Full and Model-Reduced Structure-Preserving Simulation of Incompressible Fluids

Thesis by

Gemma Mason

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2015

(Defended May 8, 2015)

© 2015
Gemma Mason
All Rights Reserved

To Hannah. You'll be here someday.

Acknowledgements

Every long road has a long list of people who helped along the way. Let me start by thanking my adviser, Mathieu Desbrun, for a great deal of advice and help, and for being such an important part of allowing me to fit in here at Caltech.

I owe a deep debt of gratitude to Christian Lessig, who provided me with support and mentorship, without which this PhD would not have been worth as nearly as much to me.

Particular thanks to Beibei Liu and Yiying Tong for their significant contributions to the results presented in Chapter 4. Thanks are also due to Dmitry Pavlov, Patrick Mullen and Dzhelil Rufat for valuable discussions and collaborations, and to Tudor Ratiu and François Gay-Balmaz, who gave advice both via email and in person.

The late Jerry Marsden was my co-adviser for a very brief time. His advice for that short time was shrewd and nurturing. I am sorry not to have been able to know him longer.

Finally, let me thank my family and my fiancé. Mum, your sympathy has been a great comfort to me. Dad, your unfailing belief in me has helped to steer me through times of trouble. Hannah, Kathryn, and Annie, I love you all and you make my life so much brighter. Paddy, I don't know how I'd have done this without you.

Abstract

This thesis outlines the construction of several types of structured integrators for incompressible fluids. We first present a vorticity integrator, which is the Hamiltonian counterpart of the existing Lagrangian-based fluid integrator [32]. We next present a model-reduced variational Eulerian integrator for incompressible fluids, which combines the efficiency gains of dimension reduction, the qualitative robustness to coarse spatial and temporal resolutions of geometric integrators, and the simplicity of homogenized boundary conditions on regular grids to deal with arbitrarily-shaped domains with sub-grid accuracy.

Both these numerical methods involve approximating the Lie group of volume-preserving diffeomorphisms by a finite-dimensional Lie-group and then restricting the resulting variational principle by means of a non-holonomic constraint. Advantages and limitations of this discretization method will be outlined. It will be seen that these derivation techniques are unable to yield symplectic integrators, but that energy conservation is easily obtained, as is a discretized version of Kelvin's circulation theorem.

Finally, we outline the basis of a spectral discrete exterior calculus, which may be a useful element in producing structured numerical methods for fluids in the future.

Contents

Chapter 1

Introduction

Imitating continuous structures in order to produce numerical integrators that preserve geometric properties of the problem being approximated is a technique that has been used in a wide variety of contexts in recent decades. Within the field of ordinary differential equations, the study of geometric integrators is very well developed [17]. In particular, the technique of variational integration [43], which imitates the action integral via a sum over discrete times, was developed in order to reliably create symplectic integrators for ODE systems.

Structured integrators for partial differential equations are a more recent and less well-developed subject, but considerable strides have been made in recent years in subjects such as electromagnetics [38] and Lagrangian field theories [41]. 2011 saw the publication of a structure-preserving integrator for incompressible fluids [32], which discretized Arnold's classic formulation of incompressible, inviscid fluids in terms of geodesics on the Lie group of volume-preserving diffeomorphisms. This paper drew on earlier techniques such as discrete exterior calculus [11] and the discrete action [43], but it also introduced the completely new technique of approximating the infinite-dimensional Lie group of volume-preserving diffeomorphisms by a finite-dimensional Lie group. Techniques such as discrete exterior calculus and the finite-dimensional Lie group method are structure-preserving discretizations in space, and may be employed either separately or together with existing variational integrators in time.

In Chapter 2 we review the existing geodesic integrator for incompressible inviscid fluids that was presented in [32]. Chapter 3 introduces a discretization of the corresponding Hamiltonian view of incompressible inviscid fluids[26], thereby giving a method based on the vorticity equation. The semidiscretization in space of a Lie-Poisson system such as that of the vorticity equation may

be achieved either variationally or by means of a discrete Lie-Poisson bracket. We also derive discrete Lie-Poisson integration in time. These developments may be viewed as analogous to the development of discrete Hamiltonian variational integrators for ordinary differential equations [43].

A notable disadvantage of these earliest geometric fluids methods [32], [14] is the large number of variables involved. These arise as a result of the discretization of the space into a large number of cells, all of which may theoretically interact with one another. In practice, we are able to restrict interactions to nearest neighbors by means of a non-holonomic constraint. However, further reduction in the number of variables and in the computational time would be desirable, particularly for applications in fluid animation. For this reason, we turn to the techniques of *model reduction*, approximating the fluid motion using only a small number of basis functions in the style of computational fluid dynamics methods such as [35, 45, 31]. Dimensionality reduction was first introduced for fluid animation by Treuille et al. [40] through Galerkin projection onto a reduced set of basis functions computed through principal component analysis of a training set of fluid motion. A number of works followed proposing the use of different bases such as Legendre polynomials [16], trigonometric functions [24], or even non-polynomial Galerkin projection [37]; eventually, Laplacian eigenvectors were pointed out by [10] to be particularly appropriate, as they guarantee divergence free flows and facilitate the conversion between vorticity and velocity, while offering a sparse advection operator for regular domains. These eigenfunctions also allow easy implementation of viscosity.

Chapter 4 presents a model-reduced variational integrator. Interestingly, despite the smaller space of variables, this method still requires a non-holonomic constraint, which will prompt some important theoretical insights. We use spectral approximation of the functional map through (cell-based) scalar and (face-based) vector Laplacian eigenvectors in order to offer model reduction without losing the variational properties of the integrator. Furthermore, we extend the embedded-boundary approach of [30] to our framework in order to compute spectral (scalar- and vector-valued) basis functions of arbitrary domains directly on *regular grids* for fast computations with sub-grid accuracy We demonstrate the efficiency of our resulting integrator through a number of examples in 2D, 3D, and curved 2D domains.

We will see in Section 4.3.1 that discretizing the Lie group picture does have some limitations. In particular, this method requires the use of a non-holonomic constraint, which makes it impossible to create a symplectic method using this style of discretization. In Chapter 5 we outline the theory

of a spectral discrete exterior calculus, which may be of use in creating a structured fluid integrator which does not employ a finite-dimensional Lie group.

Together, these new developments in the structure-preserving simulation of incompressible fluids significantly increase our theoretical understanding of the application of variational and geometric methods to the mathematics of fluid motion.

1.1 Contributions

- Chapter 2 reviews the work of Pavlov et al [32], emphasising the techniques used and outlining some of its limitations.
- Chapter 3 presents a new Hamiltonian integrator for incompressible fluids (work done with the mentorship of Christian Lessig).
- Chapter 4 presents a model-reduced integrator, using Laplacian eigenfunctions to improve the computation speed (work done in collaboration with Beibei Liu, Yiying Tong, and my adviser).
- Chapter 5 outlines a spectral discrete exterior calculus, as previously published in [33] (this work was a collaboration between Dzhelil Rufat, myself, Patrick Mullen and my adviser).

Chapter 2

Summary of Lagrangian Fluids

This chapter will outline structure-preserving integrators for inviscid, incompressible fluids as developed in [32] and [14]. We will see how to mimic the Lie group Lagrangian structure of the fluid (as discovered by Arnold in 1966 [2]) in a discretized form. This allows us to create integrators with good energy behavior, and that preserve a discrete version of Kelvin's circulation theorem. We will see in later chapters that elements of the process used here may be applied to a wide variety of structure-preserving integrators for fluids.

2.1 Continuous theory

The variational understanding of incompressible fluids is based upon the observation by Arnold [2] that the motion of an incompressible, inviscid fluid on a manifold $\mathcal{M}$ (possibly with boundary) may be derived as geodesic motion on the infinite-dimensional Lie group $\mathrm{Diff}_{\mathrm{vol}}(\mathcal{M})$ of volume-preserving diffeomorphisms on that manifold. Specifically, the position of the fluid at time t is represented by a diffeomorphism ϕ_t, which represents the change in position of each fluid particle when we change from time 0 to time t. Because the fluid is incompressible, this diffeomorphism will be volume-preserving.

This Lie group representation is Lagrangian in that it records the motions of specific particles[1]. The associated Eulerian representation, which records the velocity field at each fixed point in space, may be obtained by moving to the associated Lie algebra χ_{div}, consisting of the divergence-free

[1]The word 'Lagrangian' will inevitably arise in this dissertation in two distinct contexts: both in the sense (used here) of following specific fluid particles and also in the sense of Lagrangian (as opposed to Hamiltonian) mechanics. I apologize for any confusion caused by this unfortunate confluence of nomenclature.

vector fields on the manifold $\mathcal{M}$ which are parallel to the boundary, if $\mathcal{M}$ has a boundary. We can find the appropriate geodesic motion by solving for the time-dependent velocity field $\mathbf{v}(\mathbf{x}, t)$, which is a stationary point of the action $\int_0^t L \, \mathrm{d}t$, where the Lagrangian function L is defined by

$$L = \int_{\mathcal{M}} \|\mathbf{v}\|^2 \, \mathrm{d}V, \tag{2.1}$$

where $\|\mathbf{v}\|$ uses the standard Euclidean norm on vectors in $\mathbb{R}^n$, and the variations satisfy the *Lin constraints*:

$$\delta v = \dot{\xi} + [v, \xi]. \tag{2.2}$$

A Lagrangian stationary action problem of this type, which uses Eulerian rather than Lagrangian co-ordinates, is solved by the Euler-Poincaré equations [9], which are similar to the Euler-Lagrange equations for Lagrangian co-ordinates. In this case, the resulting equations are

$$\dot{v}^\flat + \mathcal{L}_{\mathrm{v}}(v^\flat) = -\mathrm{d}p, \tag{2.3}$$

where $v^\flat$ is the 1-form dual to the vector field $\mathbf{v}$, $\mathcal{L}_{\mathbf{v}}$ is the Lie derivative with respect to $\mathbf{v}$, and $\mathrm{d}p$ is the exterior derivative of some 0-form p.

On regions in $\mathbb{R}^2$ or $\mathbb{R}^3$, these equations are equivalent to the expected Euler equations for an incompressible, inviscid fluid:

$$\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} = -\nabla P, \tag{2.4}$$

where P is a function which accounts for the pressure.

Arnold's Lagrangian derivation can be used to show that Kelvin's theorem holds. Specifically, the circulation-preservation required by Kelvin's theorem is a special case of Noether's theorem, which states that a continuous symmetry of a Lagrangian system induces a conserved momentum. In this case, the Lagrangian is conserved under a *particle-relabeling symmetry*. That is, if we relabel each fluid particle by rearranging them according to some volume-preserving diffeomorphism $g \in \mathrm{Diff}_{\mathrm{vol}}(\mathcal{M})$, the Lagrangian function is unchanged. Kelvin's theorem follows as a result.

2.2 Discretization process

The 2011 paper of Pavlov et al [32] introduced a variational integrator for fluids, based on the continuous variational derivation of Euler fluids as geodesics on χ_{vol}. This method used several varieties of geometric discretization at once. Most notably, it introduced the idea of approximation of the infinite dimensional Lie group of volume-preserving diffeomorphisms with an appropriate finite-dimensional Lie group G. It also introduced a new variant on discrete exterior calculus [11] which interacts with the associated Lie algebra $\mathfrak{g}$. With the help of a non-holonomic constraint, these two elements can spatially discretize the motion of the fluid in a manner that respects the Lagrangian structure. It is then possible to derive an associated time-discretization using the method of a *discrete variational principle*, a geometric discretization method that has been extensively developed for use with ODEs [43].

I shall detail here the main elements of the spatial discretization.

2.2.1 The finite-dimensional Lie group method

To find a finite-dimensional Lie group that can approximate the behavior of the infinite-dimensional Lie group $\text{Diff}_{\text{vol}}(\mathcal{M})$, we begin by discretizing the fluid domain into a finite number of cells. In preparation for the rest of this thesis, the explanation given here will focus on regular Cartesian grids, with square cells of width h.

We discretize a continuous function $f(\mathbf{x})$ on the manifold $\mathcal{M}$ by taking an average (integrated) value f_i per grid cell i of the mesh, which we arrange in a vector $\mathbf{f}$. This definition of discrete functions allows us to discretize the set of possible flows $\phi_t \in \text{Diff}_{\text{vol}}(\mathcal{M})$ using a Koopman operator (or *functional map*, as it is known in computer graphics literature), which describes the action of the flow in terms of how it changes a discrete function. Specifically, the flow map is encoded as a matrix q of size the square of the number of cells, and the integrated values $\mathbf{f}$ of f per cell become $q\mathbf{f}$ once f is advected by the flow ϕ.

The constant function is unchanged when advected by the flow, so every discrete flow should take the constant function to itself. That is, for all q, we require that $q\mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ is a vector of ones (see [32] for the equivalent condition on an arbitrary mesh). This is the same as saying that the row sums of q are equal to 1, i.e., q is *signed stochastic*.

Since we are simulating an incompressible fluid, we also require that the discrete flow be volume-

preserving. This condition is achieved by asking that the discrete flow preserves the inner product of vectors, that is, q is *orthogonal*, i.e., $q^T = q^{-1}$. Thus, we find that we need to take q to be an element of the Lie group G of orthogonal, signed stochastic matrices. This matrix group represents our discrete fluid configuration, as we describe next.

We can view the finite-dimensional Lie group G as a configuration space: it encodes the space of possible "positions" for the discrete fluid, in that each element of the Lie group represents a possible way that the fluid could have evolved from its initial position. This Lie group represents a Lagrangian perspective as it identifies the fluid particles in a given cell by recording which cells they originally came from. The associated Eulerian perspective is given by the Lie algebra $\mathfrak{g}$ of matrices of the form $\dot{q} \circ q^{-1}$ for $q \in G$. It was shown in [32] that any matrix $A \in \mathfrak{g}$ of this Lie algebra is both *antisymmetric* ($A^T = -A$) and *row-null* ($A\mathbf{1} = \mathbf{0}$), and corresponds to a discrete counterpart of the Lie derivative L_v with respect to the continuous velocity field $v = \dot{\phi} \circ \phi^{-1}$. Thus, the product $A\mathbf{f}$ of such a matrix with a discrete function $\mathbf{f}$ approximates the continuous term $v \cdot \nabla f$. Furthermore, if cells i and j are nearest neighbors, then the matrix element A_{ij} represents the flux of the fluid through the face shared by cells i and j. Thus, an element of the Lie algebra $\mathfrak{g}$ of G is directly linked to the usual flux-based (Marker And Cell) discretization of vector field in fluid simulators [18].

2.2.2 Discrete exterior calculus

The Lie algebra $\mathfrak{g}$ consists of vector fields on our discrete manifold. The space Ω_1 of functionals on these vector fields may therefore be viewed as a space of discrete 1-forms, while the integrated functions $\mathbf{f}$ that we defined above would naturally form the space Ω_0 of discrete 0-forms. This suggests that we might be able to form some sort of discrete exterior calculus [11], defining discrete versions of operators such as the exterior derivative, which will be of use in describing the motion of a discrete fluid.

Throughout this thesis, the discrete versions of continuous operators will be indicated by boldface. For example, the continuous exterior derivative is written as d, and the discrete exterior derivative (defined below) will be denoted $\mathbf{d}$.

We define the space of discrete 1-forms Ω^1 as the space of antisymmetric matrices. We pair a 1-form $C^\flat$ with a vector field B using a Frobenius pairing:

$$\langle C^\flat, B \rangle = \mathrm{Tr}((C^\flat)^T B). \tag{2.5}$$

2-forms are defined in this formulation as antisymmetric 3-tensors F_{ijk}. We define the *contraction* of a 2-form F with a vector field A as

$$(\mathbf{i}_A F)_{ij} = \sum_k \left(F_{ikj} A_{ik} - F_{jki} A_{jk} \right), \tag{2.6}$$

which is constructed to be antisymmetric so as to yield a 1-form. It was shown in [32] that this definition is consistent with defining the total pairing of a 2-form F with two vector fields A, B as

$$\langle F, A, B \rangle = \sum_{i,j,k} F_{ijk} A_{ij} B_{ik}. \tag{2.7}$$

We can now define a discrete exterior derivative which takes us from a vector $\mathbf{f}$ of function values $(f_1, f_2, ...f_N)$ to an antisymmetric matrix as follows:

$$A_{ij} = (\mathbf{df})_{ij} = f_i - f_j. \tag{2.8}$$

The exterior derivative on 1-forms is similarly defined by antisymmetry as:

$$(\mathbf{d}A)_{ijk} = A_{ij} + A_{jk} + A_{ki}. \tag{2.9}$$

More generally, a k-form is given by an antisymmetric $(k+1)$-tensor, and the discrete exterior derivative on a k-form is

$$\mathbf{d}G_{i_1 i_2 \cdots i_{k+1}} = \sum_{j \in \{1,2,\cdots,k+1\}} (-1)^{j+1} G_{i_1 \cdots i_{j-1} i_{j+1} \cdots i_{k+1}}. \tag{2.10}$$

The discrete exterior derivative satisfies $\mathbf{d}^2 = 0$, as we would expect by comparison with the continuous version of the operator.

It should be noted that the version of discrete exterior calculus defined here and in [32], [14] differs in several important respects from the discrete exterior calculus that was previously defined by Desbrun et al [11]. Notably, the definitions above define elements of a 1-form corresponding to every pair of cells, rather than locating 1-forms on edges. The non-holonomic constraint defined in the next section will bring this version of DEC and the traditional version into closer alignment. One consequence of this is that we will then be able to use concepts from traditional DEC to inform

the construction of the discrete flat operator in Section 2.2.4.

2.2.3 Non-holonomic constraint

Whilst the elements A_{ij} for $A \in \mathfrak{g}$ have a clear physical interpretation in the case where i and j are nearest neighbors, this is not the case for elements representing interactions between cells that are not immediate neighbors. Similarly to a CFL condition, we prohibit fluid particles from skipping to non-neighboring cells by restricting the Lie algebra to the constrained set $\mathcal{S}$, the set of matrices A such that $A_{ij} = 0$ unless cells i and j share a face (or an edge in 2D). We require the elements of $\mathfrak{g}$ that we use to represent the fluid velocity fields to fall into this constrained set. This has the additional advantage of making the matrices sparse, dramatically decreasing the amount of memory required and the computational time, as much fewer degrees of freedom need to be updated. Moreover, a Lie algebra element in this constrained set corresponds exactly to the traditional MAC discretization with fluxes. We can also relate fluxes between nearest neighbors to the fluxes on edges that are used in discrete exterior calculus [11].

Constraining the matrices in this way requires a non-holonomic constraint, because the set $\mathcal{S}$ is not closed under the Lie bracket. That is, interactions between nearest neighbors followed by further interactions between nearest neighbors produce interactions between cells that are two-away from each other, which are therefore not inside the constrained set $\mathcal{S}$.

2.2.4 Discrete flat operator and inner product

The co-ordinate-independent nature of exterior calculus allows us to translate it to the discrete setting with the structure intact, in such a way that *no error whatsoever* is introduced by the discretization process. However, all numerical methods introduce error at some point, and with numerical methods derived from discrete exterior calculus that error comes from places in the method where the co-ordinates (and thus the local metric) must unavoidably be introduced. The method of Pavlov et al [32] does this by introducing both an inner product and a *flat operator*, which takes a vector field $A \in \mathfrak{g}$ and maps it to the 1-form $A^\flat \in \mathfrak{g}^*$ such that

$$\langle A^\flat, B \rangle = (A, B),\ \forall B \in \mathcal{S}, \tag{2.11}$$

where the angle brackets on the left hand side indicate the pairing of a 1-form with a vector field (2.5), and the curved brackets on the right hand side indicate an inner product between vector fields. We use a Frobenius inner product between vector fields as follows:

$$(A, B) = 2h^2 \mathrm{Tr}(A^T B) \tag{2.12}$$

where h is the spatial step size. This definition approximates an integral of the product of the two functions over the space $\mathcal{M}$.

One might think that condition (2.11) would be sufficient to define the flat operator. This turns out not to be the case. Equation (2.11) only defines those elements of $A^\flat$ that record interactions between cells that are nearest neighbors. However, we will also need to know the entries of $A^\flat$ for pairs of cells that share a nearest neighbor in common. To define these elements of $A^\flat$, we require that the flat operator satisfies the following condition on the total pairing of its discrete exterior derivative with two discrete vector fields:

$$\langle \mathbf{d}A^\flat, B, C \rangle \to \int_{\mathcal{M}} \mathrm{d}v^\flat(u, w)\mathrm{d}V \tag{2.13}$$

as $h \to 0$, where A, B and C converge to v, u and w, respectively, as $h \to 0$. It was shown in [32] that this may be achieved with a little help from discrete exterior calculus to determine the desired form of the vorticity elements $\mathbf{d}A^\flat$. I refer the reader to this paper for more details.

