



Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

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

Explicit, parallel Poisson integration of point vortices on the sphere

Keith W. Myerscough^{a,*}, Jason Frank^b^a *Centrum Wiskunde & Informatica, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands*^b *Mathematical Institute, Utrecht University, P.O. Box 80010, 3508 TA Utrecht, The Netherlands*

ARTICLE INFO

Article history:

Received 25 March 2015

Received in revised form 29 February 2016

Keywords:

Point vortex method
 Numerical integration
 Poisson integrator
 Parallel computing

ABSTRACT

Point vortex models are frequently encountered in conceptual studies in geophysical fluid dynamics, but also in practical applications, for instance, in aeronautics. In spherical geometry, the motion of vortex centres is governed by a dynamical system with a known Poisson structure. We construct Poisson integration methods for these dynamics by splitting the Hamiltonian into its constituent vortex pair terms. From backward error analysis, the method is formally known to provide solutions to a modified Poisson system with the correct bracket, but with a modified Hamiltonian function. Different orderings of the pairwise interactions are considered and also used for the construction of higher order methods. The energy and momentum conservation of the splitting schemes is demonstrated for several test cases. For particular orderings of the pairwise interactions, the schemes allow scalable parallelization. This results in a linear – as opposed to quadratic – scaling of computation time with system size when scaling the number of processors accordingly.

© 2016 Elsevier B.V. All rights reserved.

1. Motivation

A point vortex represents a singular measure solution to the vorticity equation for two-dimensional, incompressible fluid flow. A point vortex model consists of multiple point vortices mutually interacting. The motion of each point vortex is dictated by the flow field induced by the other vortices and by external forcing, e.g. topography. Point vortices were introduced by Helmholtz [1] and have since been the subject of much study; see for example [2–4].

Dynamical studies of point vortex systems provide insight into the (qualitative) behaviour of fluid dynamics. The series of papers by Newton et al. [4–7] discuss relative equilibria and the conditions for integrability of the dynamics. Vortex dynamics were studied extensively by Aref who compiled an extensive review on their history [8]. Newton [9] discusses the future of point vortex research in the “post-Aref era”.

In statistical fluid mechanics, the behaviour of point vortex systems has been studied as a model for two-dimensional turbulence in the limit of an infinite number of vortices. This was first done by Onsager [10], who provided an explanation for the formation of clusters of like-signed vortices in a bounded domain. This research has since been continued by, amongst others, Joyce and Montgomery [11,12], Pointin and Lundgren [13], Eyink and Spohn [14], and Lions and Majda [15]. Such results are of interest in the fields of geophysical fluid dynamics [16] and stellar dynamics [17]. Some of Onsager’s statements were tested numerically by Bühler [18].

* Corresponding author.

E-mail addresses: k.w.myerscough@cw.nl (K.W. Myerscough), j.e.frank@uu.nl (J. Frank).

Point vortices and their three-dimensional generalization, vortex filament methods, are also used as a discretization of practical fluid flows in engineering applications [19]. By using a large number of point vortices a continuous velocity field is approximated. Such techniques find practical application in the works of Chatelain et al. [20], Rossinelli et al. [21], Winckelmans et al. [22] and Rossinelli and Koumoutsakos [23] present the fast multipole, vortex-in-cell and hybrid methods that are used for computing these large systems. Sakajo [24] has developed a fast tree-code reducing the $\mathcal{O}(N^2)$ computational cost to $\mathcal{O}(N \log^3 N)$ for N interacting point vortices. Regularized approximations to the delta distributions provide more accurate representations of continuous vorticity fields, but their solutions are no longer exact, as the kernel itself ought to deform due to shearing [25,26].

It is important to develop efficient time integrators for point vortex methods for two reasons. First, the use of very large numbers of point vortices, as required for accurate approximation of continuous fluids, is hampered by the quadratic complexity of the pairwise coupling between vortices, i.e. evaluations of the vector field with N vortices requires N^2 operations. Second, the concept of numerical stability of a system of point vortices on planar geometry is not without ambiguity. Equilibria only exist for certain configurations, and are never asymptotically stable since the dynamics are Hamiltonian. The simplest nontrivial system is a pair of like-signed vortices, whose solution is periodic. If a contracting method such as backward Euler is employed, the vortices will eventually approach one another, and the derivatives grow unbounded. If an expanding method such as forward Euler is employed, the vortices will drift apart and the trajectories will grow without bounded. Hence, even for this simple configuration some degree of energy conservation is necessary to maintain a bounded solution with bounded derivative.

Recently, Vankerschaver and Leok [26] have developed a Poisson integrator for point vortex systems via the construction of a higher dimensional linear Lagrangian. The associated dynamics project down onto solutions of the point vortex equations on the sphere. The resulting integrator exactly conserves the Casimirs and momentum of the point vortex dynamics and also has good conservation of energy. The implicit definition, however, requires the use of an iterative solver.

We give an interpretation of the point vortex method in light of the approach first communicated by McLachlan [27] for discretizing Hamiltonian PDEs; namely as a scheme that discretizes the Poisson structure and Hamiltonian separately. With a vorticity field given as a sum of point vortices, the quadrature of the Hamiltonian functional is evaluated exactly as a sum of pointwise values. We do not consider regularizations of the vortices, but they could be accommodated in the quadrature scheme for the Hamiltonian. The Poisson bracket is discretized exactly for a particular class of functionals.

A numerical integrator for these dynamics follows from splitting the Hamiltonian into its constituent pairwise terms. The scheme developed is Poisson, explicit and allows scalable parallelization. It may also be applied to regularized point vortices, provided the kernel is rotation- and translation-invariant. The method requires an explicit expression for the pairwise flow map for the two-vortex system. Any regularization that maintains a pairwise Hamiltonian form will have three Poisson-commuting first integrals and is thus integrable. Both rotation of the sphere and topography introduce only decoupled, splittable terms in the Hamiltonian.

The remainder of this paper is organized as follows. Section 2 describes two-dimensional incompressible fluid flow in Hamiltonian form. Section 3 discusses the discretization according the ideas of McLachlan [27]. A Poisson integrator for the resulting point vortex description for fluids is developed in Section 4. The parallelization of this method is discussed in Section 5. Numerical results and comparisons of computation times are presented in Sections 4 and 5, respectively. Finally, in Section 6 we state conclusions and discuss the extension of the method to practical applications.

2. Continuous Hamiltonian description

The barotropic quasi-geostrophic equations on the unit sphere provide a simple model for studying geophysical fluid dynamics [28]. Point vortex representations capture much of the system's dynamics, for instance the formation of coherent vortical structures over long time [10]. This is a consequence of the existence of negative temperature states, that are possible due to the bounded domain. On a disk or on an annulus, the same behaviour can be observed, but these geometries require the inclusion of ghost vortices to maintain the boundary conditions. The boundedness of the domain also implies that solutions remain bounded for almost any initial condition when considering heterogeneous systems, i.e. systems with both positive and negative circulation vortices.

We express the barotropic quasi-geostrophic equations on the sphere [28] in terms of the stream function ψ and potential vorticity q

$$q_t + J(\psi, q) = 0 \quad (1)$$

$$q = \Delta_S \psi + 2\Omega z + h, \quad (2)$$

where Ω is the angular velocity of the sphere about the z -axis and h represents topography. The Laplace–Beltrami operator on the sphere Δ_S is defined (in spherical coordinates) as

$$\Delta_S \psi = \frac{1}{\cos \theta} \left[\frac{1}{\cos \theta} \psi_{\phi\phi} + \frac{\partial}{\partial \theta} (\cos \theta \psi_\theta) \right],$$

where ϕ is the longitude and θ the latitude. The Jacobian $J(f, g)$ is defined as

$$J(f, g) = \frac{1}{\cos \theta} (f_\phi g_\theta - g_\phi f_\theta). \quad (3)$$

On the sphere the Hamiltonian is given by:

$$\mathcal{H} = -\frac{1}{2} \int_{\mathbb{S}^2} \psi \Delta_S \psi \, dS = \frac{1}{2} \int_{\mathbb{S}^2} \nabla_S \psi \cdot \nabla_S \psi \, dS,$$

where the second equality follows from the divergence theorem. Using the rightmost expression we find the first variation of \mathcal{H}

$$\begin{aligned} \delta \mathcal{H} &= \int_{\mathbb{S}^2} \nabla_S \psi \cdot \nabla_S \delta \psi \, dS \\ &= - \int_{\mathbb{S}^2} \psi \Delta_S \delta \psi \, dS \\ &= - \int_{\mathbb{S}^2} \psi \delta (q - 2\Omega z - h) \, dS, \end{aligned}$$

and consequently

$$\frac{\delta \mathcal{H}}{\delta q} = -\psi.$$

The Poisson bracket is given by

$$\{\mathcal{F}, \mathcal{G}\} [q] = - \int_{\mathbb{S}^2} \frac{\delta \mathcal{F}}{\delta q} J \left(q, \frac{\delta \mathcal{G}}{\delta q} \right) \, dS,$$

where the Jacobian J is given by (3).

Point vortex systems are derived from Eqs. (1)–(2), assuming a potential vorticity field that can be expressed as the sum of Dirac-delta distributions, i.e.

$$q(\mathbf{x}) = \sum_{i=1}^N \Gamma_i \delta(\mathbf{x} - \mathbf{x}_i(t)).$$

We have chosen to represent the vortex centres as vectors \mathbf{x}_i embedded in \mathbb{R}^3 . This is not necessary for the development of the Poisson integrator, but it results in less cumbersome equations of motion and a clearer expression for the exact solution of the two-vortex problem. This does come at the expense of a larger memory requirement and the requirement that the vortex position remains exactly on the sphere, i.e. that $|\mathbf{x}_i(t)| \equiv 1, \forall t$. We will observe that for the method developed in this article this second requirement is satisfied by construction, without the need for a projection step.

We introduce two new streamfunctions ψ_C and ψ_h such that $2\Omega z = \Delta_S \psi_C$ and $h = \Delta_S \psi_h$. With these we rewrite (2) as

$$\Delta_S \psi = \omega = q - 2\Omega z - h = q - \Delta_S \psi_C - \Delta_S \psi_h.$$

We solve this for ψ

$$\psi = \Delta_S^{-1} q - \psi_C - \psi_h, \tag{4}$$

where $\Delta_S^{-1} q = \sum_{i=1}^N \Gamma_i G(\mathbf{x} - \mathbf{x}_i(t))$ represents the sum of Green’s functions for the Laplace equation on the sphere, given by

$$G(\mathbf{x} - \mathbf{x}_i(t)) = \frac{1}{4\pi} \ln (|\mathbf{x} - \mathbf{x}_i(t)|^2). \tag{5}$$

3. Discrete Hamiltonian representation

In this section we review the point vortex description on the sphere by interpreting it as a Hamiltonian discretization in the sense of McLachlan [27]. By discretizing the Hamiltonian and Poisson bracket individually, and ensuring that the latter defines a finite dimensional Poisson bracket, it is guaranteed that the finite dimensional approximation is again Poisson, and Poisson integrators may be employed. For point vortices in planar geometry, the bracket is canonical and hence symplectic Runge–Kutta methods are applicable. On the sphere, the Poisson bracket is nontrivial and splitting methods offer the most generic approach.

