Physica D 240 (2011) 443–458

Contents lists available at ScienceDirect

Physica D

journal homepage: www.elsevier.com/locate/physd

Structure-preserving discretization of incompressible fluids

D. Pavlov [a], P. Mullen [a], Y. Tong [b], E. Kanso [c], J.E. Marsden [a,*], M. Desbrun [a,**]

[a] Caltech, Pasadena, CA, United States
[b] Michigan State University, East Lansing, MI, United States
[c] University of Southern California, Los Angeles, CA, United States

ARTICLE INFO

Article history:
Received 19 December 2009
Received in revised form
3 August 2010
Accepted 28 October 2010
Available online 23 November 2010
Communicated by J. Bronski

Keywords:
Euler equations
Lie group variational integrators

ABSTRACT

The geometric nature of Euler fluids has been clearly identified and extensively studied over the years, culminating with Lagrangian and Hamiltonian descriptions of fluid dynamics where the configuration space is defined as the volume-preserving diffeomorphisms, and Kelvin's circulation theorem is viewed as a consequence of Noether's theorem associated with the particle relabeling symmetry of fluid mechanics. However computational approaches to fluid mechanics have been largely derived from a numerical–analytic point of view, and are rarely designed with structure preservation in mind, and often suffer from spurious numerical artifacts such as energy and circulation drift. In contrast, this paper geometrically derives discrete equations of motion for fluid dynamics from first principles in a purely Eulerian form. Our approach approximates the group of volume-preserving diffeomorphisms using a finite-dimensional Lie group, and associated discrete Euler equations are derived from a variational principle with non-holonomic constraints. The resulting discrete equations of motion yield a structure-preserving time integrator with good long-term energy behavior and for which an exact discrete Kelvin's circulation theorem holds.

1. Introduction

The geometric nature of Euler fluids has been extensively studied in the literature in works of Arnold, Ebin–Marsden and others; however the geometric-differential standpoint of these studies sharply contrasts with the numerical approaches traditionally used in Computational Fluid Dynamics (CFD). In particular, methods based on particles, vortex particles, staggered Eulerian grids, spectral elements, as well as hybrid Lagrangian–Eulerian formulations were not designed with structure preservation in mind — in fact, recent work pinpoints the loss of Lagrangian structures as a major numerical impediment of current CFD techniques [1]. In contrast, structure-preserving methods (the so-called geometric integrators) have recently become popular in the context of Lagrangian dynamics in solid mechanics. Based on discrete versions of Hamilton's principle and its variants, they have been shown to capture the dynamics of the mechanical system they discretize without traditional numerical artifacts such as loss of energy or momenta.

While the variational principles for incompressible fluid mechanics are best expressed in a Lagrangian formalism, computational efficiency often calls for an *Eulerian treatment* of fluid computations to avoid numerical issues inherent to deforming meshes. In order to circumvent these issues without giving up structure preservation, a new Eulerian formulation of discrete fluid mechanics is thus needed.

Guided by the variational integrators used in the Lagrangian setting, this paper introduces a discrete, structure-preserving theory for incompressible perfect fluids based on Hamilton–d'Alembert's principle. Such a discrete variational approach to fluid dynamics guarantees invariance under the particle relabeling group action and gives rise to a discrete form of Kelvin's circulation theorem. Due to their variational character, the resulting numerical schemes also exhibit good long-term energy behavior. In addition, the resulting schemes are not difficult to implement in practice (see Fig. 1), and we will derive particular instances of numerical update rules and provide numerical results. We will favor formalism over smoothness in the exposition of our approach in order to better elucidate the correspondences between continuous and discrete expressions.

1.1. Brief review of the continuous case

Let $M \subset \mathbb{R}^n$ be an arbitrary compact manifold, possibly with boundary (where n denotes the dimension of the domain, typically,

* Corresponding author.
** Corresponding author. Tel.: +1 626 395 6230; fax: +1 626 792 4257.
E-mail addresses: jmarsden@caltech.edu (J.E. Marsden), mathieu@caltech.edu (M. Desbrun).

0167-2789/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.physd.2010.10.012

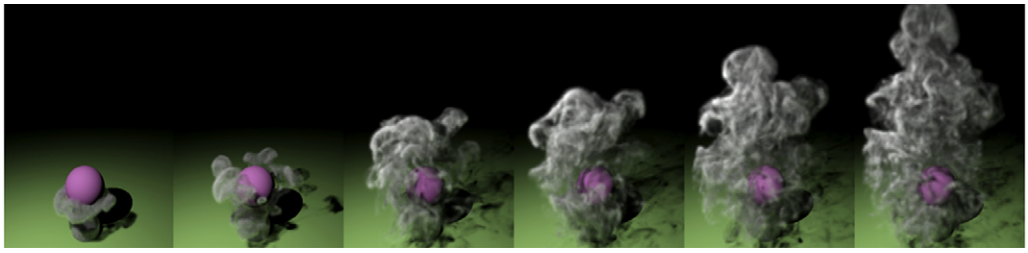

Fig. 1. Our geometric approach to discretizing the dynamics of incompressible fluids leads to discrete, structure-preserving, Lie group integrators. Here, six frames of an animation simulating heated smoke rising around a round obstacle in a closed box of incompressible fluid.

2 or 3), and SDiff(M) be the group of smooth volume-preserving diffeomorphisms on M. As was shown in [2], the motion of an ideal incompressible fluid in M may be described by a geodesic curve g_t in SDiff(M). That is, SDiff(M) serves as the configuration space — a particle located at a point $x_0 \in M$ at time $t = 0$ travels to $g_t(x_0)$ at time t. Being geodesics, the equations of motion naturally derive from Hamilton's stationary action principle:

$$\delta \int_0^1 L(g, \dot{g}) \mathrm{d}t = 0 \quad \text{where } L(g, \dot{g}) = \frac{1}{2} \int_M \|\dot{g}\|^2 \mathrm{d}V \tag{1}$$

subject to arbitrary variations δg vanishing at the endpoints. Here, the Lagrangian $L(g, \dot{g})$ is the kinetic energy of the fluid and $\mathrm{d}V$ is the standard volume element on M. As this Lagrangian is invariant under *particle relabeling* — that is, the action of SDiff(M) on itself by composition on the *right*, the principle stated in Eq. (1) can be rewritten in *reduced (Eulerian) form* in terms of the Eulerian velocity $v = \dot{g} \circ g^{-1}$:

$$\delta \int_0^1 l(v) \mathrm{d}t = 0 \quad \text{where } l(v) = \frac{1}{2} \int_{M_0} \|v\|^2 \mathrm{d}V \tag{2}$$

subject to constrained variations $\delta v = \dot{\xi} + [v, \xi]$ (called *Lin constraints*), where ξ is an arbitrary divergence-free vector field – an element of the Lie algebra of the group of volume-preserving diffeomorphisms – and $[,]$ is the Jacobi Lie bracket (or vector field commutator). There is a complex history behind this reduced variational principle which was first shown for general Lie groups by Marsden and Scheurle [3] (see also [2,4–6]). As stated above, the reduced Eulerian principle is more attractive in computations because it involves a fixed Eulerian domain (mesh); however, the constrained variations necessary in this context complicates the design of a variational Eulerian algorithm.

1.2. Overview and contributions

While time integrators for fluid mechanics are often derived by approximating equations of motion, we instead follow the geometric principles described above and discretize the configuration space of incompressible fluids in order to derive the equations of motion through the principle of stationary action. Our approach uses an Eulerian, finite-dimensional representation of volume-preserving diffeomorphisms that encodes the displacement of a fluid from its initial configuration using special orthogonal, signed stochastic matrices. From this particular discretization of the configuration space, which forms a finite-dimensional Lie group, one can derive a right invariant discrete equivalent to the Eulerian velocity through its Lie algebra, i.e., through antisymmetric matrices whose columns sum to zero. After imposing non-holonomic constraints on the velocity field to allow transfer only between neighboring cells during each time update, we apply the Lagrange–d'Alembert principle (a variant of Hamilton's principle applicable to non-holonomic systems) to obtain the discrete equations of motion for our fluid representation.

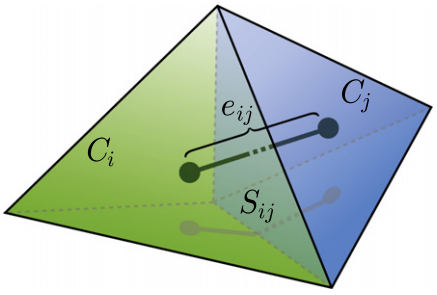

Fig. 2. Spatial Discretization: two cells C_i and C_j, with their common face $S_{ij} = C_i \cap C_j$ of area $|S_{ij}|$ and its dual edge e_{ij} of length $|e_{ij}|$.

As we will demonstrate, the resulting Eulerian variational Lie group integrator is structure-preserving, and as such, has numerous numerical properties, from momentum preservation (through a discrete Noether's theorem) to good long-term energy behavior.

1.3. Notations

The spatial discretization (mesh), either simplicial (tetrahedra) or regular (cubes), will be denoted by $\mathbb{M}$, with N being the number of n-dimensional cells $\{C_i\}_{i=1,\ldots,N}$ in $\mathbb{M}$. The ***size*** of a mesh will refer to the maximum diameter h of its cells. The Lebesgue measure will be denoted by $|.|$. Thus, $|C_i|$ is the volume of cell C_i, $|C_i \cap C_j|$ is the area of the face common to C_i and C_j, etc (see Fig. 2). The ***dual*** of $\mathbb{M}$ is the circumcentric dual cell complex [7], formed by connecting the circumcenters of each cell C_i based on the connectivity of $\mathbb{M}$. We will further assume that the mesh $\mathbb{M}$ is Delaunay with well shaped elements [8] to avoid degeneracies of its orthogonal dual as well as to simplify the exposition. We will also use the term ***regular grid*** (or Cartesian grid) to designate a mesh that consists of cells that are n-dimensional cubes of equal size. The notation $N(i)$ will denote the set of indices of cells neighboring cell C_i, that is, cell C_i shares a face with cell C_j iff $j \in N(i)$. We will say that a pair of cells C_i, C_j is ***positively oriented around an edge*** e if they share a face containing e and they are oriented such that they "turn" clockwise around the edge when viewed along the oriented edge. The same term will be used similarly for triplets of cells C_i, C_j, C_k where $i, k \in N(j)$ and all three cells contain edge e.

The notation $(.,.)$ and $\langle .,. \rangle$ will respectively refer to the L^2 inner product of vectors and the pairing of 1-forms and vector fields, while their discrete counterparts will be denoted by $((.,.))$ and $\langle\!\langle .,. \rangle\!\rangle$. Table 1 summarizes the main variables used in the remainder of this paper, along with their meaning and representation.

2. Discrete volume-preserving diffeomorphisms

We first introduce a finite-dimensional approximation to the infinite-dimensional Lie group of volume-preserving diffeomorphisms that tracks the amount of fluid transferred from one cell to another while preserving two key properties: volume and mass preservation.

Table 1
Physical/geometric meaning of the basic (continuous and discrete) variables used throughout this document.

Symbol	Meaning	Representation
M	Domain of motion	$M \subset \mathbb{R}^n$
n	Dimension of the domain	$n \in \mathbb{N}$
SDiff(M)	Configuration space of ideal fluid	Volume-preserving diffeomorphisms on M
SVect(M)	Tangent space of SDiff(M) at Id	Divergence-free vector fields on M
$\mathbb{M}$	Mesh discretizing domain M	Simplicial or regular mesh
N	Number of cells in $\mathbb{M}$	$N \in \mathbb{N}$
C_i	Cell #i of $\mathbb{M}$	Tetrahedron or cube in 3D
Ω	Discrete analog of volume form	Diagonal matrix of cell volumes, $\Omega_{ii} = \lvert C_i\rvert$
$\mathcal{D}(\mathbb{M})$	Discrete configuration space	Ω-orthogonal signed stochastic matrices
$\mathfrak{D}(\mathbb{M})$	Lie algebra of $\mathcal{D}(\mathbb{M})$	Ω-antisymmetric null-row matrices
q	Discrete configuration	Matrix $\in \mathcal{D}(\mathbb{M}) \subset GL(N) \subset \mathcal{M}^N$
A	Discrete Eulerian velocity $-\dot{q}q^{-1}$	Matrix $\in \mathfrak{D}(\mathbb{M}) \subset \mathfrak{gl}(N) = \mathcal{M}^N$
${}^kF$	Discrete k-form	N-dimensional tensor of order $(k+1)$
$\mathcal{N}$	Space of matrices with sparsity based on cell adjacency	Constrained set of matrices, with $A_{ij} \neq 0 \Rightarrow j \in N(i)$
$\mathcal{S}$	Space of sparse discrete velocities	Constrained set of velocities, $\mathcal{S} = \mathfrak{D}(\mathbb{M}) \cap \mathcal{N}$

2.1. Finite-dimensional configuration space

Suppose that the domain M is approximated by a mesh $\mathbb{M}$. Our first step in constructing a discrete representation of ideal fluids is to approximate SDiff(M) with a finite-dimensional Lie group in such a way that the elements of the corresponding Lie algebra can be considered as a discretization of divergence-free vector fields. To achieve this goal, we will *not discretize the diffeomorphism g itself*, but rather the associated operator $U_g : L_2 \to L_2$ defined by $\varphi(x) \mapsto \varphi(g^{-1}(x))$. Here $L^2 = L^2(M, \mathbb{R})$ is the space of square integrable real-valued functions on M. An important property of U_g is given by the following lemma, which follows from the change of variables formula.

Lemma 1 (*Koopman's Lemma*[1])**.** *If the diffeomorphism g is volume-preserving, then U_g is a unitary operator on L^2.*

Another important property of U_g is that it preserves constants, i.e., $U_g C = C$ for every constant function C, which can be seen as mass preservation for fluids. Next we present an approach to discretize this operator U_g while respecting its two defining properties.