2.3 Creating a variational numerical method

Using the spatial discretization outlined in the preceding section, Pavlov et al [32] showed that one can create a semidiscretized numerical method for ideal, incompressible fluids through the Euler-Poincaré equations [14] for the Lagrangian given by

$$\mathcal{L}_{\text{semidiscrete}} = \frac{1}{2}(A, A) \approx \frac{1}{2}\int v^2 \mathrm{d}x, \tag{2.14}$$

where variations of A satisfy the Lin constraint $\delta A = \dot{B} + [A, B]$, and A is subject to the non-holonomic constraint $A \in \mathcal{S}$. This gives the equations

$$\dot{A}^\flat + \mathcal{L}_A A^\flat = -\mathbf{d}p. \tag{2.15}$$

Because the semidiscretized equations constitute a Lagrangian system of ODEs, they may be discretized in time using the method of a discrete Lagrangian [43]. The resulting equations are equivalent to the Harlow-Welch Scheme [18] with a Crank-Nicolson (trapezoidal) update rule. Alternatively, we may create an energy-preserving method by updating with the midpoint rule [28].

2.4 Results and Discussion

We see that the technique of representing the Lie group structure of $\mathrm{Diff}_{\mathrm{vol}}(\mathcal{M})$ using an appropriate finite dimensional Lie group may be used, together with a version of discrete exterior calculus and a non-holonomic constraint, in order to construct a numerical method for fluids which discretizes in space in a way that preserves the variational structure. Time discretization may then be achieved either by the energy-preserving midpoint rule, or using a discrete-time variational integrator.

The resulting numerical method exhibits no numerical dissipation (see Figure 2.1), and has been empirically observed to produce good qualitative behavior over long timescales. This variational integrator also guarantees that the relabeling symmetry implies a discrete version of Kelvin's circulation theorem, i.e., circulation of velocity field (represented as a Lie algebra element) along a closed loop (represented as a 1-cycle [32]) transported along the fluid flow is invariant, which helps keep the vivid details of vorticity in the fluid simulation without resorting to additional energy-injecting measures such as vorticity confinement, as shown in [28]. However, the time integration requires a quadratic solve based on *all* the fluxes in the domain, making it inappropriate for realtime simulation. Moreover, the non-holonomic constraint means that we no longer preserve the symplectic structure.

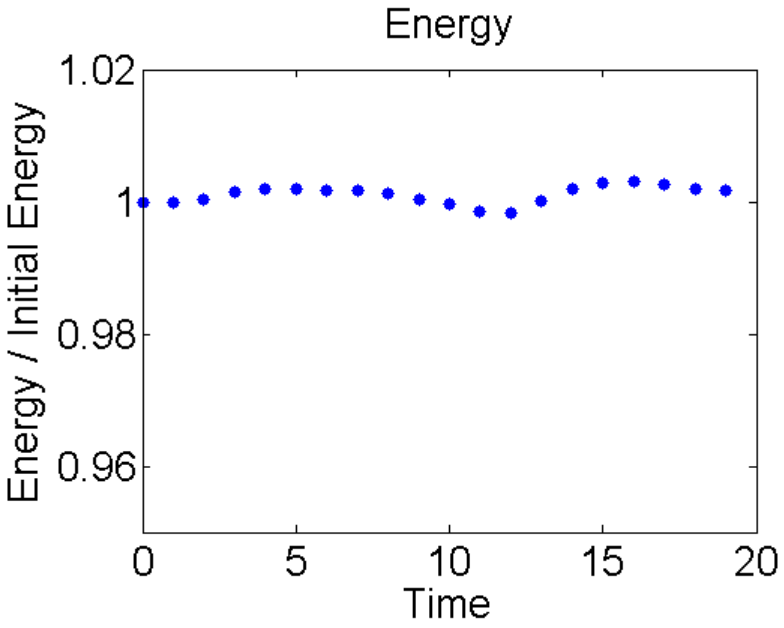

Figure 2.1: Energy behavior of the numerical method developed in [32], using a discrete-time variational integrator for the timestepping, and evaluated on a 10 by 10 Cartesian grid.

Chapter 3

Hamiltonian Fluids

Arnold's geodesic formulation of inviscid, incompressible fluids has an associated Hamiltonian formulation, as described by Marsden and Weinstein [26]. The primary variable in the Hamiltonian viewpoint is the *vorticity* rather than the velocity. This chapter describes a numerical method for fluids based on this Hamiltonian vorticity viewpoint.

Expanding our understanding of geometric numerical methods for fluids to the Hamiltonian perspective is useful from a practical perspective, in that the stream function is located on points, a property which may be useful for constructing embedded boundary methods. Specifically, if there is no fluid flow through the boundary, then the stream function is constant on the boundary. If the boundary is connected, then without loss of generality we may set the stream function to zero there. If we then wish to consider the flux through an edge that passes through an embedded boundary, we may find it by subtracting the boundary value of the stream function from the value of the stream function at one of the edge's two vertices, taking the vertex located inside the boundary.

Although the vorticity form of the Euler equations is generally used in two dimensions, the framework of differential forms does allow it to be used in three dimensions as well [26]. However, in this thesis we will consider only the two-dimensional case.

In addition to opening up new possibilities for applications, discretizing the Hamiltonian view of fluids is also interesting from a theoretical perspective. We will demonstrate two ways of doing this, one of which utilizes a variational formalism, and the other of which mimics the Poisson bracket.

3.1 Review of continuous theory

Recall from Section 2.1 that the motion of an incompressible fluid may be described by elements of $\mathrm{Diff}_{\mathrm{vol}}(\mathcal{M})$, the Lie algebra of volume-preserving diffeomorphisms, or in the Eulerian sense by elements of χ_{div}, the Lie algebra of divergence-free vector fields. The Lagrangian L is invariant under a particle-relabeling symmetry: for any $g \in \mathrm{Diff}_{\mathrm{vol}}(\mathcal{M})$, $L(g_R(\mathbf{v})) = L(\mathbf{v})$ for $\mathbf{v} \in \chi_{\mathrm{div}}$, where $g_R(\mathbf{v})$ is the right action of g on $\mathbf{v}$. Marsden and Weinstein showed in 1983 [26] that this symmetry allows us to perform a Lie-Poisson reduction in order to induce a Hamiltonian system on the dual space χ^*_{div}. This dual space consists of 1-forms modulo exact 1-forms (in the Euler equations for an incompressible fluid, the modulo over exact 1-forms is accounted for by the $\mathrm{d}p$ term). Equivalently, we may identify any element of χ^*_{div} with an exact 2-form via a bijection given by the exterior derivative[1]. Taking the exterior derivative of a 1-form is equivalent to taking the curl of the associated vector field, and this bijection takes us from a given velocity to the associated vorticity.

Since we have only kinetic (and not potential) energy, when we move to the Hamiltonian picture the Hamiltonian function is equal to our earlier Lagrangian (2.14), transformed into momentum coordinates. Given homogeneous boundary conditions (so that the codifferential δ will be the adjoint of the exterior derivative), this transformation is as follows:

$$H = (v^\flat, v^\flat) \tag{3.1a}$$
$$= (\delta \mathrm{d} \Delta^{-1} v^\flat, v^\flat) \tag{3.1b}$$
$$= (\Delta^{-1} \mathrm{d} v^\flat, \mathrm{d} v^\flat) \tag{3.1c}$$
$$= (\Delta^{-1} \omega, \omega). \tag{3.1d}$$

Note that the Laplace-deRham operator $\Delta v^\flat = (\delta \mathrm{d} + \mathrm{d}\delta) v^\flat = \delta \mathrm{d} v^\flat$ since the codifferential $\delta v^\flat = 0$ because v is divergence-free. A similar formula holds for ω, where we have $\Delta \omega = (\delta \mathrm{d} + \mathrm{d}\delta)\omega = \mathrm{d}\delta\omega$, because $\mathrm{d}\omega = 0$. I have also used the fact that d commutes with Δ^{-1}. This follows easily from the fact that d commutes with Δ, which can be verified using the definition of Δ and the nilpotence of the exterior derivative.

Marsden and Weinstein [26] use this Hamiltonian along with the following Lie Poisson bracket

[1] Strictly speaking, we need not only the exact 2-form but also the circulation around any non-contractible loops in the event that $\mathcal{M}$ is not simply connected.

(derived via momentum reduction):

$$\{F, G\}(\omega) = \int_{\mathcal{M}} \left\langle \omega, \left[\frac{\delta F}{\delta \omega}, \frac{\delta G}{\delta \omega} \right] \right\rangle \mathrm{d}x, \tag{3.2}$$

where $\frac{\delta F}{\delta \omega} \in \chi_{\mathrm{div}}$ is a functional derivative and an element of the Lie algebra. Then the Lie-Poisson equations $\dot{F} = \{F, H\}$ yield the vorticity equations:

$$\frac{\partial \omega}{\partial t} + \mathcal{L}_{\mathbf{v}} \omega = 0 \tag{3.3a}$$

$$\text{with} \quad \omega = \mathrm{d}v. \tag{3.3b}$$

More recently, Cendra et al [9] explained how to derive the Lie-Poisson equations from a variational principle. In the context of incompressible fluids, the Lie-Poisson variational principle states that the curve $(v(t), \omega(t)) \in \chi_{\mathrm{div}} \times \chi^*_{\mathrm{div}}$ is a critical point of the action

$$\delta \int_{t_0}^{t_1} \left(\langle \omega, v \rangle - H(\omega(t)) \right) \mathrm{d}t = 0 \tag{3.4}$$

for variations of the form

$$\delta v(t) = \dot{\eta}(t) - [v(t), \eta(t)], \tag{3.5}$$

where η is a curve in χ_{div} such that $\eta(t_1) = \eta(t_2) = 0$, and the variations $\delta\omega$ are arbitrary.

This variational principle is equivalent both to Hamilton's principle of critical action and to the Lie-Poisson equations, which are equivalent to the vorticity equation (3.3a) given above.

3.2 Derivation

3.2.1 Preliminary theoretical elements

By analogy with the continuous theory, we convert a velocity into a vorticity by moving from the space $\Omega_1/\mathbf{d}\Omega_0$ of discrete 1-forms modulo exact discrete 1-forms to the space Ω_2^{ex} of exact discrete 2-forms. A bijection $\beta(A) = \mathbf{d}A^\flat$ between the two spaces is given by the discrete exterior derivative. Thus, we define the vorticity 2-form $W = \mathbf{d}A^\flat$ to be the discrete exterior derivative of the velocity 1-form.

If we assume that $A \in \mathcal{S}$ satisfies the nonholonomic constraint, and use the form of the flat

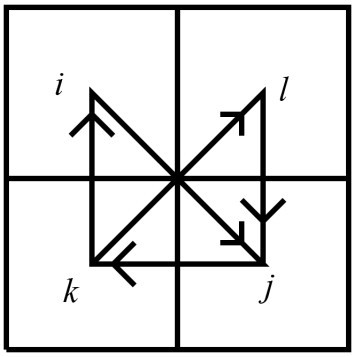

Figure 3.1: The square cells i, j, k and l share a vertex. $W_{ijk} = W_{jkl}$, because both triangles are clockwise with the given index order.

operator defined by [32], then $W_{ijk} = 0$ unless cells i, j and k share a vertex. Furthermore, if cells i, j, k and l all share a vertex, then $W_{abc} = W_{ijk}$ for any indices $a, b, c \in \{i, j, k, l\}$ such that no two indices are repeated, provided that the triangle abc is oriented in the same way as the triangle ijk (i.e., if travelling from the center of cell i to the center of cell j to the center of cell k and then back to the center of cell i represents a clockwise/anticlockwise motion, then we require that the triangle abc should also be clockwise/anticlockwise in order for the equality to hold; see Figure 3.1). This allows us to make a direct analogy between the 2-form W and an integrated 2-form on dual cells in discrete exterior calculus [11].

We define a pairing between elements of Ω_2^{ex} and vector fields by defining

$$\langle \mathbf{d}B, A \rangle = \langle B, A \rangle \tag{3.6}$$

for $\mathbf{d}B \in \Omega_2^{\text{ex}}$ and $B \in \Omega_1/\mathbf{d}\Omega_0$. That is, when we pair a vector field with an element of Ω_2^{ex}, the result is the same as if we had paired it with the corresponding element of $\Omega_1/\mathbf{d}\Omega_0$ as defined by the discrete exterior derivative bijection between the two spaces.

Previous work [32], [14] defined the metric using a flat operator, which takes us from a vector field to the 1-form which is equivalent to taking an inner product with that vector field. There are several other ways to introduce a metric, notably the *Hodge star*, denoted by $\star$. For the purposes of a derivation on vorticity, it is convenient to use a *codifferential*, which acts on a k-form α as $\delta\alpha = (-1)^{k+1} \star^{-1} \mathrm{d} \star \alpha$. When we have periodic or zero boundary conditions, the codifferential is the

adjoint of the exterior derivative. Thus, we require that the discrete codifferential $\boldsymbol{\delta}$ must satisfy

$$(\boldsymbol{\delta}F, A) = (F, \mathbf{d}A). \tag{3.7}$$

Recall that we have the following inner product on 1-forms:

$$(A, B) = 2h^2 \sum_{i,j} A_{ij} B_{ij}. \tag{3.8}$$

Define the following Frobenius inner product on 2-forms:

$$(F, G) = \frac{1}{9} \sum_{i,j,k} F_{ijk} G_{ijk}. \tag{3.9}$$

A discrete codifferential that is consistent with these inner product definitions is then defined as follows:

$$(\boldsymbol{\delta}F)_{ij} = \frac{1}{6h^2} \sum_{k} (F_{ijk} + F_{kij} + F_{jki}). \tag{3.10}$$

These definitions for the inner product and the codifferential are consistent in that they satisfy the equality (3.7). The constants have also been chosen so that, if i, j, k, and l are four cells that share a vertex, then $(\mathbf{d}\boldsymbol{\delta}F)_{ijk} + (\mathbf{d}\boldsymbol{\delta}F)_{ikl}$ gives an expression equivalent to the five-point stencil for the Laplacian. It may also be seen that, on a Cartesian grid and using the non-holonomic constraint, this definition for the codifferential is equivalent to taking appropriately-scaled differences between adjacent dual cells.

Using the discrete codifferential rather than a discrete flat operator allows us to introduce a metric to our space in a much more straightforward way. Whereas the flat operator was defined on a rather ad hoc basis as an operator satisfying the definition and the limiting equality (2.13), the discrete codifferential may be easily derived from the definition of a codifferential along with a standard discrete form of the Laplacian.

3.2.2 Variational derivation

To derive the semidiscretized vorticity equations from a variational principle, we create a non-holonomically constrained Lie-Poisson system using the Lie-Poisson variational principle (3.4). We

define the discrete Hamiltonian $H_{\mathrm{d}} = \langle \mathbf{\Delta}^{-1} W, W \rangle$, using the discrete Laplacian $\mathbf{\Delta} = \mathbf{d}\boldsymbol{\delta} + \boldsymbol{\delta}\mathbf{d}$, and consider variations of the action

$$\int \left(\langle W, A \rangle - H_{\mathrm{d}}(W) \right) dt, \quad A \in \mathcal{S} \tag{3.11}$$

for arbitrary variations δW and variations δA of the form $\delta A = \dot{B} + [B, A]$, where $B \in \mathcal{S}$ satisfies the nonholonomic constraint.

This yields

$$\frac{\delta H_{\mathrm{d}}}{\delta W} = A \tag{3.12}$$

and

$$\left\langle \dot{W} + \mathrm{ad}^*_A W, B \right\rangle = 0, \quad \text{for all } B \in \mathcal{S}. \tag{3.13}$$

Both these equations require further interpretation. The expression $\frac{\delta H_{\mathrm{d}}}{\delta W}$ is defined by

$$\mathrm{D}H_{\mathrm{d}}(\mu) \cdot \nu = \langle \nu, \frac{\delta H_{\mathrm{d}}}{\delta \mu} \rangle, \tag{3.14}$$

where $\mathrm{D}H_{\mathrm{d}} \in \mathfrak{g}$ is the derivative of F. We can calculate $\frac{\delta H_{\mathrm{d}}}{\delta W}$ as follows, using the fact that the Laplacian (and its inverse) are self-adjoint:

$$\mathrm{D}H_{\mathrm{d}}(W) \cdot \delta W = \lim_{\epsilon \to 0} \frac{1}{\epsilon} (H_{\mathrm{d}}(W + \epsilon \delta W) - H_{\mathrm{d}}(W)) \tag{3.15a}$$

$$= \lim_{\epsilon \to 0} \frac{1}{\epsilon} \big(\frac{1}{2} \langle \Delta^{-1}(W + \epsilon \delta W), W + \epsilon \delta W \rangle - \frac{1}{2} \langle \Delta^{-1} W, W \rangle \big) \tag{3.15b}$$

$$= \langle \Delta^{-1} W, \delta W \rangle. \tag{3.15c}$$

This means that the pairing of $\frac{\delta H_{\mathrm{d}}}{\delta W}$ with a 2-form $\mathbf{d}\sigma$ can be related to the following inner product on 2-forms:

$$\langle \mathbf{d}\sigma, \frac{\delta H_{\mathrm{d}}}{\delta W} \rangle = (\mathbf{d}\sigma, \Delta^{-1} W). \tag{3.16}$$

Given the definition 3.6 of the pairing of a 2-form with a vector field, we can then conclude that

$$\langle \sigma, \frac{\delta H_{\mathrm{d}}}{\delta W} \rangle = (\mathbf{d}\sigma, \Delta^{-1} W) \tag{3.17a}$$

$$= (\sigma, \delta \Delta^{-1} W). \tag{3.17b}$$

That is, if we set $A = \frac{\delta H}{\delta W}$, then we have $A^\flat = \delta\Delta^{-1}W$. When we take the discrete exterior derivative of both sides, we see that equation (3.12) implies that $W = \mathbf{d}A^\flat$, as expected.

To interpret the second equation (3.13), recall that the pairing of a 2-form with a vector field is defined (3.6) to be the same as pairing the equivalent 1-form with a vector field when we invert the bijection β between the two spaces. Equation (3.13) thus implies that

$$\beta^{-1}(\dot{W} + \mathrm{ad}^*_A W) \in \mathcal{S}^\perp, \tag{3.18}$$

where $\mathcal{S}^\perp$ is the space perpendicular to $\mathcal{S}$, which was shown in [32] to be the space of exact 1-forms $\mathrm{d}P$ for some 0-form P. Thus,

$$\beta^{-1}(\dot{W} + \mathrm{ad}^*_A W) = \mathbf{d}P \tag{3.19}$$

$$\implies \dot{W} + \mathrm{ad}^*_A W = 0, \tag{3.20}$$

since the bijection $\beta : \Omega_1/\mathbf{d}\Omega_0 \to \Omega_2^{\mathrm{ex}}$ is given by the discrete exterior derivative. Given the definition of the adjoint action ad^*, this is equivalent to

$$\dot{W} + \mathcal{L}_A W = 0 \tag{3.21}$$

This is a semidiscrete version of the vorticity equation (3.3a).

3.2.3 Derivation from a discrete version of the Lie-Poisson bracket

We can construct an alternate derivation based on the fact that the vorticity equations can be derived from the Lie-Poisson equations $\dot{F} = \{F, H\}$ on χ^*_{vol}. Following the continuous derivation given in [26], we can construct a similar derivation on our discrete Lie algebra dual $\mathfrak{g}^*$.

The Lie-Poisson bracket is defined in the continuous case by

$$\{F, G\}(\omega) = \int \langle \omega, [\frac{\delta F}{\delta \omega}, \frac{\delta G}{\delta \omega}] \rangle \mathrm{d}x. \tag{3.22}$$

In the discrete case, we drop the integral as it is effectively included in the pairing. This gives:

$$\{F, G\}(W) = \langle W, [\frac{\delta F}{\delta W}, \frac{\delta G}{\delta W}] \rangle. \tag{3.23}$$

This is already well-defined as an expression. We have a Lie bracket on $\mathfrak{g}$ and a discrete pairing with the dual. It is worth noting, however, that if we constrain W to be in the set $\delta A^\flat$ for $A \in \mathcal{S}$, then the introduction of the non-holonomic constraint will mean that this bracket no longer satisfies the Jacobi identity [5].