With the assumption that the vorticity field is a sum of Dirac delta distributions, the integration of \mathcal{H} reduces to a sum over the values of the integrand at the vortex centres

$$\mathcal{H} = -\frac{1}{2} \int_{\mathbb{S}^2} \psi q \, dS = -\frac{1}{2} \sum_{i=1}^N \Gamma_i \psi(\mathbf{x}_i).$$

Substituting the inverse Laplacian of (4) with Green’s function (5), we find

$$\mathcal{H} = \frac{1}{2} \sum_{i=1}^N \Gamma_i \left(\psi_C(\mathbf{x}_i) + \psi_h(\mathbf{x}_i) - 2 \sum_{j<i} \Gamma_j \frac{1}{4\pi} \ln (|\mathbf{x}_j - \mathbf{x}_i|^2) \right) =: H. \tag{6}$$

The Hamiltonian can thus be expressed discretely in terms of only the positions of the vortex centres \mathbf{x}_i . This discrete representation of the dynamics is exact if the point vortices are singular, and hence the discrete H can be defined equal to the functional \mathcal{H} . It is assumed that the stream functions associated with the Coriolis and topography terms are known explicitly. For regularized systems, the Green’s function may still be known, but the quadrature of \mathcal{H} can no longer be performed exactly and the discretized Hamiltonian will no longer be exact.

The Poisson bracket is discretized separately. First of all, it is useful to rewrite (3), because we have defined the point vortex positions as vectors in \mathbb{R}^3 rather than in spherical coordinates. For any $\mathbf{x} \in \mathbb{R}^3 : |\mathbf{x}| = 1$, (3) is equivalent to

$$J(f, g, \mathbf{x}) = (\mathbf{x} \times \nabla f) \cdot \nabla g.$$

The Poisson bracket then follows as

$$\begin{aligned} \{\mathcal{F}, \mathcal{G}\} &= \int \frac{\delta \mathcal{F}}{\delta q} \nabla q \cdot \mathbf{x} \times \left(\nabla \frac{\delta \mathcal{G}}{\delta q} \right) dS \\ &= - \int q \nabla \cdot \left(\frac{\delta \mathcal{F}}{\delta q} \mathbf{x} \times \nabla \frac{\delta \mathcal{G}}{\delta q} \right) dS \\ &= - \int q \nabla \left(\frac{\delta \mathcal{F}}{\delta q} \right) \cdot \left(\mathbf{x} \times \nabla \frac{\delta \mathcal{G}}{\delta q} \right) dS, \end{aligned} \tag{7}$$

using first the divergence theorem and then the fact that the divergence of the curl equals zero. The discrete form of the functional $\mathcal{F} = \int f(\mathbf{x}) dS$ is given by

$$F = \sum_{i=1}^N f(\mathbf{x}_i) = \int f(\mathbf{x}) \left(\sum_{i=1}^N \delta(\mathbf{x} - \mathbf{x}_i) \right) dS.$$

We assume there exists a field λ_q for the vorticity field q with the properties:

$$\begin{aligned} \lambda_q(\mathbf{x}_i) &= \Gamma_i^{-1}, \\ \nabla \lambda_q(\mathbf{x})|_{\mathbf{x}=\mathbf{x}_i} &= 0, \\ \lim_{\varepsilon \rightarrow 0} \frac{\lambda_{q+\varepsilon v} - \lambda_q}{\varepsilon} &= 0. \end{aligned}$$

With this we write

$$F = \int f(\mathbf{x}) \lambda(\mathbf{x}) \left(\sum_{i=1}^N \Gamma_i \delta(\mathbf{x} - \mathbf{x}_i) \right) dS = \int f(\mathbf{x}) \lambda(\mathbf{x}) q(\mathbf{x}) dS,$$

from which the variational derivative follows

$$\begin{aligned} \int \frac{\delta F}{\delta q} v dS &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left(\int f \lambda_{q+\varepsilon v} (q + \varepsilon v) dS - \int f \lambda_q q dS \right) \\ &= \int f \lambda_q v dS. \end{aligned}$$

Substitution of this form for the functional \mathcal{F} and \mathcal{G} in (7) leads to the discrete form of the Poisson bracket

$$\begin{aligned} \{F, G\} &= - \int q \nabla \frac{\delta F}{\delta q} \cdot \mathbf{x} \times \nabla \frac{\delta G}{\delta q} dS \\ &= - \int q \nabla (f \lambda_q) \cdot \mathbf{x} \times \nabla (g \lambda_q) dS \\ &= - \int q \lambda_q^2 \nabla f \cdot \mathbf{x} \times \nabla g dS \\ &= - \sum_{i=1}^N \Gamma_i^{-1} \nabla f(\mathbf{x}_i) \cdot \mathbf{x}_i \times \nabla g(\mathbf{x}_i). \end{aligned}$$

This is a generalization of the well-known Poisson bracket for rigid body rotation [29], also used in models for ferromagnetism [30]. The bracket is in fact equivalent to the bracket for a Heisenberg spin chain [31].

We introduce the vector $\mathbf{y} \in \mathbb{R}^{3N}$ as the concatenation of the $\mathbf{x}_i \in \mathbb{R}^3, i = 1, 2, \dots, N$. The dynamics are then

$$\dot{\mathbf{y}} = B(\mathbf{y}) \nabla H(\mathbf{y}), \tag{8}$$

Table 1
Summary of discrete identities compared to continuous identities.

Continuous representation	Discrete representation
$\mathcal{F}(t) = \int f(\mathbf{x}) \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{x}_i(t)) \, dS$ $= \int f(\mathbf{x}) \lambda(\mathbf{x}) q(\mathbf{x}, t) \, dS$	$F(t) = \sum_{i=1}^N f(\mathbf{x}_i(t))$
$\mathcal{H}(t) = \int q \psi \, dS$	$H(t) = \frac{1}{2} \sum_{i=1}^N \Gamma_i \left(\psi_c(\mathbf{x}_i(t)) + \psi_h(\mathbf{x}_i(t)) - 2 \sum_{j<i} \Gamma_j \frac{1}{4\pi} \ln(\mathbf{x}_j(t) - \mathbf{x}_i(t) ^2) \right)$
$\mathcal{J}(t) = \int \mathbf{x} q(\mathbf{x}, t) \, dS$	$J(t) = \sum_{i=1}^N \Gamma_i \mathbf{x}_i(t)$
$\mathcal{C}_j(t) = \int \mathbf{x} ^2 \delta(\mathbf{x} - \mathbf{x}_j(t)) \, dS$	$C_j(t) = \mathbf{x}_j(t) ^2$
$\{\mathcal{F}, \mathcal{G}\} = - \int q \nabla \left(\frac{\delta \mathcal{F}}{\delta q} \right) \cdot \left(\mathbf{x} \times \nabla \frac{\delta \mathcal{G}}{\delta q} \right) \, dS$	$\{F, G\} = - \sum_{i=1}^N \Gamma_i^{-1} \nabla f(\mathbf{x}_i(t)) \cdot \mathbf{x}_i \times \nabla g(\mathbf{x}_i(t))$
$\{q, \mathcal{H}\}$	$\{x_i^{(1)}, H\}, \{x_i^{(2)}, H\}, \{x_i^{(3)}, H\}$

with the block-diagonal structure matrix

$$B(\mathbf{y}) = \begin{bmatrix} \Gamma_1^{-1} \widehat{\mathbf{x}}_1 & & & 0 \\ & \Gamma_2^{-1} \widehat{\mathbf{x}}_2 & & \\ & & \ddots & \\ 0 & & & \Gamma_N^{-1} \widehat{\mathbf{x}}_N \end{bmatrix},$$

where $\widehat{\mathbf{x}}$ is the 3×3 skew-symmetric matrix such that $\widehat{\mathbf{x}}\mathbf{u} = \mathbf{x} \times \mathbf{u} \, \forall \mathbf{u} \in \mathbb{R}^3$. The vortex position radii $C_i = |\mathbf{x}_i|$ are Casimirs of the Poisson bracket associated with structure matrix $B(\mathbf{y})$. That is, for any function $F(\mathbf{y})$ and any C_i , one has $\{F, C_i\} \equiv 0$. This property is important as it implies that if the vortex positions initially satisfy $|\mathbf{x}_i| = 1$, this is maintained throughout the simulation, ensuring the point vortices remain on the sphere. The numerical integration scheme developed below respects this property inherently, without the need of a projection step.

Due to the rotational symmetries of the sphere, the dynamics exhibit three Noether momenta given in vector form as $\mathcal{J} = \int_{\mathbb{S}^2} \mathbf{x} q \, dS$. In the point vortex discretization, these momenta persist as

$$\mathbf{J} = \mathcal{J} = \int_{\mathbb{S}^2} \mathbf{x} q \, dS = \sum_{i=1}^N \Gamma_i \mathbf{x}_i.$$

We summarize the discrete Hamiltonian representation in [Table 1](#).

4. Poisson integrator

For Poisson systems such as the point vortex system it is essential to employ a numerical time integrator that maintains the structure of the underlying ordinary differential equations. Standard numerical integrators do not conserve Casimirs. Hence Runge–Kutta or multistep methods will result in point vortices drifting from the sphere. This can be corrected with projections, but as is known from the rigid body equation, doing so can introduce artificial stable equilibria.

Integrators that conserve the geometric structure are of special importance when one is interested in the statistics of long simulations. In geophysical fluid dynamics the long time mean vorticity field and streamfunction, as well as the pointwise statistics, depend heavily on the geometric properties of the numerical integrator [32–34].

Patrick [35] suggests applying a Poisson splitting method to point vortex dynamics, but does not detail the method. A Poisson integrator preserves Casimirs by definition. We will see that the splitting also preserves the Noether momenta exactly, and the Hamiltonian approximately in the sense of backward error analysis, as detailed in Section 4.2.

We expand system (8) with Hamiltonian (6) as

$$\dot{\mathbf{y}} = \sum_{i=1}^N B(\mathbf{y}) \nabla H_i + \sum_{j<i} B(\mathbf{y}) \nabla H_{ij},$$

where

$$H_i = \frac{1}{2} \Gamma_i (\psi_c(\mathbf{x}_i) + \psi_h(\mathbf{x}_i)), \quad \text{and}$$

$$H_{ij} = \frac{1}{4\pi} \ln(|\mathbf{x}_j - \mathbf{x}_i|^2) = \frac{1}{4\pi} \ln(2 - 2\mathbf{x}_i \cdot \mathbf{x}_j).$$

We treat the dynamics for each of these terms separately. The time- τ flow map associated with each of the H_i terms will be denoted by ϕ_τ^i . The dynamics associated with $\dot{\mathbf{y}} = B(\mathbf{y}) \nabla H_{ij}$ is just that of a two vortex system with time- τ flow map denoted by ϕ_τ^{ij} . This flow is known explicitly as detailed in Section 4.1. A splitting method is a composition of the flow maps of all the individual terms in the dynamics.

We initially restrict ourselves to Lie–Trotter splittings and Strang splittings [36], respectively of the form

$$\Phi_\tau^{\text{LT}} = \prod_{(i,j) \in C_N} \phi_\tau^{ij} \circ \prod_{i=1}^N \phi_\tau^i \quad \text{and} \quad \Phi_\tau^{\text{S}} = \prod_{(i,j) \in C_N} \phi_{\tau/2}^{ij} \circ \prod_{i=1}^N \phi_\tau^i \circ \prod_{(i,j) \in C_N^*} \phi_{\tau/2}^{ij}. \tag{9}$$

For both cases C_N represents an ordering of all the possible pairs (i, j) , $i \neq j$. The Strang splitting also uses the reverse ordering, labelled C_N^* , to create a symmetric method. The symmetry results in a cancellation of first order error terms, yielding a method that is second order accurate. The Strang splitting can subsequently be used in the construction of higher order methods [37,38]. In the remainder we will ignore the effects of the Coriolis force and topography. The corresponding flows ϕ_τ^i are perfectly parallelizable and their evaluation represents an ever smaller fraction of the total workload as the number of vortices increases.