Discrete functions. To discretize the operator U_g we first need to discretize the space on which U_g acts. Since the mesh $\mathbb{M}_h$ splits the domain of motion M into N cells C_i of maximum diameter h, a function $\varphi \in C^0(M; \mathbb{R})$ can be approximated by a step function $\bar{\varphi}$, constant within each cell of the mesh, through a map $R_{\mathbb{M}_h} : C^0(M; \mathbb{R}) \to$ step functions, which averages φ per cell:

$$R_{\mathbb{M}_h} : \varphi \mapsto \bar{\varphi}, \qquad \bar{\varphi} = \sum_i \left[\frac{1}{\Omega_i}\int_{C_i} \varphi\right] \chi_{C_i},$$

where χ_{C_i} is the indicator function for the cell C_i, and $\Omega_i = \lvert C_i\rvert$ is the volume of cell C_i. Since the space of all step functions on $\mathbb{M}_h$ is isomorphic to $\mathbb{R}^N$, we can consider the step functions as vectors: using the map $P_{\mathbb{M}_h} : L_2 \to \mathbb{R}^N$ defined by

$$(P_{\mathbb{M}_h}\varphi)_i = \frac{1}{\Omega_i}\int_{C_i} \varphi, \tag{3}$$

we can define a vector $\varphi_h = P_{\mathbb{M}_h}\varphi$ of size N to represent the step function $\bar{\varphi}$. To reconstruct a step function from an arbitrary vector $\varphi_h \in \mathbb{R}^N$ we define an operator $S_{\mathbb{M}_h} : \mathbb{R}^N \to L_2$ by

$$(S_{\mathbb{M}_h}\varphi_h)(x) = (\varphi_h)_i, \quad \text{if } x \in C_i. \tag{4}$$

[1] Many dynamical properties of g, such as ergodicity, mixing etc., can be studied using spectral properties of U_g. The idea of using methods of Hilbert spaces to study dynamical systems was first suggested by Koopman [9] and is usually called Koopmanism; it is closely related to the Perron–Frobenius methodology.

Thus, the operators $R_{\mathbb{M}_h}$, $P_{\mathbb{M}_h}$ and $S_{\mathbb{M}_h}$ are related through $R_{\mathbb{M}_h} = S_{\mathbb{M}_h}P_{\mathbb{M}_h}$.

The vector φ_h will be called a ***discrete function*** as it provides an approximation of a continuous function φ: when $h \to 0$,

$$S_{\mathbb{M}_h}\varphi_h \xrightarrow{C^0} \varphi, \quad \text{i.e., } \|S_{\mathbb{M}_h}\varphi_h - \varphi\|_{C^0} \to 0.$$

We also introduce a discrete approximation of the continuous L^2 inner product of functions $\langle\varphi, \psi\rangle = \int_M \varphi\psi$ through

$$\langle\varphi_h, \psi_h\rangle = \sum_i \Omega_i(\varphi_h)_i(\psi_h)_i. \tag{5}$$

Discrete diffeomorphisms. Let $\mathcal{M}^N$ be the space of real-valued $N \times N$ matrices. For a domain M discretized with N_h cells of maximum diameter h, we will say that a matrix $q_h \in \mathcal{M}^{N_h}$, acting on a vector φ_h, approximates U_g if $S_{\mathbb{M}_h}(q_h\varphi_h)$ is close to $U_g\varphi$. More precisely, we can define a notion of weak convergence of matrices to diffeomorphisms based on the C^0 norm:

Definition 1. Consider a family of meshes $\{\mathbb{M}_h\}$, with each mesh $\mathbb{M}_h$ discretizing a domain M and consisting of N_h cells of maximum size h. We will say that a family $\{q_h\}$ of matrices in $\mathcal{M}^{N_h}$, acting on discrete functions on $\mathbb{M}_h$, approximates a diffeomorphism $g \in$ SDiff(M) (and denote this property as $q_h \rightsquigarrow g$) if the following is true:

$$S_{\mathbb{M}_h}(q_hP_{\mathbb{M}_h}\varphi) \xrightarrow[h\to 0]{C^0} U_g\varphi \quad \text{for every } \varphi \in C(M; \mathbb{R}).$$

Note that although the space of matrices $\mathcal{M}^{N_h}$ changes as h approaches zero (as the dimension N_h of $\mathbb{M}_h$ increases), the operators P and S defined in Eqs. (3) and (4) are used to map the action of our discrete diffeomorphisms q_h on continuous functions back to the space of piecewise constant functions where we can properly define our notion of weak convergence with the continuous C^0 norm.

In order to better respect the continuous structures at play, we further enforce that our discrete configuration space of diffeomorphisms satisfies two key properties of U_g: volume preservation, reflecting the fact that U_g is unitary, and total mass preservation, as U_g preserves constants. We will thus only consider matrices q that

- *preserve the discrete L^2 inner product of functions*, i.e.,

  $\langle q\varphi_h, q\psi_h\rangle = \langle\varphi_h, \psi_h\rangle$,

  where the inner product of discrete functions is defined by Eq. (5). Denoting

$$\Omega = \begin{pmatrix} |C_1| & 0 & \dots & 0 \\ 0 & |C_2| & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & |C_N| \end{pmatrix},$$

note that this discrete notion of volume preservation directly implies that for our mesh $\mathbb{M}_h$ a volume-preserving matrix q satisfies

$q^T \Omega q = \Omega$.

The matrix q is thus *Ω-orthogonal*, restricted to matrices of determinant 1.

- *preserve constant vectors* (i.e., vectors having all coordinates equal) as well:

$$q\mathbf{1} = \mathbf{1}, \quad \text{where: } \mathbf{1} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}.$$

The matrix q must thus be *signed stochastic* as well.

Consequently, the finite-dimensional space of matrices we will use to discretize volume-preserving diffeomorphisms has the following definition.

Definition 2. Let $\mathbb{M}$ be a mesh consisting of cells C_i, $i = 1, \ldots, N$ and Ω be the diagonal matrix consisting of volumes of the cells, i.e., $\Omega_{ii} = |C_i|$ and $\Omega_{ij} = 0$ when $i \neq j$ (we will abusively use the shorter notation Ω_i to denote a diagonal element of Ω for simplicity in what follows). We will call a matrix $q \in \mathcal{M}^N$ ***volume-preserving*** and ***constant-preserving*** with respect to the mesh $\mathbb{M}$ if, for all i in $\{1, \ldots, N\}$,

$$q^T \Omega q = \Omega \tag{6}$$

and

$$\sum_j q_{ij} = 1. \tag{7}$$

The set of all such Ω-orthogonal, signed stochastic matrices of determinant 1 will be denoted by $\mathcal{D}(\mathbb{M})$, and will be used as a discretization of the configuration space SDiff(M).

Our finite-dimensional configuration space $\mathcal{D}(\mathbb{M})$ for fluid dynamics is thus the intersection of two Lie groups: the Ω-orthogonal group, and the group of invertible stochastic matrices; therefore, it is a *Lie group*. Note that if all cells of $\mathbb{M}$ have the same volume, i.e., $\Omega = \Omega_0 \mathrm{Id}$, then a matrix $q \in \mathcal{D}(\mathbb{M})$ is *orthogonal* in the usual sense and the equality (6) implies $\sum_i q_{ij} = 1$. For such meshes (which include Cartesian grids), the matrix q is *signed doubly stochastic*.

Remark. An alternate, arguably more intuitive way to discretize a diffeomorphism $g \in$ SDiff(M) on a mesh $\mathbb{M}$ would be to define a matrix T as

$$T_{ij}(g) \equiv \frac{|g^{-1}(C_j) \cap C_i|}{|C_i|}.$$

This discretization also satisfies by definition a discrete preservation of mass and a (different) notion of volume preservation. While it has the added benefit of enforcing that T has no negative terms (therefore respecting the positivity of U_g), the class of matrices it generates is, unfortunately, only a *semi-group*, which would be an impediment for establishing a variational treatment of fluids as an inverse map will be needed in the Eulerian formulation. So instead, we take the *orthogonal* part of this matrix T as our configuration q (which can be obtained in practice through the *polar decomposition*). (Notice that polar factorization has often been proposed in the context of fluids (see, e.g., [10]), albeit for more general nonlinear Hodge-like decomposition.) The negative coefficients happening in our finite-dimensional approximations, while not strictly physical, are the price to pay to capture the group structure. Conveniently, transporting functions instead of measures will not be an impediment, neither in theory nor numerically. Nevertheless, we will discuss in Section 6.3 that one could construct non-negative matrices T_k from any existing discrete geodesic path q_k.

2.2. Discrete velocity field

Now that we have established a finite-dimensional configuration space $\mathcal{D}(\mathbb{M})$, we describe its associated Lie algebra, and show that elements of this Lie algebra provide a discretization of divergence-free vector fields SVect(M). We will assume continuous time for simplicity, but a fully discrete treatment of space and time will be introduced in Section 5.

Consider a smooth path in the space of volume-preserving diffeomorphisms $g_t \in$ SDiff(M) with $g_0 = \mathrm{Id}$, and let $q_h(t)$ be an approximation of g_t, i.e., for any piecewise constant function φ_h^0 approximating a smooth function $\varphi^0 \in C^1(M, \mathbb{R})$, a discrete version of $\varphi^0 \circ g_t^{-1} = \varphi(t)$ is given by

$$\varphi_h(t) = q_h(t)\varphi_h^0.$$

Assuming $q_h(t)$ is smooth in time, we define its ***Eulerian velocity*** $A_h(t)$ to be

$$A_h(t) = -\dot{q}_h(t) q_h^{-1}(t),$$

thus yielding

$$\dot{\varphi}_h(t) = -A_h(t)\varphi_h(t).$$

Since $\frac{\mathrm{d}}{\mathrm{d}t}(\varphi^0 \circ g_t^{-1}) = -\langle \mathbf{d}\varphi(t), v_t \rangle = -\mathbf{L}_{v_t}\varphi$, where $v_t = \dot{g}_t \circ g_t^{-1}$ and $\mathbf{L}_{v_t}$ is the Lie derivative, the matrix $A_h(t)$ represents an approximation of the Eulerian velocity field v_t, which motivates the following definition:

Definition 3. Consider a one-parameter family of volume-preserving diffeomorphisms $g_t \in$ SDiff(M) and the associated time-dependent vector field $v_t = \dot{g}_t \circ g_t^{-1} \in$ SVect(M). Consider a family of meshes $\mathbb{M}_h$ of size h consisting of cells C_i^h and an operator $P_{\mathbb{M}_h} : C(M; \mathbb{R}) \to \mathbb{R}^{N_h}$ defined by Eq. (3).

We will say that a family of matrices $A_h(t) \in \mathcal{M}^{N_h}$ approximates a vector field v_t (denoted by $A_h(t) \rightsquigarrow v_t$) if the following statement is true:

$$S_{\mathbb{M}_h}(A_h(t) P_{\mathbb{M}_h}\varphi) \xrightarrow{C^0} \mathbf{L}_{v_t}\varphi \quad \text{for every } \varphi \in C^\infty(M; \mathbb{R}).$$

Remark. The choice of the minus sign in the definition of $A_h(t)$ stems from the fact that $q_h(t)$ represents U_g (thus, g^{-1} in essence). Since $\mathbf{L}_v = -\dot{U}_g U_g^{-1}$, we picked the sign to make $A_h(t)$ represents $\mathbf{L}_v$, consistent with the continuous case.

If a curve of matrices $q(t)$ belongs to the configuration space $\mathcal{D}(\mathbb{M})$ (i.e., if $q(t)$ is Ω-orthogonal signed stochastic), then its associated A belongs to its Lie algebra that we denote as $\mathfrak{D}(\mathbb{M})$. Matrices from this Lie algebra inherit the properties that *their rows must sum to zero*:

$$\sum_j A_{ij} = 0 \quad \text{(preservation of mass)},$$

and *they are Ω-antisymmetric*:

$$A^T\Omega + \Omega A = 0 \quad \text{(preservation of volume)}.$$

These two properties can be intuitively understood as discrete statements that A represents an advection, and the vector field representing this advection is divergence-free. Lie algebra elements for arbitrary simplicial meshes will be called *null-row Ω-antisymmetric matrices*. Note that if the mesh is regular ($\Omega = \Omega_0 \mathrm{Id}$), q belongs to the orthogonal group and the matrix A has to be *antisymmetric* with both its rows and columns summing to zero ("doubly null").

The link between convergence of $A_h(t)$ to $\mathbf{L}_{v_t}$ and convergence of $q_h(t)$ to U_{g_t} is described by the following lemma.

Lemma 2. *Consider the setup of Definition 3 and suppose a family of matrices $A_h(t) \in \mathfrak{D}(\mathbb{M}_h)$ approximates the Lie derivative $\mathbf{L}_{v_t}$ (in the sense of Definition 3) uniformly in t when $t \in [0, T]$ for some $T > 0$.*

Then there is a family of matrices $q_h(t) \in \mathcal{D}_{\mathbb{M}_h}$ such that $A_h(t) = -\dot{q}_h(t)q_h(t)^{-1}$ and $q_h(t)$ approximates g_t (in the sense of Definition 1).

Proof. Consider a family of smooth functions $\varphi(t, x)$ satisfying the advection equation

$$\dot{\varphi}(t, x) = -\mathbf{L}_{v_t}\varphi(t, x).$$

Suppose that $\bar{\varphi}(0, x) = S_{\mathbb{M}_h}P_{\mathbb{M}_h}\varphi(0, x)$ is an approximation to $\varphi(0, x)$ with

$$\sup_{x\in M} |\bar{\varphi}(0, x) - \varphi(0, x)| < \epsilon_1$$

and that $\varphi_h(t) = P_{\mathbb{M}_h}\bar{\varphi}(t, x)$ satisfies the discrete advection equation

$$\dot{\varphi}_h(t) = -A_h(t)\varphi_h(t).$$

Since $A_h(t)$ approximates $\mathbf{L}_{v_t}$, given $\varepsilon_2 > 0$, we can choose h such that

$$\|S_{\mathbb{M}_h}(A_h(t)\varphi_h(t)) - \mathbf{L}_{v_t}\varphi\| < \epsilon_2, \quad \text{for all } t \in [0, T].$$

Therefore,

$$\|S_{\mathbb{M}_h}(\dot{\varphi}_h(t))(x) - \dot{\varphi}(t, x)\| < \epsilon_2, \quad \text{for all } t \in [0, T]$$

and

$$\|S_{\mathbb{M}_h}(\varphi_h(t))(x) - \varphi(t, x)\| < \epsilon_1 + \epsilon_2 t.$$

Thus, we have shown that $\varphi_h(t, x)$ approximates $\varphi(t, x)$. However, $\varphi(t, x)$ satisfies

$$\varphi(t, x) = U_{g_t}\varphi(0, x),$$

and φ_h satisfies

$$\varphi_h(t) = q(t)\varphi_h(0),$$

where $q(t)$ is the matrix satisfying the equation

$$\dot{q}(t) = -A_h(t)q(t).$$

Therefore, we see that $q(t)\varphi(0)$ approximates $U_{g_t}\varphi(0, x)$. Thus, $A_h(t) \rightsquigarrow v_t$ implies that $q(t) \rightsquigarrow g_t$. □

2.3. Discrete commutator

A space-discrete flow that approximates a continuous flow $g(t) \in \mathrm{SDiff}(M)$ is defined to be a smooth path $q(t) \in \mathcal{D}(\mathbb{M})$ in the space of Ω-orthogonal signed stochastic matrices, such that $q(t) \rightsquigarrow g(t) \in \mathrm{SDiff}(M)$ (see Definition 2) and $A(t) = -\dot{q}(t)q^{-1}(t) \rightsquigarrow v_t = \dot{g}_t(g_t^{-1})$ (see Definition 3). It is straightforward to show that the Lie algebra structure of the space of divergence-free vector fields is preserved by our discretization. Indeed, if two matrices A and B approximate vector fields u and v then their commutator $[A, B]$ approximates the commutator of the Lie derivative operators:

$$[A, B] \to \mathbf{L}_u\mathbf{L}_v - \mathbf{L}_v\mathbf{L}_u.$$

Since $\mathbf{L}_u\mathbf{L}_v - \mathbf{L}_v\mathbf{L}_u = \mathbf{L}_{[u,v]}$, we obtain $[A, B] \rightsquigarrow [u, v]$, where $[.,.]$ denotes both the commutator of vector fields and the commutator of matrices. This property will be very useful to deal with Lin constraints later on.