We can now derive the vorticity equations from the Lie-Poisson equations as follows. For an arbitrary function F, we know that:

$$\dot{F} = \{F, H_\mathrm{d}\}, \tag{3.24}$$

where $H_\mathrm{d} = \langle \delta^{-1} W, W \rangle$ is the discrete Hamiltonian as before. Considering the left hand side, we have:

$$\frac{\mathrm{d}F}{\mathrm{d}t} = \mathrm{D}F(W) \cdot \dot{W} = \langle \dot{W}, \frac{\delta F}{\delta W} \rangle. \tag{3.25}$$

On the right hand side, we have:

$$\{F, H_\mathrm{d}\} = \langle W, \big[\frac{\delta F}{\delta W}, \frac{\delta H_\mathrm{d}}{\delta W}\big] \rangle \tag{3.26a}$$

$$= \langle W, \big[\frac{\delta F}{\delta W}, A\big] \rangle \tag{3.26b}$$

$$= \langle W, -\mathrm{ad}_A \frac{\delta F}{\delta W} \rangle \tag{3.26c}$$

$$= \langle -\mathrm{ad}^*_A W, \frac{\delta F}{\delta W} \rangle \tag{3.26d}$$

$$= \langle -\mathcal{L}_A W, \frac{\delta F}{\delta W} \rangle. \tag{3.26e}$$

Given that F is arbitrary, we have thus proved that

$$\dot{W} + \mathcal{L}_A W = 0 \tag{3.27}$$

in the discrete case.

Equations (3.2.2) to (3.17b) apply here, just as they did in the previous derivation, to show that $W = \mathbf{d}A^\flat$.

3.2.4 Geometric properties of the semidiscretization

We can use the Lie-Poisson derivation of the semidiscretized equations of motion to show that these equations preserve the Hamiltonian. We can see that the Hamiltonian is preserved, because the

derivation shows that equation (3.27) holds if and only if $\dot{F} = \{F, H_{\mathrm{d}}\}$ for all functions F. Thus, if we take $F = H_{\mathrm{d}}$,

$$\frac{\mathrm{d}H_{\mathrm{d}}}{\mathrm{d}t} = \{H_{\mathrm{d}}, H_{\mathrm{d}}\} = 0. \tag{3.28}$$

Note that, as with the Lagrangian method developed in [32] and outlined in Chapter 2, we expect that the semidiscretized equations cannot be symplectic, due to the use of a nonholonomic constraint. We will, however, still satisfy a discrete version of Kelvin's theorem, in the same style as in Chapter 2. This may be seen from the equivalence between equation (3.19) and equation (2.15). Since the latter satisfies the discrete Kelvin's theorem property, it follows that the former will, also.

3.2.5 Time discretization

Having obtained the semidiscretization (3.21), we now wish to discretize in time. To retain good energy behavior and avoid numerical viscosity, we will use a time-symmetric discretization [17]. The lowest order time-symmetric discretizations are of order two; one may use either the trapezoidal rule or the midpoint rule. The midpoint rule has the advantage that it preserves quadratic invariants [17], including the energy of the fluid.

The trapezoidal rule, which may be derived from a variational method in the style of [43], would ordinarily give a symplectic method. However, as a rule, systems derived from a non-holonomic constraint are not symplectic unless the constraint can be rewritten to be holonomic [5]. As this is not the case for the constraint that we are using, our semidiscretization (3.21) is already not a symplectic system. Nonetheless, for completeness, we detail here how to use a discrete Lie-Poisson action to variationally derive a trapezoidal time discretization. Variational Lie-Poisson integrators have been developed before [25], but not in the presence of a non-holonomic constraint, so we derive the time integration from first principles.

The discrete action for timesteps of size h is defined to be

$$S_h = \sum_{k=0}^{N} \big[\langle W_k, A_k\rangle - H(W_k)\big]h. \tag{3.29}$$

We place no restrictions on variations of W_k, and we ask that variations of A_k satisfy a *discrete Lin constraint*. These discrete Lin constraints were developed in [32], and may vary in form depending

on how we posit the group element q_k (denoting position of the fluid in the Lagrangian sense) to update based on the current fluid velocity A_k. The simplest of these, which corresponds to an explicit Euler method update of the fluid position, is

$$\delta_k A_k = -\frac{B_k}{h} + A_k B_k \tag{3.30a}$$

$$\text{with} \quad \delta_{k+1} A_k = \frac{B_{k+1}}{h} - B_{k+1} A_k, \tag{3.30b}$$

where $\delta_k A_k$ denotes the variations of A_k with respect to δq_k, and $B_k \in \mathcal{S}$.

Equating the variation of the action (3.29) with respect to W_k to zero yields

$$A_k = \frac{\delta H}{\delta W}(W_k), \tag{3.31}$$

which, as we have seen before, is equivalent to the statement $W_k = \mathbf{d}A_k$.

Equating the variations of the action with respect to δq_k to zero for $k = 1...N-1$ yields

$$\langle W_{k-1}, \delta_k A_{k-1} \rangle + \langle W_k, \delta_k A_k \rangle = 0. \tag{3.32}$$

Now we substitute in the discrete Lin constraints to get

$$\langle W_{k-1}, \frac{B_k}{h} - B_k A_{k-1} \rangle + \langle W_k, -\frac{B_k}{h} + A_k B_k \rangle = 0 \tag{3.33}$$

$$\langle W_{k-1} - W_k, \frac{B_k}{h} \rangle + \langle W_k, A_k B_k \rangle - \langle W_{k-1}, B_k A_{k-1} \rangle = 0 \tag{3.34}$$

$$\left\langle \frac{W_k - W_{k-1}}{h}, B_k \right\rangle - \left\langle W_k, \frac{A_k - A_k^T}{2} B_k \right\rangle - \left\langle W_{k-1}, B_k \frac{A_{k-1} - A_{k-1}^T}{2} \right\rangle = 0, \tag{3.35}$$

where we have used the fact that B_k is antisymmetric. However, as it happens A_k is also antisymmetric, which means $A_k^T B_k = -(A_k^T B_k)^T = -B_k^T A_k = B_k A_k$. A similar expression of course holds

for A_{k-1}, and thus,

$$\left\langle \frac{W_k - W_{k-1}}{h}, B_k \right\rangle - \left\langle W_k, \frac{A_k B_k - B_k A_k}{2} \right\rangle - \left\langle W_{k-1}, \frac{B_k A_{k-1} - A_{k-1} B_k}{2} \right\rangle = 0 \tag{3.36}$$

$$\left\langle \frac{W_k - W_{k-1}}{h}, B_k \right\rangle + \left\langle W_k, \frac{1}{2}[A_k, B_k] \right\rangle - \left\langle W_{k-1}, \frac{1}{2}[A_{k-1}, B_k] \right\rangle = 0 \tag{3.37}$$

$$\left\langle \frac{W_k - W_{k-1}}{h}, B_k \right\rangle + \left\langle \frac{1}{2}(\mathrm{ad}^*_{A_k} W_k + \mathrm{ad}^*_{A_{k-1}} W_{k-1}), B_k \right\rangle = 0. \tag{3.38}$$

$B_k \in \mathcal{S}$ is arbitrary, and so we find that, by the definition (3.6) of the pairing of a 2-form with a vector field,

$$\beta^{-1}\left(\frac{W_k - W_{k-1}}{h} + \frac{1}{2}(\mathrm{ad}^*_{A_k} W_k + \mathrm{ad}^*_{A_{k-1}} W_{k-1}) \right) \in \mathcal{S}^\perp \tag{3.39}$$

$$\beta^{-1}\left(\frac{W_k - W_{k-1}}{h} + \frac{1}{2}(\mathrm{ad}^*_{A_k} W_k + \mathrm{ad}^*_{A_{k-1}} W_{k-1}) \right) = \mathbf{d}P \tag{3.40}$$

$$\frac{W_k - W_{k-1}}{h} + \frac{1}{2}(\mathrm{ad}^*_{A_k} W_k + \mathrm{ad}^*_{A_{k-1}} W_{k-1}) = 0, \tag{3.41}$$

as expected.

3.2.6 Explicit numerical formulas

Recall that the semidiscretized equations were

$$\dot{W} + \mathcal{L}_A W = 0 \tag{3.42a}$$

$$W = \mathbf{d}A. \tag{3.42b}$$

Let us examine in more detail the Lie derivative term $\mathcal{L}_A W$. We expand this term using the discrete version of Cartan's "magic" formula, which states that

$$\mathcal{L}_A W = \mathbf{d}\mathbf{i}_A W. \tag{3.43}$$

Using the definition of the contraction (2.6), we have that, for a 2-form W and a 1-form A,

$$(\mathbf{i}_A W)_{ij} = \sum_m (W_{imj} A_{im} - W_{jmi} A_{jm}). \tag{3.44}$$

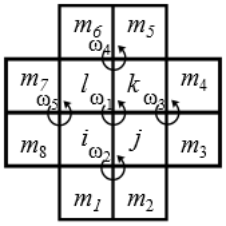

Figure 3.2: A grid for calculating the (discrete) Lie derivative.

Consider cells i and j in Figure 3.2. Then, bearing in mind the non-holonomic constraint, we find

$$(\mathbf{i}_A W)_{ij} = \sum_m (W_{imj} A_{im} - W_{jmi} A_{jm}) \tag{3.45a}$$

$$= W_{ilj} A_{il} + W_{im_1 j} A_{im_1} - W_{jki} A_{jk} - W_{jm_2 i} A_{jm_2}. \tag{3.45b}$$

Recall that if the 2-form W is equal to $\mathbf{d}B^\flat$ for some $B \in \mathcal{S}$, then $W_{im_8 j} = W_{jm_3 i} = 0$ (since the cells are in a straight line), so those terms are not included above. Furthermore, this constraint on W also implies that $W_{ilj} = -W_{jki}$. Write $W_{jki} = \omega_1$, and label the vorticities on other nodes similarly (see Figure 3.2 for details). This results in the equation

$$(\mathbf{i}_A W)_{ij} = -\omega_1 (A_{il} + A_{jk}) + \omega_2 (A_{im_1} + A_{jm_2}). \tag{3.46}$$

Similarly, it can be calculated that

$$(\mathbf{i}_A W)_{jk} = -\omega_1 (A_{ji} + A_{kl}) + \omega_3 (A_{jm_3} + A_{km_4}) \tag{3.47}$$

$$(\mathbf{i}_A W)_{ki} = \omega_1 (A_{il} + A_{kl} - A_{kj} - A_{ij}). \tag{3.48}$$

Putting these together, the discrete form of Cartan's formula is

$$(\mathbf{d}\mathbf{i}_A W)_{ijk} = (\mathbf{i}_A W)_{ij} + (\mathbf{i}_A W)_{jk} + (\mathbf{i}_A W)_{ki} \tag{3.49a}$$

$$= \omega_2 (A_{im_1} + A_{jm_2}) + \omega_3 (A_{jm_3} + A_{km_4}). \tag{3.49b}$$

Notice that, while $W_{ijk} = W_{kli}$ whenever $W = \mathbf{d}B^\flat$ for some $B \in \Omega_1$, it is *not* generally true in this case that $(\mathbf{d}\mathbf{i}_A W)_{ijk} = (\mathbf{d}\mathbf{i}_A W)_{kli}$. This means that while the vorticity values over ijk and

kli will be the same, the update equations for these quantities will still differ. Knowing this, how can we be sure that $W_{ijk} = W_{kli}$ will still hold at the next timestep? The answer comes from the other equation governing the motion of the fluid, which tells us that $W = \mathbf{d}A^\flat$. Given that A has been restricted to satisfy the non-holonomic constraint, this places the desired restriction on the structure of W.

We can further streamline our equations by defining the value $G_{ijkl} = W_{ijk} + W_{kli}$, located on the square dual cell that overlaps with cells i, j, k and l. Physically speaking, this is as if we have added the vorticity integrated over the triangle ijk to the vorticity integrated over the triangle kli to obtain a value integrated over the dual square $ijkl$. It is then natural to use linearity of the exterior derivative and define

$$(\mathbf{di}_A G)_{ijkl} = (\mathbf{di}_A W)_{ijk} + (\mathbf{di}_A W)_{kli} \tag{3.50a}$$

$$= \omega_2(A_{im_1} + A_{jm_2}) + \omega_3(A_{jm_3} + A_{km_4}) \tag{3.50b}$$

$$+ \omega_4(A_{km_5} + A_{lm_6}) + \omega_5(A_{lm_7} + A_{im_8}). \tag{3.50c}$$

This formula is nicely symmetric, and is the form that was used for the numerical results below.

3.3 Results

To show the good qualitative behavior of the method outlined in this chapter, we used it to simulate the behavior of two co-rotating vortices. Each vortex was given by a Gaussian vorticity distribution $\omega = -a \cdot e^{\frac{1}{2}\frac{1-(x^2+y^2)}{a^2}}$, appropriately shifted away from the origin so that the two vortices start at a distance $d = 0.8$ apart. The width factor a was set to be 0.3. These parameters give a system that is very close to a bifurcation point. If the vortices were ever so slightly closer, they would merge. With the parameters given, however, the correct behavior should be that they move apart. Figure 3.4 shows the behavior of corotating vortex simulations using several other methods. It can be seen that methods with good energy behavior generally perform well at this test, whereas methods that introduce numerical dissipation generally cannot resolve the bifurcation point.

In Figure 3.3, we see the behavior obtained from the Hamiltonian semidiscretization, using the midpoint rule in time, on a 50 by 50 grid. It can be seen that the correct behavior is obtained, even for this low resolution.

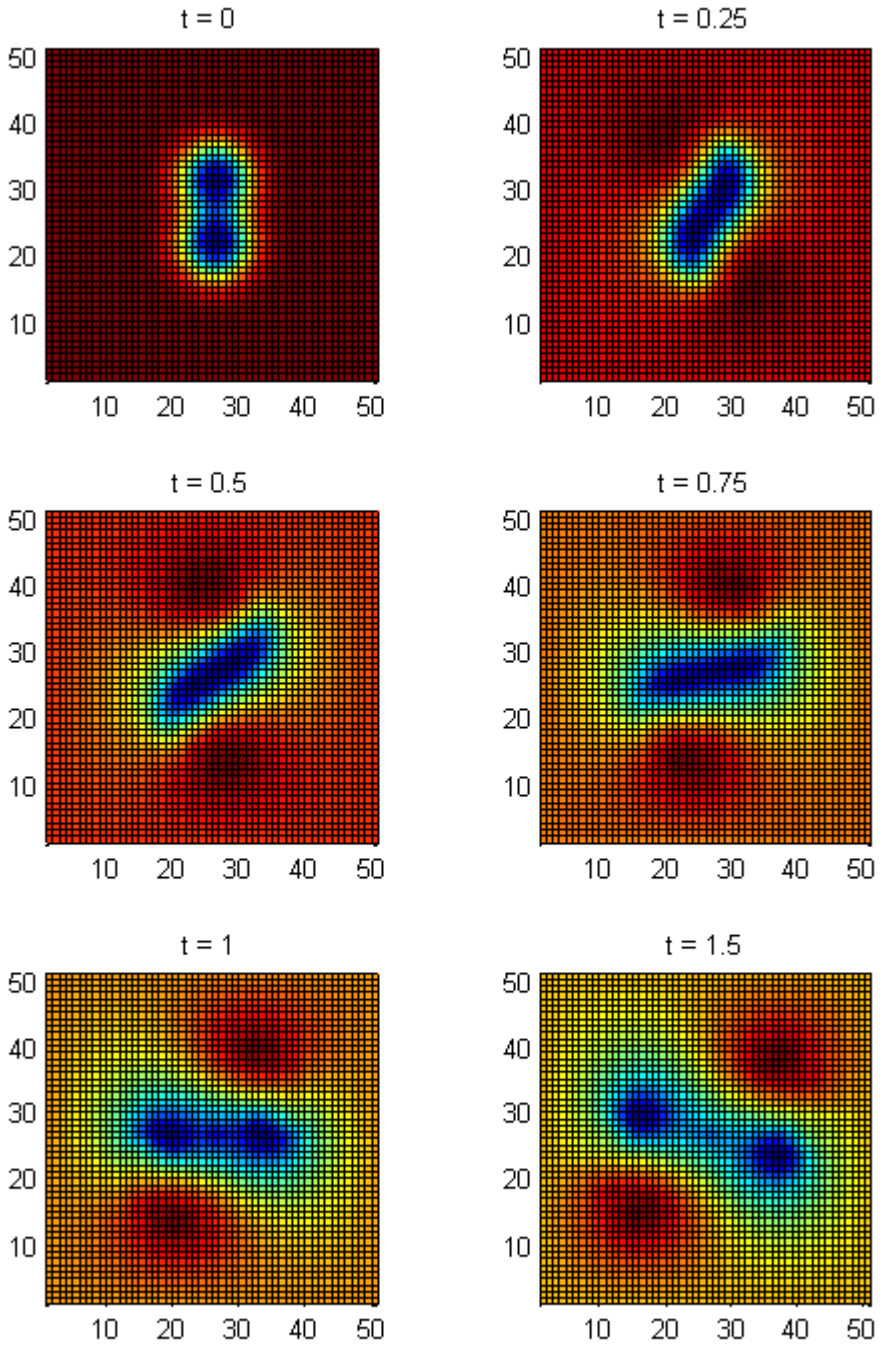

Figure 3.3: The Hamiltonian midpoint method on a 50 by 50 grid captures the behavior of two co-rotating vortices which almost merge, then move apart.

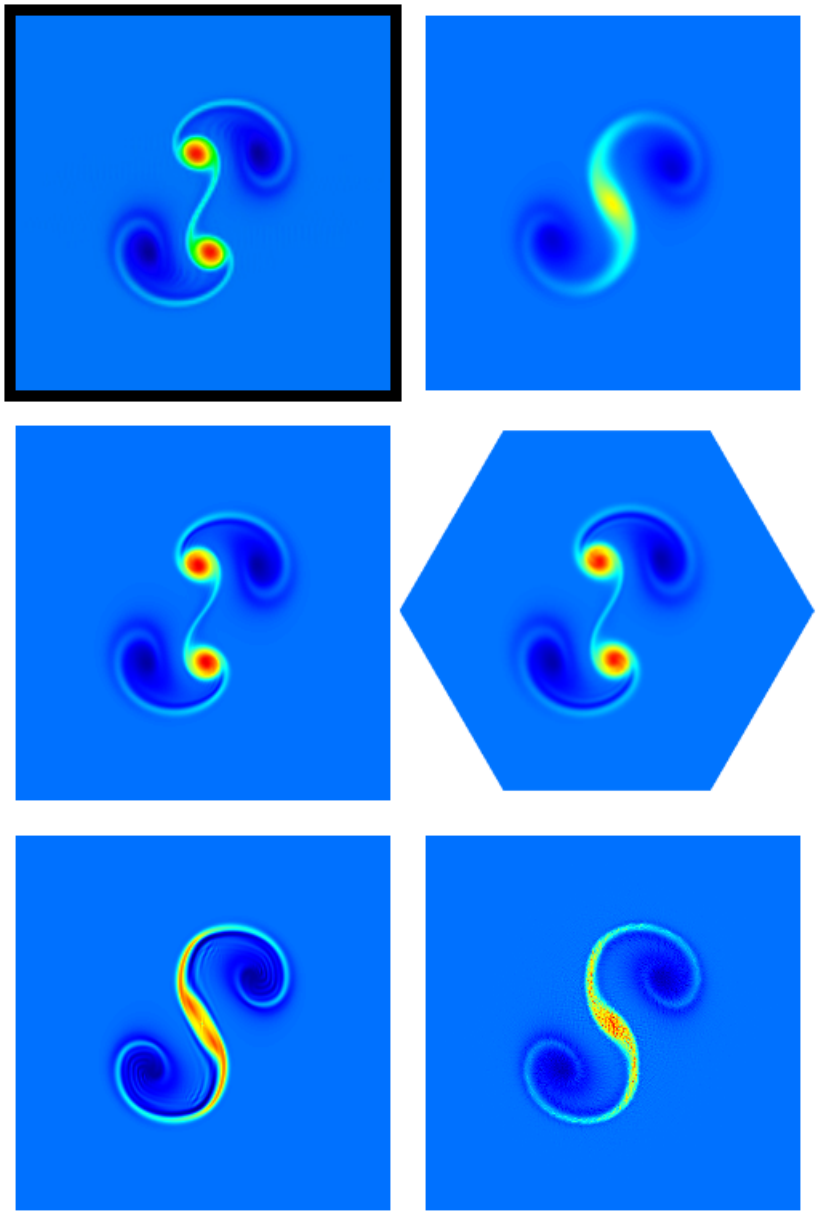

Figure 3.4: Performance of some other methods on the problem of two co-rotating vortices, taken from [28]. From left to right, top to bottom: Reference solution; Stable fluids [36]; energy-preserving scheme (Harlow-Welch [18] with midpoint time discretization); a simplicial energy-preserving scheme [28]; a MacCormack scheme [34]; FLIP [46]. All results were computed on grids of around 2^{16} cells or triangles.