Because each of the pairwise interactions in the splittings in (9) is the exact solution to a Poisson dynamical system, each ϕ_τ^{ij} is a Poisson map with respect to the bracket $\{F, G\} = \nabla_{\mathbf{y}} F^T B(\mathbf{y}) \nabla_{\mathbf{y}} G$. As the composition of Poisson maps is again Poisson, the splitting schemes are also Poisson maps with respect to the same bracket [36].

4.1. Integration of the two vortex system

The dynamics $\dot{\mathbf{y}} = B(\mathbf{y}) \nabla H_{ij}$ (with flow map ϕ_τ^{ij}) affects only vortices i and j and can thus be expressed as

$$\begin{aligned} \dot{\mathbf{x}}_i &= \Gamma_i^{-1} \mathbf{x}_i \times \nabla_i H_{ij} = \frac{-\Gamma_j}{4\pi} \frac{\mathbf{x}_j \times \mathbf{x}_i}{1 - \mathbf{x}_i \cdot \mathbf{x}_j}, \\ \dot{\mathbf{x}}_j &= \Gamma_j^{-1} \mathbf{x}_j \times \nabla_j H_{ij} = \frac{-\Gamma_i}{4\pi} \frac{\mathbf{x}_i \times \mathbf{x}_j}{1 - \mathbf{x}_i \cdot \mathbf{x}_j}. \end{aligned}$$

This two-vortex system has Noether momenta expressed by the vector

$$\mathbf{J}_{ij} = \Gamma_i \mathbf{x}_i + \Gamma_j \mathbf{x}_j.$$

Using the Noether momenta we find

$$\dot{\mathbf{x}}_i = \frac{-1}{4\pi} \frac{\mathbf{J}_{ij}}{1 - \mathbf{x}_i \cdot \mathbf{x}_j} \times \mathbf{x}_i =: \mathbf{a} \times \mathbf{x}_i, \tag{10}$$

$$\dot{\mathbf{x}}_j = \frac{-1}{4\pi} \frac{\mathbf{J}_{ij}}{1 - \mathbf{x}_i \cdot \mathbf{x}_j} \times \mathbf{x}_j =: \mathbf{a} \times \mathbf{x}_j. \tag{11}$$

Conservation of the Hamiltonian H_{ij} implies the denominators in (10)–(11) are constant. This implies the vector \mathbf{a} is invariant under the two-vortex dynamics.

Using Rodrigues’ formula [39], the solution to Eq. (10) is given by

$$\begin{aligned} \mathbf{x}_i(\tau) &= \exp(\widehat{\mathbf{a}}\tau) \mathbf{x}_i(0) = \mathbf{x}_i(0) + \frac{\sin a\tau}{a} \widehat{\mathbf{a}} \mathbf{x}_i(0) + \frac{1 - \cos a\tau}{a^2} \widehat{\mathbf{a}}^2 \mathbf{x}_i(0) \\ &= \mathbf{x}_i(0) + \sin a\tau \tilde{\mathbf{a}} \times \mathbf{x}_i(0) + (1 - \cos a\tau) (\tilde{\mathbf{a}}(\tilde{\mathbf{a}} \cdot \mathbf{x}_i(0)) - \mathbf{x}_i(0)), \end{aligned}$$

where $\widehat{\mathbf{a}}$ is the matrix such that $\widehat{\mathbf{a}} \mathbf{x} = \mathbf{a} \times \mathbf{x}$, $a = |\mathbf{a}|$ and $\tilde{\mathbf{a}} = \mathbf{a}/a$. The solution to (11) follows by substituting \mathbf{x}_j for \mathbf{x}_i .

This flow map presents an explicit formulation of the exact solution to the two-vortex system of vortices i and j . This pairwise solution is therefore a Poisson system with the same bracket as the N -vortex problem that also preserves the pairwise Hamiltonian and momenta exactly. A splitting composed of Poisson flows with identical brackets respects the Casimirs of that bracket. The Noether momenta of the N -vortex system may be written as

$$\mathbf{J} = \sum_{i=1}^N \Gamma_i \mathbf{x}_i = \Gamma_i \mathbf{x}_i + \Gamma_j \mathbf{x}_j + \sum_{k \neq i,j} \Gamma_k \mathbf{x}_k = \mathbf{J}_{ij} + \sum_{k \neq i,j} \Gamma_k \mathbf{x}_k.$$

This implies that the total momenta \mathbf{J} are conserved, because the pairwise flows preserve the pairwise momenta \mathbf{J}_{ij} and do not modify the other vortices. The Hamiltonian is not conserved exactly, as the evaluation of pair (i, j) perturbs the values of the Hamiltonian terms H_{ik} and H_{jk} for $k \neq i, j$. This is considered in more detail in the following section.

4.2. Modified Hamiltonian

For splitting schemes consisting of exactly integrated Poisson flows with the same bracket, the combined map approximates, to an exponentially high order, a Poisson system with the same bracket, but a modified Hamiltonian. Studying the modified problem in this case will lead to two conclusions. First, we will show that there exists no ordering of pairwise

interactions that would lead to a cancellation of error terms.¹ Second, it will become explicitly clear that the energy error is dominated by close approaches between vortices.

Before considering the modified Hamiltonian for the point-vortex system, let us recall the simpler problem where the dynamics is given by

$$\dot{y} = B(y) (H^{[1]}(y) + H^{[2]}(y)). \tag{12}$$

For the symmetric Strang splitting

$$\Phi_\tau = \phi_{\tau/2}^{[1]} \circ \phi_\tau^{[2]} \circ \phi_{\tau/2}^{[1]}$$

the modified dynamics read

$$\dot{\tilde{y}} = B(\tilde{y}) (H^{[1]}(\tilde{y}) + H^{[2]}(\tilde{y}) + \tau^2 H_3(\tilde{y}) + \tau^4 H_5(\tilde{y}) + \dots).$$

Throughout the present work we will only consider the first correction term, $H_3(\tilde{y})$, corresponding to an $\mathcal{O}(\tau^2)$ modification to the Hamiltonian. For the Strang splitting of (12) this term is given by [36, p. 299]

$$H_3 = -\frac{1}{24} \{ \{ H^{[2]}, H^{[1]} \}, H^{[1]} \} + \frac{1}{12} \{ \{ H^{[1]}, H^{[2]} \}, H^{[2]} \}. \tag{13}$$

When the splitting contains more than two different flow maps, the modified Hamiltonian is constructed by applying (13) repeatedly “from the inside out”, as illustrated by the following three-vortex example.

Consider a system consisting of three vortices that are integrated according to the Strang splitting Φ_τ^S of (9). We consider an ordering of pairwise interactions such that the splitting reads

$$\Phi_\tau^{S,3} = \phi_{\tau/2}^{12} \circ \phi_{\tau/2}^{02} \circ \phi_\tau^{01} \circ \phi_{\tau/2}^{02} \circ \phi_{\tau/2}^{12}.$$

Note that the innermost map ϕ_τ^{01} is a composition of two successive maps in the definition of (9). We first construct the modified Hamiltonian for the inner map $\phi_\tau^{\text{inner}} = \phi_{\tau/2}^{02} \circ \phi_\tau^{01} \circ \phi_{\tau/2}^{02}$ using (13), resulting in

$$H^{\text{inner}} = H_{02} + H_{01} + \tau^2 H_3^{\text{inner}} + \dots,$$

with

$$H_3^{\text{inner}} = \frac{-1}{24} \{ \{ H_{01}, H_{02} \}, H_{02} \} + \frac{1}{12} \{ \{ H_{02}, H_{01} \}, H_{01} \}.$$

The modified Hamiltonian of the full step is found by applying (13) to $\Phi_\tau^{S,3} = \phi_{\tau/2}^{12} \circ \phi_\tau^{\text{inner}} \circ \phi_{\tau/2}^{12}$, which results in

$$H^{[S]} = H_{12} + H_{02} + H_{01} + \tau^2 H_3^{\text{inner}} + \tau^2 H_3^{\text{outer}} + \dots, \tag{14}$$

with

$$H_3^{\text{outer}} = \frac{-1}{24} \{ \{ H_{01} + H_{02}, H_{12} \}, H_{12} \} + \frac{1}{12} \{ \{ H_{12}, H_{01} + H_{02} \}, H_{01} + H_{02} \} + \dots.$$

Terms arising from the Poisson bracket of H_3^{inner} with H_{12} are of order τ^4 and are subsequently neglected. Combining both second order corrections we find

$$\begin{aligned} H_3 &= H_3^{\text{inner}} + H_3^{\text{outer}} \\ &= \frac{-1}{24} (\{ \{ H_{01}, H_{02} \}, H_{02} \} + \{ \{ H_{01}, H_{12} \}, H_{12} \} + \{ \{ H_{02}, H_{12} \}, H_{12} \}) \\ &\quad + \frac{1}{12} (\{ \{ H_{02}, H_{01} \}, H_{01} \} + \{ \{ H_{12}, H_{01} \}, H_{01} \} + \{ \{ H_{12}, H_{02} \}, H_{02} \}) \\ &\quad + \frac{1}{12} (\{ \{ H_{12}, H_{01} \}, H_{02} \} + \{ \{ H_{12}, H_{02} \}, H_{01} \}). \end{aligned} \tag{15}$$

The first two lines of (15) consist of Poisson brackets that are all of the form

$$T_{k\ell m}^I := \{ \{ H_{k\ell}, H_{km} \}, H_{k\ell} \} = - \{ \{ H_{km}, H_{k\ell} \}, H_{k\ell} \}. \tag{16}$$

The last line consists of brackets of the form

$$T_{k\ell m}^{II} := \{ \{ H_{k\ell}, H_{km} \}, H_{\ell m} \} = - \{ \{ H_{km}, H_{k\ell} \}, H_{\ell m} \} = -T_{k\ell m}^{II}. \tag{17}$$

¹ At least no such ordering exists for general systems, in the highly specialized test cases of Section 4.3, different orderings may yield very different results.

With these definitions Eq. (15) is expressed more compactly as

$$H_3 = \frac{-1}{24} (-T_{021}^I - T_{120}^I - T_{210}^I) + \frac{1}{12} (-T_{012}^I - T_{102}^I - T_{201}^I + T_{120}^{II} + T_{210}^{II}).$$

We illustrate the accuracy of this correction to the Hamiltonian by simulating a three-vortex system using different time steps. Fig. 1 compares the L_1 norm of errors in the Hamiltonian and the modified Hamiltonian. We simulate 10 time units starting from arbitrary initial conditions. The second-order convergence of the original Hamiltonian is visible for sufficiently small time steps. The modified Hamiltonian is preserved to fourth order as expected.