2.4. Non-holonomic constraints (NHC)

For a smooth path $q(t)$, the matrix $A(t)$ describes the infinitesimal exchanges of fluid particles between any pair of cells C_i and C_j. We will thus assume that A_{ij} is non-zero *only* if cells C_i and C_j share a common boundary, i.e., are immediate neighbors. This sparsity will be numerically advantageous later on to reduce the computational complexity of the resulting integration schemes. We thus choose to *restrict discrete paths $\{q(t)\}$ on $\mathcal{D}(\mathbb{M})$ to those for which $A(t)$ satisfies this constraint*.[2] In other words, we *only* consider null-row Ω-antisymmetric matrices satisfying the constraints as valid discrete vector fields. The non-zero elements A_{ij} of these matrices correspond to boundaries between adjacent cells C_i and C_j, and can be interpreted as *directional transfer densities (per second) from C_i to C_j* — they could abusively be called "fluxes" on regular grids; but we will make the proper link with the integrals of the velocity field over mesh faces in the next section.

More formally, we define the ***constrained set*** $S_q \subset T_q\mathcal{D}(\mathbb{M})$ as the set of matrices corresponding to exchanges between neighboring cells only, i.e., $\dot{q} \in S_q$ if and only if $(\dot{q}q^{-1})_{ij} \neq 0$ implies that the cells C_i and C_j are neighbors. In this case the matrix A is defined by a set of non-zero values A_{ij} defined on faces between adjacent cells C_i and C_j. As mentioned previously, to indicate their adjacency, we will write that $j \in N(i)$ and $i \in N(j)$, where $N(k)$ refers to the set of indices of adjacent cells to cell C_k in the mesh $\mathbb{M}$. We will say that a matrix A belongs to the class $\mathcal{N}$ if $A_{ij} \neq 0$ implies $j \in N(i)$. Finally, we will denote by $\mathcal{S} \equiv S_{\mathrm{Id}} = \mathfrak{D}(\mathbb{M}) \cap \mathcal{N}$, the constrained set at the identity. Consequently, our treatment of fluid dynamics will only consider matrices A in $\mathcal{S} \subset \mathfrak{D}(\mathbb{M})$, i.e., matrices in $\mathfrak{D}(\mathbb{M})$ satisfying the sparsity constraints.

Note that if two matrices A and B both satisfy the constraints, their commutator need *not*: while the element of the commutator corresponding to any pair of cells which are more than two cells away is zero, the element $[A, B]_{ij}$ may be non-zero when cells C_i and C_j are "two cells away" from each other since

$$[A, B]_{ij} = \sum_k (A_{ik}B_{kj} - B_{ik}A_{kj}).$$

Notice that the commutator is zero for neighboring cells since $A_{kk} = B_{kk} = 0$ due to their Ω-antisymmetry. Writing $[\mathcal{S}, \mathcal{S}] = \{[A, B] \mid A, B \in \mathcal{S}\}$, one sees that $\mathcal{S} \cap [\mathcal{S}, \mathcal{S}] = \{\mathbf{0}\}$, where $\mathbf{0}$ is the zero matrix. Therefore, the constraints we just defined are *non-holonomic*.

Remark. When a discrete vector field A is in $\mathcal{S}$, the non-zero values $\Omega_i A_{ij}$ of the antisymmetric matrix ΩA can be understood as dual 1-chains, i.e., 1D chains on the dual of $\mathbb{M}$ [11]. This connection with 1-chains will become crucial later when dealing with advection of curves to derive a discrete Kelvin's theorem in Section 4.3.

2.5. Relation between elements of A and fluxes

Suppose we have a family of discrete flows $q_h(t)$ which approximates a flow $g_t \in \mathrm{SDiff}(M)$ such that $A_h(t) = -\dot{q}_h(t)q_h(t)^{-1}$ approximates $\mathbf{L}_{v_t}$ and satisfies the NHC. Let us see how individual elements $(A_h)_{ij}(t)$ of $A_h(t)$ are related to spatial values of v_t. Recall that

$$\dot{\varphi}_h(t) = -A_h(t)\varphi_h(t)$$

is a discrete version of the advection equation

$$\dot{\varphi} = -\mathbf{L}_{v_t}\varphi$$

and $A_h(t)\varphi_h(t) \to \mathbf{L}_{v_t}\varphi$ in the C^0 norm. But it also means that $(\Omega A_h(t)\varphi_h(t))_i$ is an approximation to the integral $\int_{C_i} \mathbf{L}_{v_t}\varphi_t$, i.e.,

$$\sum_{j\in N(i)} \Omega_i(A_h)_{ij}(t)\varphi_j(t) \approx \int_{C_i} \mathbf{L}_{v_t}\varphi_t \overset{\nabla\cdot v_t=0}{=} \int_{\partial C_i} \varphi_t(v_t, \vec{n}) \tag{8}$$

[2] Although we will adopt this sparsest form of the velocity in this paper, there may be advantages in considering larger non-zero neighborhoods in future work.

where $\vec{n}$ is the normal vector to the boundary of C_i and $(.,.)$ denotes the inner product of vectors. However,

$$\int_{\partial C_i} \varphi_t(v_t, \vec{n}) \approx \sum_{j\in N(i)} \frac{1}{2}(\varphi_i + \varphi_j) \int_{S_{ij}} (v_t, \vec{n}_{ij}) \overset{\nabla\cdot v_t = 0}{=} \sum_{j\in N(i)} \frac{1}{2}\varphi_j \int_{S_{ij}} (v_t, \vec{n}_{ij})$$

where S_{ij} is the face shared by cells C_i and C_j, and $\vec{n}_{ij}$ the normal vector to S_{ij} oriented from C_i to C_j. By comparing this result to Eq. (8), it is clear that an element $\Omega_i(A_h)_{ij}(t)$ can be considered (up to a constant) as an approximation to the *flux* of a vector field $v(t)$ through S_{ij}:

$$\Omega_i A_{ij}(t) \approx \frac{1}{2}\int_{S_{ij}} (v_t, \vec{n}_{ij}).$$

We know that $\int_{S_{ij}} (v_t, \vec{n}) \approx (v_t(x_{ij}), \vec{n}_{ij})S_{ij} + O(h^2)$, where x_{ij} is the barycenter of the boundary S_{ij} and $|S_{ij}|$ is the area of S_{ij}. Therefore, we obtain that, up to a constant dependent on local mesh measures, $(A_h)_{ij}$ approximates the flux through the boundary between C_i and C_j, i.e.,

$$(A_h)_{ij}(t) \approx (v_t(x_{ij}), \vec{n}_{ij})\frac{|S_{ij}|}{2\Omega_i}.$$

In the case of a Cartesian grid of size h this formula simplifies to:

$$(A_h)_{ij}(t) \approx \frac{(v_t(x_{ij}), \vec{n}_{ij})}{2h}.$$

2.6. Towards Lagrangian dynamics with non-holonomic constraints

One of the goals of this paper is to approximate geodesic flows on SDiff(M) by Lagrangian flows on $\mathcal{D}(\mathbb{M})$. To achieve this goal, we first need to define a Lagrangian $\mathcal{L}_h(q, \dot{q})$ such that

$$\mathcal{L}_h(q, \dot{q}) \to \int_M \frac{1}{2}\|v\|^2 \mathrm{d}V \quad \text{when } -\dot{q}q^{-1} \rightsquigarrow v \tag{9}$$

and

$$\delta\mathcal{L}_h(q, \dot{q}) \to \delta\int_M \frac{1}{2}\|v\|^2 \mathrm{d}V \quad \text{when } -\dot{q}q^{-1} \rightsquigarrow v$$
$$\text{and } \delta(-\dot{q}q^{-1}) \rightsquigarrow \delta v. \tag{10}$$

Such a Lagrangian, depending only on $A = -\dot{q}q^{-1}$ to mimic the continuous case, can then be used to formulate fluid dynamics through a discrete *Lagrange–d'Alembert principle* (to account for the non-holonomic constraint we impose on the sparsity of our Eulerian velocity approximation):

$$\delta\int_0^1 \mathcal{L}_h(q, \dot{q})\mathrm{d}t = 0 \quad \text{with } \begin{cases}\delta q \in S_q \\ \delta q(0) = \delta q(1) = 0.\end{cases}$$

Note that the constraint on the variations of q will induce a constraint on the variations of A, giving rise to a discrete version of the well-known Lin constraints of the form $\delta A = \dot{B} + [A, B]$, with $B = -\delta q q^{-1}$ (see Section 4.2).

However, we will show in later sections that coming up with a proper Lagrangian will require great care. As is typical with non-holonomic systems, the dynamics on $\mathcal{D}(\mathcal{M})$ will depend strongly on the values of $\partial\mathcal{L}_h/\partial A$ (i.e., the matrix with $\partial\mathcal{L}_h/\partial A_{ij}$ as its (i, j) element) *outside* of the constraint set $\mathcal{S}$ because of the commutator present in the Lin constraints. In particular, a conventional discretization of the kinetic energy via the sum of all the squared fluxes on the grid would lead to a matrix $\partial\mathcal{L}_h/\partial A$ with only values on pairs of adjacent cells, resulting in no dynamics. Instead, the Lagrangian *must* depend on values A_{ij} where $i \notin N(j)$.

To satisfy properties (9) and (10), we will look for a Lagrangian $\mathcal{L}_h$ of the form

$$\mathcal{L}_h(A) = \frac{1}{2}((A, A)),$$

where the discrete L^2-inner-product $((\cdot,\cdot))$ will be defined to satisfy the following properties (where $(\cdot,\cdot)$ denotes the continuous inner product of vector fields): for all $A, B \in \mathcal{S}$,

$$((A, B)) = ((B, A)) \to \int_M (u, v)\mathrm{d}V, \quad \text{when } A \rightsquigarrow u, B \rightsquigarrow v$$

and for all $A, B, C \in \mathcal{S}$

$$((A, [B, C])) \to \int_M (u, [v, w])\mathrm{d}V = \int_M -\mathbf{d}u^\flat(v, w)\mathrm{d}V,$$
$$\text{when } \begin{cases}A \rightsquigarrow u \\ B \rightsquigarrow v \\ C \rightsquigarrow w\end{cases} \tag{11}$$

where $\flat$ is the continuous flat operator (see for instance [12]). These properties will guarantee that conditions (9) and (10) are satisfied, and will lead to the proper dynamics. In the next section we will present a discretization of differential forms and a few operators acting on them to help us construct the discrete L^2-inner product (or equivalently, the discrete flat operator $\flat$).

3. Structure-preserving spatial field discretization

We now introduce a discrete calculus consistent with our discretization of vector fields. Unlike previous discrete exterior calculus approaches, mostly based on chains and cochains (see [13–15] and the references therein), we clearly distinguish between discrete vector fields and discrete forms acting on them. Moreover, our notion of forms will need to act not only on vector fields satisfying the NHC (being thus very reminiscent of the chain/cochain approach), but also on vector fields resulting from a commutator as imposed by the Lin constraints. We also introduce a discrete contraction operator $\mathbf{i}_v$ and a discrete Lie derivative $\mathbf{L}_v$ to complete our set of spatial operators — we will later show that the algebraic definition of our Lie derivative matches its dynamic counterpart as expected. We will not make any distinction in symbols between the discrete and continuous exterior calculus operators ($\mathbf{i}_v, \mathbf{L}_v, \mathbf{d}, \flat$, etc.) as the context will make their meaning clear.

3.1. Discrete 0-forms

In our context, a discrete 0-form is a function ${}^0F$ that is piecewise constant per cell as previously defined in Section 2. Note that its representation is a vector of N cell values,

$${}^0F = ({}^0F_1, {}^0F_2, \ldots, {}^0F_N)^T,$$

where ${}^0F_i$ represents the value of the function ${}^0F$ in cell C_i. Also, the volume integral of such a discrete 0-form is obtained by weighting the value of each cell by the Lebesgue measure of this cell, and summing all contributions:

$$\int_M {}^0F\mathrm{d}V = \sum_{i=1}^N \Omega_i {}^0F_i.$$

Remark. Our definition of 0-forms is no different from dual 0-cochains in dimension n as used extensively in, e.g., [11,13]. They naturally pair with dual 0-chains (i.e., linear combinations of cell circumcenters).

3.2. Discrete 1-forms

As the space $\mathcal{D}(\mathbb{M})$ of matrices is used to discretize vector fields, a natural way to discretize 1-forms is to also use matrices

to respect the duality between these two entities. Moreover, it is in line with the previous definition for 0-forms that were encoded as a 1-tensor: 1-forms will now be encoded by a 2-tensor. Notice that this is also reminiscent of the approximation $TM \approx M \times M$ used in discrete mechanics [16].

Discrete contraction. We define the contraction operator by a discrete vector field A, acting on a discrete 1-form ${}^1F$ to return a discrete 0-form, as

$$\mathbf{i}_A {}^1F \equiv \mathrm{diag}(A {}^1F^T) \stackrel{\text{def}}{=} ((A {}^1F^T)_{11}, \ldots, (A {}^1F^T)_{NN})^T. \tag{12}$$

Notice the metric-independence of this definition, and that if the discrete vector field contains only non-zero terms for neighboring cells, any term $({}^1F)_{ij}$ where cell C_i and cell C_j are not neighbors does not contribute to the contraction. In this case, the value of the resulting 0-form for cell C_i is thus: $(\mathbf{i}_A {}^1F)_i = \sum_{j\in N(i)} A_{ij} {}^1F_{ij}$, which is a local sum of the natural pairings of ${}^1F$ and A on each face of cell C_i.

Discrete total pairing. With this contraction defined, we derive a total pairing between a discrete 1-form and a discrete vector field as

$$\langle\!\langle {}^1F, A \rangle\!\rangle \equiv \mathrm{Tr}(\Omega A {}^1F^T).$$

This definition satisfies the following connection with the contraction defined in Eq. (12): indeed, for all $A \in \mathfrak{D}(\mathbb{M})$,

$$\int_M \mathbf{i}_A {}^1F \mathrm{d}V = \langle\!\langle {}^1F, A \rangle\!\rangle.$$

Note that the volume form Ω is needed to integrate the piecewise constant 0-form $\mathbf{i}_A {}^1F$ over the entire domain as explained in Section 3.1. Finally, since the matrix ΩA is antisymmetric, the symmetric component of ${}^1F$ does not play any role in the pairing.