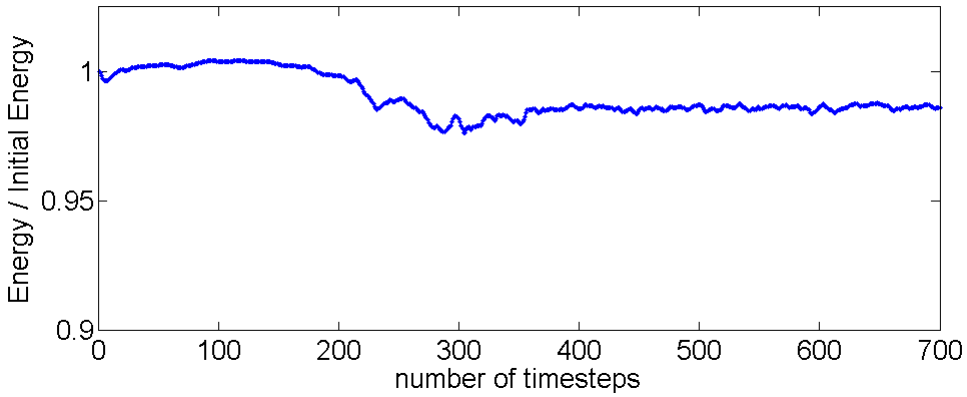

Figure 3.5: Energy divided by initial energy for a long-time simulation using the Hamiltonian trapezoidal method.

We can also see empirically (Figure 3.5) that the Hamiltonian trapezoidal method does not have any innate tendency to gain or lose energy over time. Over 700 timesteps, with a spatial resolution of $h = 0.08$ and timesteps of length 0.05, the energy drifts by less than three percent. Unlike the results from a symplectic method, here there is nothing to stop the energy from eventually drifting arbitrarily far from the start point. However, the energy behavior of this method is stable enough that this drift would take a very long time to occur. If we instead use the midpoint rule, we attain exact energy conservation as expected.

3.4 Discussion

The Hamiltonian methods detailed here expand our theoretical and practical understanding of the structure-preserving simulation of incompressible fluids. The spatial semi-discretization has been shown to be possible either using a continuous-time variational formalism, or using a discrete version of the Lie-Poisson bracket. We have also seen how to construct a variational Lie-Poisson time integrator in the presence of a non-holonomic constraint. We have created an energy and circulation preserving vorticity method.

In this case, because of the non-holonomic constraint, the discrete Lie-Poisson bracket does not satisfy the Jacobi identity. However, we have seen that it is nevertheless useful in the creation of a numerical method, particularly since the conservation of energy relies only on the antisymmetry of the bracket, a property which we still keep.

In the course of creating these methods, we have included the discrete codifferential, which was previously proposed by Bochev et al [6]. In combining the Hodge star with a derivative, the codifferential is sometimes more elegant than a direct Hodge star, particularly since the simplest version of the Hodge star used in DEC is the diagonal Hodge star, which is, strictly speaking, a zeroth-order approximation, whereas the codifferential *based* on the diagonal Hodge star may be created so as to be second-order on a regular grid, as it is essentially a central difference method.

Future work includes the extension of this work to three dimensions, and the potential application of the stream function formalism to embedded boundaries. The work done here is also potentially applicable to other Lie-Poisson systems besides that of fluids, such as the Poisson-Vlasov system, which describes plasma [27].

Chapter 4

Model-Reduced Lagrangian Fluids

So far, we have seen that the technique of using a finite-dimensional Lie group to approximate the infinite-dimensional Lie group $\mathrm{Diff}_{\mathrm{vol}}(\mathcal{M})$ is useful for creating a semidiscretization that preserves the variational structure of the equations, resulting in energy-preservation (with an appropriate time discretization), a discrete version of Kelvin's circulation theorem, and good qualitative behavior in resolving the bifurcation point between merging and separation of two co-rotating vortices. However, the resulting implicit methods are not especially fast, making them impractical for use in, for example, computer graphics. Moreover, the introduction of the non-holonomic constraint was a computational necessity in order to avoid dense matrices whose number of entries would be on the order of $(\Delta x)^{-2}\times(\Delta x)^{-2}=(\Delta x)^{-4}$, where Δx is the spatial step size. This non-holonomic constraint then made it impossible to use this technique to construct a symplectic numerical method. Thus, we aim to improve the speed of computation, with an eye to perhaps making the non-holonomic constraint unnecessary by reducing the number of variables. In collaboration with Beibei Liu and Yiying Tong, we consider how to introduce *model reduction* to a variational fluids method. That is, instead of describing the fluid by recording the value of a function at many different points is space, we instead describe functions by projecting them onto a partial basis for the function space. This allows us to use fewer variables while still retaining considerable accuracy. Interestingly, it does *not* remove the need for a non-holonomic constraint, and the insights gained from the recurrence of this necessity will be detailed at the end of the chapter.

4.1 Derivation

We will first define the discrete, reduced scalar and velocity fields on which our functional map Lie group will act. Extending what was advocated in [10], we expand functions and vector fields using the orthonormal bases for 2-forms and 3-forms given by the eigenfunctions of the (deRham-)Laplacian operators on an arbitrary discrete mesh $\mathcal{M}$. These are calculated using the discrete operators of Finite-Element Exterior Calculus (FEEC [1]) and Discrete Exterior Calculus (DEC [11]), allowing us to leverage the large literature on their implementation and structure-preserving properties with respect to topology. Although there are several sets of orthogonal eigenfunctions that we could use, this set also has the advantage of simplifying our equations in the event that we wish to add vorticity. This set of basis functions can efficiently encode through reduced coordinates the full-space fields typically used in the MAC scheme, i.e., fluxes through cell boundaries (discrete 2-forms) to represent velocity fields, and densities integrated in each cell (discrete 3-forms) to represent scalar fields (such as smoke density).

4.1.1 Spectral Bases

Choice of eigenfunctions. We consider a reduced space of volume forms given by the span of the first N Laplacian eigenfunctions on our space. To find these Laplacian eigenfunctions, we follow the approach of [10]. That is, we either calculate Laplacian eigenfunctions explicitly (in the case of a simple space like a circle or rectangle), or, more likely, we use a mesh, which need not be regular, and calculate eigenfunctions of a discrete Laplacian operator on that mesh.

We denote the i-th eigenfunction of 3-form Laplacian Δ_3 as Φ_i with associated eigenvalue $-\mu_i^2$,

$$\Delta_3 \Phi_i = -\mu_i^2 \Phi_i.$$

The eigenfunctions corresponding to the M_3+1 *smallest* eigenvalues μ_i can be assembled into a low-frequency basis:

$$\{\Phi_0, ..., \Phi_{M_3}\}.$$

A general volume form ρ will then be represented by a vector $(\rho_0, ..., \rho_{M_3})$ of (M_3+1) values, each representing the inner product of ρ with an element of this reduced orthonormal basis, so that $\rho = \sum_{i=0}^{M_3} \rho_i \Phi_i$. Note that, depending on the boundary condition, $\mu_0 = 0$ may correspond to more

than one harmonic function, but these are not influenced by divergence-free velocity fields with zero flux across the boundary and are omitted in our discussion.

Similarly, we denote the i-th eigenfunction of 2-form Laplacian Δ_2 as Ψ_i, with its associated eigenvalue $-\kappa_i^2$:

$$\Delta_2 \Psi_i = -\kappa_i^2 \Psi_i.$$

We also assemble the first M_2 eigenvector fields (corresponding to the M_2 smallest eigenvalues) into a finite dimensional low-frequency basis:

$$\{\Psi_1, ...\Psi_{M_2}\}.$$

Some of the 2-form eigenfunctions are not divergence-free, and these eigenfunctions can be identified as gradient fields, $\nabla\Phi_i/\mu_i$. Thus, we can reorder the eigenfunctions of Δ_2 into

$$\{h_1, ..., h_{\beta_1}, \frac{\nabla\Phi_1}{\mu_1}, ..., \frac{\nabla\Phi_{M_3}}{\mu_{M_3}}, \Psi_1, ...\Psi_{M_C}\},$$

where h_i are harmonic 2-forms (corresponding to frequency $\kappa_i = 0$) with β_1 is the first Betti number determined by the topology of the domain (basically, the number of tunnels plus the number of connected components of the boundary minus one), and $M_C = M_2 - M_3 - \beta_1$ denoting the number of non-harmonic but divergence-free basis functions.

Lie Group action. In similar fashion to the other derivations outlined in this thesis, we represent the action of a volume-preserving diffeomorphism on the space using the approach of Koopmanism, by looking at how a given diffeomorphism acts on functions. We encode the fluid motion through a time-varying Lie group element $q(t)$ that represents a functional map induced by the fluid flow ϕ_t, mapping a function $f(\mathbf{x}) = \sum_i f_i \Phi_i(\mathbf{x})$ linearly to another function $g(\mathbf{x}) = \sum_i g_i \Phi_i(\mathbf{x})$ such that $g(\mathbf{x}) = f \circ \phi^{-1}(\mathbf{x})$. Since a volume form $f\mathrm{d}V$ is represented by a vector of $(M_3 + 1)$ values, a diffeomorphism q can be encoded by a $(M_3 + 1)$-by-$(M_3 + 1)$ matrix.

The volume-preserving property of the flow still implies the orthogonality of the matrix q, i.e, $q^t q = \mathrm{Id}$. So we are looking for a subgroup of $O(M_3+1)$, or, more accurately, of $SO(M_3+1)$, since we wish to describe gradual changes from the identity.

The condition that constant functions are mapped to themselves in this low-frequency Lie group

becomes $q_{0i} = \delta_{0i}$ and $q_{i0} = \delta_{i0}$, where δ_{ij} is the Kronecker symbol, since 0-th frequency represents the constant function. This effectively removes one dimension of possible changes that an element of the Lie group could make. The Lie group that we are using is thus isomorphic to $SO(M_3)$. This is much smaller than the full Lie group used for the spatial representation [32], which had a dimension proportional to the square of the number of cells of the mesh—a potential reduction of several orders of magnitude.

Lie Algebra. We identify each velocity eigenfunction Ψ_m with an element of the Lie algebra of the above Lie group as follows. We take the Lie derivative along the velocity field Ψ_m of a scalar eigenfunction Φ_i projected onto another scalar eigenfunction Φ_j, which is a matrix A_m for each velocity eigenfunction Ψ_m, with entries

$$A_{m,ij} = \int_{\mathcal{M}} \Phi_i (\Psi_m \cdot \nabla \Phi_j). \tag{4.1}$$

As in the non-spectral case, the divergence-free condition leads to the antisymmetry of these matrices:

$$\begin{aligned} A_{m,ij} + A_{m,ji} &= \int_{\mathcal{M}} (\Phi_i \Psi_m \cdot \nabla \Phi_j + \Phi_j \Psi_m \cdot \nabla \Phi_i) && \text{(4.2a)} \\ &= \int_{\mathcal{M}} \Psi_m \cdot \nabla (\Phi_i \Phi_j) && \text{(4.2b)} \\ &= -\int_{\mathcal{M}} \Phi_i \Phi_j \nabla \cdot \Psi_m && \text{(4.2c)} \\ &= 0. && \text{(4.2d)} \end{aligned}$$

This is expected, since the Lie algebra $\mathfrak{so}(M_3)$ of $SO(M_3)$ contains only antisymmetric matrices. The Lie algebra has a Lie bracket given by the usual matrix commutator $[A_m, A_n] = A_m A_n - A_n A_m$.

4.1.2 Non-holonomic constraint.

There are $M_3(M_3 - 1)/2$ degrees of freedom in the Lie algebra of antisymmetric matrices $\mathfrak{so}(M_3)$. However, we cannot simply use the first $M_3(M_3 - 1)/2$ elements of the Laplacian basis of vector fields, because most of these elements will be of such high frequency that when we calculate its effect on a volume form basis element Φ_j according to the definition (4.1), we will find that the result is

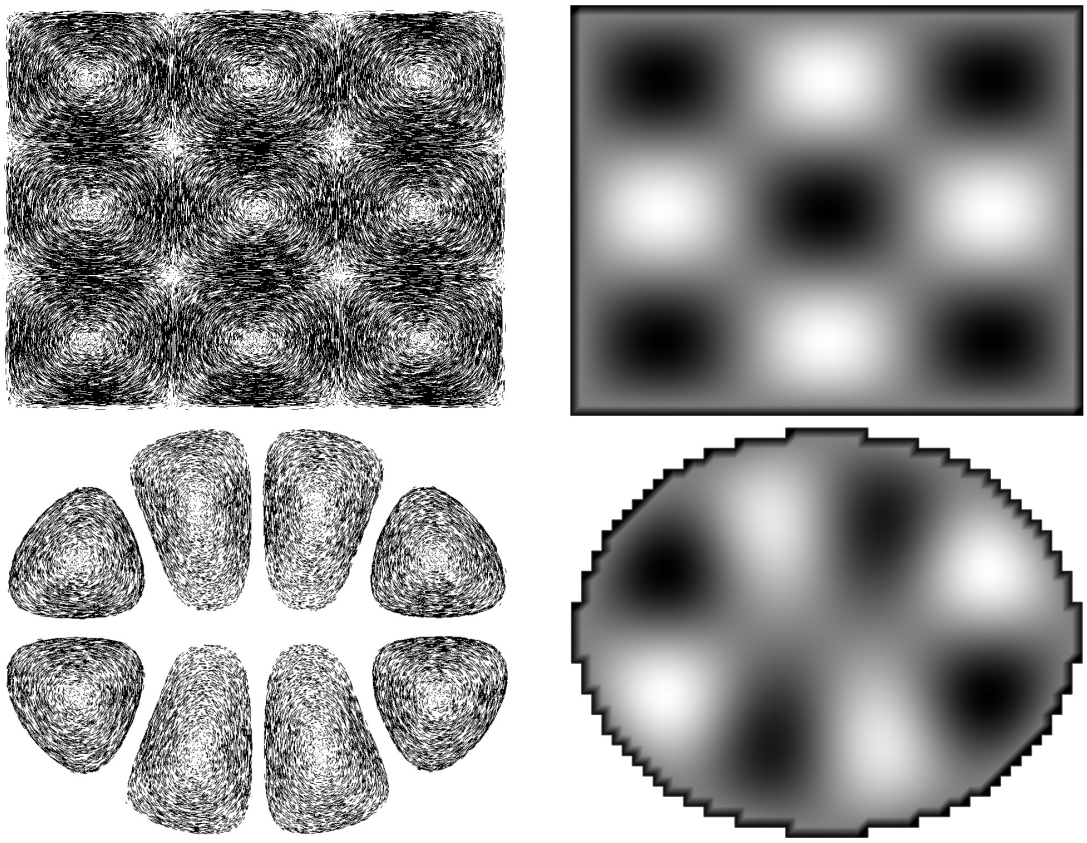

Figure 4.1: **Effect of shape on spectral bases**: The Laplacian eigenvectors depends heavily on the domain Ω. Here, rectangle (top) vs. ellipse (bottom) domains (both computed on 2D rectangular grid of size 120^2) exhibit very different eigenvectors Ψ_{10} and Φ_{10}.

still of such high frequency that its inner product with each basis element Φ_i gives zero. The space of divergence-free vector fields acting on the reduced basis of volume forms is of lower dimension than the Lie algebra. It follows that some elements of the Lie algebra must not correspond to a vector field at all.

We force the dynamics on the Lie algebra to remain within the domain of physically-sensible elements using the following non-holonomic constraint, which keeps the velocity within the space spanned by the lowest frequency $M_C+\beta_1$ divergence-free 2-form basis fields:

$$A = \sum_{i=1}^{M_C+\beta_1} v_i \, A_i \tag{4.3}$$

where v_i is a coefficient for A_i representing the modal amplitude of frequency κ_i. This linear condition can be seen as an intuitive extension of the one-away spatial constraint on Lie algebra elements that was used in [32] and explained in Section 2.2.3. Instead of constraining interactions to be local in space, we have constrained them to the part of the Lie algebra that corresponds to lower-frequency basis functions[1].

4.1.3 Implementation

Computing our spectral bases requires a proper discretization of the Laplacian operators and of boundary conditions. Both topics are well studied, and many implementations can be leveraged [12, 4]. A detailed guide to discretization on *arbitrary unstructured meshes* may be found in Appendix A of [23] to explicate how to enforce no-transfer and free-slip conditions (corresponding, respectively, to $v_n|_{\partial\mathcal{M}}=0$ and $\partial v_t/\partial n\,|_{\partial\mathcal{M}}=0$ if the continuous velocity field is decomposed at the boundary into its normal and tangential components, $v=v_n+v_t$). Note that only two operators are required: the exterior derivative d and the Hodge star $\star$. The first operator is purely topological, while the second is just a scaling operation per edge, face, and cell. Moreover, as explained in the next paragraph, this latter operator can be trivially modified to handle arbitrary fluid domains without having to use anything else but a regular grid. From these two operators, both Laplacians are easily assembled, and low-frequency eigenfields are found via Lanczos iterations since each operator is symmetric by construction.

[1] Note that the rest of the Lie algebra does not actually correspond to higher-frequency basis functions (if it did, we could keep them). Rather, it corresponds to elements with no clear physical interpretation.

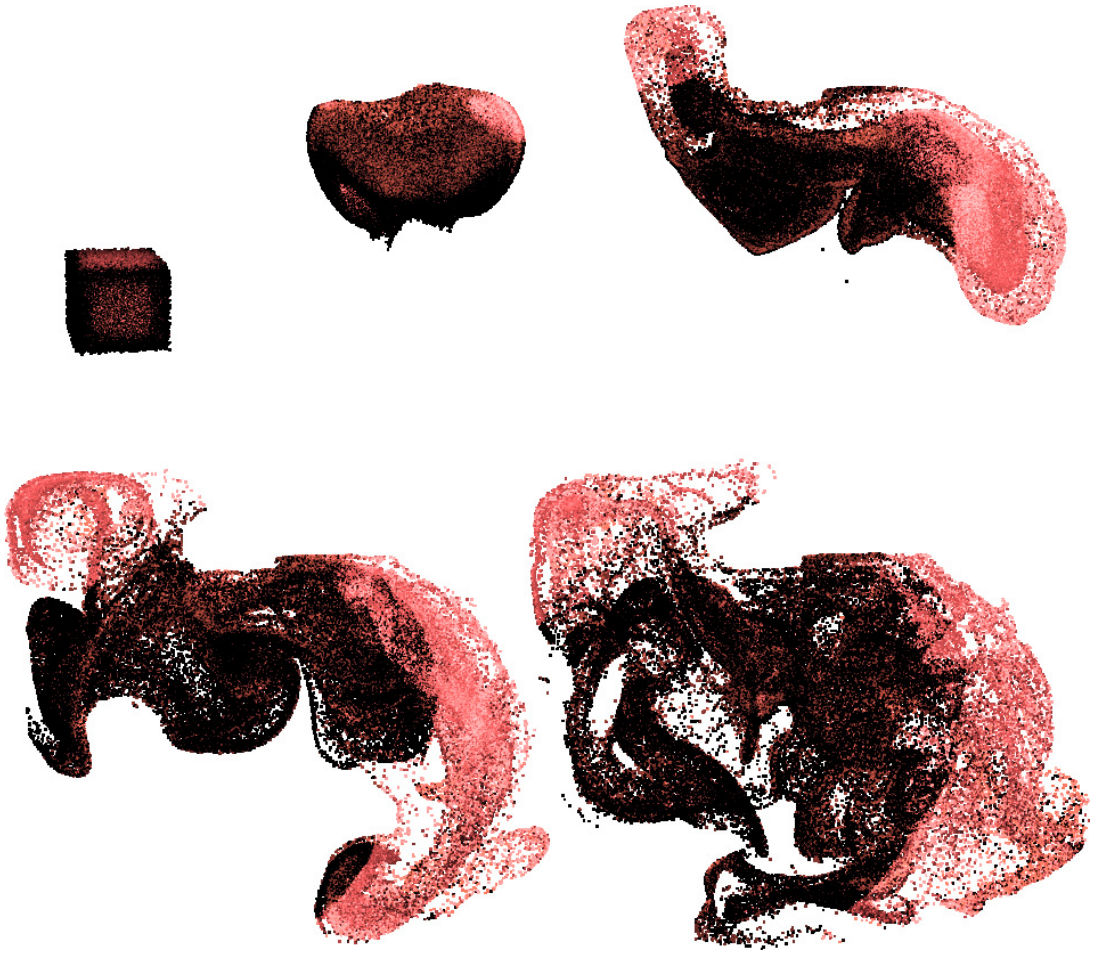

Figure 4.2: **3D bunny buoyancy test**: A hot cube of air initially located at the center of a 3D bunny-shaped domain is advected through buoyancy. Computations were performed using a modified Hodge star on a $42\times42\times32$ grid, with only 100 modes.