In the general N -vortex case, the modified Hamiltonian follows from computing the modified Hamiltonian of the innermost composition and repeatedly working outwards. The ordering C_N consists of $R = \frac{1}{2}N(N - 1)$ pairs; we label these $C_N^{[r]}$, with $r \in \{1, \dots, R\}$. Let $H^{[r]}$ denote the Hamiltonian corresponding to the pair $C_N^{[r]}$. Following the same procedure as for the three-vortex system, only now for a more general number of steps, we find the second-order correction to the Hamiltonian to be

$$H_3 = \sum_{r=1}^{R-1} \sum_{s=r+1}^R \left[\frac{-1}{24} \{ \{ H^{[s]}, H^{[r]} \}, H^{[r]} \} + \sum_{t=r+1}^R \frac{1}{12} \{ \{ H^{[r]}, H^{[s]} \}, H^{[t]} \} \right]. \tag{18}$$

These commutators will be zero if either the two pairs inner commutators share no vortex, or if the outer vortex pair shares no vortex with either of the inner vortex pairs. Consequently, for large systems there will only be a small number of non-trivial terms in (18). The other terms in the sum are all distinct, because $r < s \leq t$. This means it is not possible to construct orderings where terms cancel, at least not without evaluating vortex pairs multiple times.

When the terms in (18) involve only three distinct vortices, they are of a form as in (16) or (17). Nontrivial terms that involve four distinct vortices follow either of the following two patterns

$$T_{k\ell mn}^{III} := \{ \{ H_{k\ell}, H_{km} \}, H_{kn} \} = - \{ \{ H_{km}, H_{k\ell} \}, H_{kn} \} \tag{19}$$

$$T_{k\ell mn}^{IV} := \{ \{ H_{k\ell}, H_{km} \}, H_{\ell n} \} = - \{ \{ H_{km}, H_{k\ell} \}, H_{mn} \}. \tag{20}$$

Without presenting the derivations, we state that (16)–(17) and (19)–(20) are computed as

$$\begin{aligned} T_{k\ell m}^I &= A_{k\ell m} B_{k\ell}^2 B_{km} [-\Gamma_\ell B_{km} C_{\ell km}^2 + (\Gamma_\ell \mathbf{x}_m \times \mathbf{x}_\ell + \Gamma_k \mathbf{x}_k \times \mathbf{x}_m) \cdot (\mathbf{x}_k \times \mathbf{x}_\ell)] \\ T_{k\ell m}^{II} &= A_{k\ell m} B_{k\ell} B_{km} B_{\ell m} [(\Gamma_\ell B_{km} - \Gamma_m B_{k\ell}) C_{\ell km}^2 + (\Gamma_m \mathbf{x}_m + \Gamma_\ell \mathbf{x}_\ell) \times \mathbf{x}_k \cdot (\mathbf{x}_m \times \mathbf{x}_\ell)] \\ T_{k\ell mn}^{III} &= A_{k\ell m} B_{k\ell} B_{km} B_{kn} \Gamma_n [B_{k\ell} C_{\ell km} C_{\ell kn} + B_{km} C_{\ell km} C_{mkn} + (\mathbf{x}_m \times \mathbf{x}_\ell) \cdot (\mathbf{x}_k \times \mathbf{x}_n)] \\ T_{k\ell mn}^{IV} &= A_{k\ell m} B_{k\ell} B_{km} B_{\ell n} \Gamma_n [B_{k\ell} C_{\ell km} C_{k\ell n} + (\mathbf{x}_k \times \mathbf{x}_m) \cdot (\mathbf{x}_\ell \times \mathbf{x}_n)], \end{aligned}$$

where

$$\begin{aligned} A_{k\ell m} &= \frac{-\Gamma_k \Gamma_\ell \Gamma_m}{(4\pi)^3} \\ B_{k\ell} &= \frac{1}{1 - \mathbf{x}_k \cdot \mathbf{x}_\ell} \\ C_{k\ell m} &= \mathbf{x}_k \cdot \mathbf{x}_\ell \times \mathbf{x}_m. \end{aligned}$$

From this it follows that the error in the Hamiltonian is dominated by close approaches between vortices, as there the denominator in $B_{k\ell}$ approaches zero. The magnitude of both $A_{k\ell m}$ and $C_{k\ell m}$ are obviously bounded.

Long time conservation of the Hamiltonian by symplectic methods can be rigorously shown in the case of analytic Hamiltonian H , but is often observed in practice for more general Hamiltonians. For point vortices, the Hamiltonian has singularities when two vortices coincide. The motion of a single pair could inadvertently place one vortex in close proximity to another, effectively a “numerical collision”. In practice we have not encountered this. This is only problematic if that vortex pair is evaluated before either of the coincident vortices are moved by a different vortex pair interaction.

4.3. Numerical results

Our primary purpose for developing the explicit and symplectic integrator outlined in the present work is to allow performing numerical experiments on moderate to large vortex systems to verify hypothesized statistical behaviour. These simulations must be run over a long time period to yield meaningful statistics. In this section we demonstrate the approximate energy conservation and exact momentum conservation of the splitting schemes in different settings. We compare the results of the Strang splitting to results for a symplectic, implicit method developed by Vankerschaver and Leok [26] (labelled “Hopf” in the figures as by their own convention) and to a fourth order explicit Runge–Kutta scheme. Both implementations can be found at https://github.com/jvkersch/hopf_vortices along with two other implicit schemes.

We investigate the conservation of energy and momentum in a number of different configurations. For each case we compare energy and momentum errors given by $\varepsilon_H(t) = |H(t) - H(0)|$ and $\varepsilon_j(t) = \|\mathbf{J}(t) - \mathbf{J}(0)\|_2$. We have repeated




Fig. 1. Error convergence for the Hamiltonian of a three vortex system in solid red. In dash-dot blue is the convergence of the error for the modified Hamiltonian given by (14). For reference two dashed lines with slopes τ^2 and τ^4 are given.




Fig. 2. Left: initial conditions for the collapsing vortices, colours indicate vortex strength ranging from negative unit circulation in dark red to positive unit circulation in pale yellow. Middle: trajectories leading up to the collapse event, spiralling inwards. Right: trajectories after the collapse event, spiralling out.

the experiments of Vankerschaver and Leok [26], but observe that in most of these configurations symmetries play an important role. This symmetry is broken by the splitting methods, resulting in a poorer performance than methods that maintain the symmetry. We also compare results for an arbitrary initial condition at a given energy level in Section 4.3.4. This is a more practical test case for engineering applications and statistical mechanics; it is in fact the setting in which we use this integrator in other work [40].

4.3.1. Collapsing vortices

For certain initial conditions three or more vortices will collapse onto a single point in finite time, while the energy remains bounded. Such initial conditions with three vortices have been studied by Kidambi and Newton [41] and with four vortices by Sakajo [42]. We simulate the same three vortex system as Vankerschaver and Leok [26].

The vortex circulations are $\Gamma_1 = \Gamma_2 = 1$ and $\Gamma_3 = -\frac{1}{2}$. The vortices start at the vertices of a triangle with lengths $l_{12} = \frac{1}{2}\sqrt{2}$, $l_{23} = \frac{1}{2}\sqrt{2}$ and $l_{13} = 1$. These initial conditions result in a collapse of the three vortices onto a single point at $T^- = 4\pi(\sqrt{23} - \sqrt{17}) \approx 8.4537$. At this time the exact equations of motion become undefined. The numerical methods will not collapse exactly, due to the numerical error—and in the case of the methods in [26], due to regularization of the dynamics. Instead there will be a moment that the vortices approach each other closely. This event is repeated periodically. The left hand panel of Fig. 2 depicts the initial conditions, the middle panel shows the simulation trajectories leading up to the near-collapse, and the right panel shows trajectories after this event.

We have performed this simulation with time steps $\tau = 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}$ to illustrate how the behaviour changes. Conservation of the Hamiltonian over a short time – enough to show the first collapse event – is illustrated in the left-hand panel of Fig. 3. As the time step is reduced, the magnitude of the maximum energy error during this part of the simulation does not change, but the length of time over which there is a significant error is greatly reduced by using a smaller time step. In this pathological configuration, the lack of analyticity of the Hamiltonian negates backward error analysis and its prediction of second-order convergence.

The moment of the closest approach is indicated by the largest error in the Hamiltonian. With smaller time steps this instant approaches the correct time of the true collapse event.

The right-hand panel of Fig. 3 demonstrates how the energy changes over longer time, including four more near-collapse events. Between the near-collapses, the energy consistently returns close to its initial value. Note also that the results presented by Vankerschaver and Leok [26] include some regularization. This in fact slows down the dynamics around the




Fig. 3. Energy error for a system of $N = 3$ collapsing vortices for various time step sizes. Initial conditions that lead to a collapse of the three vortices onto a single point (in finite time) are given by Kidambi and Newton [41]. The exact time of this collapse is indicated by a black dotted line.




Fig. 4. Left: initial conditions for the stable vortex ring, all vortices have equal circulation $\Gamma = \frac{1}{6}$. Right: z -coordinate of vortices in the Strang splitting simulation, indicating a steady oscillation.

near-collapses so much that the fifth event does not occur before the end of the simulation at $T = 500$. We should note that the behaviour of this system with a repetition of near-collapses does not occur despite numerical error, but because of it. The exact solution becomes undefined at the (first) collapse. When the error is reduced by using a higher order method (results not shown), the dynamics leading up to the first event are more accurate, resulting in a much closer approach between the vortices. This causes a larger energy error after the event than in the Strang splitting.

4.3.2. Stable vortex ring

A ring of N equidistantly placed vortices of equal strength rotates stably around its centre provided $N \leq 7$ and provided that the latitude of the ring (assumed parallel to the equator) is above a certain critical value [43].

We simulate a stable configuration with $N = 6$, $\Gamma = \frac{1}{N}$ and latitude corresponding to $z = 0.92$ for 1000 time units. All simulations use a time step of 0.05. The initial conditions and the z -coordinate of the vortices in the resulting simulation are depicted in Fig. 4.

Fig. 5 compares the energy and momentum errors against those for the Hopf and Runge–Kutta integration methods. The splitting method only approximately conserves energy, but the error is bounded. The momentum is conserved to machine precision throughout the simulation. Due to the rotational symmetry of this configuration, a method that respects this symmetry will easily exhibit energy conservation. The splitting method does not respect the symmetry due to the influence of the ordering of pairs. Consequently its energy conserving quality in this rather specialized test case is inferior to that of the Hopf integrator. We have also performed the experiment with McLachlan’s 6th order composition [38] of the Strang splitting. This shows energy conservation to the same degree as the Hopf method.

4.3.3. Von Kármán vortex streets

Another relative equilibrium is that of the Von Kármán vortex streets presented by Chamoun et al. [44]. This configuration consists of two staggered rings of N_r vortices placed at latitudes $\theta = \pm\theta_0$ plus one vortex at each pole.

We take $N_r = 5$ vortices per ring, each with $\Gamma = \pm 1$, placed at $\theta = \pm\frac{\pi}{6}$, respectively. The polar vortices satisfy $\Gamma_n = -\Gamma_s = \frac{1}{2}$. This configuration rotates about the z -axis with a period of $T = 10.85$. We simulate this system with a time step of $\tau = 0.5$ for 10 000 time units. Fig. 6 depicts the initial conditions and the z -coordinate of the vortices in the ensuing simulation.

The splitting scheme and Hopf method both conserve the momenta exactly by construction, and this is reflected in the simulation results shown in Fig. 7. The error in the energy remains bounded throughout the simulation at an accuracy that is somewhat better than that of the Hopf integrator.