Therefore, we will assume hereafter that a *discrete 1-form* ${}^1F$ *is defined by an antisymmetric matrix*: ${}^1F \in \mathfrak{so}(N)$.

Remark 1. When viewed as acting on vector fields in the NHC space $\mathcal{S}$, our representation of discrete 1-forms coincides with the use of 1-cochains on the dual of $\mathbb{M}$ [13]: the value ${}^1F_{ij}$ (resp., ${}^1F_{ji}$) can be understood as the integral of a continuous 1-form ${}^1f$ on the oriented dual edge going from cell C_i to cell C_j (resp., from C_j to C_i). However, our use of antisymmetric matrices extends this cochain interpretation. This will become particularly useful when 1-forms need to be paired with vector fields that have the form of the commutator $[A, B]$ of two vector fields A and B both in $\mathcal{S}$ as in Eq. (11).

Remark 2. Notice finally that we can also define the notion of contraction of the volume form Ω by a discrete vector field A using $\mathbf{i}_A \Omega = 2\Omega A$. The resulting matrix can be thought of as a discrete 2-form encoding the *flux* of A over each mesh face as derived in Section 2.5. In the notation convention of [7], this would be called a "primal" 2-form, while the 2-forms we will work with in this paper are "dual" 2-forms. We will not discuss these primal 2-forms further in this paper (as the construction of a consistent discrete calculus of forms and tensors is a subject on its own), but it is clear that they naturally pair with dual 1-forms ${}^1F$ (forming a discrete wedge product between primal 2- and dual 1-forms), numerically resulting in the same value as the discrete pairing $\langle\!\langle {}^1F, A \rangle\!\rangle$.

3.3. Discrete 2-forms

We extend our definition of 1-forms to 2-forms in a similar fashion: discrete 2-forms will be encoded as 3-tensors ${}^3F_{ijk}$ that are completely antisymmetric, i.e., antisymmetric with respect to any pair of indices.

Discrete contraction. Contraction of a 2-form ${}^2F$ by a vector field A is defined as

$$(\mathbf{i}_A {}^2F)_{ij} = \sum_k ({}^2F_{ikj} A_{ik} - {}^2F_{jki} A_{jk}).$$

Notice again here that the resulting discrete 1-form is indeed an antisymmetric matrix (by construction), and that if $A \in \mathcal{S}$, many of the terms in the sum vanish.

Discrete total pairing. The total pairing of a discrete 2-form ${}^2F$ by two discrete vector fields A and B, the discrete equivalent of $\int_M {}^2f(a, b)\mathrm{d}V$, will be defined as

$$\langle\!\langle {}^2F, A, B \rangle\!\rangle \equiv 2 \sum_{i,j,k} \Omega_i {}^2F_{ijk} A_{ij} B_{ik}. \tag{13}$$

This definition satisfies the expected property linking contraction and pairing: for all $B \in \mathcal{S}$,

$$\langle\!\langle \mathbf{i}_A {}^2F, B \rangle\!\rangle = \langle\!\langle {}^2F, A, B \rangle\!\rangle.$$

Indeed, using our previous definitions, we have

$$\begin{aligned}
\langle\!\langle \mathbf{i}_A {}^2F, B \rangle\!\rangle &= \mathrm{Tr}(\Omega B (\mathbf{i}_A {}^2F)^T) = \sum_{i,j} \Omega_i B_{ij} (\mathbf{i}_A {}^2F)_{ij} \\
&= \sum_{i,j,k} \Omega_i ({}^2F_{ikj} A_{ik} - {}^2F_{jki} A_{jk}) B_{ij} \\
&\left.\begin{array}{l}\text{using } {}^2F_{ijk} = -{}^2F_{ikj} \\ \text{and } \Omega_i B_{ij} = -\Omega_j B_{ji}\end{array}\right\} \\
&= \sum_i \left(\sum_{j,k} (-\Omega_i {}^2F_{ijk} A_{ik} B_{ij} + \Omega_j {}^2F_{jki} A_{jk} B_{ji}) \right) \\
&= -2 \sum_i \left(\sum_{j,k} \Omega_i {}^2F_{ijk} A_{ik} B_{ij} \right) = -\langle\!\langle {}^2F, B, A \rangle\!\rangle \\
&= \langle\!\langle {}^2F, A, B \rangle\!\rangle.
\end{aligned}$$

3.4. Other operators on discrete forms

A few more operators acting on 0-, 1-, or 2-forms will be valuable to our discretization of incompressible fluids.

Discrete exterior derivative. We can easily define a discrete version $\mathbf{d}$ of the exterior derivative. For a discrete 0-form ${}^0F$, the 1-form $\mathbf{d}{}^0F$ is defined as

$$(\mathbf{d}{}^0F)_{ij} = {}^0F_j - {}^0F_i.$$

Similarly, if ${}^1F$ is a discrete 1-form then we can define

$$(\mathbf{d}{}^1F)_{ijk} = {}^1F_{ij} + {}^1F_{jk} + {}^1F_{ki}.$$

More generally, we define our operator $\mathbf{d}$ as acting on a k-form ${}^kF$ through:

$$(\mathbf{d}{}^kF)_{i_1 i_2 \ldots i_{k+1}} = \sum_{j\in[1..k+1]} (-1)^{j+1} {}^kF_{i_1 \ldots \widehat{i_j} \ldots i_{k+1}}$$

where $\widehat{\cdot}$ indicates the omission of a term. This expression respects the antisymmetry of our discrete form representation.

Remark. Notice here again that when the circumcenters of cells $C_{i_1}, C_{i_2}, \ldots, C_{i_{k+1}}$ form a k-simplex on the dual of mesh $\mathbb{M}$, our definition of $\mathbf{d}$ simply enforces Stokes' theorem and thus coincides with the discrete exterior derivative widely used in the literature [13]. Our discrete exterior derivative extends this simple geometric property to arbitrary $(k+1)$-tuples of cells, while trivially enforcing that $\mathbf{d} \circ \mathbf{d} = 0$ on the discrete level as well.

Discrete Lie derivative. Now that we have defined contraction and derivatives on discrete 1-forms we can define the Lie derivative using Cartan's "magic" formula in the continuous setting $\mathbf{L}_v = \mathbf{i}_v \mathbf{d} + \mathbf{d}\mathbf{i}_v$.

Definition 4. Let A be a discrete vector field satisfying the NHC and ${}^1F$ be a discrete 1-form. Then the discrete Lie derivative of ${}^1F$ along A is defined as

$$\mathbf{L}_A {}^1F = \mathbf{i}_A \mathbf{d} {}^1F + \mathbf{d}\mathbf{i}_A {}^1F.$$

Lemma 3. *For a vector field represented through an Ω-antisymmetric and null-row A, and a discrete closed 1-form represented as a null-row and antisymmetric ${}^1F$:*

$$\mathbf{L}_A {}^1F = [A, {}^1F\Omega]\Omega^{-1} = A^1F - (A^1F)^T. \tag{14}$$

Proof. As A is null-row, we have

$$\sum_k {}^1F_{ij}A_{ik} = {}^1F_{ij}\sum_k A_{ik} = 0.$$

Therefore,

$$\begin{aligned}(\mathbf{i}_A\mathbf{d}^1F)_{ij} &= \sum_k (((\mathbf{d}^1F)_{ikj}A_{ik} - (\mathbf{d}^1F)_{jki}A_{jk})\\ &= \sum_k ({}^1F_{ik} + {}^1F_{kj} + {}^1F_{ji})A_{ik}\\ &\quad - \sum_k ({}^1F_{jk} + {}^1F_{ki} + {}^1F_{ij})A_{jk}\\ &= (A^1F)_{ij} + ({}^1FA^T)_{ii} - (A^1F)_{ji} - ({}^1FA^T)_{jj}.\end{aligned}$$

Now, since $A^T = -\Omega A\Omega^{-1}$ and ${}^1F^T = -{}^1F$, we can write:

$$(A^1F)_{ji} = ((A^1F)^T)_{ij} = ({}^1F\Omega A\Omega^{-1})_{ij}$$

and, therefore,

$$(\mathbf{i}_A\mathbf{d}^1F)_{ij} = ([A, {}^1F\Omega]\Omega^{-1})_{ij} + ({}^1FA^T)_{ii} - ({}^1FA^T)_{jj}.$$

Also, one has

$$\mathbf{i}_A {}^1F = \operatorname{diag}({}^1FA^T),$$

therefore,

$$(\mathbf{d}\mathbf{i}_A {}^1F)_{ij} = ({}^1FA^T)_{jj} - ({}^1FA^T)_{ii},$$

which implies the result. □

Note that the resulting formula corresponds to an antisymmetrization of A applied to ${}^1F$ — leading, up to the volume form Ω, to the commutator of A and ${}^1F$.

3.5. Discrete L^2-inner product and discrete flat operator

The Lagrangian for incompressible, inviscid fluid dynamics is the squared L^2-norm of the velocity field. Hence, we wish to define a discrete L^2-inner product between two discrete vector fields. Since we require spatial sparsity (NHC condition) of the velocity field A, and Lin constraints for its variation $\delta A = \dot{B} + [A, B]$, we are only concerned with vector fields in $\mathcal{s} \cup [\mathcal{s}, \mathcal{s}]$.

Recall that the continuous flat of a vector field v is a 1-form $v^\flat$ such that

$$\langle v^\flat, w\rangle = (v, w), \quad \text{for every vector field } w,$$

where (v, w) is the L^2-inner product of vector fields. Since the discrete total pairing is essentially a Frobenius inner product, discretizing the L^2 inner product for vector fields is equivalent to discretizing the flat operator $\flat : A \mapsto A^\flat$ such that the pairing of matrices $\langle\langle A^\flat, B\rangle\rangle$ approximates the inner product of vector fields integrated on M:

$$\begin{aligned}((A, B)) &= \langle\langle A^\flat, B\rangle\rangle = \operatorname{Tr}(\Omega B(A^\flat)^T)\\ &\underset{h\to 0}{\to} \int_M (v, w)\mathrm{d}V, \quad \text{if } A \rightsquigarrow v \text{ and } B \rightsquigarrow w.\end{aligned}$$

Looking ahead, we will only use the L^2 inner product of the type $((A + \delta A, A + \delta A))$ when taking variations of the Lagrangian. Therefore, we need only to define L^2 inner products of the form $((A, B))$, $((A, [B, C]))$ and $(([B, C], A))$, for any $A, B, C \in \mathcal{s}$ (equivalently, $\langle\langle A^\flat, B\rangle\rangle$, $\langle\langle A^\flat, [B, C]\rangle\rangle$, and $\langle\langle [B, C]^\flat, A\rangle\rangle$). As our discrete L^2 inner product will be symmetric, we only need to focus on inner products of the form $((A, \cdot))$ (resp., $\langle\langle A^\flat, \cdot\rangle\rangle$) for $A \in \mathcal{s}$. Note that this discrete L^2 inner product can *not* be trivial: indeed, for any matrices $A, B, C \in \mathcal{s}$, we have $\operatorname{Tr}(A[B, C]) = 0$ because $\mathcal{s} \cap [\mathcal{s}, \mathcal{s}] = \{\mathbf{0}\}$; but we could choose A, B and C that approximate vector fields v, u and w such that $\int_M (v, [u, w]) \neq 0$. As we now introduce, we define our discrete symmetric L^2 inner product in a matter that satisfies a discrete version of the continuous identity $\int_M (v, [u, w]) = -\int_M \mathrm{d}v^\flat(u, w)$, which holds for divergence-free vector fields.

Definition 5. Consider a family of meshes $\mathbb{M}_h$ of size h. An operator

$$\flat_h : \mathcal{s} \to \mathfrak{D}(\mathbb{M}_h)$$

is called a discrete flat operator if the following two conditions are satisfied:

$$\begin{aligned}&\langle\langle A_h^{\flat_h}, B_h\rangle\rangle \to \int_M (v(x), u(x))\mathrm{d}x, \quad \text{when } h \to 0,\\ &\quad \text{for every } A_h, B_h \in \mathcal{s}, A_h \to \mathbf{L}_v, B_h \to \mathbf{L}_u\end{aligned} \tag{15}$$

$$\begin{aligned}&\langle\langle A_h^{\flat_h}, [B_h, C_h]\rangle\rangle \to \int_M (v(x), [u, w](x))\mathrm{d}x, \quad \text{when } h \to 0,\\ &\quad \text{for every } A_h, B_h, C_h \in \mathcal{s}, A_h \to \mathbf{L}_v,\\ &\quad B_h \to \mathbf{L}_u, C_h \to \mathbf{L}_w.\end{aligned} \tag{16}$$

Note that in this definition, $\langle\langle A^{\flat_h}, X\rangle\rangle$ approximates the continuous inner product both when $X \in \mathcal{s}$ *and* when $X \in [\mathcal{s}, \mathcal{s}]$.

The next lemma introduces a necessary and sufficient condition to guarantee the validity of a discrete flat operator. This particular condition will be very useful when we study the dynamics of discrete fluids, as it involves the vorticity $\omega = \nabla \times u$ of a vector field:

Lemma 4. *A family of operators $\flat_h$ satisfies condition (16) if and only if for every $A_h, B_h, C_h \in \mathcal{s}$ approximating vector fields $v, u, w \in \operatorname{SVect}(M)$ respectively we have*

$$\langle\langle \mathbf{d}A_h^{\flat_h}, B_h, C_h\rangle\rangle \to \int_M \omega(u, w)\mathrm{d}V, \quad \text{where } \omega = \mathbf{d}v^\flat.$$

Proof. First, let us show that for any $u, v, w \in \operatorname{SVect}(M)$

$$\int_M (v, [u, w])\mathrm{d}x = \int_M -\mathbf{d}v^\flat(u, w)\mathrm{d}x.$$

Indeed, since

$$\int_M (v, [u, w])\mathrm{d}x = \int_M \mathbf{i}_{[u,w]}v^\flat$$

and (see [6])

$$\mathbf{i}_{[u,w]}v^\flat = \mathbf{L}_u\mathbf{i}_w v^\flat - \mathbf{i}_w\mathbf{L}_u v^\flat,$$

we have

$$\int_M (v, [u, w])\mathrm{d}x = \int_M \mathbf{L}_u\mathbf{i}_w v^\flat - \mathbf{i}_w\mathbf{L}_u v^\flat \overset{\nabla\cdot u=0}{=} -\int_M \mathbf{i}_w\mathbf{L}_u v^\flat.$$

But, by Cartan's formula $\mathbf{L}_u v^\flat = \mathbf{i}_u\mathbf{d}v^\flat + \mathbf{d}\mathbf{i}_u v^\flat$. Therefore,

$$\int_M \mathbf{i}_w\mathbf{L}_u v^\flat = \int_M \mathbf{i}_w\mathbf{i}_u\mathbf{d}v^\flat = \int_M \mathbf{d}v^\flat(u, w),$$

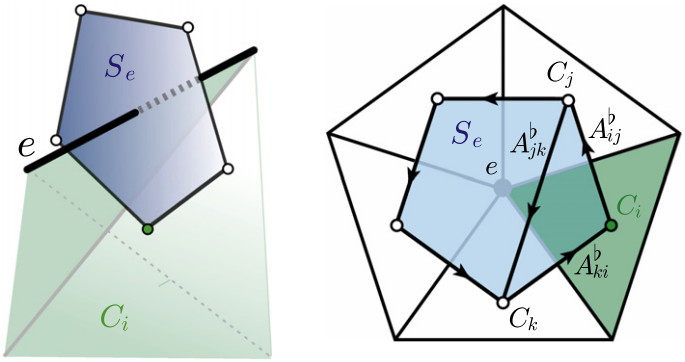

Fig. 3. Flat Operator: schematic representation of $A^\flat_{ij}$ as a part of the cell S_e dual to edge e in 3D (left) and a view of the dual cell seen straight along the edge (right). This last figure can also be seen as the 2D schematic version of the flat operator, where e is now a vertex and S_e is its associated dual Voronoi face.

where we used the fact that $w \in \mathrm{SVect}(M)$ and therefore $\int_M \mathbf{i}_w \mathbf{d}\mathbf{i}_u v^\flat = 0$.