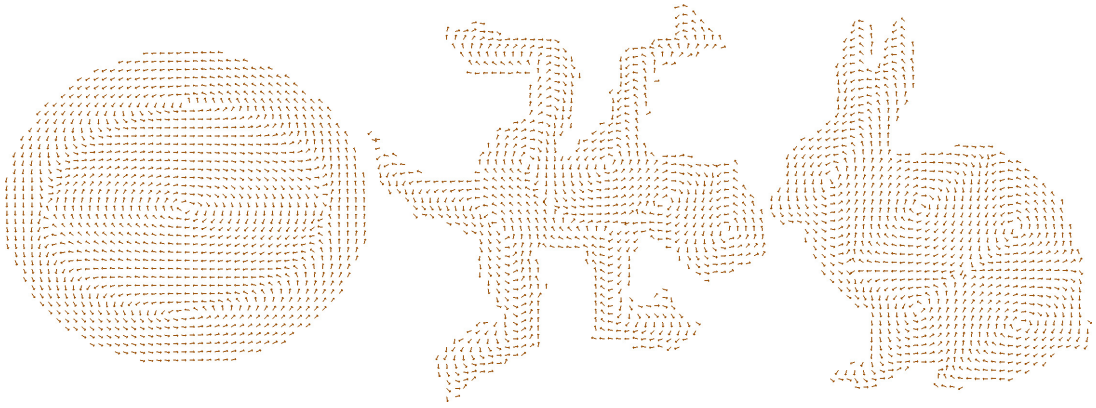

Figure 4.3: **Domain-altered Hodge stars**: Hedge-hog visualization of Ψ_5 on a 256^2 grid for three different 2D domain shapes, obtained through a simple alteration of the Hodge star $\star$ operator.

To implement embedded boundaries we propose a simple extension of the technique of [30] to compute k-form Laplacians of an arbitrary domain *while still using a regular grid*. This renders the implementation of Laplacians and their boundary conditions quite trivial, and removes the arduous task of tetrahedralizing arbitrary domains. This idea was introduced in [3] for their pressure-based projection, and a simple alteration proposed by [30] made the approach robust and convergent. We leverage this latter work by noticing that the modification of the Laplacian Δ_3 that they proposed amounts to a local change to the Hodge star operator $\star_2$.

More precisely, consider a continuous domain Ω, e.g., defined implicitly by a function χ via $\Omega = \{\mathbf{x}|\chi(\mathbf{x}) \geq 0\}$ (see Figure 4.4). Recall that the diagonal Hodge stars on a mesh $\mathcal{M}$ are all expressed using local ratios of measurements (edge lengths, face areas, cell volumes) on both the primal elements of $\mathcal{M}$ and its dual elements [11]. The changes to the Laplacian operator Δ_3 that Ng et al. [30] introduced can be reexpressed by an alteration of the Hodge star $\star_2$, where *each primal area measurement only counts the part of the primal face that is inside* Ω, but dual edge lengths are kept unchanged. We extend this simple observation (which amounts to a local, numerical homogenization to capture sub-grid resolution) by computing modified Hodge stars $\hat{\star}_1, \hat{\star}_2$, and $\hat{\star}_3$ where only the parts of the primal elements (partial lengths, areas, or volumes) that are within the domain Ω are counted (see inset). Note that changing directly the Hodge stars does not affect the symmetry and positive-definiteness of the Laplacians, and thus incurs no additional cost for our method. This straightforward extension allows us to compute our spectral bases on regular grid for arbitrary domains Ω as illustrated in Fig. 4.3 for a basis element of vector fields. We also show the

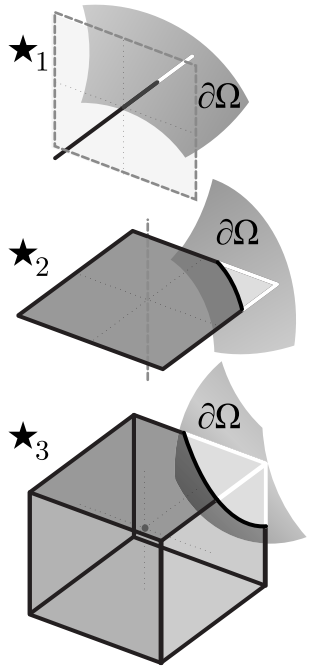

Figure 4.4: Diagram of the Hodge star near an embedded boundary.

behavior of this Hodge star modification under refinement of the regular grid for a given continuous elliptic domain Ω, resulting in very good approximations of the eigenvectors.

4.1.4 Spectral variational integrator

The Lagrangian of the fluid motion (i.e., its kinetic energy in the case of Euler fluids) can be written as $\mathcal{L}_{\text{Euler}} = \frac{1}{2}\langle A(t), A(t)\rangle$ as we reviewed in Sec. 2.3. Thus, the equation of motion can be derived from Hamilton's (least action) principle:

$$\int \langle A(t), \delta A\rangle dt = 0 \tag{4.4}$$

under the Lin constraint [14] restricting the variations of A to those of the form

$$\delta A = \dot{B} + [A, B], \tag{4.5}$$

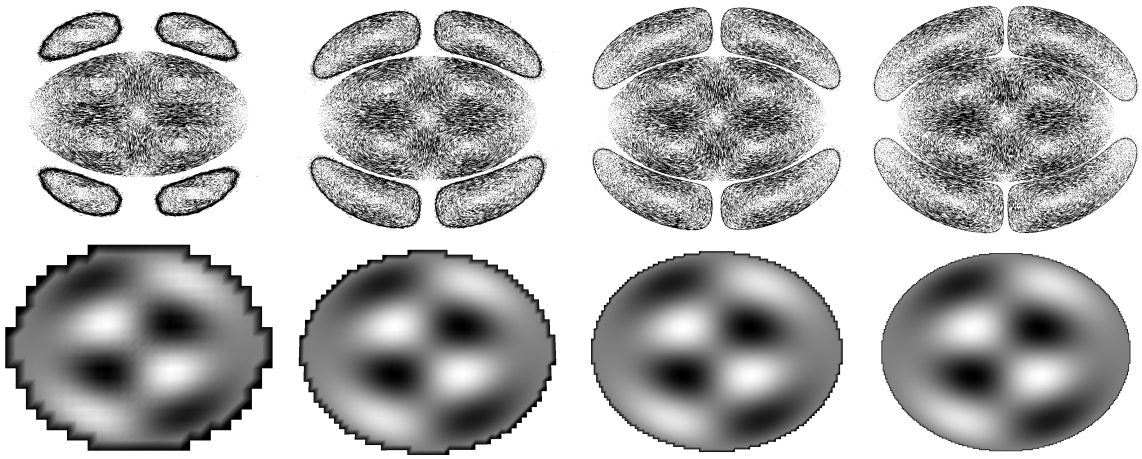

Figure 4.5: **Convergence of Laplacians**: Our discretization of the two Laplacians creates (vector and scalar) eigenfields that converge under refinement of the regular grid used to compute them, extending the linear convergence proved in [30]. Here, particle-tracing visualization of the 15^{th} eigenbasis for vector fields on the ellipse (top) at resolution 30^2, 60^2, 120^2, and 240^2, and 15^{th} eigen function (bottom) at the same resolutions.

where $B = \sum_i \xi_i A_i$ is an arbitrary element of the Lie algebra with coordinates $\{\xi_i\}_i$ in the 2-form basis. Substituting Eq.(4.5) into Eq.(4.4), we then have

$$\begin{aligned} 0 &= \int \langle A, \delta A \rangle \, dt \\ &= \int \sum_{i,k} v_k \dot{\xi}_k \langle A_i, A_k \rangle + \sum_{i,j,k} v_i v_j \xi_k \langle A_i, [A_j, A_k] \rangle \, dt \\ &= \int \sum_k \left(- \sum_i \dot{v}_k \langle A_i, A_k \rangle + \sum_{i,j} v_i v_j \langle A_i, [A_j, A_k] \rangle \right) \xi_k \, dt. \end{aligned}$$

Since this last equation must be valid for any ξ_k, the update rule for the velocity field has to be

$$\dot{v}_k = \sum_{i,j} v_i v_j \langle A_i, [A_j, A_k] \rangle \equiv \mathbf{v}^T \mathbf{C}_k \mathbf{v}, \tag{4.6}$$

where $\mathbf{v}$ is the column vector storing the coefficients v_i of the discrete velocity A (Eq. 4.3), and $\mathbf{C}_k$ is the square matrix with components

$$C_{k,ij} = \langle A_i, [A_j, A_k] \rangle = \int_{\mathcal{M}} (\nabla \times \Psi_i) \cdot (\Psi_k \times \Psi_j). \tag{4.7}$$

Time integrator. The continuous-time update in Eq. 4.6 is then discretized via either a midpoint rule (which will lead to an *energy-preserving* model-reduced variant of [28]) or a trapezoidal rule (which corresponds to a model-reduced variant of the variational method of [32]). Specifically, the midpoint rule is implemented as

$$v_k^{t+h} - v_k^t = h \sum_{i,j} C_{k,ij} \frac{v_i^t + v_i^{t+h}}{2} \frac{v_j^t + v_j^{t+h}}{2}. \tag{4.8}$$

The energy preservation can be easily verified by multiplying $v_k^t + v_k^{t+h}$ on both sides of the above equation, summing over k, and invoking the property of coefficients $C_{k,ij} = -C_{j,ik}$. The trapezoidal rule can, instead, be implemented as

$$v_k^{t+h} - v_k^t = \frac{h}{2} \sum_{i,j} C_{k,ij} (v_i^t v_j^t + v_i^{t+h} v_j^{t+h}), \tag{4.9}$$

which is derived from a temporal discretization of the action with variation of $(\delta q)q^{-1}$ for q along the path to be in the restricted Lie algebra set (to enforce Lin constraints). Both the energy-preserving and trapezoidal variational method are time-reversible implicit methods solved through a simple quadratic set of equations with a small number of variables. Note that an explicit forward Euler integration can also be used for small time steps, but with no guarantee of good behavior over long periods of time.

4.1.5 Kelvin's circulation theorem

The model-reduced method obeys a form of Kelvin's theorem as follows. We can define generalized curves as spectral dual 1-chains (also called 1-currents [11]) of the form:

$$\Gamma = \sum_i \gamma_i \star_2 A_i. \tag{4.10}$$

Although generalized curves of this form do not necessarily have an obvious relation to a one-dimensional curve in space, the above expression represents a closed curve in the sense of chains, because each A_i corresponds to a closed (divergence-free) 2-form, which means that its dual 1-chain

is boundaryless. A pairing between a 2-form and a generalized loop is defined as expected:

$$\langle A, \Gamma \rangle = \left\langle \sum_i v_i A_i, \sum_j \gamma_j \star_2 A_j \right\rangle = \sum_i v_i \gamma_i. \tag{4.11}$$

The Lie advection of the generalized curve along the velocity field

$$\dot{\Gamma} = -[A, \Gamma] \tag{4.12}$$

indicates that the coefficients $\{\gamma_i\}_i$ must evolve such that

$$\begin{aligned} \dot{\gamma_k} &= -\sum_{i,j} \gamma_i v_j \int_{\mathcal{M}} \Psi_k \cdot (\nabla \times (\Psi_i \times \Psi_j)) \\ &= \sum_{i,j} \gamma_i v_j \int_{\mathcal{M}} (\nabla \times \Psi_k) \cdot (\Psi_j \times \Psi_i) = \sum_{i,j} \gamma_i v_j C_{j,ki}. \end{aligned}$$

Thus, the spectral version of Kelvin's theorem holds since

$$\begin{aligned} \frac{d}{dt} \langle A, \Gamma \rangle &= \sum_i (\dot{v_i} \gamma_i + v_i \dot{\gamma_i}) \\ &= \sum_i \mathbf{v}^t \mathbf{C}_i \mathbf{v} \gamma_i + \sum_i v_i \sum_{j,k} v_k \gamma_j C_{k,ij} \\ &= \sum_{i,j,k} v_j C_{i,jk} v_k \gamma_i - \sum_{i,j,k} v_j v_k \gamma_i C_{i,jk} = 0. \end{aligned}$$

In the above derivation, dummy index variables are swapped and the identity $C_{k,ij} = -C_{j,ik}$ is used.

4.2 Results

All these results are taken from [23], with thanks to Beibei Liu and Julian Hodgson for the images, and were generated on an Intel i7 laptop with 12GB RAM.

Reduced vs. full simulation. In order to check that validity of our reduced approach, we performed a stress test in a periodic 2D domain to visualize how the increase in the number of bases used in our spectral integrator impacts the simulation over time. We selected a band-limited

initial velocity field at time $t=0$ that only contains non-zero components for the first 120 frequencies. We then advected the fluid using our integrator, with fluid markers initially set as two colored disks near the center. Because of the propensity of vorticity to go to higher scales, our reduced approach does not lead to the exact same position of the fluid markers after 12s of simulation if one uses only 120 bases. However, as the number of bases increases to 300 or 500, the simulation quickly captures the same dynamical behavior as a full variational integrator with 256^2 degrees of freedom (see Figure 4.6).

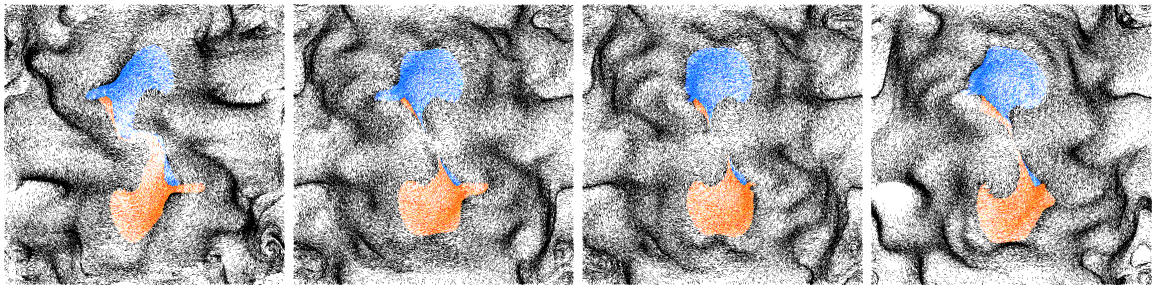

Figure 4.6: **Convergence of simulation**: A flow in a periodic domain is initialized with a band-limited velocity fields with 120 wave number vectors. Fluid markers (forming a blue and red circle) are added for visualization. After 12s of simulation, the results of our reduced approach (from the left: with 120, 300, 500 modes) vs. the full 256^2 dynamics (right) are qualitatively similar.

Arbitrary domains. We also show in Figs. 4.1, 4.3, and 4.8 for 2D and Figs. 4.7, 4.2, 4.10, and 4.11 for 3D that the use of a homogenized boundary condition on regular grids leads to the expected visual behavior near the domain boundaries, eliminating the staircase artifacts of traditional immersed-grid methods. Our homogenized boundary treatment obtains results similar to those of unstructured meshes while using only calculations that are directly performed on regular grids—thus requiring significantly simpler, smaller, and more efficient data structures. Our fluid dynamics is also consistent across a wide range of temporal and spatial discretizations (see Figure 4.8). In addition, the regular grid structure also simplifies the interaction with immersed solid objects as in the case demonstrated in Fig. 4.10 through a flow induced by a scripted car turning around a corner; interactive fluid stirring by a paddle manipulated by the user is also easily achieved as shown in Fig. 4.11. Finally, our spectral integrator can be carried out in the same fashion on curved domains as well, since the eigenvectors of the Laplace(-Beltrami) operator are no more difficult to compute on a triangulated surface; Fig. 4.9 shows a simple laminar flow on the *surface* of the bunny model.

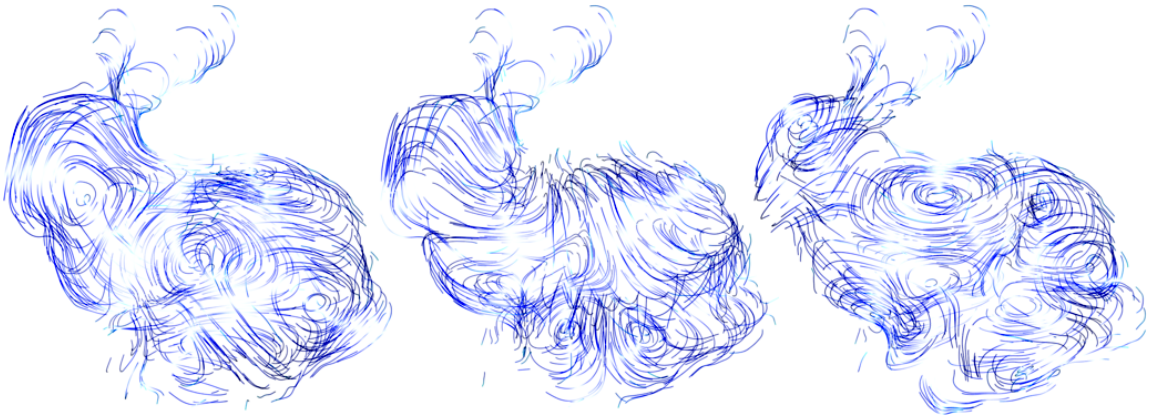

Figure 4.7: **Model-reduced fluids on regular grids.** Our energy-preserving approach integrates a fluid flow variationally using a small number of divergence-free velocity field bases over an arbitrary domain (visualized here are the 5^{th}, 10^{th}, and 15^{th} eigenvectors of the 2-form Laplacian) computed with subgrid accuracy on a regular grid (here, a $42\times42\times32$ grid). Our integrator is versatile: it can be used for realtime fluid animation, magnetohydrodynamics, and turbulence models, with either explicit or implicit integration.

Advanced fluid models. We also extended our method to the LANS-α turbulence model to better capture the spectral energy distribution with a small number of modes. On a 3D regular grid, we performed a simulation as described in [13] by holding the low wave number components v_i fixed for $|\kappa_i|<2$ to act as a forcing term, and running the simulation until $t=100$. We then extracted the average spectral energy distribution present between $t=33$ to $t=100$. We show in Fig. 4.12 that the Kolmogorov "$-5/3$ law" is much better captured than with the usual Navier-Stokes model, even for the low number of modes used in our spectral context: the α-model produces a decay rate at high wave numbers much steeper than a Navier-Stokes simulation, allowing us to cut off the higher frequencies at a lower threshold without significant deviation from the spectral distribution. This indicates that our approach consisting in a simple scaling of the structural coefficients helps improving fluid simulation on coarse grids.

Computational efficiency. Our use of model reduction via Laplacian eigenbases provides a significantly more efficient alternative to full simulators, obviously. Due to our variational treatment of time integration, we also prevent many shortcomings of the previous reduced models as we ensure consistency of the results over a large spectrum of spatial and temporal discretization rates, and maintain a qualitatively correct behavior even on coarse grids. The efficiency gain compared to the full variational simulation is apparent, be it in 2D, curved 2D, or 3D. For instance, a full-

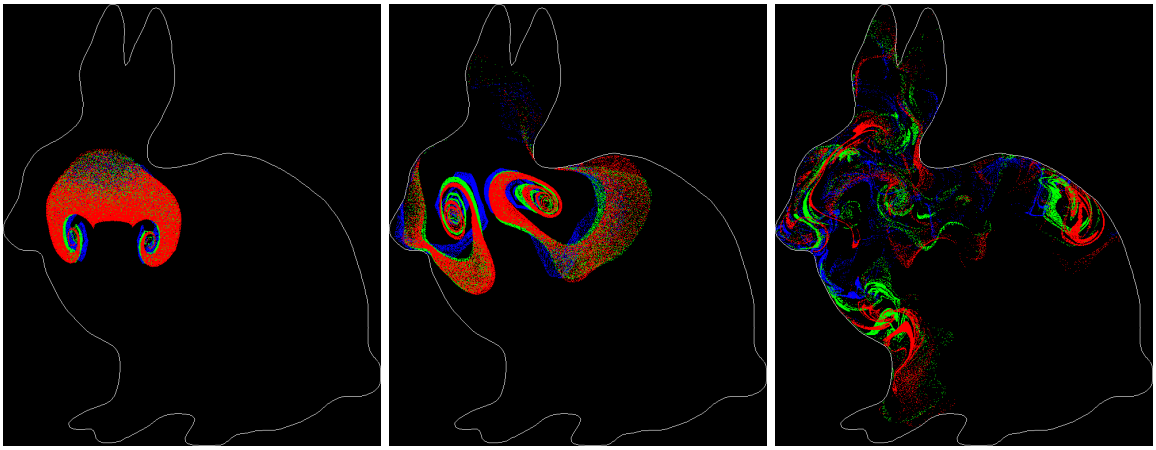

Figure 4.8: **Robustness to resolution**: With the homogenized boundary condition on grids of resolution 40^2 (blue), 80^2 (green), and 160^2 (red), no staircase artifacts are observed, and the simulation results are consistent across resolutions.

blown 128^2 grid takes around 50s for the variational integrator to update one step (through a Newton solver) in a typical simulation using the trapezoidal rule update, while a 50-mode (resp., 100-mode and 200-mode) simulation with our integrator takes only 0.098s (resp., 0.65s and 2.0s) for complex boundaries (i.e., with dense structural coefficients), and 0.026s (resp., 0.070s, 0.28s) for simple box domains (with sparse coefficients). Our Newton solver normally converges in a couple of iterations depending on the time step size (which determines the quality of the initial guess); for instance, the average in our 3D bunny buoyancy test in Fig. 4.2 is below 3 iterations. Depending on how many modes the user is willing to discard (and replace by wavelet noise or dynamical texture for efficiency), the computational gains can thus total several orders of magnitude, and this allows us to simulate flows at interactive rates—or even in realtime for periodic 3D domains if we compute the eigenbases in closed form as shown in Fig. 4.11. Note however that our model-reduced integrator suffers from the usual limitation of model reduction: the complexity is actually growing quadratically (resp., cubically) with the number of modes for sparse (resp., dense) structural coefficients. So our integrator is numerically efficient only for relatively low mode counts.