Fig. 5. Energy (left) and momentum (right) error for the stable vortex ring of Polvani and Dritschel [43] with $N = 6$ vortices. The energy oscillates about a fixed mean, with bounded error. The momentum is conserved to machine precision.

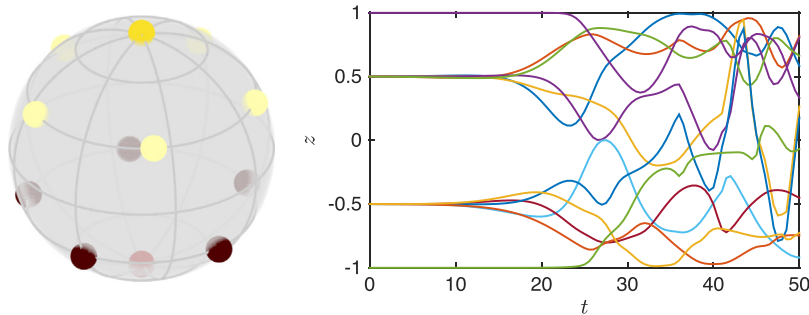


Fig. 6. Left: initial conditions for the Von Kármán vortex streets, colours indicate vortex strength ranging from negative unit circulation in dark red to positive unit circulation in pale yellow. Right: z -coordinate of vortices in the Strang splitting simulation, showing the onset of instability after some time.

This configuration is believed to be inherently unstable [26], making the symmetry of vital importance. The splitting scheme breaks this symmetry, yet the error remains bounded.

4.3.4. Generic initial conditions

In this final test case we consider a system with 48 vortices, eight with circulation $\Gamma = \pm 1$ and 40 with $\Gamma = \pm \frac{1}{5}$ with equal numbers positive and negative. The initial configuration is drawn randomly from the set of all states with a given energy level and zero angular momentum. We use the same initial condition for the different methods. We consider both large negative energy ($H = -2$), resulting in a strong clustering of like-signed vortices [10,45], and large positive energy ($H = 2$), leading to a well-mixed configuration with close approaches between opposite signed-vortices. These configurations, depicted graphically in Fig. 8, are extreme in the sense that the specified energy levels lie close to the tails in the distribution of all attainable energy states for this number of vortices with these circulations.

Fig. 9 compares the energy error over a short time for the Strang splitting method against the Hopf method [26], a fourth order Runge–Kutta scheme and an implicit midpoint method [36], all with time step $\tau = 0.0001$. There is a remarkable difference in the performance of the integrators at positive and negative energies. At negative energies, the Runge–Kutta scheme conserves energy accurately while it exhibits rapid error growth in the positive energy simulation. The inaccuracy at positive energy is a consequence of the formation of many and long-lived like-signed vortex pairs. These pairs orbit each other rapidly and this circular motion is poorly captured by the Runge–Kutta scheme, where each time step consists of a number of rectilinear extrapolations with a projection onto the sphere. The Hopf integrator performs poorly at negative energy, and fails to converge at positive energy, even for this modest time step size. In the negative energy simulation, the Strang splitting and implicit midpoint method have roughly the same accuracy. For positive energy, the implicit midpoint is more accurate.

We also compare the results over a longer time period with an increased time step of $\tau = 0.1$. All three implicit methods considered by Vankerschaver and Leok [26] fail to converge with this time step. Therefore Fig. 10 shows only the results for the Strang splitting and the fourth order Runge–Kutta scheme. In both cases the Strang splitting shows a smaller energy error. The momentum error (not shown) is within machine precision for the Strang splitting and of order one for the Runge–Kutta approach.

We attribute the higher accuracy of the splitting scheme in the positive energy case to the formation of long lived coherent structures of like-signed vortices. It was concluded in Section 4.2 that the error is dominated by close approaches between vortices. Conservation of energy implies that a close approach of a like-signed (e.g. “+/+)” vortex pair can only occur with a simultaneous close approach of an opposite-signed (e.g. “+/-”) vortex pair. Such encounters occur less frequently in




Fig. 7. Energy (left) and momentum (right) error for the Von Kármán vortex street with two rings of 5 vortices and a vortex at either pole, giving $N = 12$ vortices in total. The energy oscillates rapidly about a slowly varying mean. The momentum is initially conserved to machine precision, but accumulation of arithmetic errors eventually leads to a small drift.




Fig. 8. Generic initial conditions with $H = -2$ (left) and $H = 2$ (right), colours indicate vortex strength ranging from negative unit circulation in dark red to positive unit circulation in pale yellow. The left hand panel shows positive and negative vortices are mixed, whereas the right hand panel shows a clear clustering of like-signed vortices, for example a number of negative vortices close to the “north pole”.

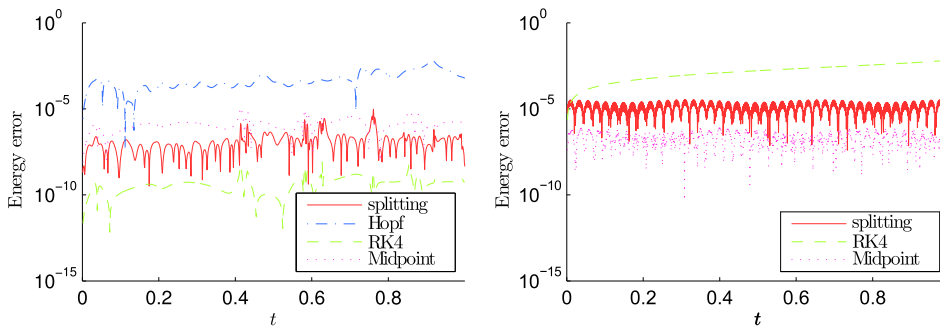


Fig. 9. Energy error in simulations with negative (left) and positive (right) energies. The Strang splitting is compared against a fourth order Runge–Kutta (RK4) scheme. The Hopf fibration method fails to converge in the positive energy case.




Fig. 10. Energy error in simulations with negative (left) and positive (right) energies. The Strang splitting is compared against a fourth order Runge–Kutta (RK4) scheme.

Table 2
Round-robin scheme for ordering vortex pairs in a four-vortex system.

Round	0	1	2
Pair	0–3 1–2	1–3 2–0	2–3 0–1

the positive energy scenario, as the fluid develops large coherent patches of positive and negative vorticity, relegating opposite-signed vortex collisions (+/−) to the boundaries of these patches.

4.4. Error comparison against computational cost

The previous sections compared the conservation properties of the splitting integrator against those of various other integrators, but the parameter choices for the different methods do not result in an equal workload for all of the simulations. In this section, we present the error against workload in terms of vector field evaluations. For the implicit methods used by Vankerschaver and Leok [26], this measure varies per timestep depending on the convergence of the linear solver. Therefore we restrict the comparison to the second- and fourth order Runge–Kutta schemes (labelled “RK2” and “RK4”) and various splitting schemes: the Lie–Trotter (“LT”) and Strang (“S”) splittings discussed above, and four higher order methods. These higher order methods follow from a composition of a number of Strang steps of different sizes [36]. We study both fourth and sixth order methods, considering the methods with the minimal number of stages presented by Yoshida [37] (“Y4” and “Y6”), as well as the methods of same order but with smaller error coefficients found by McLachlan [38] (“M4” and “M6”). The workload is viewed as the number of force evaluations for the Runge–Kutta schemes and as the number of Lie–Trotter (and adjoint) splitting evaluations for the splitting methods. These are of approximately the same cost.

The first three panels of Fig. 11 show the energy error against computational workload for the first three test cases respectively, i.e. the collapsing vortices, the stable vortex ring and the Von Kármán vortex streets. In the case of the collapsing vortices, the Runge–Kutta schemes outperform most of the splitting methods. We attribute this to the inclusion of regularization in the Runge–Kutta implementation developed by Vankerschaver and Leok [26]. The stable vortex ring and the Von Kármán vortex streets both show that higher order methods are only advantageous in terms of computational cost if a high accuracy is required.

The fourth panel of Fig. 11 compares the error in position for the different methods in the stable vortex ring case. The splitting methods, which are developed for accuracy in conservation properties, show poorer results in this metric. The Runge–Kutta schemes respect the symmetry of this test case, resulting in a far better accuracy in the limit of a small time step, i.e. $\tau \rightarrow 0$. If there is only a moderate requirement for accuracy the apparent difference in performance is inconclusive, considering that the chosen definition of workload is not entirely accurate. The other test cases do not exhibit a clear convergence in position due to the chaotic nature of their motion.

5. Parallelization

In this section we will discuss several different parallelization algorithms for point vortex dynamics. All of them are based on rearranging the pairwise flow maps of (9) and grouping together subsets of the composition. By choosing groups that consist of disjoint pairs, the order of evaluation does not affect the result and the pairs may be evaluated in parallel, without loss of accuracy.

As an illustration, let us consider a system with just 4 vortices, labelled² $i = 0, 1, 2, 3$. The ordering C must contain all 6 possible vortex pairings once. This is similar to round robin tournament scheduling, where all competitors play each other once. Such a scheme is presented in Table 2, each round is performed from top to bottom before moving onto the next round. The ordering is constructed by fixing the last vortex, number 3, in position, while the other vortices rotate over the remaining positions. Note that during each “round” each vortex occurs only once, this will be important later on when discussing parallelizations.

The Strang splitting for this ordering is written as

$$\Phi_{2\tau}^S = \phi_\tau^{03} \circ \phi_\tau^{12} \circ \phi_\tau^{13} \circ \phi_\tau^{20} \circ \phi_\tau^{23} \circ \phi_\tau^{01} \circ \phi_\tau^{01} \circ \phi_\tau^{23} \circ \phi_\tau^{20} \circ \phi_\tau^{13} \circ \phi_\tau^{12} \circ \phi_\tau^{03}.$$

For notational convenience we use a time step of 2τ . Because this ordering is symmetrical and each of the ϕ_τ^{ij} is self-adjoint, the resulting method $\Phi_{2\tau}^S$ is a second order accurate method.

The more general ordering of N interacting vortices can be arranged in the same way. We write the s th pair of round r in the ordering C_N as $C_N^{sr} = (A_{sr}, B_{sr})$. Following the same pattern as for the four-vortex system results in the expressions

$$A_{sr} = \text{mod}(s + r, N - 1) \quad \text{for } r = 0, 1, \dots, N - 2, s = 0, 1, \dots, P - 1$$

$$B_{sr} = \begin{cases} N - 1 & \text{for } r = 0, 1, \dots, N - 2, s = 0 \\ \text{mod}(N - 1 - s + r, N - 1) & \text{for } r = 0, 1, \dots, N - 2, s = 1, 2, \dots, P - 1, \end{cases}$$

² For convenience in the ensuing modular arithmetic, we switch to indexing from zero.