Now, let us show that $\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle = -\langle\langle A^\flat, [B, C]\rangle\rangle$. Using properties of the trace operator we have

$$\begin{aligned}\langle\langle A^\flat, [B, C]\rangle\rangle &= \mathrm{Tr}(\Omega[B, C](A^\flat)^T) = -\mathrm{Tr}(A^\flat \Omega[B, C])\\ &= -\mathrm{Tr}([A^\flat \Omega, B]C).\end{aligned}$$

By Lemma 3, $[A^\flat \Omega, B]\Omega^{-1} = -\mathbf{L}_B A^\flat$. Thus,

$$\begin{aligned}-\mathrm{Tr}([A^\flat \Omega, B]C) &= \mathrm{Tr}((\mathbf{L}_B A^\flat)\Omega C) = -\mathrm{Tr}(\Omega C(\mathbf{i}_B \mathbf{d}A^\flat)^T)\\ &= -\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle,\end{aligned}$$

where we used that $\mathrm{Tr}(\Omega C(\mathbf{d}\mathbf{i}_B A^\flat)^T) = 0$ because C is divergence-free. □

Discrete vorticity in the sense of DEC. As our derivation relies on having a predefined notion of discrete vorticity, we first provide a definition used in [7,17] (we will refer to it as the DEC vorticity, as it was derived from a Discrete Exterior Calculus [13]):

$$\omega_{\mathrm{DEC}}(e) = \sum_{\substack{(i,j)\\ e\subset(C_i\cap C_j)}} 2\Omega_i \frac{|e_{ij}|}{|S_{ij}|} A_{ij} s_{ij}, \tag{17}$$

where $s_{ij} = 1$ if the cells C_i and C_j are positively oriented around e and $s_{ij} = -1$ otherwise. Notice that this represents the integral of the vector field A along dual edges e_{ij} around the edge e: by Stokes' theorem, $\omega_{\mathrm{DEC}}(e)$ is thus *the vorticity of A integrated over the dual Voronoi face to e* (see Fig. 3, left). More importantly, it has been established that this approximation does converge (as long the mesh does not get degenerate) to the notion of vorticity in the limit of refinement [18].

A flat operator on a 3D mesh. From the previous lemma, we can derive a construction of a flat operator on a 3D simplicial mesh. Given a matrix A, we need to find a matrix $A^\flat$ which satisfies the properties

$$\langle\langle A^\flat, A\rangle\rangle = \langle\langle A, A^\flat\rangle\rangle \to \int_M \|v\|^2 \mathrm{d}V,$$

and

$$\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle \to \int_M \omega(u, w)\mathrm{d}V.$$

To satisfy the first property we simply define the values of $A^\flat$ for *immediate* neighbors as

$$A^\flat_{ij} = A_{ij}\Omega_i \frac{2|e_{ij}|}{|S_{ij}|} \quad \text{for } j \in N(i). \tag{18}$$

Notice that it corresponds to the flux $2\Omega_i A_{ij}$ of the velocity field, further multiplied by the diagonal Hodge star of 2-forms for the face S_{ij} (see, e.g., [18]) to make $A^\flat$ a 1-form on the dual edge between C_i and C_j.

Enforcing the second property of the flat operator is more difficult; our construction will use the fact that in the limit, one must have

$$\int_M \omega(u, w)\mathrm{d}V = \int_M *\omega \wedge u^\flat \wedge w^\flat.$$

Let us assume that the values of $A^\flat$ for adjacent cells are defined by Eq. (18), and that the values of $A^\flat$ for *non-adjacent* pairs of cells C_j and C_k are *defined* by

$$(\mathbf{d}A^\flat)_{ijk} = A^\flat_{ij} + A^\flat_{jk} + A^\flat_{ki} = K_{ijk}\omega_{\mathrm{DEC}}(e_{ijk}), \tag{19}$$

where C_i is adjacent to both C_j and C_k (see Fig. 3 (right) for a schematic depiction), e_{ijk} is the primal edge common to the cells C_i, C_j, C_k, and K_{ijk} is a coefficient of proportionality whose exact expression will be provided later on. In other words, we assume that the flat operator allows us to evaluate vorticity not only on dual (Voronoi) faces as in the DEC sense, but on any triplet of cells C_i, C_j, C_k as depicted in Fig. 3 (right); this will give us values of vorticity on *subparts* of Voronoi faces as well.

Then the pairing $\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle$ can be written (see Def. (13)) as

$$\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle = 2\sum_{i,j,k} \Omega_i K_{ijk}\omega_{\mathrm{DEC}}(e_{ijk})B_{ij}C_{ik},$$

or, if one uses the flat of both vector fields B and C,

$$\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle = \frac{1}{2}\sum_{i,j,k} \Omega_i \widetilde{K}_{ijk}\omega_{\mathrm{DEC}}(e_{ijk})B^\flat_{ij}C^\flat_{ik}, \tag{20}$$

where

$$B^\flat_{ij} = \Omega_i B_{ij}\frac{2|e_{ij}|}{|S_{ij}|}, \qquad C^\flat_{ik} = \Omega_i C_{ik}\frac{2|e_{ik}|}{|S_{ik}|}, \quad \text{and}$$

$$\widetilde{K}_{ijk} = K_{ijk}\frac{1}{\Omega_i^2}\frac{|S_{ij}|}{|e_{ij}|}\frac{|S_{ik}|}{|e_{ik}|}.$$

Now, suppose we have a discrete version of the wedge product (e.g., [7,19]) between two dual 1-forms, written with given weights W_{ijk} as

$$(B^\flat \wedge C^\flat)_{S_{e_{ijk}}} = \sum_{\substack{i,j,k\\ e_{ijk}=C_i\cap C_j\cap C_k}} W_{ijk}B^\flat_{ij}C^\flat_{ik},$$

where $S_{e_{ijk}}$ is the 2D face dual to the primal edge e_{ijk} and the sum is taken as before over all consecutive cells i, j and k which have e_{ijk} as a common edge. If we further define

$$\widetilde{K}_{ijk} = 2W_{ijk}\frac{|e|}{\Omega_i|S_{e_{ijk}}|}$$

(where, as usual, $|e_{ijk}|$ denotes the length of the edge e_{ijk} and $|S_{e_{ijk}}|$ is the area of the dual face $S_{e_{ijk}}$), we can reexpress equation (20) by summing over all edges e_{ijk} to find a simple wedge-product-based version of the total pairing of the vorticity with two vector fields:

$$\langle\langle \mathbf{d}A^\flat, B, C\rangle\rangle = \sum_e \omega_{\mathrm{DEC}}(e)\frac{|e|}{|S_e|}(B^\flat \wedge C^\flat)_{S_e} \approx \int_M (*\omega)\wedge u^\flat \wedge w^\flat.$$

Thus, we can derive the flat operator $\flat$ once a set of coefficients W_{ijk} is known: given a vector field $A \in \mathcal{S}$, $A^\flat$ for adjacent cells is defined using Eq. (18), while the rest of its non-zero values are defined such that

$$A^\flat_{ij} + A^\flat_{jk} + A^\flat_{ki} = K_{ijk}\omega_{\mathrm{DEC}}(e_{ijk}),$$

where

$$K_{ijk} = 2W_{ijk}\Omega_i \frac{|e_{ij}|}{|S_{ij}|}\frac{|e_{ik}|}{|S_{ik}|}\frac{|e_{ijk}|}{|S_{e_{ijk}}|} \quad \text{for } e = C_i \cap C_j \cap C_k.$$

A concrete expression of W_{ijk} can be used by extending the definition of the primal–primal wedge product given in [7] (Definition 7.1.1) to the dual in a straightforward fashion to make it exact for constant volume 2-forms through:

$$W_{ijk} = s_{ijk}\frac{|S_{e_{ijk}} \cap C_i|}{|\Delta_{ijk}|},$$

where Δ_{ijk} is a triangle with vertices at the (circum)centers of the cells C_i, C_j and C_k, and $s_{ijk} = 1$ if the triplet of cells C_i, C_j, C_k is positively oriented around e and $s_{ijk} = -1$ otherwise.

To simplify the expression for K_{ijk} we use the equality $|\Delta_{ijk}| = \frac{1}{2}|e_{ij}|\,|e_{ik}|\sin\alpha_{ijk}$, where α_{ijk} is the angle between dual edges, yielding:

$$K_{ijk} = 4s_{ijk}\Omega_i|e|\frac{1}{\sin\alpha_{ijk}}\frac{1}{|S_{ij}|\,|S_{ik}|}\frac{|S_{e_{ijk}} \cap C_i|}{|S_{e_{ijk}}|}.$$

Now, applying the generalized law of sines for the volume of a tetrahedron yields

$$\Omega_i = \frac{2}{3|e_{ijk}|}|S_{ij}|\,|S_{ik}|\sin\alpha_{ijk}$$

and thus

$$K_{ijk} = \frac{8}{3}s_{ijk}\frac{|S_{e_{ijk}} \cap C_i|}{|S_{e_{ijk}}|}.$$

This formula was used in the implementation of our method as described in [20] (note that the wedge product was rewritten as a function of the flux $F_{ij} = 2\Omega_i A_{ij}$).

Flat operator on regular grids in 2D. Our construction of the flat operator is particularly simple for regular (Cartesian) grids as we now review for completeness.

Lemma 5. *For a domain represented with a Cartesian grid of size h, let A be an antisymmetric doubly null matrix satisfying the NHC. The operator* $\flat : A \mapsto A^\flat$ *defined as*

$$A^\flat_{ij} = 2h^2 A_{ij}, \quad \text{for } i \in N(j),$$
$$A^\flat_{ij} = h^2 \sum_{k\in N(i)\cap N(j)} (A_{ik} + A_{kj}), \quad \text{for } i \neq N(j)$$

is a discrete flat operator.

Note that while A satisfies the NHC, $A^\flat$ has non-zero elements for neighboring cells and for cells that share a common neighbor (i.e., two cells away). Now, let i, k, j, l be four cells on a regular mesh sharing a common node x, oriented counter-clockwise (see Fig. 4). Then it is easy to see that for $A^\flat$ defined above we have

$$\mathbf{d}A^\flat_{ijk} = h^2(A_{ik} + A_{kj} + A_{jl} + A_{li}) = \frac{\omega_{\mathrm{DEC}}(x)}{2},$$

where $\omega_{\mathrm{DEC}}(x)$ is the discrete vorticity in the sense of Discrete Exterior Calculus integrated over the dual cell of node x. Since ω_d converges to vorticity the condition of Lemma 4 is satisfied, just as in the simplicial mesh case.

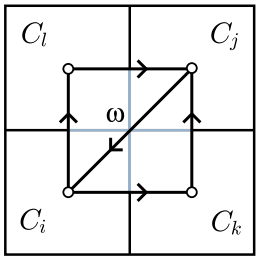

Fig. 4. Flat Operator on a Regular Grid in 2D: our definition of the flat operator is particularly simple when the spatial discretization is a regular mesh.

4. Dynamics on the group of Ω-orthogonal stochastic matrices

We now focus on defining a Lagrangian on the tangent bundle of the group $\mathcal{D}(\mathbb{M})$ of Ω-orthogonal, signed stochastic matrices and studying the corresponding variational principle with non-holonomic constraints. We will first assume a discrete-space/continuous-time setup before presenting a fully discrete version.

4.1. Variational principle and symmetry

We wish to study dynamics on the Lie group $\mathcal{D}(\mathbb{M})$ of Ω-orthogonal, signed stochastic matrices representing volume-preserving diffeomorphisms on a mesh $\mathbb{M}$. While the group's Lie algebra $\mathfrak{D}(\mathbb{M})$ consists of null-row Ω-antisymmetric matrices, we restrict the Eulerian velocity $A = -\dot{q}q^{-1}$ to lie in the NHC space $\mathcal{S}$, i.e., with fluid transfer happening only between adjacent cells (see Section 2.4).

We first establish a discrete Lagrangian $\mathcal{L}_h(q, \dot{q})$ on $T\mathcal{D}(\mathbb{M})$ with the property that $\mathcal{L}_h \xrightarrow[h\to 0]{} \frac{1}{2}\int \|v\|^2$ for $A \underset{h\to 0}{\rightsquigarrow} v$ by defining:

$$\mathcal{L}_h(A) = \frac{1}{2}\langle\langle A^\flat, A\rangle\rangle \equiv \frac{1}{2}\mathrm{Tr}(\Omega A(A^\flat)^T).$$

When A satisfies the NHC, it was shown in Section 3.5 that $\langle\langle A^\flat, A\rangle\rangle \to \int(v, v)$; thus the discrete Lagrangian is a proper approximation to the L^2-norm of the velocity field in this case. Note also that it is trivially right invariant as in the continuous case, since one can compose q by a discrete diffeomorphism η without changing the Eulerian velocity $A = -(\dot{q}\eta)(q\eta)^{-1} = -\dot{q}q^{-1}$. Our discrete setup thus respects *particle relabeling symmetry*.

4.2. Computing variations

To compute the variation of $A(t)$, we assume that q depends on a parameter s, we denote $q' = \frac{dq}{ds}$ and $\dot{q} = \frac{dq}{dt}$, and we differentiate the Eulerian velocity:

$$\frac{\mathrm{d}}{\mathrm{d}s}A(s,t) = -\dot{q}'q^{-1} + \dot{q}q^{-1}q'q^{-1}.$$

If we denote by B the vector field satisfying $B = -q'q^{-1}$, we directly get the well-known *Lin constraints*:

$$\frac{\mathrm{d}}{\mathrm{d}s}A(s,t) = \dot{B} + [A, B], \tag{21}$$

where $[A, B] = AB - BA$ is the commutator of matrices.