4.2.1 Generalization to other bases

While we provided detail on the construction of a variational model-reduced integrator for fluid simulation using Laplace eigenvectors, one can easily adapt our approach to arbitrary basis functions, even those extracted from a training set of fluid motions. Suppose that we are given a set

of scalar basis elements Φ_i (orthonormalized through the Gram-Schmidt procedure) and a set of velocity basis elements Ψ_i. The Lie derivative matrix A will still be antisymmetric as long as Ψ_i's are divergence-free. This means that one can use existing finite element basis functions instead of our Laplace eigenvectors—or even wavelet bases of $H(\mathrm{div}, \Omega)$ (see for instance [42]) if spatially localized basis functions are sought after to get a sparser advection. The key to the numerical benefits of our variational approach is to ensure the anti-commutativity of the Lie bracket in the evaluation of $\langle A_i, [A_j, A_k] \rangle$ (Eq. (4.7)) and the advection of other fields by the exponential map (or approximations thereof) of the matrix representing the Lie derivative (as done in MHD and complex fluids [14]). In a way, the original non-spectral variational integrators can be seen as a special case of our framework where Whitney basis functions are used. However, viscosity can no longer be handled as easily in this case as the Laplacian is not diagonal in general bases. Moreover, the required number of degrees of freedom to produce smooth flows may be high if the bases are arbitrary.

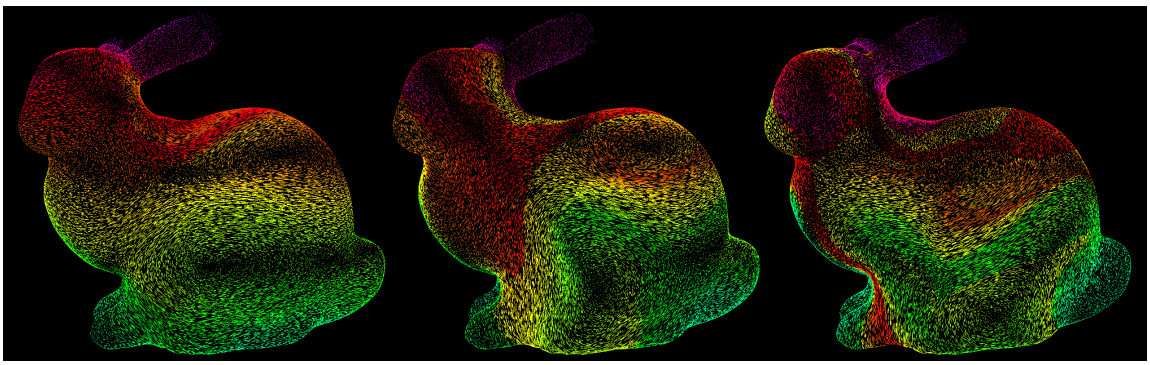

Figure 4.9: **Curved domain**: While all our other results were achieved on a regular grid, our approach applies to arbitrary domains, here on the *surface* of a triangulated domain; a simple laminar flow with initial horizonal velocity smoothly varying along the vertical direction quickly develops vortical structures on this complex surface.

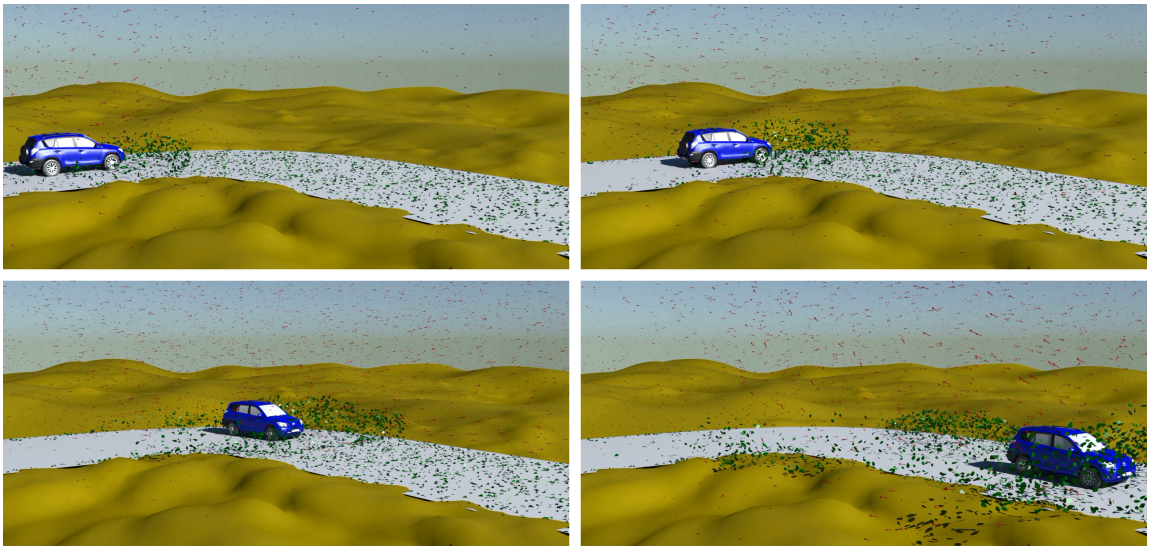

Figure 4.10: **Immersed moving objects.** As the car makes a right turn, the low frequency motion of the air displaced around it lifts the dead leaves. The velocity field above is visualized through arrows.

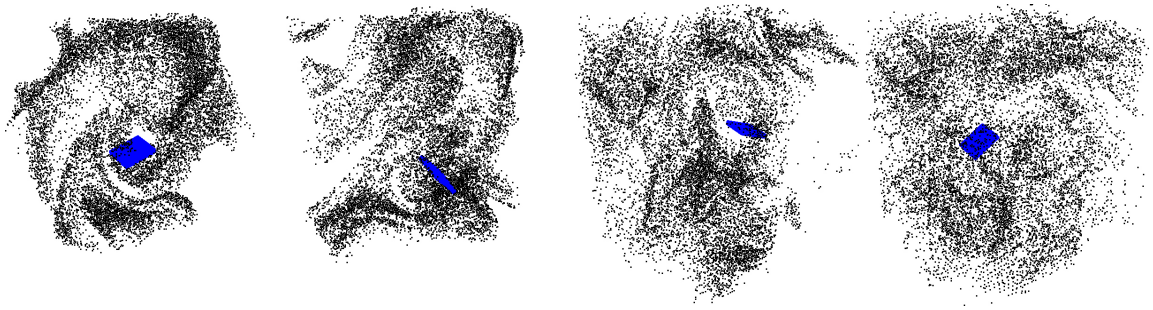

Figure 4.11: **Interactivity**: We can also use the analytic expressions for Ψ_k and $C_{k,ij}$ in a periodic 3D domain to handle a large number of modes without even calculating the spectral bases. The explicit update rule exhibits no artificial damping of the energy as expected, but offers realtime flows.

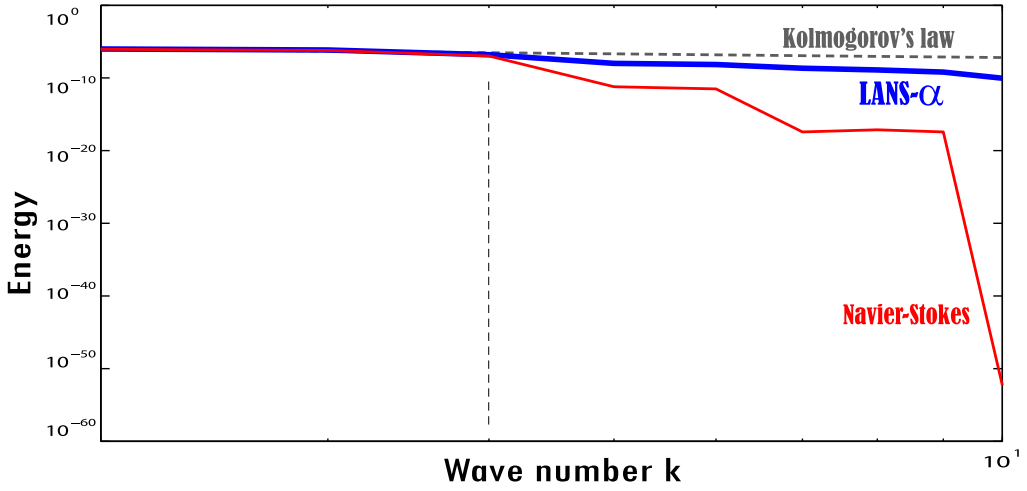

Figure 4.12: **Spectral energy distribution**: With forcing terms keeping the low wave number amplitudes fixed [13], our 3D reduced model applied to the LANS-α model of turbulence produces an average spectral energy distribution (blue) much closer to the expected Kolmogorov distribution (black) than with the usual Navier-Stokes equations (red).

4.3 Discussion

We have introduced a variational integrator for fluid simulation in reduced coordinates. By restricting the variations in Hamilton's principle to a low-dimensional space spanned by low-frequency divergence-free velocity fields, our method exhibits the properties of variational integrators in capturing the qualitatively correct behavior of ideal incompressible fluids (such as Kelvin's circulation and energy preservation) while greatly reducing the computational cost. The resulting method is versatile, energy-preserving, and computationally efficient.

While any application targeting low computational complexity of fluid simulation will benefit from this reduced space approach, one possible limitation of our method in some cases is the lack of spatial locality in our bases. For this reason, it should be noted that our integrator is not restricted to a particular set of basis functions. Future work could explore the use of wavelets for vorticity to offer optimal sparsity in the structural coefficients. Another possible future extension is to incorporate free surface boundary conditions through our modified Hodge star, combined with wavelet representations for the volume of fluid per cell.

4.3.1 Persistence of non-holonomic constraints

Among the initial aims for this method was the hope that the use of a reduced basis would allow us to forgo the non-holonomic constraint used in earlier methods [32]. This proved impossible, because the mapping that we chose between Laplacian basis functions for the vector field and the Lie algebra $\mathfrak{g}$ was not surjective, and thus there were elements of the Lie algebra that did not correspond to any vector field.

A related problem is that the Lie bracket is not preserved by this form of discretization. Instead, taking Lie brackets in the continuous space leads us to higher-frequency functions, which would require more basis functions if we wished to resolve them. This tendency of the Lie bracket to lead us outside our reduced basis is related to the cascade of energy to higher-frequency motions, as famously recognized by Kolmogorov [20]. In theory, if we could find a set of finite-dimensional Lie algebra discretizations which *did* preserve the Lie bracket, this problem could be solved. In particular, this would require a Lie algebra isomorphic to a suitable finite-dimensional Lie subalgebra of χ_{vol}. I have not found in the literature any complete classification of finite-dimensional Lie subalgebras of χ_{vol}. In practice, however, if so perfect a closure of the energy cascade existed, it

would probably have been noted before now.

We therefore see that the problem of creating a symplectic method for incompressible fluids is, perhaps surprisingly, intimately related to the problem of closing the turbulent energy cascade.

Chapter 5

Spectral Discrete Exterior Calculus

We have seen that we can create a Lie group based method for incompressible fluids using Laplacian eigenfunctions, and this method may be easily extended to other sets of orthogonal eigenfunctions. We have also seen that the Lie group method for fluids has limitations, because it requires the non-holonomic constraint, which breaks the symplecticity of the method. One possible avenue for future work is to try to find a structured integrator for fluids using a field theoretic perspective rather than the Lie group perspective. For this purpose, it may be useful to build on work I did with Dzhelil Rufat, Patrick Mullen, and my adviser on a spectral version of discrete exterior calculus.

Discrete exterior calculus [11] represents k-forms by integrating them on k-dimensional elements of a mesh. Thus, 0-forms are represented by their values at points, 1-forms are represented by quantities integrated along edges, 2-forms on faces, and 3-forms on volumes. One particularly nice aspect of this style of discretization is that we can use Stokes' theorem to define a discrete exterior derivative which introduces *no error whatsoever* – see Section 5.2.1. Along with this discrete exterior derivative, we define operators of discrete exterior calculus such as the wedge product and the hodge star. These operators *do* introduce error. In particular, existing versions are sometimes of very low order. For example, the diagonal hodge star is zeroth order [11].

To obtain a discrete operators of higher accuracy, we introduce a *spectral* discrete exterior calculus, obtaining faster-than-polynomial convergence, and using the Fast Fourier Transform for calculational speed. For this, we will consider reconstruction maps (Section 5.1.1) which interpolate, or *histopolate*, discrete k-forms using either the Fourier or Chebyshev basis. Histopolation allows

us to take into account the fact that we are dealing with integrated quantities, which is an essential element in preserving the perfect accuracy of the discrete exterior derivative. We may then apply the corresponding continuous operator to the interpolated or histopolated function, before discretizing again using a *reduction* operator. In this way we are able to create a spectrally-accurate discrete exterior calculus.

5.1 Basic Spectral Tools

Before delving into the design of spectrally-accurate discrete operators, we must define a series of basic tools and conventions which will be particularly useful for our task. We start by introducing the general notions of reduction and reconstruction maps. The reduction map will allow us to represent a continuous form in terms of a finite number of basis functions. The reconstruction map will then allow us to reconstruct the continuous form again from this information.

The reduction and reconstruction maps will first be explained in terms of general basis functions. We will then proceed to define the spectral basis functions on periodic regular grids which our specific reduction and reconstruction maps will use, before moving on to the case of Chebyshev grids, which allow us to consider regions with boundaries.

5.1.1 Reduction and Reconstruction Maps

The reduction and reconstruction maps provide a way to go back and forth between continuous forms and their discrete realizations.

Reduction. The reduction map (also called the de Rham map) is a linear operator $\mathcal{P}$ that projects a continuous form to its discrete realization on the grid through integration over mesh elements:

$$\begin{aligned} \mathcal{P}: \quad & \Lambda^k \to \bar{\Lambda}^k \\ & \omega^k \to \bar{\omega}^k \quad \text{with} \quad \bar{\omega}_i^k = (\mathcal{P}\omega^k)_i \equiv \int_{\sigma_i^k} \omega^k. \end{aligned}$$

We will denote by $\widetilde{\mathcal{P}}$ the analogous operator mapping continuous forms to their *dual* discrete counterparts in a similar fashion:

$$\begin{aligned} \widetilde{\mathcal{P}}: \quad & \Lambda^k \to \widetilde{\Lambda}^k \\ & \omega^k \to \widetilde{\omega}^k, \quad \text{with} \quad \widetilde{\omega}^k_i = (\widetilde{\mathcal{P}}\omega^k)_i \equiv \int_{\widetilde{\sigma}^k_i} \omega^k. \end{aligned}$$

Note that this definition of reduction extends the notion of *point sampling*: while the reduction of a 0-form is found by simply point-sampling its value at each vertex of the grid, the reduction of a general k-form is its evaluation (i.e., integral) on all the k-dimensional elements (vertices for $k=0$, edges for $k=1$, faces for $k=2$, etc) of the grid.

Reconstruction. Conversely, the reconstruction map ($\mathcal{R}$) is a map which reconstructs a continuous k-form from its discrete realization by interpolation for $k=0$, and by histopolation otherwise:

$$\begin{aligned} \mathcal{R}: \quad & \bar{\Lambda}^k \to \Lambda^k \\ & \bar{\omega}^k \to \omega^k. \end{aligned}$$

We will denote by $\widetilde{\mathcal{R}}$ the analogous operator mapping *dual* discrete forms to continuous forms in a similar fashion. Note that we will sometimes omit the dual sign $\widetilde{\cdot}$ for clarity, as which reconstruction operator is meant is unambiguously implied by the (primal or dual) nature of the discrete form it is applied on.

If one has a set of basis functions $\{\phi^k_0(x), \phi^k_1(x), \dots\}$ for k-forms that satisfy the property

$$\forall i, j, \quad \int_{\sigma^k_i} \phi^k_j = \delta_{ij}, \tag{5.1}$$

then, $\mathcal{R}$ can trivially be defined as

$$\omega^k = \mathcal{R}\bar{\omega}^k \equiv \sum_i \bar{\omega}^k_i \phi^k_i.$$

One can readily verifies that this reduction map is a left inverse of the reconstruction map:

$$\mathcal{P}\mathcal{R} = \mathrm{Id}.$$

However, the converse is not true; the reconstruction map is only approximately the right inverse of the reduction map, with equality in the limit when the mesh element size h approaches 0:

$$\|\omega - \mathcal{R}\mathcal{P}\omega\| \xrightarrow[h\to 0]{} 0,$$

with a rate of convergence determined by the chosen norm on forms and the degree of the basis functions. While Whitney first introduced a one-sided inverse of the de Rham map using what amounts to piecewise linear basis functions [44], we will instead use global basis functions satisfying Eq. (5.1) to provide spectrally-accurate reconstructions.

5.1.2 1D Periodic Interpolator & Histopolator Functions

To build our spectral wedge and Hodge star operators to work on discrete forms of arbitrary degree, we will need not only spectral interpolating basis functions, but also spectral *histopolating* basis functions, i.e., basis functions which integrate to one over assigned intervals [7]. To this end, we consider a one-dimensional periodic domain of width 2π with N *regularly-spaced nodes*, and define over this canonical domain two scalar functions α_N and β_N—that we will respectively call *interpolator* and *histopolator*—as follows:

$$\boxed{\alpha_N(x) = \frac{1}{N}\begin{cases} \cot\frac{x}{2}\,\sin\frac{Nx}{2} & \text{if N even,} \\ \csc\frac{x}{2}\,\sin\frac{Nx}{2} & \text{if N odd,} \end{cases}} \tag{5.2}$$

and

$$\boxed{\beta_N(x) = \begin{cases} \frac{1}{2\pi} - \frac{1}{4}\cos\frac{Nx}{2} + \frac{1}{N}\sum_{n=1}^{N/2}\frac{n\cos nx}{\sin\frac{n\pi}{N}} & \text{if N even,} \\ \frac{1}{2\pi} + \frac{1}{N}\sum_{n=1}^{(N-1)/2}\frac{n\cos nx}{\sin\frac{n\pi}{N}} & \text{if N odd.} \end{cases}} \tag{5.3}$$

For convenience, let us denote any translation of the above functions using the notation below

$$\alpha_{N,n}(x) = \alpha_N(x - hn), \quad \text{and} \quad \beta_{N,n}(x) = \beta_N(x - hn)$$

where $h = 2\pi/N$ is the mesh's edge element size and the nodes are enumerated by $x_n = nh$. For any given number N of nodes, these two functions satisfy the following important properties mentioned in Eq. (5.1) (where δ_{mn} refers to the Kronecker delta):

$$\alpha_{N,n}(x_m) = \delta_{mn}, \quad \text{and} \quad \int_{x_m - \frac{h}{2}}^{x_m + \frac{h}{2}} \beta_{N,n}(x)\, dx = \delta_{mn}.$$

In other words, α_N provides a smooth interpolation of a discrete function with 1 at node x_0 and 0 at every other nodes, while β_N integrates to 1 over the interval $[x_0 - \frac{h}{2}, x_0 + \frac{h}{2}]$, and to 0 over all other intervals (see Figure 5.1). Notice that these functions are the only Fourier series with N sinusoidal components satisfying the point-wise (resp., interval-wise) constraints. Note also that α_N and β_N are related to each other:

$$\int_{x - \frac{h}{2}}^{x + \frac{h}{2}} \beta'_N(\xi) d\xi = \alpha_N(x), \quad \text{or equivalently,} \quad \beta_N(x + \frac{h}{2}) - \beta_N(x - \frac{h}{2}) = \alpha'_N(x).$$

Next, we show that these two functions provide building blocks to construct basis functions for arbitrary forms on regular, periodic grids—from which we will derive basis functions for bounded domains via pushforward.

5.1.3 Spectral Basis Functions on Regular Grids

Basis functions for periodic grids in arbitrary dimensions can be easily built through tensor products of (translated) interpolators α_N and histopolators β_N, where the number of β_N used in the tensor product is equal to the degree of the form. To make this point clear, we will now introduce our notation for 1-, 2-, and 3-dimensional basis functions, with the extension to higher dimensions being straightforward. We will denote the basis functions as $\phi^{p[\text{comp}]}_{[\text{ind}]}$ with superscripts used to indicate the degree p, followed by the component of the form when appropriate (e.g., x, xy, etc), and subscripts used for grid indices; for instance, $\phi^{1y}_{mnk}(.)$ is the function of $\mathbb{R}^3$ representing the $\mathbf{d}y$-component of the 1-form basis, located on the edge parallel to the y axis indexed by (m, n, k).