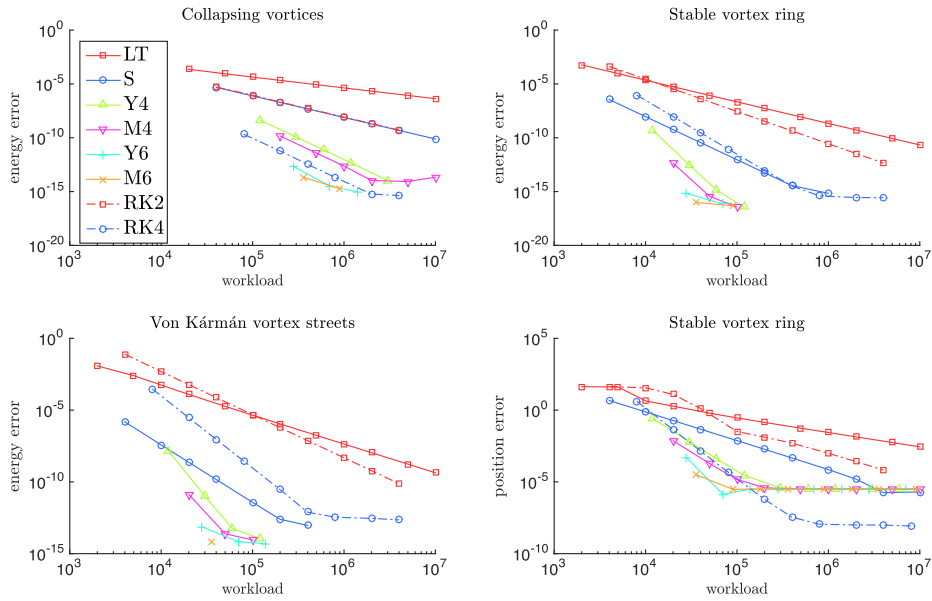


Fig. 11. A comparison of the error plotted against computational workload for various methods. Top left shows the energy error for the collapsing vortices case, top right that for the stable vortex ring, lower left for the Von Kármán vortex streets; lower right shows the *position* error for the stable vortex ring. All errors are time averages over the simulation lengths as described in the text.

Table 3
Round-robin scheme for ordering vortex pairs in a ten-vortex system.

Round	0	1	2	3	4	5	6	7	8
Pair	0–9 1–8 2–7 3–6 4–5	1–9 2–0 3–8 4–7 5–6	2–9 3–1 4–0 5–8 6–7	3–9 4–2 5–1 6–0 7–8	4–9 5–3 6–2 7–1 8–0	5–9 6–4 7–3 8–2 0–1	6–9 7–5 8–4 0–3 1–2	7–9 8–6 0–5 1–4 2–3	8–9 0–7 1–6 2–5 3–4

where P is the number of parallel operations and $N = 2P$ is the number of vortices. An odd number of vortices can be simulated by including a dummy vortex of zero circulation. A ten-vortex system exemplifies this ordering in Table 3.

5.1. Complete parallelization

The arrangement of pairwise vortex interactions into “rounds” in Table 2 is not just convenient for notation. Let us denote the composition for each round by $\psi_\tau^{ij,kl} = \phi_\tau^{ij} \circ \phi_\tau^{kl}$, for i, j, k, ℓ all different. In this notation the first round of Table 2 is the composition $\psi_\tau^{03,12} = \phi_\tau^{03} \circ \phi_\tau^{12}$. But as the evolution of vortex pair (0, 3) is independent of that of pair (1, 2), the order of the operations is irrelevant, i.e. ϕ_τ^{03} and ϕ_τ^{12} commute. This means that while evolving the system, these two vortex pairs can be evaluated simultaneously, in parallel. Let us stress this fact by using the notation $\chi_\tau^{ij,kl}$ to denote the time- τ flow map of the evolution of vortex pairs (i, j) and (k, l) in *either* order. The Lie–Trotter splitting method where both pairs for each round are evaluated in parallel is then written as

$$\Phi_\tau^{LT||} = \chi_\tau^{03,12} \circ \chi_\tau^{13,20} \circ \chi_\tau^{23,01},$$

and the Strang splitting

$$\Phi_{2\tau}^{S||} = \chi_\tau^{03,12} \circ \chi_\tau^{13,20} \circ \chi_\tau^{23,01} \circ \chi_\tau^{23,01} \circ \chi_\tau^{13,20} \circ \chi_\tau^{03,12}.$$

Because each of the $\chi_\tau^{ij,kl}$ compositions is self-adjoint, this method is again second-order accurate. Note that the order of the underlying pairwise interactions is now no longer necessarily symmetric.

When integrating a system with four vortex, each round of the round-robin scheme contains two pairwise interactions that can be performed in parallel. More generally, this scheme allows P processors to evaluate a system with $2P$ vortices. But each process only evaluates a single vortex pair interaction per round, meaning there is a lot of communication relative to the amount of work done each round.

Table 4
Reduced communication round-robin scheme for ordering vortex pairs in a ten-vortex system.

Round	0	1	2	3	4	5	6	7	8
Pair	0–9	4–9	8–9	3–9	7–9	2–9	6–9	1–9	5–9
	1–8	0–8	0–7	8–7	8–6	7–6	7–5	6–5	6–4
	2–7	1–7	1–6	0–6	0–5	8–5	8–4	7–4	7–3
	3–6	2–6	2–5	1–5	1–4	0–4	0–3	8–3	8–2
	4–5	3–5	3–4	2–4	2–3	1–3	1–2	0–2	0–1

5.2. Reducing communication

It is not necessary for each round to be finished completely before starting evaluating the next. For the s th vortex pair of round r , (A_{sr}, B_{sr}) , to be evaluated, it is only necessary to wait for these two vortices to have been evaluated in the previous round $r - 1$. The evaluation of the remaining vortex pairs commutes with the evaluation of (A_{sr}, B_{sr}) .

If thread s in round r has to wait only for vortices A_{sr} and B_{sr} to be done in round $r - 1$, it is beneficial to choose the ordering such that one of the two vortices is evaluated on the same thread as in the previous round. This means that each thread has to wait for only one other vortex pair of the previous round. We construct such an ordering based on the construction used in the previous section. Again the vortex $2P - 1$ is kept fixed in place, but now the other vortices rotate through $P - 1$ positions. In the ten-vortex example, this can be seen as jumping straight to round 4 in Table 3 after the first round. In doing so, all vortices on the right under the fixed vortex move to the left, but in reversed order. By subsequently reversing the order of these pairs in all odd rounds, each thread needs to wait only for one other vortex during each round. The resulting ordering $\tilde{C}_N^{sr} = (\tilde{A}_{sr}, \tilde{B}_{sr})$ is given by

$$(\tilde{A}_{sr}, \tilde{B}_{sr}) = \begin{cases} (A_{s\tilde{r}}, B_{s\tilde{r}}) & \text{for } r \text{ even} \\ (A_{s\tilde{r}}, B_{s\tilde{r}}) & \text{for } r \text{ odd, } s = 0 \\ (B_{s\tilde{r}}, A_{s\tilde{r}}) & \text{for } r \text{ odd, } s \neq 0, \end{cases}$$

where $\tilde{r} = \text{mod}(r(P - 1), N - 1)$ and $\tilde{s} = P - s$ (recall that $N = 2P$). An example ordering with 10 vortices is presented in Table 4.

5.3. Hierarchical parallelization

An efficient parallelization finds a balance between reducing the work load per processor and reducing the time required for communication between threads. Without parallelization there is no communication time, but the workload per thread is largest. The parallelization scheme in Section 5.1 represents the other extreme – where communication time dominates the total computation time. In this section we develop a scheme that has adjustable parallelization. Both the “complete parallelization” and the non-parallel computation are special cases of this scheme.

The method we develop here is constructed as a hierarchy of methods that act on a hierarchical system of *leagues* of vortices. Given a number of parallel threads P and a number of vortices N to be evaluated, we place $L = \frac{N}{2P}$ vortices in each lowest-level league in the hierarchy. If the fraction $\frac{N}{2P}$ is not integer, we introduce a number of dummy vortices with zero circulation to increase N such that it is a multiple of $2P$.

We also define a factorization $P = p_1 \times p_2 \times \dots \times p_n$. We will use the prime factorization, but depending on system architecture a different factorization may be desirable. Given this factorization, we construct a hierarchy of leagues as follows: there are $2p_1$ level-one leagues, each consisting of p_2 level-two leagues, each consisting of p_3 level-three leagues and so on. The level- n leagues are the lowest level and consist of L vortices each.

We denote a level- m league in this tree by \mathcal{L}_m^k , where the vector $k = (k_1, k_2, \dots, k_m) \in \mathbb{R}^m$ denotes the ancestry of the league. In other words, k_m denotes the current child of parent league k_{m-1} of grandparent k_{m-2} and so forth. Each k_i is in the range $\{0, \dots, p_i - 1\}$, determined by the chosen factorization of P .

The hierarchical splitting uses the fact that each vortex pair is either an interaction *between* two leagues, with one vortex from each of the two leagues, or *within* a league, with both vortices from that same league. Looking at the top level first, we see that this means that we have interactions between the $2p_1$ groups, and interactions within each of these groups. One way to do this would be by first evaluating all $p_2(2p_2 - 1)$ possible combinations of level-one leagues according to an ordering as in Section 5.2 and then evaluating the interactions within each level-one league separately. This can, however, have the unfortunate effect of wasted computation time if one of the factors of the factorization is odd. Instead we evaluate the interaction between level-one leagues according to the ordering of Section 5.2, but we omit the final round. The interactions between these omitted leagues are combined with their interior interactions. For the top level, this forms a system of N/p_1 vortices that has to be evaluated completely.

Let $\Lambda_1[(k, \ell)]$ denote the interaction between all vortices in level-one league \mathcal{L}_1^k and those in \mathcal{L}_1^ℓ and let the composition $\Phi_1[(k, \ell)]$ denote the evaluation of *all* interactions within or between level-one leagues \mathcal{L}_1^k and \mathcal{L}_1^ℓ . The Lie–Trotter splitting




Fig. 12. Tree representation of the hierarchical parallelization with $L = 2$ and $P = 4$, using prime factorization: $p_1 = p_2 = 2$. Only a few branches are expanded for clarity.

Table 5

Vortex interaction pairs between vortices $\{0, \dots, 3\}$ and $\{4, \dots, 7\}$.

Round	0	1	2	3
Pair	0–4 1–5 2–6 3–7	0–5 1–6 2–7 3–4	0–6 1–7 2–4 3–5	0–7 1–4 2–5 3–6

is then given by (we drop the subscript indicating the time step for notational convenience)

$$\Phi_0^{LT} = \prod_{r=0}^{2p_1-3} \prod_{s=0}^{p_1-1} \Lambda_1[\tilde{C}_{2p_1}^{sr}] \circ \prod_{s=0}^{p_1-1} \Phi_1[\tilde{C}_{2p_1}^{s,2p_1-2}]. \tag{21}$$

This composition is represented graphically as the top level of the tree diagram in Fig. 12. The compositions $\Lambda_1[(k, \ell)]$ and $\Phi_1[(k, \ell)]$ both follow a recursive definition detailed below.

5.3.1. Interactions between leagues

The flow map $\Lambda_q[(\mathbf{k}, \ell)]$ represents all possible interactions with one vortex in league $\mathcal{L}_q^{\mathbf{k}}$ and one in league \mathcal{L}_q^{ℓ} . Using the hierarchical ordering of the vortices, we evaluate all such pairings by evaluating all possible combination between leagues one level down. This is represented graphically by the branches on the left in Fig. 12.

The ordering D_{2p} contains all possible pairs (i, j) with $i \in \{0, \dots, P - 1\}$ and $j \in \{P, \dots, 2P - 1\}$.

$$D_{2p}^{sr} = (s, P + \text{mod}(s + r, P)), \quad s, r \in \{0, \dots, P - 1\}.$$

An example with $P = 4$ is presented in Table 5.