Now remember that the dynamics of systems with nonholonomic constraints can be derived from the *Lagrange–d'Alembert principle*:

$$\delta \int_0^1 \mathcal{L}_h(A)\mathrm{d}t = 0, \quad \delta q \in S_q,\ A \in S,$$
$$\delta q(0) = \delta q(1) = 0. \tag{22}$$

Since $\delta q \in S_q$, the vector field B must be in $\mathcal{S}$, i.e., $B_{ij} = 0$ except for neighboring cells C_i and C_j. We can then compute $\delta\mathcal{L}_h$:

$$\delta\mathcal{L}_h(A) = \frac{1}{2}(\langle\langle \delta A^\flat, A\rangle\rangle + \langle\langle A^\flat, \delta A\rangle\rangle) = \langle\langle A^\flat, \delta A\rangle\rangle.$$

As we restrict A to lie in the NHC subspace $\mathcal{S}$, the Lin constraints in Eq. (21) imply

$$\delta\mathcal{L}_h(A) = \langle\langle A^\flat, \dot{B} + [A, B]\rangle\rangle.$$

Recall that if A approximates $\mathbf{L}_v$ and B approximates $\mathbf{L}_\xi$, then by the definition of $\flat$,

$$\langle\langle A^\flat, \dot{B}\rangle\rangle \to \int_M (v, \dot{\xi})\mathrm{d}V$$

and

$$\langle\langle A^\flat, [A, B]\rangle\rangle \to \int_M (v, [v, \xi])\mathrm{d}V.$$

Thus,

$$\delta\mathcal{L}_h(A) \to \int_M (v, \dot{\xi} + [v, \xi])\mathrm{d}V = \delta l(v),$$

so the discrete Lagrangian (resp., its variation) is an approximation of the continuous Lagrangian (resp., its variation).

Since $A^\flat$ is antisymmetric and $\mathrm{Tr}(A[B, C]) = \mathrm{Tr}([A, B]C)$ for any matrices A, B, C, we get

$$\delta\mathcal{L}_h(A) = \mathrm{Tr}(-A^\flat\Omega\dot{B} - [A^\flat\Omega, A]B).$$

After integration by parts and because variations are zero at each extremity of the time interval [0, 1], the discrete Euler–Lagrange equations of Eq. (22) are

$$\forall B \in \mathcal{S}, \quad \mathrm{Tr}((\dot{A}^\flat\Omega + [A, A^\flat\Omega])B) = 0. \tag{23}$$

To express the resulting equations in a more intuitive fashion, we introduce the following lemma.

Lemma 6. *If matrix $Z \in \mathcal{M}^N$ is antisymmetric, with $\mathrm{Tr}(ZY^T) = 0$ for every $Y \in \mathcal{S}$ then there exists a discrete pressure field, i.e., a vector $P = (p_1, \ldots, p_N)$ such that*

$$Z_{ij} = p_j - p_i, \quad \text{where } j \in N(i).$$

Proof. Since $Y \in \mathcal{S}$, the inner product of matrices $\mathrm{Tr}(ZY^T)$ does not depend on Z_{ij} when i and j are not direct neighbors. We can thus assume that $Z \in \mathcal{N}$. The space $\mathcal{S}$ has codimension $N - 1$ in the space $\mathcal{N}$. Indeed, it is defined by a system of $N - 1$ independent equations:

$$\sum_{j\in N(i)} Y_{ij} = 0, \quad 1 \le i \le N - 1,$$

the last equation for $i = N$ being automatically enforced by the others.

Moreover, the space of discrete gradients (i.e., matrices $M \in \mathcal{N}$ such as $M_{ij} = p_j - p_i$) is orthogonal to the space of null-row antisymmetric matrices w.r.t. the Frobenius inner product $M \cdot Y = \mathrm{Tr}(MY^T)$ and has dimension $N - 1$. Therefore, the orthogonal complement to $\mathcal{S}$ in $\mathcal{N}$ coincides with the space of discrete gradients. □

We directly deduce our main theorem:

Theorem 1. *Consider the discrete-space/continuous-time Lagrangian on $T\mathcal{D}(\mathbb{M})$:*

$$\mathcal{L}_h(A) = \frac{1}{2}\langle\langle A^\flat, A\rangle\rangle,$$

where $A = -\dot{q}q^{-1} \in \mathcal{S}$ is a sparse, null-row, and Ω-antisymmetric matrix, $q \in \mathcal{D}(\mathbb{M})$ is a signed stochastic Ω-orthogonal matrix, and $A^\flat$ is the discrete flat operator defined in Section 3.5 applied to A. Then the Lagrange–d'Alembert principle

$$\delta \int_0^1 \mathcal{L}_h(A)\mathrm{d}t = 0, \quad \delta q \in S_q,\ A \in S, \qquad \delta q(0) = \delta q(1) = 0$$

implies

$$(\dot{A}^\flat + \mathbf{L}_A A^\flat + \mathbf{d}p)_{ij} = 0, \quad \text{for } j \in N(i), \tag{24}$$

or, equivalently,

$$\dot{A}_{ij} + \frac{|S_{ij}|}{2\Omega_i\Omega_j|e_{ij}|}[A, A^\flat\Omega]_{ij} + \frac{|S_{ij}|}{2\Omega_i|e_{ij}|}(p_j - p_i) = 0,$$
$$\text{for } j \in N(i) \tag{25}$$

where p is a discrete pressure field to enforce $A \in \mathcal{S}$.

Proof. Apply Lemma 6 (for $Z = \dot{A}^\flat + [A, A^\flat\Omega]\Omega^{-1}$ and $Y = \Omega B$) to Eq. (23) and substitute the definition of discrete Lie derivative given in Eq. (14). □

The resulting discrete Euler–Lagrange (DEL) equations we obtained represent a *weak form* of the continuous Euler equations expressed as

$$\dot{v}^\flat + \mathbf{L}_v v^\flat + \mathbf{d}p = 0.$$

Furthermore, these equations of motion can be reexpressed in various ways, mimicking different forms of Euler equations: for instance, the discrete equations of motion written as

$$(\dot{A}^\flat + i_A\mathbf{d}A^\flat + \mathbf{d}\tilde{p})_{ij} = 0$$

for all $j \in N(i)$, which are equivalent to

$$\dot{v} + v \times \omega + \nabla\tilde{p} = 0$$

(where $\tilde{p}$ is the dynamic pressure), while taking the exterior derivative of these same equations leads to

$$((\mathbf{d}\dot{A}^\flat) + \mathbf{L}_A(\mathbf{d}A^\flat))_{ij} = 0 \quad \forall j \in N(i),$$

a discrete version of

$$\dot{\omega} + \mathbf{L}_v\omega = 0.$$

4.3. Discrete Kelvin's theorem

This section presents a discrete version of Kelvin's theorem that the discrete Euler–Lagrange equations fulfill. Through replacing the continuous notions of a curve and its advection by discrete Eulerian counterparts, the proof of this discrete Kelvin's theorem will be essentially the same as in the continuous case, which we will describe first for completeness.

Kelvin's theorem: the continuous case. Kelvin's theorem states that the circulation along a closed curve stays constant as the curve is *advected* with the flow. Let γ_t be a closed curve and $C_{\gamma_t} v_t$ be the circulation of v_t along γ_t, i.e.:

$$C_{\gamma_t} v_t = \oint_{\gamma_t} v_t \cdot \mathrm{d}s.$$

Consider a divergence-free vector field γ_0^ε representing a "narrow current" of width ε flowing along γ_0, with unit flux when integrated over transversal sections of the curve. This *current* can be thought of as an ε-spreading (akin to a convolution by a smoothed Dirac function) of the tangent vector field to the immediate surroundings of the curve γ_t, forming a smoothed notion of a curve. Let γ_t^ε be the field γ_0^ε advected by the flow v_t, i.e., it satisfies

$$\dot{\gamma}_t^\varepsilon + \mathbf{L}_{v_t}\gamma_t^\varepsilon = 0. \tag{26}$$

Note that this equation encodes the notion of advection of a curve seen from a current point of view, hence without the need for a parameterization of the curve; see [2]. Then, as $\varepsilon \to 0$,

$$\langle v_t^\flat, \gamma_t^\varepsilon\rangle \to C_{\gamma_t} v_t,$$

so the pairing $\langle v_t^\flat, \gamma_t^\varepsilon\rangle$ can be considered as an approximate circulation converging to the real circulation as $\varepsilon \to 0$. We can compute its derivative:

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\langle v_t^\flat, \gamma_t^\varepsilon\rangle &= \langle \dot{v}_t^\flat, \gamma_t^\varepsilon\rangle + \langle v_t^\flat, \dot{\gamma}_t^\varepsilon\rangle \\ &= -\langle \mathbf{L}_{v_t} v_t^\flat, \gamma_t^\varepsilon\rangle - \langle v_t^\flat, \mathbf{L}_{v_t}\gamma_t^\varepsilon\rangle = 0.\end{aligned}$$

And since this pairing represents the circulation along the ε-smoothed curve for any ϵ, the circulation itself stays constant.

Remark. A current is formally the dual of a 1-form (in the sense of vector space duality), i.e., it is a linear map that takes a 1-form to $\mathbb{R}$. When the space is equipped with a metric, one can think of a current as a vector field as described above. While a metric-independent treatment is possible as well, we will stick to the vector field point of view for simplicity in this paper.

Achieving the goal of finding a discrete Kelvin's theorem first requires a definition of *discrete curves* and their advection, for which we will borrow the concept of *1-chains* used in algebraic topology and demonstrate that curves and vector fields satisfying the non-holonomic constraints share the same representation; that is, *the discretization of a curve $\gamma(s)$ will be thought of as a discretization of the narrow current γ^ε*. Since we already have established a discrete analog to the Lie derivative (based on the commutator of matrices), we will be able to define how to *advect a discrete curve along a discrete vector field*. We will find that, just like Kelvin's circulation theorem in the continuous case, for any discrete curve γ_t advected by a discrete vector field $A(t)$ satisfying the discrete Euler equations, the circulation of $A(t)$ along γ_t remains constant.

Discrete curves. A discrete curve in our Eulerian setup can be nicely defined using the concept of 1-chains. Let us recall that *dual 1-chains* are linear combinations of dual edges (linking two adjacent cells), converging to 1-manifolds as the mesh gets finer (see [11] for a thorough exposition of chains and simplicial homology, and [18] for applications in electromagnetism). In our context, in order to consider curves as "currents" (i.e., localized vector fields) as in the continuous description above, we will be using a linear combination of primal fluxes instead, exploiting the well-known isomorphism (from the Poincaré duality theorem) between dual 1-chains and primal 2-forms in 3D (i.e., between dual 1-chains and primal $(n-1)$-cochains in dimension n; see [11]). In other words, an Ω-antisymmetric matrix will be used to describe a discrete curve as it was used to describe a 2-form. We start by defining a *simple* curve:

Definition 6. A ***simple discrete curve*** is a discrete path from cell C_{i_1} to cell $C_{i_2}, \ldots,$ to cell C_{i_n} with C_{i_k} adjacent to $C_{i_{k+1}}$ and such that $(i_k, i_{k+1}) \neq (i_j, i_{j+1})$ for $k \neq j$. It is represented by an Ω-antisymmetric matrix Γ whose entries Γ_{ij} satisfy

$$\Omega_{i_k}\Gamma_{i_k i_{k+1}} = -\Gamma_{i_{k+1} i_k}\Omega_{i_{k+1}} = \frac{1}{2},$$

and

$$\Gamma_{ij} = 0, \quad \text{for } (i,j) \neq (i_k, i_{k+1})\ \forall k.$$

The matrix Γ representing a simple discrete curve $\gamma(s)$ that exactly follows dual edges can be considered as a discrete current induced by the tangent field $\mathrm{d}\gamma(s)/\mathrm{d}s$. Moreover, one can extend the notion of discrete curves to encompass arbitrary dual 1-chains. In our work, we will be focusing on *closed* discrete curves described as discrete *divergence-free* currents:

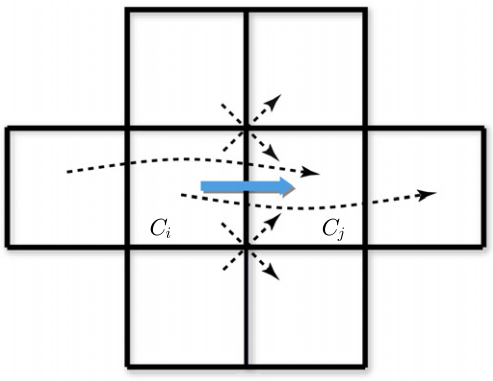

Fig. 5. Projection on Regular Grids: our projection of $[A, B]$ onto the subspace of non-holonomic constraints accumulates on the common boundary of C_i and C_j all the two-cell-away transfers going through this boundary. In this figure, the transfers in dotted lines are summed up (with a weight of $\frac{1}{2}$ for the diagonal ones) and assigned to the blue one-away transfer.

Definition 7. A ***discrete closed curve*** is a simple discrete curve that closes (i.e., a discrete path from cell C_{i_1} to cell $C_{i_2}, \ldots,$ to cell C_{i_n}, and back to cell C_{i_1}). It is represented by a null-row Ω-antisymmetric matrix Γ such that $\Gamma_{ij} = 0$ when $(i,j) \neq (i_k, i_{(k+1) \bmod n})$ for some k.

Since our discrete representation of a 1-manifold coincides with our definition of discrete Eulerian velocities in the NHC, we will no longer distinguish between discrete curves and discrete velocities.

Discrete circulation. Due to the duality between discrete curves and discrete fluxes of a vector field, the circulation of a vector field along a curve is trivially computed using the same pairing of matrices we used earlier.

Definition 8. The circulation $C_\Gamma A$ of a discrete vector field A along a discrete curve Γ is defined as

$$C_\Gamma A \equiv \langle\langle A^\flat, \Gamma\rangle\rangle.$$

We finally need to define a discrete notion of advection, which should be an approximation to $\mathbf{L}_v\gamma^\varepsilon$. We use a matrix $A \in \mathcal{S}$ to discretize v and a matrix $\Gamma \in \mathcal{S}$ to discretize γ^ε, so their commutator $[A, \Gamma]$ is a discretization of $\mathbf{L}_v\gamma^\varepsilon$. However, $[A, \Gamma] \notin \mathcal{S}$. So instead, we can only consider the elements of $[A, \Gamma]$ that satisfy the constraints to define our *weak* notion of curve advection.