See Figure 5.2 for an illustration of the 2D notations.

Primal and dual nodes. On a 1D regular grid the N primal nodes are equally spaced, and the dual nodes are placed in the middle between two adjacent primal nodes. Consequently, the primal node are: $x_n = 2\pi n/N$, and the dual nodes are at: $\widetilde{x}_n = 2\pi(n + \frac{1}{2})/N$.

Basis Functions for Forms in 1D. Because of their interpolation (resp., histopolation) property, translated versions of the functions α_N and β_N can directly be used as basis functions for 0- and 1-forms respectively as follows:

$$\phi^0_{N,n}(x) = \alpha_{N,n}(x), \qquad \text{and} \qquad \phi^1_{N,n}(x) = \beta_{N,n+\frac{1}{2}}(x)\,\mathbf{d}x.$$

where $n \in \{0, 1, \ldots, N-1\}$. Indeed, a discrete 0-form $\bar{\omega}^0$ (resp., a discrete 1-form $\bar{\omega}^1$) can be reconstructed as a smooth form through $\omega^0 = \sum_i \bar{\omega}^0_i \phi^0_i$ (resp., $\omega^1 = \sum_i \bar{\omega}^1_i \phi^1_i$); the reconstructed form then satisfies $\mathcal{P}\omega^0 = \bar{\omega}^0$ (resp., $\mathcal{P}\omega^1 = \bar{\omega}^1$). Similarly, dual basis functions are easily designed as well through:

$$\tilde{\phi}^0_{N,n}(x) = \alpha_{N,n+\frac{1}{2}}(x) \qquad \text{and} \qquad \tilde{\phi}^1_{N,n}(x) = \beta_{N,n}(x)\,\mathbf{d}x.$$

Basis Functions for Forms in 2D. In two dimensions, tensor products of the one-dimensional bases provide basis functions for 0-, 1-, and 2-forms on a regular $M \times N$ grid. These functions are expressed as follows:

$$\begin{aligned}
\phi^0_{MN,mn}(x,y) &= \phi^0_{M,m}(x) \wedge \phi^0_{N,n}(y) \\
\phi^{1x}_{MN,mn}(x,y) &= \phi^1_{M,m}(x) \wedge \phi^0_{N,n}(y) \\
\phi^{1y}_{MN,mn}(x,y) &= \phi^0_{M,m}(x) \wedge \phi^1_{N,n}(y) \\
\phi^2_{MN,mn}(x,y) &= \phi^1_{M,m}(x) \wedge \phi^1_{N,n}(y)
\end{aligned}$$

One can easily check that these functions are 1 on their associated degree of freedom and 0 on all others (vertices for 0-forms, edges for 1-forms, and faces for 2-forms), thus offering a proper set

of bases for smooth reconstructions of discrete forms. Formulas for the dual basis functions are strictly analogous, where $\widetilde{\phi}$ is used in lieu of ϕ.

Basis Functions for Forms in 3D. The same construction can be used in three or higher dimensions. For completeness, we describe the three-dimensional basis functions for primal forms:

$$\begin{aligned}
\phi^{0}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{0}_{K,k}(z) \\
\phi^{1x}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{0}_{K,k}(z) \\
\phi^{1y}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{0}_{K,k}(z) \\
\phi^{1z}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{1}_{K,k}(z) \\
\phi^{2xy}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{0}_{K,k}(z) \\
\phi^{2xz}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{1}_{K,k}(z) \\
\phi^{2yz}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{1}_{K,k}(z) \\
\phi^{3}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{1}_{K,k}(z)
\end{aligned}$$

Here again, it is easy to check that these functions provide smooth reconstructions from values of vertices, edges, faces, or volumes on a regular grid indexed by m, n, and k, in a periodic domain. Formulas for the dual basis functions are strictly analogous, where $\widetilde{\phi}$ is used in lieu of ϕ.

5.1.4 Chebyshev Grids over Bounded Domains

For bounded domains, a popular choice of spatial discretiation in spectral methods is the use of Chebyshev computational grids [8]. It is well known that Chebyshev polynomials can be derived as the pullback of Fourier basis functions from a circle onto its diameter; see Fig. 5.3. We can, in fact, perform the same pullback of our regular-grid basis functions of forms to obtain new spectral bases of forms applicable for Chebyshev grids over non-periodic domains.

Denote by φ the map from the interval between -1 and 1 on the real line (with Cartesian

coordinate x) to the unit semicircle (with polar angle θ, see Figure 5.3):

$$\begin{aligned} \varphi: \quad [-1,1] \subset \mathbb{R} &\to [0,\pi] \subset \mathcal{S}^1 \\ x &\mapsto \theta = \arccos(-x) \end{aligned}$$

Primal and dual nodes. Let the 1D grid consist of N points (including the boundaries) in the interval $[-1,1]$. The corresponding unit circle will then have $2N-2$ points (see Figure 5.3). The primal nodes $\{x_n\}$ ($n = 0..N-1$) and dual nodes $\{\widetilde{x}_n\}$ ($n = 0..N-2$) of the Chebyshev grid are thus

$$x_n = -\cos\frac{n\pi}{N-1} \quad \text{and} \quad \widetilde{x}_n = -\cos\frac{(n+\frac{1}{2})\pi}{N-1}.$$

Basis functions. To design our basis functions, we first define functions on the semi-circle ($\theta \in [0,\pi]$) by mirroring/antimirroring the regular basis functions ϕ on the whole circle (see [39] for the usual case of primal 0-forms), in order to satisfy the interpolation/histopolation properties (Eq. (5.1)) on the semi circle. The resulting functions κ are expressed as:

$$\kappa^0_{N,n} = \begin{cases} \phi^0_{2N-2,n}, & n = 0 \text{ or } n = N-1 \text{ (endpoints)} \\ \phi^0_{2N-2,n} + \phi^0_{2N-2,2N-2-n}, & n \in \{1,\dots,N-2\} \text{ (midpoints)} \end{cases}$$

for primal 0-forms,

$$\kappa^1_{N,n} = \phi^1_{2N-2,n} - \phi^1_{2N-2,2N-3-n}, \quad n \in \{0,\dots,N-2\}$$

for primal 1-forms,

$$\widetilde{\kappa}^0_{N,n} = \widetilde{\phi}^0_{2N-2,n} + \widetilde{\phi}^0_{2N-2,2N-3-n} \quad n \in \{0,\dots,N-2\}$$

for dual 0-forms, and

$$\widetilde{\kappa}^1_{N,n}(\theta) = \begin{cases} \delta_N(\theta) & n = 0 \\ \left[\widetilde{\phi}^0_{2N-2,n} - \widetilde{\phi}^0_{2N-2,2N-2-n} - \rho_{N,n}\right](\theta) & n \in \{1,\dots,N-2\} \\ \delta_N(\pi-\theta) & n = N-1 \end{cases}$$

for dual 1-forms, where

$$\delta_N(\theta) = \left((N-1)^2 \alpha_{2N-2,0}(\theta) + \frac{1}{2}\cos\left((N-1)\theta\right)\right)\sin\theta\,\mathbf{d}\theta,$$

$$\text{and} \quad \rho_{N,n}(\theta) = 2\left(\gamma_{2N-2,n}\delta_N(\theta) + \gamma_{2N-2,N-n-1}\delta_N(\pi-\theta)\right)\mathbf{d}\theta,$$

$$\text{with: } \gamma_{N,n} = \int_0^{\frac{\pi}{N}} \beta_{N,n}(\theta)\,d\theta$$

$$= \begin{cases} \frac{1-(-1)^k}{2N} + \frac{1}{N}\sum_{n=1}^{N/2} \frac{\sin(2kn\pi/N) - \sin((2k-1)n\pi/N)}{\sin\frac{n\pi}{N}} & \text{if N even,} \\ \frac{1}{2N} + \frac{1}{N}\sum_{n=1}^{(N-1)/2} \frac{\sin(2kn\pi/N) - \sin((2k-1)n\pi/N)}{\sin\frac{n\pi}{N}} & \text{if N odd} \end{cases}$$

The function δ is used to deal with the special case of the two boundary dual (half-)edges, and the function ρ adds contributions to intermediate basis functions so that they integrate to zero at both boundary dual edges. Finally, we pull back the functions κ by φ to obtain the form basis functions ψ on the Chebyshev 1D grid ($\psi = \varphi^*\kappa$):

$$\begin{aligned} \psi^0_{N,n}(x) &= \kappa^0_{N,n}(\arccos(-x)) \\ \psi^1_{N,n}(x) &= \kappa^1_{N,n}(\arccos(-x))\frac{\mathbf{d}x}{\sqrt{1-x^2}} \\ \tilde{\psi}^0_{N,n}(x) &= \tilde{\kappa}^0_{N,n}(\arccos(-x)) \\ \tilde{\psi}^1_{N,n}(x) &= \tilde{\kappa}^1_{N,n}(\arccos(-x))\frac{\mathbf{d}x}{\sqrt{1-x^2}} \end{aligned}$$

One can easily check that the functions above satisfy the property in Eq. (5.1) required for basis functions, as the functions κ were designed to satisfy these properties, and the pullback φ^* commutes with integration.

Note that the basis functions for primal zero-forms turns out (unsurprisingly) to be the Lagrange polynomials of order N, i.e.,

$$\psi^0_{N,n}(x) \quad = \quad \prod_{\substack{m=0 \\ m \neq n}}^{N-1} \frac{x - x_m}{x_n - x_m}$$

where $\{x_n\}_{n=0,\ldots,N-1}$ represent the coordinates of the primal points. All other basis functions of forms are *also* polynomials for any choice of N. Examples for the primal and dual basis functions for 0- and 1-forms are provided in Figures 5.4 and 5.5 for $N = 7$.

Finally, basis functions in higher dimensions for arbitrary forms are derived using tensor products of these two primal and two dual $1D$ basis functions, just as we explained in Section 5.1.3.

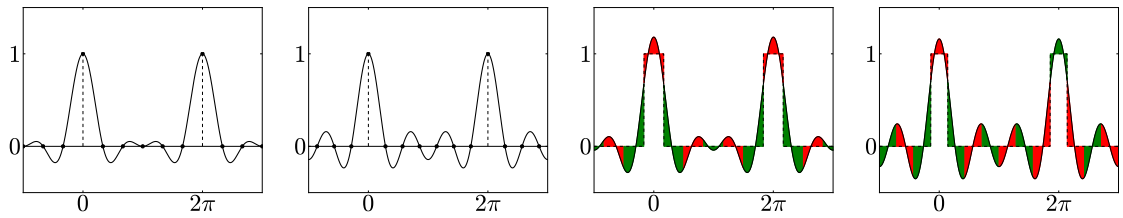

(a) $\alpha_6(x)$ (b) $\alpha_7(x)$ (c) $h\beta_6(x)$ (d) $h\beta_7(x)$

Figure 5.1: Examples of periodic interpolators for $N = 6$ (a) and 7 (b), and corresponding periodic histopolators (c) and (d), scaled by $h = 2\pi/N$ for clarity. While the interpolator α_N satisfies $\alpha_N(nh \mod 2\pi) = \delta_{0n}\ \forall n$, the histopolator β_N integrates to 1 over the dual cell straddling $x = 0$, and to 0 over other dual cells in the range $[0, 2\pi]$. Note that the alternating red and green colors are used to mark out dual cells, and to illustrate that the integral of β_N over each of these dual cells sums to zero or one.

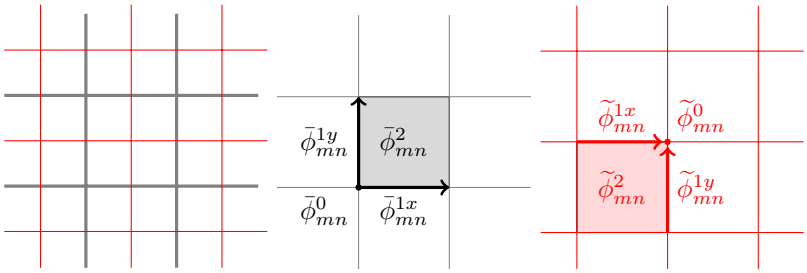

(a) Primal/Dual grid (b) Primal grid (c) Dual grid

Figure 5.2: Illustration of a regular 2D grid (in black) along with its dual (in red), along with our mesh orientation and index convention.

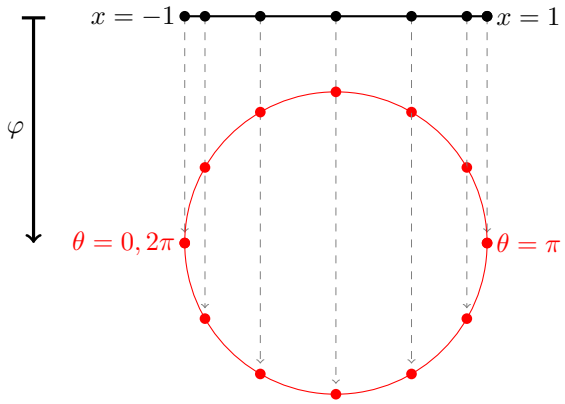

Figure 5.3: The map φ mapping the canonical interval $[-1, 1]$ to the bottom unit hemicircle ($0 \leq \theta \leq \pi$).

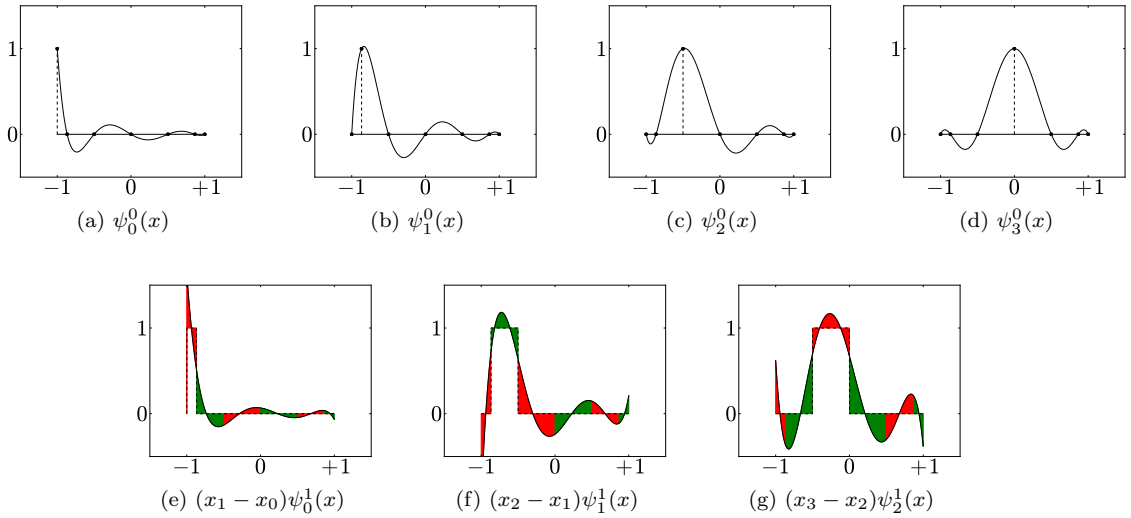

(a) $\psi_0^0(x)$ (b) $\psi_1^0(x)$ (c) $\psi_2^0(x)$ (d) $\psi_3^0(x)$

(e) $(x_1 - x_0)\psi_0^1(x)$ (f) $(x_2 - x_1)\psi_1^1(x)$ (g) $(x_3 - x_2)\psi_2^1(x)$

Figure 5.4: Chebyshev primal basis functions for a grid with $N = 7$. We normalize the one-form basis functions by $x_n - x_{n-1}$ to have approximately the same scale in our visualizations.

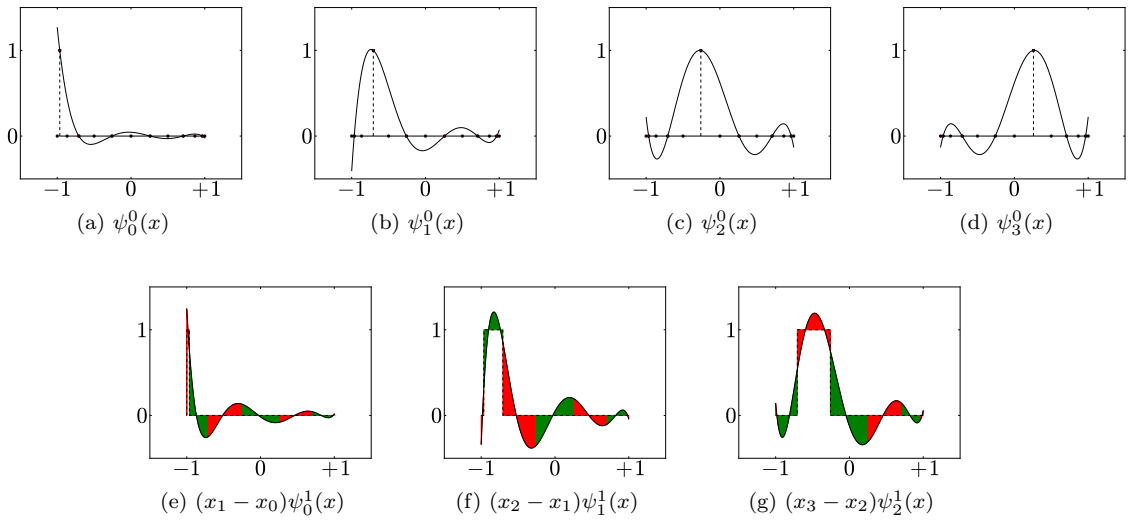

(a) $\psi_0^0(x)$ (b) $\psi_1^0(x)$ (c) $\psi_2^0(x)$ (d) $\psi_3^0(x)$

(e) $(x_1 - x_0)\psi_0^1(x)$ (f) $(x_2 - x_1)\psi_1^1(x)$ (g) $(x_3 - x_2)\psi_2^1(x)$

Figure 5.5: Chebyshev dual basis functions for a grid with $N = 7$

5.2 Spectrally Accurate Discrete Operators

Equipped with these basis functions on regular and Chebyshev grids, we can now derive the discrete, spectrally accurate version of the exterior derivative $\mathbf{d}$, the wedge product $\wedge$, and the Hodge star $\star$. These operators may be quickly calculated using the Fast Fourier Transform. More thorough implementation details are given in [33]; here we focus on the theoretical construction.

5.2.1 Discrete Exterior Derivative D

The most common discrete realization of the exterior derivative based on algebraic topology [29] using chains and cochains is a linear operator $\mathbf{D}$:

$$\mathbf{D}: \quad \bar{\Lambda}^k \to \bar{\Lambda}^{k+1}.$$

which is the (signed) incidence matrix between $(k+1)$ elements and k elements of the grid, with a sign determined by the relative orientation of the elements. In other words,

$$\boxed{\mathbf{D} = \boldsymbol{\partial}^t,}$$

where $\boldsymbol{\partial}$ refers to the *boundary operator* acting on chains (see [29, 11]). Note that the operator $\mathbf{D}$ is thus implemented via a sparse matrix whose non-zeros elements are only $+1$ and -1 values to indicate incidence between mesh cells. It also satisfies $\mathbf{D}^2 = 0$ like its continuous equivalent (since the boundary of a boundary is always the empty set), and that a *discrete Stokes' theorem* is also enforced on chains since the very definition of $\mathbf{D}$ is tantamount to enforcing

$$\int_\sigma \mathbf{d}\omega = \int_{\partial\sigma} \omega$$

for every mesh element σ. This discrete realization of the exterior derivative is exact, in the sense that the operator $\mathbf{D}$ commutes with the reduction operator:

$$\mathbf{d}\mathcal{P} = \mathcal{P}\mathbf{D}. \tag{5.4}$$

This implies that the following diagram fully commutes:

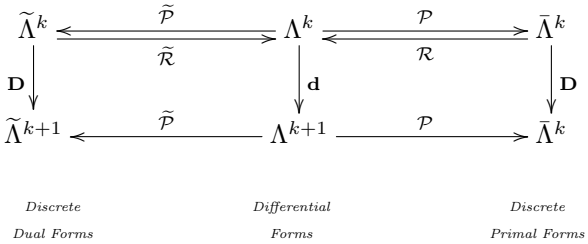

Therefore, as mentioned in the introduction to this chapter, the classic discrete exterior derivative used in mimetic methods, DEC, or finite-dimensional exterior calculus needs no special treatment to obtain spectral accuracy. However, its use for spectral computations represents a clear departure from the conventional spectral methods [7], since all operators of classical vector calculus (divergence, gradient, and curl) can be achieved exactly via $\mathbf{D}$ [15]. In our framework, we also use the exterior derivative on dual forms:

$$\widetilde{\mathbf{D}}: \quad \widetilde{\Lambda}^k \to \widetilde{\Lambda}^{k+1}.$$

Its implementation and properties are no different from its primal version, since the adjacency of the dual mesh is directly derived from the adjacency on the primal:

$$\boxed{\widetilde{\mathbf{D}} = \widetilde{\boldsymbol{\partial}}^t}.$$

We now turn to the definition of a discrete wedge product and discrete Hodge star operator.