With this, we define the recursive definition for the interaction between leagues

$$\Lambda_{q-1}[(\mathbf{k}, \ell)] = \prod_{r=0}^{p_q-1} \prod_{s=0}^{p_q-1} \Lambda_q[I_{\mathbf{k},\ell}^{p_q}(D_{2p_q}^{sr})], \tag{22}$$

where the mapping

$$I_{\mathbf{k},\ell}^{p_q}(k) = \begin{cases} (\mathbf{k}, k) & \text{for } k < p_q \\ (\mathbf{l}, k - p_q) & \text{for } k \geq p_q \end{cases} \tag{23}$$

associates each of the vortices from an ordering with $2p_q$ vortices with children in the groups $\mathcal{L}_{q-1}^{\mathbf{k}}$ and \mathcal{L}_{q-1}^{ℓ} . The lowest level interaction between leagues is given by

$$\Lambda_n^{\mathbf{k},\ell} = \prod_{r=0}^{L-1} \prod_{s=0}^{L-1} \phi_{\mathbf{k}\ell}^{L(D_{2L}^{sr})}. \tag{24}$$

Note that we also use the ancestry to enumerate the vortices, rather than linear indexing.

5.3.2. Interactions within leagues

The definition of $\Phi_1[(k, \ell)]$ is given recursively, so it follows the same pattern as $\Phi_{q-1}[(\mathbf{k}, \ell)]$ with $\mathbf{k}, \ell \in \mathbb{R}^{q-1}$, that is

$$\Phi_{q-1}[(\mathbf{k}, \ell)] = \prod_{r=0}^{2p_q-3} \prod_{s=0}^{p_q-1} \Lambda_q[I_{\mathbf{k}\ell}^{p_q}(\tilde{C}_{2p_q}^{sr})] \circ \prod_{s=0}^{p_q-1} \Phi_q[I_{\mathbf{k}\ell}^L(\tilde{C}_{2p_q}^{s,2p_q-2})], \tag{25}$$

for $q = 1, 2, \dots, n$. The mapping $I_{\mathbf{k}\ell}^{p_q}$ is the same as in (23). Note that the definition of Φ_0^{LT} in (21) is in fact equivalent to (25) for $q = 1$.

At the lowest level, the league \mathcal{L}_n^k no longer consists of leagues, but of L point vortices. So when considering the interaction of all vortices within two lowest-level leagues \mathcal{L}_n^k and \mathcal{L}_n^ℓ , this constitutes evaluating a $2L$ point vortex system.

$$\Phi_n[\mathbf{k}, \ell] = \prod_{r=0}^{2L-2} \prod_{s=0}^{L-1} \phi_{\mathbf{k}\ell}^{p_n}(C_L^{sr}). \tag{26}$$

By (recursively) substituting Eqs. (22)–(26) into (21) we find a Lie–Trotter splitting for the system with $N = 2LP$ vortices.

5.3.3. Symmetric splitting

To construct a symmetric splitting we need to compose the splitting Φ_0^{LT} of (21) with its adjoint $\Phi_0^{LT,*}$ as

$$\Phi_0^S = \Phi_0^{LT} \circ \Phi_0^{LT,*}.$$

The adjoint of the Lie–Trotter splitting follows from a reversal of the order of the operators – insofar as this is necessary – and taking the adjoint of each of the interior operators.

$$\Phi_0^{LT,*} = \prod_{s=0}^{p_1-1} \Phi_1^*[\tilde{C}_{2p_1}^{s,2p_1-2}], \circ \prod_{r=2p_1-3}^0 \prod_{s=0}^{p_1-1} \Lambda_1^*[\tilde{C}_{2p_1}^{sr}].$$

The adjoint of $\Phi_{q-1}[\mathbf{k}, \ell]$ for $q \in \{2, \dots, n\}$ of (25) is taken in a similar fashion to be

$$\Phi_{q-1}^*[\mathbf{k}, \ell] = \prod_{s=0}^{p_q-1} \Phi_q^*[I_{\mathbf{k}\ell}^L(\tilde{C}_{2p_q}^{s,2p_q-2})], \circ \prod_{r=2p_q-3}^0 \prod_{s=0}^{p_q-1} \Lambda_q^*[I_{\mathbf{k}\ell}^{p_n}(\tilde{C}_{2p_q}^{sr})].$$

For the interactions between leagues given by $\Lambda_q[\mathbf{k}, \ell]$ in (22) we need only reverse the order of the outer product and take the adjoint of the Λ_{q+1} inside. The adjoint of the lowest level operators in (24) and (26) is achieved by reversing the order of the outside product over r .

The complete parallelization of Section 5.1 follows from the choice $L = 1, N = 2P = 2p_1$, whereas the non-parallel scheme follows from $P = 1, N = 2L$.

5.4. Implementation details

All simulations were performed on a desktop Macintosh MacPro running OS X 10.9.5. The system has two Intel Xeon 2.93 GHz processors with six cores each. The system has 32 GB shared memory, 12 MB L3 cache (per processor) and 256 kB L2 cache (per core). The programming code was written in C, compiled into stand-alone applications using Matlab’s mex with llvm-gcc-4.2. The motivation for this compiling strategy is to allow for easier transfer of data to Matlab, which was used for all post-processing and data-analysis purposes. All source files are available at <https://github.com/KeithWM/Poissonpv>.

5.5. Timing experiments

To investigate the practical use for the different splitting schemes and vortex orderings we perform several experiment measuring the required time for different simulations. In all cases the configuration consists of eight strong vortices, four positive ($\Gamma = 1$) and four negative ($\Gamma = -1$), with the remaining vortices weaker with circulation $\pm \frac{1}{5}$ in equal numbers. The initial conditions are chosen such that the total energy and momentum are zero. All timing experiments are performed five times independently, to confirm that the results are not influenced by external factors.

The efficiency of the parallelized splitting method is best represented by studying the scaling of the method to large number of vortices while linearly increasing the number of threads. This implies each thread always operates on the same number of vortices and consequently the time taken for each round does not change. The number of rounds does change when the number of vortices is increased, but this is only a linear increase.

For these simulations we use a Strang splitting with a time step of $\tau = 0.001$ and simulate up to only $T = 0.01$. This short time makes the timing results for small systems somewhat noisy, but it means that simulating a large system remains feasible, even when using only a few threads. We present results for the splitting as detailed in Section 5.3 as well as results for the same principle with a small modification to reduce communication.

In Fig. 13 we display the time required to simulate a systems with 64 (left) and 1024 (right) vortices per thread, when varying the number of threads from 1 to 12. There is indeed only a linear increase in the workload. In the case with 64 vortices per thread, the benefits of parallelization outweigh the costs roughly from a system size of 192 onwards. With 1024 vortices per thread there is an immediate benefit to parallelization.

The different orderings appear of little to no effect on the speed of computation. We attribute this to the fact that the intended improvement by rearranging the ordering would arise only in situations where there is a significant difference in




Fig. 13. Wall clock time required plotted against number of vortices N , keeping the number of vortices per thread fixed (64 on left, 1024 on right). A solid black line represents the time needed to perform the simulations on a single core and dashed black lines represent $\mathcal{O}(N)$ and $\mathcal{O}(N^2)$ scaling.




Fig. 14. Wall clock time plotted against number of vortices N , using different numbers of threads. The times are normalized by N^2 , thus estimating the workload per vortex pair. All simulations are performed using a Strang splitting.

the time required to evaluate the different pairwise interactions. This is certainly not the case when each thread is evaluating many pairwise interactions each round.

The same set-up is used to compare the speeds using different combinations of system size and thread counts (N and P respectively). The number of vortices per group is then chosen to be $L = \lceil \frac{N}{2P} \rceil$, where $\lceil \cdot \rceil$ indicates the ceiling function. This rounding up may lead to the introduction of dummy vortices, but this will have minimal effect on computation time in large systems.

Simulation times for system sizes ranging from $N = 24$ to $N = 24\,000$ and thread counts from one to twelve are compared in Fig. 14. The times are normalized by N^2 representing an estimation of the simulation time per vortex interaction. For increasing system size this quantity approaches a fixed constant for a fixed number of threads. This is clearly visible for the smaller thread counts. A general trend that larger systems are evaluated fastest using more threads is apparent, but there are some notable exceptions. Most standing out is the speed of the 12 thread computation of a system with 180 vortices. This is probably a result of the efficiency of evaluating 8 vortices per group due to memory management.

We investigate the error convergence for the different orderings in a large system with 360 vortices – 8 strong and 352 weak as before. We run a short simulation up to $T = 0.1$ time units. The time step used is varied from $\tau = 10^{-2}$ down to $\tau = 10^{-5}$. We consider the energy error $\varepsilon_H(t) = |H(t) - H(0)|$ and take the mean $\bar{\varepsilon}_H = \frac{\tau}{T} \sum_{i=1}^l \varepsilon_H(i\tau)$ over the simulation interval.

The mean energy error is plotted in Fig. 15; it is compared against the time step in the left-hand panel, and against the simulation time on the right. As expected the Lie–Trotter splittings show first order convergence and the Strang splittings second order. The heuristic modification made to the separated ordering has little effect on accuracy or speed.

Over the range of time steps considered, Strang splitting outperforms Lie–Trotter splitting not only in terms of accuracy for a given time step, but also in terms of accuracy against computational cost. This suggests higher order methods could be even more efficient. This is investigated by comparing fourth and sixth order methods against the first order Lie–Trotter and second order Strang splittings for the same problem. The higher order methods follow from a composition of a number of Strang steps of different sizes [36]. For both fourth and sixth order we consider the methods with the minimal number

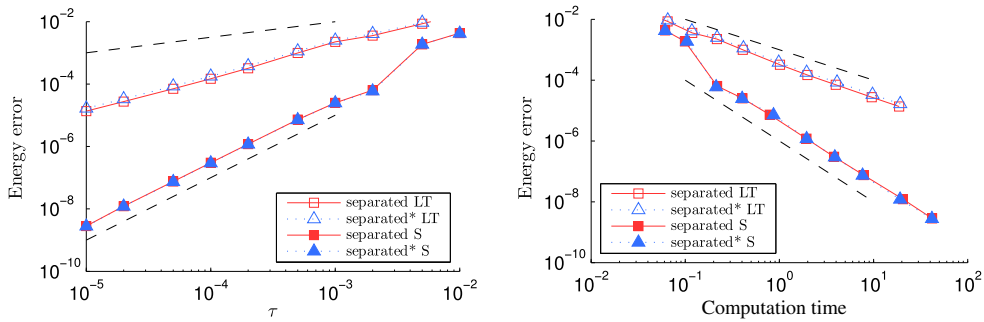


Fig. 15. Comparison between different orderings in the Lie–Trotter (LT) and Strang (S) splitting schemes for a system with 360 vortices. On the left the error in energy is compared against different time steps, on the right it is compared against simulation time required. All simulations use eight parallel threads. Black dashed lines represent first and second order convergence.