Definition 9. Let $\Gamma_t \in \mathcal{S}$ be a family of discrete curves evolving in time and A_t is a (time-dependent) discrete vector field. We say that $\Gamma_t \in \mathcal{S}$ is advected by A_t if Γ_t satisfies the advection equation

$$\langle\langle X^\flat, \dot{\Gamma} + [A, \Gamma]\rangle\rangle = 0, \quad \text{for any } X \in \mathcal{S}. \tag{27}$$

Note that this definition defines a projection of the commutator $[A, \Gamma]$ onto the subspace $\mathcal{S}$ of non-holonomic constraints, and Fig. 5 depicts this projection for the case of a regular grid.

Now, let us prove that if $\dot{\Gamma}$ satisfies Eq. (27), it is a discrete (weak) approximation of Eq. (26). Indeed, if $X \rightsquigarrow w, A \rightsquigarrow v, \Gamma \rightsquigarrow \gamma$, then, by the definition of the discrete operator $\flat$,

$$\langle\langle X^\flat, \dot{\Gamma}\rangle\rangle \to \int_M (w, \dot{\gamma})$$

and

$$\langle\langle X^\flat, [A, \Gamma]\rangle\rangle \to \int_M (w, [v, \gamma]).$$

Thus, if Eq. (27) is satisfied, $\dot{\gamma}$ has to satisfy

$$\int_M (w, \dot{\gamma} + [v, \gamma]) = 0$$

for every $w \in \mathrm{SVect}(M)$. Since $\gamma \in \mathrm{SVect}(M)$, this last equation is a weak form of $\dot{\gamma} = -[v, \gamma] = -\mathbf{L}_v \gamma$.

Discrete Kelvin's theorem. We are now ready to give a discrete analog of Kelvin's circulation theorem satisfied by our discrete Euler equations.

Theorem 2. *If A_t satisfies the DEL equations (24) and Γ_0 is an arbitrary discrete curve, then the circulation of A along Γ_t stays constant:*

$$C_{\Gamma_t} A_t = C_{\Gamma_0} A_0,$$

where Γ_t is the curve Γ_0 advected by A_t.

Proof. The time derivative of the circulation $C_{\Gamma_t} A_t$ is expressed as

$$\frac{\mathrm{d}}{\mathrm{d}t} C_{\Gamma_t} A_t = \frac{\mathrm{d}}{\mathrm{d}t} \langle\langle A_t^\flat, \Gamma_t \rangle\rangle = \langle\langle \dot{A}_t^\flat, \Gamma_t \rangle\rangle + \langle\langle A_t^\flat, \dot{\Gamma}_t \rangle\rangle.$$

Since $A_t^\flat$ satisfies the DEL equations $(\dot{A}_t^\flat + [A_t^\flat \Omega, A_t]\Omega^{-1} + \mathbf{d}\bar{p}_t)_{ij} = 0$ for i and j representing two neighboring cells' indices, and as $\mathrm{Tr}((\mathbf{d}\bar{p}_t)\Gamma_t) = 0$, we have

$$\begin{aligned} \langle\langle \dot{A}_t^\flat, \Gamma_t \rangle\rangle &= -\langle\langle [A_t^\flat \Omega, A_t]\Omega^{-1}, \Gamma_t \rangle\rangle = -\mathrm{Tr}([A_t^\flat \Omega, A_t]\Gamma_t) \\ &= -\mathrm{Tr}(A_t^\flat \Omega [A_t, \Gamma_t]) = -\langle\langle A_t^\flat, [A_t, \Gamma_t] \rangle\rangle. \end{aligned}$$

But since Γ_t is advected by A, we get

$$\frac{\mathrm{d}}{\mathrm{d}t} C_{\Gamma_t} A_t = \langle\langle \dot{A}_t^\flat, \Gamma_t \rangle\rangle + \langle\langle A_t^\flat, \dot{\Gamma}_t \rangle\rangle = \langle\langle A_t^\flat, \dot{\Gamma}_t + [A_t, \Gamma_t] \rangle\rangle = 0. \quad \square$$

Remark. In the continuous case, Kelvin's circulation theorem can be derived from Noether's theorem using right-invariance of the metric on SDiff (particle relabeling symmetry). In the discrete case, the Lagrangian is also right invariant, but the presence of the non-holonomic constraints prevents us from using Noether's theorem directly: in a system with non-holonomic constraints a momentum is no longer expected to be conserved in general. However, we can still use the symmetry to obtain the momentum equation, i.e. the rate of change of the momentum in time. Doing so for our discrete fluid model we also get our discrete circulation theorem.

5. Fluid evolution in discrete time

In this section, we revisit our discrete version of the variational principle discussed above by making time discrete instead of continuous. We assume that the fully discrete fluid motion is given as a discrete path $q_0, q_1, \ldots, q_K$ in the space of Ω-orthogonal signed stochastic matrices, where the motion has been sampled at regular time $t_k = k\tau$ for $k \in \{0, 1, \ldots, K\}$, τ being referred to as the time step size.

5.1. Discrete velocity

Given a pair q_k, q_{k+1} of consecutive configurations in time, we can compute a discrete-time analog of Eulerian velocity using, e.g., one of the following classical formulas:

- $q_{k+1} = q_k - \tau A_k q_k$, (explicit Euler)
- $q_{k+1} = q_k - \tau A_k q_{k+1}$, (implicit Euler)
- $q_{k+1} = q_k - \tau A_k (q_k + q_{k+1})/2$ (midpoint rule).

Note that the midpoint rule preserves the Lie group structure of the configuration space. Note also that many other discretizations could be used, but we will restrict our explanations to the first two cases as they suffice to illustrate how our continuous-time procedure can be adapted to the purely discrete case.

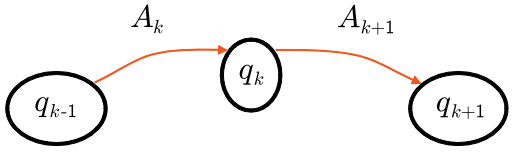

Fig. 6. Three consecutive configurations q_{k-1}, q_k, q_{k+1} of a fluid in time, with Eulerian velocities A_k and A_{k+1} in between.

5.2. Discrete Lagrangian and action

We define the discrete-space/discrete-time Lagrangian $\mathcal{L}_d(q_k, q_{k+1})$ as

$$\mathcal{L}_d(q_k, q_{k+1}) = \mathcal{L}_h(A_k).$$

The discrete action $\mathcal{A}_d$ along a discrete path is then simply the sum of all pairwise discrete Lagrangians:

$$\mathcal{A}_d(q_0, \ldots, q_K) = \sum_{k=0}^{K-1} \mathcal{L}_d(q_k, q_{k+1}).$$

We can now use the Lagrange–d'Alembert principle that states that $\delta \mathcal{A}_d = 0$ for all variations of the q_k (for $k = 1, \ldots, K-1$, with q_0 and q_K being fixed) in S_q while A_k is restricted to $\mathcal{S}$ (Fig. 6).

Variations. The variations of A_k can be easily derived.

- *Explicit Euler.* In this case, $A_k = -(q_{k+1} - q_k)/\tau q_k^{-1}$. The variation $\delta_k A_k$ and $\delta_{k+1} A_k$ with respect to q_k and q_{k+1} respectively become

  $$\delta_k A_k = \frac{1}{\tau} \delta q_k q_k^{-1} + \frac{q_{k+1} - q_k}{\tau} q_k^{-1} \delta q_k q_k^{-1},$$

  $$\delta_{k+1} A_k = -\frac{1}{\tau} \delta q_{k+1} q_k^{-1}.$$

  If we denote, similar to the continuous case, $B_k = -\delta q_k q_k^{-1}$, we get

  $$\delta_k A_k = -\frac{B_k}{\tau} + A_k B_k$$

  and

  $$\delta_{k+1} A_k = \frac{B_{k+1}}{\tau} - B_{k+1} A_k.$$

- *Implicit Euler.* In this case $A_k = -(q_{k+1} - q_k)/\tau q_{k+1}^{-1}$. It yields

  $$\delta_k A_k = \frac{1}{\tau} \delta q_k q_{k+1}^{-1}$$

  and

  $$\delta_{k+1} A_k = -\frac{1}{\tau} \delta q_{k+1} q_{k+1}^{-1} + \frac{q_{k+1} - q_k}{\tau} q_{k+1}^{-1} \delta q_{k+1} q_{k+1}^{-1}.$$

  Similarly to the previous case we now obtain

  $$\delta_k A_k = -\frac{B_k}{\tau} - B_k A_k,$$

  and

  $$\delta_{k+1} A_k = \frac{B_{k+1}}{\tau} + A_k B_{k+1}.$$

5.3. Discrete Euler–Lagrange equations

Equating the variations of the action $\mathcal{A}_d$ with respect to δq_k to zero for $k \in [1, K-1]$ yields

$$\delta_k \langle\langle A_{k-1}^\flat, A_{k-1} \rangle\rangle + \delta_k \langle\langle A_k^\flat, A_k \rangle\rangle = 0. \tag{28}$$

Thus we obtain

$$\mathrm{Tr}[A_{k-1}^\flat \Omega (\delta_k A_{k-1}) + A_k^\flat \Omega (\delta_k A_k)] = 0.$$

Now, let us solve it for A_k in the explicit Euler case. Substituting the expressions for $\delta_k A_k$ and $\delta_k A_{k-1}$ yields

$$\mathrm{Tr}[A^\flat_{k-1}\Omega(B_k - \tau B_k A_{k-1}) + A^\flat_k \Omega(-B_k + \tau A_k B_k)] = 0.$$

Denoting $\dot{A}^\flat_k = (A^\flat_k - A^\flat_{k-1})/\tau$ we can rewrite the last equation as

$$\mathrm{Tr}[(\dot{A}^\flat_k \Omega + A_{k-1}A^\flat_{k-1}\Omega - A^\flat_k \Omega A_k)B_k] = 0.$$

Therefore, we get the following discrete Euler–Lagrange equations in the explicit Euler case:

$$\dot{A}^\flat_k - \left(\frac{A^\flat_k \Omega A_k \Omega^{-1} - A_{k-1}A^\flat_{k-1}}{2} - \frac{(A^\flat_k \Omega A_k \Omega^{-1} - A_{k-1}A^\flat_{k-1})^T}{2}\right) + \mathbf{d}p_k = 0.$$

As $(A^\flat_k \Omega A_k \Omega^{-1})^T = A_k A^\flat_k$ and $(A_{k-1}A^\flat_{k-1})^T = A^\flat_{k-1}\Omega A_{k-1}\Omega^{-1}$, this last expression is equivalent to

$$\dot{A}^\flat_k + \frac{1}{2}([A_{k-1}, A^\flat_{k-1}\Omega]\Omega^{-1} + [A_k, A^\flat_k \Omega]\Omega^{-1}) + \mathbf{d}p_k = 0, \qquad (29)$$

corresponding to the discrete-time version of Eq. (23).

Using the implicit Euler formula for A_k instead of the explicit Euler one leads to the exact same equation, we thus omit the computations here.

5.4. Update rule for regular grids in 2D

The discrete Euler equation we derived above turns out to be particularly simple when applied to a regular grid. Indeed, let us consider a regular grid of size h, on a 2D domain and with continuous time for simplicity. Then the discrete Euler equation (25) becomes

$$2h^2\dot{A}_{ij} + [A, A^\flat]_{ij} + (p_j - p_i) = 0, \quad \text{for } j \in N(i).$$

Now let us fix i and j and expand $[A, A^\flat]_{ij}$. Since $A \in \mathcal{S}$ we have

$$[A, A^\flat]_{ij} = \sum_{l\in N(i)} A_{il}A^\flat_{lj} - \sum_{k\in N(j)} A^\flat_{ik}A_{kj}.$$

From the definition of $A^\flat$ (see Lemma 5) we get2

$$A^\flat_{ik} = -\frac{1}{2}\omega_{ik}s_{ijk} + 2h^2(A_{ij} + A_{jk}), \quad \text{for } k \in N(j) \text{ and } k \notin N(i)$$

and

$$A^\flat_{lj} = -\frac{1}{2}\omega_{lj}s_{jil} + 2h^2(A_{ij} + A_{li}), \quad \text{for } l \in N(i) \text{ and } l \notin N(j),$$

where $\omega_{i_1 i_2}$ is the vorticity in the DEC sense computed at the common node of cells i_1 and i_2 if i_1 and i_2 have a common node (see Fig. 4), and 0 otherwise; also $s_{i_1 i_2 i_3} = 1$ if the triplet of cells i_1, i_2, i_3 is oriented counter-clockwise and $s_{i_1 i_2 i_3} = -1$ otherwise.

Now the equations for the commutator $[A, A^\flat]$ become

$$[A, A^\flat]_{ij} = \frac{1}{2}\sum_{k\in N(j)} A_{kj}\omega_{ik}s_{ikj} - \frac{1}{2}\sum_{l\in N(i)} A_{il}\omega_{lj}s_{jil} + 2h^2\sum_{k\in N(j)} A^2_{jk} - 2h^2\sum_{l\in N(i)} A^2_{il}.$$

If $k \in N(l)$ then $\omega_{ik} = \omega_{lj}$, so only two ω's are present in the expression above. Let us denote them by ω_- and ω_+ as depicted in Fig. 7, and write

$$[A, A^\flat]_{ij} = -\omega_- \frac{A_{jk_1} + A_{il_1}}{2} - \omega_+ \frac{A_{k_2 j} + A_{l_2 i}}{2} + Q_j - Q_i,$$

where $Q_i = 2h^2\sum_{l\in N(i)} A^2_{il}$.

As we know, ω_-/h^2 and ω_+/h^2 approximate the values $\omega(x_-)$, $\omega(x_+)$ of vorticity at the corresponding nodes. Also, $A_{ij} \approx -v_{ij}/2h$.

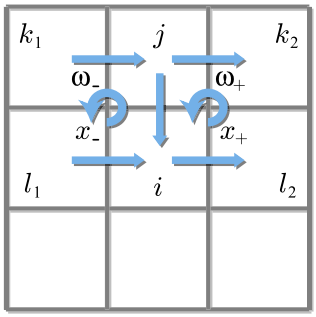

Fig. 7. Notations used to rewrite the discrete Euler equation on a regular mesh as a function of local velocities and vorticities.

Now suppose the pair of cells C_i and C_j is oriented along the y direction (see Fig. 7 again) and $v = (v_1, v_2)$. Let us denote $A_{ij} \approx -v_2/2h$ and

$$2hA_{ik_1} = -v_1^{--} \qquad 2hA_{k_2 i} = -v_1^{+-}$$
$$2hA_{jk_1} = -v_1^{-+} \qquad 2hA_{k_2 j} = -v_1^{++}.$$

Now, the discrete Euler equation implies

$$\dot{v}_2 + \frac{1}{4}(\omega(x_-)(v_1^{--} + v_1^{-+}) + \omega(x_+)(v_1^{+-} + v_1^{++})) + P_j - P_i = 0,$$

where P is some discrete function, playing the role of pressure. This equation, together with the equations for every pair i and j, is a discrete version of the 2D Euler equations written in the form

$$\begin{aligned} &\dot{v}_1 - \omega v_2 + P_x = 0,\\ &\dot{v}_2 + \omega v_1 + P_y = 0,\end{aligned} \quad \mathrm{div} v = 0, \ \omega = \frac{\partial v_2}{\partial x} - \frac{\partial v_1}{\partial y}.$$

The discretization of the Euler equations that we have obtained on the regular grid coincides with the Harlow–Welsh scheme (see [21]), and Eq. (29) is a Crank–Nicolson (trapezoidal) time update. Therefore, our variational scheme can be seen as an extension of this approach to arbitrary grids, offering the added bonus of providing a geometric picture to these numerical update rules.