5.2.2 Discrete Wedge Product W

Various discrete definitions of the wedge product have been proposed, sometimes mixing primal and dual elements [19]. We instead follow Bochev and Hyman's treatment [6]: by utilizing the de Rham and reconstruction maps, one can define a discrete wedge product $\mathbf{W}(\bar{\alpha}, \bar{\beta})$ in a general way that applies to any pair of discrete (primal or dual) forms $\bar{\alpha}$ and $\bar{\beta}$ through:

$$\boxed{\mathbf{W}(\bar{\alpha}, \bar{\beta}) = \mathcal{P}(\mathcal{R}\bar{\alpha} \wedge \mathcal{R}\bar{\beta})} \tag{5.5}$$

That is, we first reconstruct the two discrete forms with our trigonometric or Chebyshev basis functions, we then apply the wedge product to the two continuous forms that we have created, and

finally, integrate (or, in the case of 0-forms, point-sample) the result to get the final discrete form $\mathbf{W}(\bar{\alpha}, \bar{\beta})$. The resulting operator satisfies a discrete Leibniz rule:

$$\mathbf{DW}(\bar{\alpha}, \bar{\beta}) = \mathbf{W}(\mathbf{D}(\bar{\alpha}), \bar{\beta}) + (-1)^k \mathbf{W}(\bar{\alpha}, \mathbf{D}(\bar{\beta})),$$

where $\bar{\alpha}$ is a discrete k-form, and $\bar{\beta}$ is an arbitrary discrete form. Indeed, because the exterior derivative commutes with the reduction and reconstruction maps, this discrete Leibniz rule can be derived directly from the continuous Leibniz rule. Our spectrally-accurate definition of the discrete wedge product is bilinear and anticommutative, but it is *not* associative. It is, however, associative in the limit as the mesh size approaches zero. That is, the associator $\mathbf{W}(\bar{\alpha}, \mathbf{W}(\bar{\beta}, \bar{\gamma})) - \mathbf{W}(\mathbf{W}(\bar{\alpha}, \bar{\beta}), \bar{\gamma})$ (i.e., the measure of nonassociativity) will approach zero exponentially fast as the step size is reduced. The difficulty of creating an associative discrete wedge product has been previously noted by, e.g., Kotiuga [21], who refers to it as the "commutative cochain problem" and discusses some deeper topological reasons behind it. The spectral accuracy of our wedge operator will mitigate this lack of associativity exponentially fast under mesh refinement.

Note that the discrete wedge product of a (primal or dual) p-form and a (primal or dual) q-form can also be seen as a three tensor

$$\mathbf{W}^{p,q}_{ijk} = \int_{\sigma_i^{p+k}} \phi^p_j \wedge \phi^q_k,$$

since we can write the value of the wedge product on element σ_i^{p+q} as :

$$\left[\mathbf{W}(\bar{\alpha}, \bar{\beta})\right]_i = \sum_{j,k} \mathbf{W}_{ijk} \bar{\alpha}_j \bar{\beta}_k$$

5.2.3 Discrete Hodge Star H

As explained earlier, our discrete Hodge star **H** exploits the notion of mesh duality in that the discrete Hodge star of a primal k-form is a dual $(d-k)$-form, and vice-versa. The discrete Hodge star for a discrete form is realized conceptually by first reconstructing the continuous form with our primal (resp., dual) spectral bases, applying the continuous Hodge star to this form, and then

projecting this form back to the dual (resp., primal) grid. In our notation, this can be written as

$$\mathbf{H}^k = \widetilde{\mathcal{P}} \star^k \mathcal{R}, \tag{5.6}$$

and is easily understood from the following diagram:

$$\begin{array}{ccc} \Lambda^k & \underset{\mathcal{R}}{\overset{\mathcal{P}}{\rightleftarrows}} & \bar{\Lambda}^k \\ \downarrow{\scriptstyle \star^k} & & \vdots{\scriptstyle \mathbf{H}^k} \\ \Lambda^{n-k} & \xrightarrow{\widetilde{\mathcal{P}}} & \widetilde{\Lambda}^{n-k} \end{array}$$

The operator $\mathbf{H}$ can be built as a matrix with easily precomputed entries. For instance, for k-forms, this matrix contains the terms

$$\mathbf{H}^k_{ij} = \int_{\widetilde{\sigma}_i^{n-k}} \star \phi_j^k. \tag{5.7}$$

Similarly, the dual Hodge star operator is given by:

$$\widetilde{\mathbf{H}}^k_{ij} = \int_{\sigma_i^{n-k}} \star \widetilde{\phi}_j^k \tag{5.8}$$

We require our discrete Hodge stars to satisfy the discrete equivalent of the continuous requirement that $\star^k \star^{d-k} = (-1)^{k(d-k)\cdot}$ That is, we need:

$$\mathbf{H}^k \widetilde{\mathbf{H}}^{n-k} = \widetilde{\mathbf{H}}^{n-k} \mathbf{H}^k = (-1)^{k(n-k)} \mathbf{Id}. \tag{5.9}$$

This imposes a constraint on dual basis functions. Indeed, for Eq. (5.9) to hold true, they must be a linear combination of the primal basis functions:

$$\tilde{\phi}_i^k = \sum_j [(\mathbf{H}^{n-k})^{-1}]_{ji} \star \phi_j^{n-k}. \tag{5.10}$$

Both the regular (ϕ) and Chebyshev (ψ) basis functions that we defined above satisfy this constraint.

Finally, our Hodge matrices are dense; however, each matrix is *circulant* for a regular grid on periodic domains (due to the invariance of the grid under translation), and *centrosymmetric* for a Chebyshev grid (due to mirror symmetry around its center). While circulant matrices are special cases of Toeplitz matrices, for which multiplications by vectors can be done $\mathcal{O}(N \log N)$, we will

show that *all* discrete Hodge stars for arbitrary dimensions can be efficiently implemented in Fourier space in $\mathcal{O}(N \log N)$ time complexity where N is the total number of spatial points.

5.3 Discussion

This chapter outlines the foundations for a spectral instance of Discrete Exterior Calculus. Using chains and cochains as discretization for differential forms, we have provided a discrete exterior derivative, a discrete Hodge star, and a discrete wedge product that are spectrally accurate. One can use the resulting calculus as a chain collocation method to solve differential equations. Figure 5.6 gives an example of the expected spectral convergence. Further analysis of convergence and implementaton may be found in [33].

While the three exterior operators we treated cover already a large range of differential operators (allowing a number of computational tasks including Hodge decomposition), contractions (exterior product) and Lie derivatives still need to be derived in a similar geometric manner. Note that defining the contraction operator would be enough, since one can then define the Lie derivative algebraically with Cartan's "magic formula". This may be useful in the construction of numerical methods for incompressible fluids in the future.

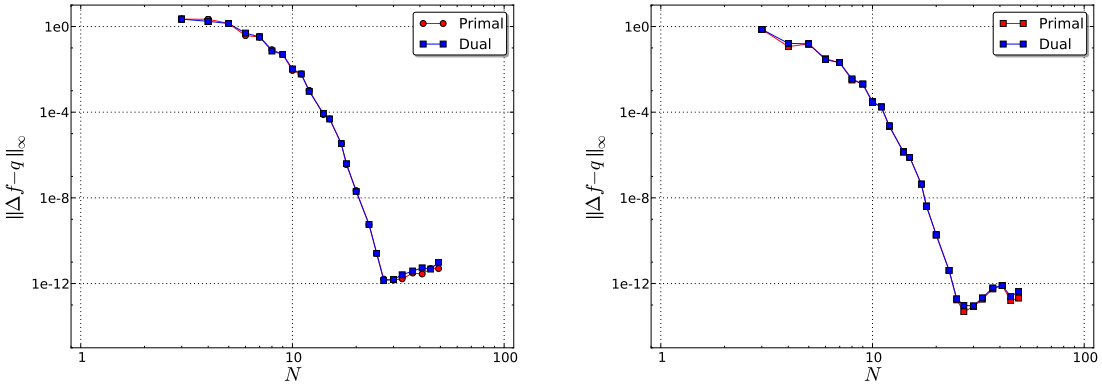

(a) Periodic 0-forms

(b) Periodic 1-forms

(c) Chebyshev 0-forms

(d) Chebyshev 1-forms

Figure 5.6: Convergence graphs for a 2D Poisson equation, using spectral discrete exterior calculus to calculate the Laplacian operator; (a) we solve $\Delta f = e^{\sin x}(\cos^2 x - \sin x) + e^{\sin y}(\cos^2 y - \sin y)$ on a periodic domain for either a primal, or dual 0-form f; (b) Now for $\Delta f = e^{\sin x}(\cos^2 x - \sin x)\mathbf{d}x + e^{\sin y}(\cos^2 y - \sin y)\mathbf{d}y$; (c) we solve $\Delta f = e^x + e^y$ on a Chebyshev grid, for either a primal 0-form with Dirichlet boundary conditions $f(x, y) = e^x + e^y$, or for a dual 0-form with Neumann boundary conditions $\nabla f(x) \cdot \mathbf{n} = (e^x \ e^y)^t \ \mathbf{n}$; (d) Now for $\Delta f = e^x \mathbf{d}x + e^y \mathbf{d}y$. All of our results exhibit spectral convergence (measured through the L^∞ error $\|\Delta f - q\|_\infty$), with the conventional plateau in accuracy for fine meshes.

Chapter 6

Conclusion

The structure-preserving simulation of incompressible fluids is a field that is still very much under development. This thesis has significantly broadened the tools at our disposal. We have seen how to introduce model-reduction using a given set of orthogonal eigenfunctions in a manner that allows us to spatially discretize while retaining the Lagrangian structure of the fluid. This significantly reduces the number of variables necessary in order to represent fluid motion in a way which respects the structure, and decreases the time necessary for computation. We have also seen how to discretize the Hamiltonian structure associated with the Lagrangian structure of the fluid. Together, these provide two very different expansions of our knowledge of how to geometrically simulate an incompressible fluid, which may be used in future work on conservative fluid-structure interactions, moving boundaries, and fast but visually sensible simulations of fluids for the purpose of computer graphics.

The toolbox of discretization techniques that were introduced in [32] has been extended and reapplied in new contexts to create a wide variety of numerical methods for fluids with useful geometric and computational properties. However, we have also seen that some of these techniques are perhaps best *not* used together. For example, if we wish to use a finite-dimensional Lie algebra to mimic the behavior of χ_{div}, we will need to use a non-holonomic constraint. This then means that the time-discrete Lagrangians of [43] will no longer yield symplectic integrators as expected. On the other hand, because the energy of the fluid is a quadratic invariant, we may use the midpoint rule to obtain energy-preserving integrators. Energy-preserving methods which are of higher order in time would be equally easy to obtain, since the property of preserving quadratic invariants is shared by all Gauss collocation methods, as detailed in [17].

An important development is the relationship between the non-holonomic constraint and the closure problem for the turbulence cascade. We have seen that the development of a symplectic method for incompressible fluids is unexpectedly difficult, or, to put it more optimistically, if we *could* find a truly symplectic fluids method, we might find ourselves with a better understanding of the turbulence cascade phenomenon as a result. A better understanding of *viscous* structure-preserving incompressible fluids would be important to this endeavor. The obvious way to include viscosity is with a discrete Laplace-deRham operator $\boldsymbol{\Delta} = \boldsymbol{\delta}\mathbf{d} + \mathbf{d}\boldsymbol{\delta}$. However, further work is needed in order to determine how best to preserve viscous energy behavior. We would hope to obtain similar results to those found for dissipative ordinary differential equations in [22].

These theoretical and practical developments have expanded our ability to efficiently simulate incompressible fluids in a manner that supplies good qualitative behavior. We now have methods that are significantly faster than earlier structure-preserving fluid simulators, along with a broader theoretical basis on which to build. This will enable the development of structure-preserving fluid integrators for new applications such as fluid-structure interactions and free surfaces in the future.

Bibliography

[1] Douglas N Arnold, Richard S Falk, and Ragnar Winther. Finite element exterior calculus, homological techniques, and applications. *Acta numerica*, 15:1–155, 2006.

[2] Vladimir Arnold. Sur la géométrie différentielle des groupes de lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. *Annales de l'institut Fourier*, 16(1):319–361, 1966.

[3] Christopher Batty, Florence Bertails, and Robert Bridson. A fast variational framework for accurate solid-fluid coupling. *ACM Trans. Graph.*, 26(3):100:1–100:8, July 2007.

[4] Nathan Bell and Anil N. Hirani. PyDEC: A Python library for Discrete Exterior Calculus, 2008. Google Code project at: `http://code.google.com/p/pydec/`.

[5] Anthony M Bloch. *Nonholonomic Mechanics and Control*. Springer, 2003.

[6] Pavel B. Bochev and James M. Hyman. Principles of mimetic discretizations of differential operators. In *Compatible Spatial Discretizations, volume 142 of The IMA Volumes in Mathematics and its Applications*, pages 89–119. Springer, 2006.

[7] Mick Bouman, Artur Palha, Jasper Kreeft, and Marc Gerritsma. A conservative spectral element method for curvilinear domains. In T. J. Barth et al., editor, *Spectral and High Order Methods for Partial Differential Equations*, volume 76 of *Lecture Notes in Computational Science and Engineering*, pages 111–119. Springer Berlin Heidelberg, 2011.

[8] C.G. Canuto, M.Y. Hussaini, A. Quarteroni, and Th.A. Zang. *Spectral Methods: Fundamentals in Single Domains*. Scientific Computation Series. Springer, 2006.

[9] Hernán Cendra, Jerrold E. Marsden, Sergey Pekarsky, Tudor, and S. Ratiu. Variational Principles for Lie–Poisson and Hamilton–Poincaré Equations, 2002.

[10] Tyler De Witt, Christian Lessig, and Eugene Fiume. Fluid Simulation Using Laplacian Eigenfunctions. *ACM Trans. Graph.*, 31(1):10:1–10:11, February 2012.

[11] M. Desbrun, E. Kanso, and Y. Tong. Discrete differential forms for computational modeling. In Bobenko and Alexander et al., editors, *Discrete Differential Geometry*, volume 38 of *Oberwolfach Seminars*, pages 287–324. Birkhäuser Basel, 2008.

[12] Sharif Elcott and Peter Schröder. Building your own DEC at home. In *Discrete Differential Geometry*, ACM SIGGRAPH Courses, pages 55–59, 2006.

[13] Ciprian Foias, DarrylD. Holm, and EdrissS. Titi. The three dimensional viscous Camassa–Holm equations, and their relation to the Navier–Stokes equations and turbulence theory. *Journal of Dynamics and Differential Equations*, 14(1):1–35, 2002.

[14] E. S. Gawlik, P. Mullen, D. Pavlov, J. E. Marsden, and M. Desbrun. Geometric, Variational Discretization of Continuum Theories. *Physica D*, 240:1724–1760, October 2011.

[15] Marc Gerritsma. Edge functions for spectral element methods. In Jan S. Hesthaven et al., editor, *Spectral and High Order Methods for Partial Differential Equations*, volume 76 of *Lecture Notes in Computational Science and Engineering*, pages 199–207. Springer Berlin Heidelberg, 2011.

[16] Mohit Gupta and Srinivasa G. Narasimhan. Legendre fluids: A unified framework for analytic reduced space modeling and rendering of participating media. In *Symposium on Computer Animation*, pages 17–25, 2007.

[17] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer series in computational mathematics. Springer, 2006.

[18] Francis H. Harlow and J. Eddie Welch. Numerical calculation of time?dependent viscous incompressible flow of fluid with free surface. *Physics of Fluids (1958-1988)*, 8(12), 1965.

[19] A. N. Hirani. *Discrete Exterior Calculus*. PhD thesis, California Institute of Technology, Pasadena, CA, 2003.

[20] A. Kolmogorov. The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds' Numbers. *Akademiia Nauk SSSR Doklady*, 30:301–305, 1941.

[21] P R Kotiuga. Theoretical limitations of discrete exterior calculus in the context of computational electromagnetics. *IEEE Transactions on Magnetics*, 44(6):1162–1165, 2008.

[22] A. Lew, J. E. Marsden, M. Ortiz, and M. West. Variational time integrators. *International Journal for Numerical Methods in Engineering*, 60(1):153–212, 2004.

[23] Beibei Liu, Gemma Mason, Julian Hodgson, Yiying Tong, and Mathieu Desbrun. Model-reduced variational fluid simulation. 2015. Preprint submitted to the 2015 SIGGRAPH Annual Conference.

[24] Benjamin Long and Erik Reinhard. Real-time fluid simulation using discrete sine/cosine transforms. In *Symposium on Interactive 3D Graphics and Games*, pages 99–106, 2009.

[25] Z. Ma and C. Rowley. Lie-Poisson integrators: a Hamiltonian, variational approach. *International Journal for Numerical Methods in Engineering*, 82:1609–1644, June 2010.

[26] J. E. Marsden and A. Weinstein. Coadjoint Orbits, Vortices, and Clebsch Variables for Incompressible Fluids. *Physica D: Nonlinear Phenomena*, 7(1-3):305–323, May 1983.

[27] J.E. Marsden and T.S. Ratiu. *Introduction to mechanics and symmetry: a basic exposition of classical mechanical systems*. Texts in applied mathematics. Springer, 1999.

[28] Patrick Mullen, Keenan Crane, Dmitry Pavlov, Yiying Tong, and Mathieu Desbrun. Energy-preserving integrators for fluid animation. *ACM Trans. Graph.*, 28(3):38:1–38:8, July 2009.

[29] James R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Menlo Park, CA, 1984.

[30] Yen Ting Ng, Chohong Min, and Frédéric Gibou. An efficient fluid-solid coupling algorithm for single-phase flows. *J. Comput. Phys.*, 228(23):8807–8829, December 2009.

[31] Steven A. Orszag. Numerical methods for the simulation of turbulence. *Physics of Fluids*, 12:II 250–257, 1969.

[32] D. Pavlov, P. Mullen, Y. Tong, E. Kanso, J. E. Marsden, and M. Desbrun. Structure-preserving discretizations of incompressible fluids. *Physica D*, 240:443–458, 2011.

[33] Dzhelil Rufat, Gemma Mason, Patrick Mullen, and Mathieu Desbrun. The chain collocation method: A spectrally accurate calculus of forms. *Journal of Computational Physics*, 257, Part B(0):1352 – 1372, 2014.

[34] Andrew Selle, Ronald Fedkiw, ByungMoon Kim, Yingjie Liu, and Jarek Rossignac. An Unconditionally Stable MacCormack Method. *Journal of Scientific Computing*, 35(2-3):350–371, 2008.

[35] Isadore Silberman. Planetary waves in the atmosphere. *J. Meteor.*, 11:27–34, 1954.

[36] Jos Stam. Stable fluids. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, pages 121–128, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[37] Matt Stanton, Yu Sheng, Martin Wicke, Federico Perazzi, Amos Yuen, Srinivasa Narasimhan, and Adrien Treuille. Non-polynomial Galerkin projection on deforming meshes. *ACM Trans. Graph.*, 32(4):86:1–86:14, July 2013.

[38] Ari Stern, Yiying Tong, Mathieu Desbrun, and Jerrold E. Marsden. Geometric computational electrodynamics with variational integrators and discrete differential forms. In Dong Eui Chang, Darryl D. Holm, George Patrick, and Tudor Ratiu, editors, *Geometry, Mechanics, and Dynamics*, volume 73 of *Fields Institute Communications*, pages 437–475. Springer New York, 2015.

[39] Lloyd N. Trefethen. *Spectral Methods in MATLAB*. SIAM, Philadelphia, 2000.

[40] Adrien Treuille, Andrew Lewis, and Zoran Popović. Model reduction for real-time fluids. *ACM Trans. Graph.*, 25(3):826–834, July 2006.

[41] Tomasz Michal Tyranowski. *Geometric integration applied to moving mesh methods and degenerate Lagrangians*. PhD thesis, California Institute of Technology, 2014.

[42] Karsten Urban. Wavelet bases for $H(\mathrm{div})$ and $H(\mathrm{curl})$. In *Wavelets in Numerical Simulation*, volume 22 of *Lecture Notes in Computational Science and Engineering*, pages 83–107. 2002.

[43] Matthew West. *Variational Integrators*. PhD thesis, California Institute of Technology, 2004.

[44] H. Whitney. *Geometric integration theory*. Princeton University Press, 1957.

[45] V.I. Yudovich. Non-stationary flow of an ideal incompressible liquid. *USSR Computational Mathematics and Mathematical Physics*, 3(6):1407–1456, 1963.

[46] Yongning Zhu. Animating sand as a fluid. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 2005.