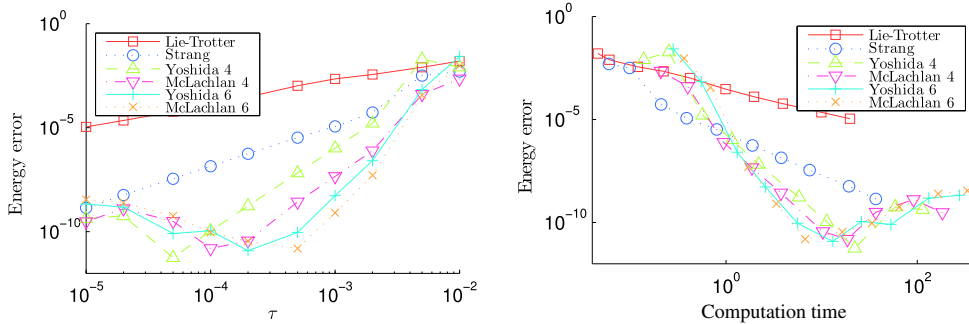


Fig. 16. Comparison between different splitting schemes for a system with 360 vortices. On the left the error in energy is compared against different time steps, on the right it is compared against simulation time required. All simulations are performed on eight parallel threads.

of stages presented by Yoshida [37], as well as the methods of same order but with smaller error coefficients found by McLachlan [38].

The results are shown in Fig. 16, again with energy error versus time step on the left and energy error versus computational time on the right. All methods exhibit the expected error convergence, albeit for a limited range of time steps. The methods due to McLachlan [38] have considerably smaller error for the same time step. When considering the benefit for the same computational load the difference is much reduced, but still in favour of McLachlan’s schemes.

6. Conclusion and outlook

Solutions to ideal fluid flow with a singular measure vorticity field result in a Poisson system describing the motion of the vortex centres. By splitting the Hamiltonian of such a point vortex system into the interactions of individual vortex pairs we construct a splitting method. By composing the basic Lie–Trotter splitting with its adjoint (the same method with reversed ordering) a symmetric Strang splitting with second order accuracy is constructed. Solution trajectories from these schemes provide exact solutions to modified Poisson problems with the original bracket, thereby respecting the Casimirs of the original dynamics. The modified Hamiltonian is studied with the use of backward error analysis, showing that inaccuracies in the Hamiltonian occur with close vortex interactions. The conservation properties are studied in a number of test cases, including those considered by Vankerschaver and Leok [26].

The splitting method can also be rearranged into the interactions between groups of vortices, allowing parallelization of the workload. This reduces the natural quadratic scaling of computation time with system size to linear scaling when the number of processors is increased accordingly. The ordering of the pairwise evaluations can be modified to reduce communication overhead. The method therefore extends well to distributed memory implementations for large systems, allowing the method to be used for engineering applications or for studying statistical mechanics of point vortex dynamics.

The Strang splitting is also used as a basis for constructing higher order methods following Yoshida [37] and McLachlan [38]. When higher accuracy is desired, these methods are more efficient in terms of computational power than Lie–Trotter or Strang splitting.

In the statistical mechanics study of point vortex systems it is usually assumed that the vortex strength decreases as the number of vortices is increased. This follows immediately if it is desired that the enstrophy $Z = \int_{\mathbb{S}^2} \omega^2 dS = \sum_{i=1}^N \Gamma_i^2$ matches the enstrophy of some real fluid. The same can be expected of point vortex models used as predictive models. This means that when the number of vortices is increased, the time step need not be decreased, thus preserving the linear increase in computational load.

Acknowledgement

The work of the first author was supported by a grant financed by the Netherlands Organisation for Scientific Research (NWO), as part of research programme 613.001.009.

References

- [1] H. Helmholtz, On integrals of the hydrodynamical equations which express vortex motion, *Phil. Mag.* 33 (1867) 485–512.
- [2] H. Lamb, *Hydrodynamics*, sixth ed., Dover Publications, New York, 1932.
- [3] P.G. Saffman, *Vortex Dynamics*, Cambridge University Press, 1992.
- [4] P.K. Newton, *The N-Vortex Problem: Analytical Techniques*, first ed., Springer-Verlag, New York, 2001.
- [5] M.I. Jamalooden, P.K. Newton, The N -vortex problem on a rotating sphere. II. heterogeneous platonic solid equilibria, in: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 462, The Royal Society, 2006, pp. 3277–3299.
- [6] P.K. Newton, T. Sakajo, The N -vortex problem on a rotating sphere. III. Ring configurations coupled to a background field, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* 463 (2080) (2007) 961–977.
- [7] P.K. Newton, T. Sakajo, The N -vortex problem on a rotating sphere: IV. Ring configurations coupled to a background field, *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* 463 (2080) (2007) 961–977.
- [8] H. Aref, 150 years of vortex dynamics, *Theor. Comput. Fluid Dyn.* (ISSN: 0935-4964) 24 (1–4) (2010) 1–7. <http://dx.doi.org/10.1007/s00162-009-0178-6>.
- [9] P.K. Newton, Point vortex dynamics in the post-Aref era, *Fluid Dynam. Res.* 46 (3) (2014) 031401.
- [10] L. Onsager, Statistical hydrodynamics, *Nuovo Cimento* (1943–1954) 6 (1949) 279–287.
- [11] G. Joyce, D. Montgomery, Negative temperature states for the two-dimensional guiding-centre plasma, *J. Plasma Phys.* 10 (01) (1973) 107–121.
- [12] G. Joyce, D. Montgomery, Simulation of the “negative temperature” instability for line vortices, *Phys. Lett. A* 39 (5) (1972) 371–372.
- [13] Y.B. Pointin, T.S. Lundgren, Statistical mechanics of two-dimensional vortices in a bounded container, *Phys. Fluids* 19 (10) (1976) 1459–1470. <http://dx.doi.org/10.1063/1.861347>.
- [14] G. Eyink, H. Spohn, Negative-temperature states and large-scale, long-lived vortices in two-dimensional turbulence, *J. Stat. Phys.* 70 (3–4) (1993) 833–886.
- [15] P.-L. Lions, A. Majda, Equilibrium statistical theory for nearly parallel vortex filaments, *Comm. Pure Appl. Math.* 53 (1) (2000) 76.
- [16] M.J. Grote, A.J. Majda, Crude closure dynamics through large scale statistical theories, *Phys. Fluids* (1994-present) 9 (11) (1997) 3431–3442.
- [17] P.-H. Chavanis, J. Sommeria, R. Robert, Statistical mechanics of two-dimensional vortices and collisionless stellar systems, *Astrophys. J.* 471 (1) (1996) 385.
- [18] O. Bühler, Statistical mechanics of strong and weak point vortices in a cylinder, *Phys. Fluids* 14 (7) (2002) 2139–2149.
- [19] G.-H. Cottet, P.D. Koumoutsakos, *Vortex Methods: Theory and Practice*, Cambridge University Press, 2000.
- [20] P. Chatelain, A. Curioni, M. Bergdorf, D. Rossinelli, W. Andreoni, P. Koumoutsakos, Billion vortex particle direct numerical simulations of aircraft wakes, *Comput. Methods Appl. Mech. Engrg.* 197 (13) (2008) 1296–1304.
- [21] D. Rossinelli, M. Bergdorf, G.-H. Cottet, P. Koumoutsakos, GPU accelerated simulations of bluff body flows using vortex particle methods, *J. Comput. Phys.* 229 (9) (2010) 3316–3333.
- [22] G. Winckelmans, R. Cocle, L. Dufresne, R. Capart, Vortex methods and their application to trailing wake vortex simulations, *C.R. Phys.* 6 (4) (2005) 467–486.
- [23] D. Rossinelli, P. Koumoutsakos, Vortex methods for incompressible flow simulations on the GPU, *Vis. Comput.* 24 (7–9) (2008) 699–708.
- [24] T. Sakajo, An extension of Draghicescu’s fast tree-code algorithm to the vortex method on a sphere, *J. Comput. Appl. Math.* 225 (1) (2009) 158–171.
- [25] J.P. Boyd, C. Zhou, Three ways to solve the Poisson equation on a sphere with Gaussian forcing, *J. Comput. Phys.* 228 (13) (2009) 4702–4713.
- [26] J. Vankerschaver, M. Leok, A novel formulation of point vortex dynamics on the sphere: geometrical and numerical aspects, *J. Nonlinear Sci.* 24 (1) (2014) 1–37.
- [27] R.I. McLachlan, Symplectic integration of Hamiltonian wave equations, *Numer. Math.* 66 (1) (1993) 465–492.
- [28] A.J. Majda, X. Wang, *Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows*, first ed., Camb. Univ. Press, Camb., UK, 2006.
- [29] D.D. Holm, *Geometric Mechanics*, Vol. 2, Imperial College Press, London, 2008.
- [30] C.-S. Liu, K.-C. Chen, C.-S. Yeh, A mathematical revision of the Landau–Lifshitz equation, *J. Mar. Sci. Technol.* 17 (3) (2009) 228–237.
- [31] L. Faddeev, L. Takhtajan, *Hamiltonian Methods in the Theory of Solitons*, in: *Springer Series in Soviet Mathematics*, Springer-Verlag, New York, 1987.
- [32] R.V. Abramov, A.J. Majda, Statistically relevant conserved quantities for truncated quasigeostrophic flow, *Proc. Natl. Acad. Sci.* 100 (7) (2003) 3841–3846.
- [33] S. Dubinkina, J. Frank, Statistical mechanics of Arakawa’s discretizations, *J. Comput. Phys.* 227 (2) (2007) 1286–1305.
- [34] S. Dubinkina, J. Frank, Statistical relevance of vorticity conservation in the Hamiltonian particle-mesh method, *J. Comput. Phys.* 229 (7) (2010) 2634–2648.
- [35] G. Patrick, Dynamics of perturbed relative equilibria of point vortices on the sphere or plane, *J. Nonlinear Sci.* 10 (3) (2000) 401–415.
- [36] E. Hairer, C. Lubich, G. Wanner, *Geometric Numerical Integration: Structure-preserving Algorithms for Ordinary Differential Equations*, Vol. 31, Springer, 2006.
- [37] H. Yoshida, Construction of higher order symplectic integrators, *Phys. Lett. A* 150 (5) (1990) 262–268.
- [38] R.I. McLachlan, On the numerical integration of ordinary differential equations by symmetric composition methods, *SIAM J. Sci. Comput.* 16 (1) (1995) 151–168.
- [39] B. Leimkuhler, S. Reich, *Simulating Hamiltonian Dynamics*, first ed., in: *Camb. Monogr. on Appl. and Comput. Math.*, Camb. Univ. Press, Camb., UK, 2004.
- [40] K.W. Myerscough, J. Frank, B. Leimkuhler, Least-biased correction of extended dynamical systems using observational data, 2014. arXiv Preprint arXiv:1411.6011.
- [41] R. Kidambi, P.K. Newton, Motion of three point vortices on a sphere, *Physica D* 116 (1) (1998) 143–175.
- [42] T. Sakajo, Non-self-similar, partial, and robust collapse of four point vortices on a sphere, *Phys. Rev. E* 78 (1) (2008) 016312.
- [43] L.M. Polvani, D.G. Dritschel, Wave and vortex dynamics on the surface of a sphere, *J. Fluid Mech.* 255 (1993) 35–64.
- [44] G. Chamoun, E. Kanso, P.K. Newton, Von Kármán vortex streets on the sphere, *Phys. Fluids* (1994-present) 21 (11) (2009) 116603.
- [45] O. Bühler, A Brief Introduction to Classical, Statistical and Quantum Mechanics, first ed., in: *Courant Lect. Notes*, Am. Math. Soc./Courant Inst. of Math. Sci., Providence, RI, 2006.