6. Conclusions and discussions

The discrete geometric derivation of Euler equations we presented above differs sharply from previous geometric approaches. First and foremost, our work derives the fluid mechanics equations from the least action principle, while many previous techniques are based on finite volume, finite difference, or finite element methods applied to Euler equations (see [21–23] and references therein). Second, our derivation does not presume or design a Lie derivative or a Poisson bracket in the manner of geometric approaches such as [24]. Finally, we discretize the volume-preserving diffeomorphism group, offering a purely Eulerian alternative to the inverse map approach proposed in [25,26]. The resulting scheme does, however, have most of the numerical properties sought after, including energy conservation over long simulations [22], time reversibility [27], and circulation preservation [17]. We now go over some of the computational details and present a few results, before discussing possible extensions.

6.1. Computational details

Implementing the discrete update rules derived in Section 5.3 is straightforward: from the sparse matrix $A_k \in \mathcal{S}$ describing the velocity field at time t_k, a new sparse matrix $A_{k+1} \in \mathcal{S}$ is computed by solving the discrete Euler–Lagrange equations through repeated Newton steps until convergence. Notice that the configuration state q_k is not needed in the computations, making

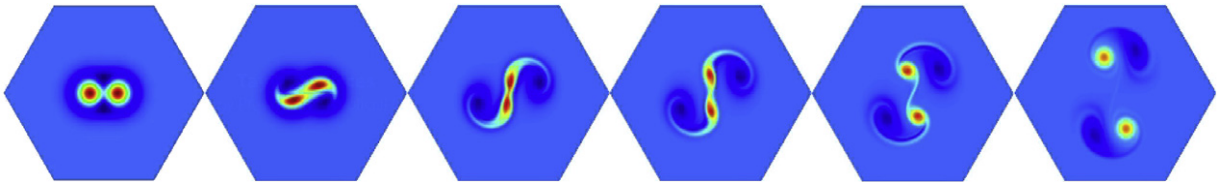

Fig. 8. Taylor Vortices Separating: two like-signed Taylor vortices (with a finite vorticity core) will merge when their distance of separation is smaller than some critical value. In this example, the two vortices in a domain discretized with 55296 triangles were initialized at a distance slightly above this critical value, leading to a separation.

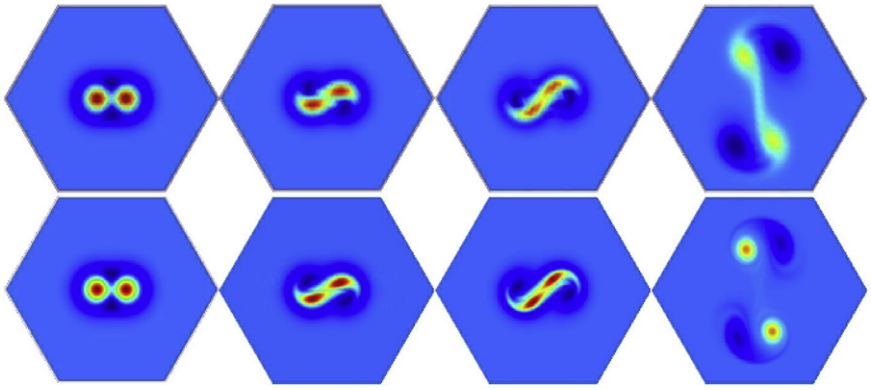

Fig. 9. 2D fluid simulation of Taylor vortices separating: even at two very different resolutions, our variational scheme leads to similar results (top: 4056 triangles; bottom: 55296 triangles; same continuous initial conditions discretized on both grids).

the numerical scheme capable of directly computing A_{k+1} from A_k. Moreover, the update rules can be rewritten entirely as a function of fluxes $F_{ij} = 2\Omega_i A_{ij}$, rendering the assembly of the advection operator simple. Finally, the Newton steps can also be made more efficient by approximating the Jacobian matrix involved in solving by only its diagonal terms. Further details on the computational procedure can be found in [20], including linearizations of the discrete Euler–Lagrange equations that tie our method with [28]. It is also worth mentioning that our simulator applies to arbitrary topology, and that viscosity is easily added by incorporating a term proportional to the Laplacian of the velocity field [20].

6.2. Numerical tests and results

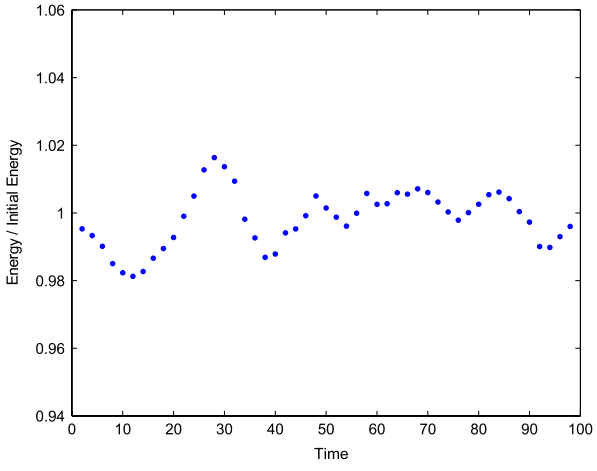

As expected from the time reversibility of the resulting discrete Euler–Lagrange equations, our fully Eulerian scheme demonstrates excellent energy behavior over long simulations, even for very low thresholds on the Newton solver. This numerical property was well known for the Harlow–Welsh discretization over regular grids when using a trapezoidal time integration scheme; our approach extends this scheme and its properties to arbitrary mesh discretization. Fig. 8 shows the results of this geometric integrator on a common test used in CFD, where a periodic 2D domain is initialized with two Taylor vortex distributions of same sign placed at a distance close to a critical bifurcation in the dynamics: as expected, the two vortices eventually separates, and our integrator keeps the energy close to the initial energy over extended simulation time (see inset). Fig. 9 demonstrates the robustness of the integrator to grid size: the same dynamics of the vortices is still captured even on a number of triangles thirteen times smaller. Finally, Fig. 1 shows frames of a simulation of a 3D fluid on a tetrahedral mesh.

6.3. Extensions

The results of this paper are rich in possible extensions. For instance, convergence and stability of our resulting schemes should be studied. Generalizing our approach to higher-order integrators is another obvious research direction; a midpoint approximation of the Eulerian velocity between q_k and q_{k+1} preserves the Lie group structure of the configuration space, but leads to additional cubic terms in the variation δA, thus requiring a flat operator valid for three-away cells as well. Finding a systematic approach to deriving such higher-order updates is the subject of future work.

We could also investigate alternative expressions for the discrete Lagrangian. One possibility is to notice that in the continuous case, the Lagrangian can be written as

$$\mathcal{L}(q, \dot{q}) = \frac{1}{2}\sum_{i=1}^{n}(\mathbf{L}_v x_i, \mathbf{L}_v x_i)$$

where x_i represents the ith coordinates in $\mathbb{R}^n$. Its discrete equivalent in 2D could therefore be written as $\mathrm{Tr}(A^2(X+Y))$ where the matrix $X \in \mathcal{S}$ (resp., $Y \in \mathcal{S}0$) is expressed as $X_{ij} = x_{C_i} \cdot x_{C_j}$ (resp., $Y_{ij} = y_{C_i} \cdot y_{C_j}$), with x_{C_k} (resp., y_{C_k}) represents the x-coordinate (resp., y-coordinate) of the circumcenter of cell C_k. Taking variation would lead to $\mathrm{Tr}(A\delta A(X+Y)) = \mathrm{Tr}(((X+Y)\dot{A} - [(X+Y)A, A])B) = 0\ \forall B \in \mathcal{S}$. We see that this alternate definition of the Lagrangian defines another flat operator (albeit, in a less geometric way).

Similarly, one may change the sparsity requirement of the NHC by defining the space $\mathcal{S}$ to be the sparsity induced by adjacency

through *vertices*. It would require a new Lie bracket which is not directly the Lie bracket for the matrices representing the vector fields, but the sparsity constraint would no longer be non-holonomic.

It would be a natural extension to look at how the energy of our discrete simulator cascades at lower scales. More generally, understanding what this geometric picture of fluid flows brings compared to traditional Large Eddy Simulation or Reynolds-Averaged Navier–Stokes methods would be interesting, as our structure-preserving approach is also based on local averages (or, equivalently, integrated values) of the velocity field.

We also wish to investigate the use of an "upwind" version of the discrete velocity A, possessing only positive fluxes as often used in the discretization of hyperbolic partial differential equations [29]. This trivial modification of the velocity matrix would allow the reconstruction, via exponentiation, of non-negative transport matrices T_k for any given discrete path $\{q_k\}$. These resulting matrices could then be formally considered as transition matrices of a Markov chain. However, this positive discretization breaks the structure leveraged in our approach as mentioned in Section 2, and we do not know at present how to derive a structure-preserving integrator *directly* with this discretization.

Finally, we note that the geometric understanding developed here should offer good foundations to tackle related problems, such as magnetohydrodynamics, variable density fluids, or Burgers' equations. Our initial results using an extension to systems with semi-direct product group structure show promise.

Acknowledgements

We thank Daryl Holm and Yann Brénier for helpful early discussions and input, Evan Gawlik for generating the energy plots, and Keenan Crane for generating 2D tests. This research was partially supported by NSF grants CMMI-0757106, CCF-0811373 & 1011944, IIS-0953096, and DMS-0453145.

References

[1] T.Y. Hou, Z. Lei, On the stabilizing effect of convection in 3D incompressible flow, Commun. Pure Appl. Math. 62 (4) (2009) 501–564.
[2] V.I. Arnold, Sur la géométrie différentielle des groupes de Lie de dimenson infinie et ses applications à l'hydrodynamique des fluides parfaits, Ann. Inst. Fourier, Grenoble 16 (1966) 319–361.
[3] J.E. Marsden, J. Scheurle, The reduced Euler–Lagrange equations, Fields Institute Comm. 1 (1993) 139–164.
[4] D.G. Ebin, J.E. Marsden, Groups of diffeomorphisms and the motion of an incompressible fluid, Ann. of Math. 92 (1970) 102–163.
[5] V.I. Arnold, B. Khesin, Topological methods in hydrodynamics, Ann. Rev. Fluid Mech. 24 (1992) 145–166.
[6] J.E. Marsden, T.S. Ratiu, Introduction to Mechanics and Symmetry, in: Texts in Applied Mathematics, vol. 17, Springer-Verlag, 1994. Second edition, 1999.
[7] A. Hirani, Discrete Exterior Calculus, Ph.D. Thesis, California Institute of Technology, 2003.
[8] J. Tournois, C. Wormser, P. Alliez, M. Desbrun, Interleaving delaunay refinement and optimization for practical isotropic tetrahedron mesh generation, ACM Trans. Graph. 28 (3) (2009) Art. 75.
[9] B.O. Koopman, Hamiltonian systems and transformations in Hilbert spaces, Proc. Nat. Acad. Sci. (USA) 17 (1931) 315–318.
[10] Y. Brenier, Polar factorization and monotone rearrangement of vector-valued functions, Comm. Pure Appl. Math. 44 (4) (1991) 375–417.
[11] J.R. Munkres, Elements of Algebraic Topology, Addison-Wesley, Menlo Park, CA, 1984.
[12] R. Abraham, J.E. Marsden, T. Ratiu, Manifolds, Tensor Analysis, and Applications, in: Applied Mathematical Sciences, vol. 75, Springer, 1988.
[13] M. Desbrun, E. Kanso, Y. Tong, Discrete differential forms for computational modeling, Chapter in ACM SIGGRAPH Course Notes on Discrete Differential Geometry (2005).
[14] P.B. Bochev, J.M. Hyman, Principles of mimetic discretizations of differential operators, 2005, Preprint LA-UR-05- 4244.
[15] D.N. Arnold, R.S. Falk, R. Winther, Finite element exterior calculus, homological techniques, and applications, Acta Numer. 15 (2006) 1–155.
[16] J.E. Marsden, M. West, Discrete mechanics and variational integrators, Acta Numer. 10 (2001) 357–515.
[17] S. Elcott, Y. Tong, E. Kanso, P. Schröder, M. Desbrun, Stable, circulation-preserving simplicial fluids, ACM Trans. Graph. 26 (1) (2007) Art. 4.
[18] A. Bossavit, Computational Electromagnetism, Academic Press, Boston, 1998.
[19] S. Sen, S. Sen, J.C. Sexton, D.H. Adams, Geometric discretization scheme applied to the abelian Chern–Simons theory, Phys. Rev. E 61 (3) (2000) 3174–3185.
[20] P. Mullen, K. Crane, D. Pavlov, Y. Tong, M. Desbrun, Energy-preserving integrators for fluid simulation, ACM Trans. Graph. 28 (3) (2009) Art. 38.
[21] F.H. Harlow, J.E. Welch, Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface, Phys. Fluids 8 (12) (1965) 2182–2189.
[22] B. Perot, Conservation properties of unstructured staggered mesh schemes, J. Comput. Phys. 159 1 (2000) 58–89.
[23] P.M. Gresho, R.L. Sani, Incompressible Flow and the Finite Element Method, J. Wiley & Sons, 2000.
[24] R. Salmon, Poisson–Bracket approach to the construction of energy and potential-enstrophy conserving algorithms for the shallow-water equations, J. Atmospheric Sci. 61 (2004) 2016–2036.
[25] C.J. Cotter, D.D. Holm, P.E. Hydon, Multisymplectic formulation of fluid dynamics using the inverse map, Proc. Roy. Soc. A 463 (2007) 2671–2688.
[26] C.J. Cotter, D.D. Holm, Continuous and discrete Clebsch variational principles, Found. Comput. Math. 9 (2) (2009) 221–242.
[27] M. Duponcheel, P. Orlandi, G. Winckelmans, Time-reversibility of the Euler equations as a benchmark for energy conserving schemes, J. Comput. Phys. 227 19 (2008) 8736–8752.
[28] J.C. Simo, F. Armero, Unconditional stability and long-term behavior of transient algorithms for the incompressible Navier–Stokes and Euler equations, Comput. Methods Appl. Mech. Eng. 111 (2000) 111–154.
[29] R.J. LeVeque, Finite Volume Methods for Hyperbolic Problems, Cambridge University Press, 2002.