Fast and accurate evaluation of Biot-Savart integrals over spatial curves

Juan Ignacio Polanco

FAST AND ACCURATE EVALUATION OF BIOT–SAVART INTEGRALS OVER SPATIAL CURVES

JUAN IGNACIO POLANCO*

Abstract. The Biot–Savart law is relevant in physical contexts including electromagnetism and fluid dynamics. In the latter case, when the rotation of a fluid is confined to a set of very thin vortex filaments, this law describes the velocity field induced by the spatial arrangement of these objects. The Biot–Savart law is at the core of vortex methods used in the simulation of classical and quantum fluid flows. Naïve methods are inefficient when dealing with large numbers of vortex elements, which makes them inadequate for simulating turbulent vortex flows. Here we exploit a direct analogy between the Biot–Savart law and electrostatics to adapt Ewald summation methods, routinely used in molecular dynamics simulations, to vortex filament simulations in three-dimensional periodic domains. In this context, the basic idea is to split the induced velocity onto (i) a coarse-grained velocity generated by a Gaussian-filtered vorticity field, and (ii) a short-range correction accounting for near-singular behaviour near the vortices. The former component can be accurately and efficiently evaluated using the nonuniform fast Fourier transform algorithm. Analytical accuracy estimates are provided as a function of the parameters entering the method. We also discuss how to properly account for the finite vortex core size in kinetic energy estimations. Using numerical experiments, we verify the accuracy and the conservation properties of the proposed approach. Moreover, we demonstrate the $O(N \log N)$ complexity of the method over a wide range of problem sizes N, considerably better than the $O(N^2)$ cost of a naïve approach.

Key words. Biot–Savart, Ewald summation, nonuniform fast Fourier transform, vortex filament model, quantum vortices

MSC codes. 65D07, 65L05, 70H05, 76B47, 76M23, 76Y05

1. Introduction.

The Biot–Savart law is well known for describing the magnetic field generated by a steady electric current. It is also very relevant in fluid dynamics, where it allows to obtain the fluid velocity induced by a known vorticity field. We are interested here in the specific case where the electric current or the vorticity field are confined to spatial curves in three-dimensional space. This is a commonly encountered problem in undergraduate electromagnetism lectures, where the question is what is the magnetic field induced by a current flowing through a thin conducting wire. In the fluid dynamics case, the equivalent would be a very thin vortex filament, which can be a reasonable idealised model for vortices found in the flow of viscous fluids such as water or air. In fact, it is also a very accurate description of vortices in certain superfluids such as liquid helium-4 near the absolute zero, where rotational motion is confined to so-called quantum vortices of atomic-size thickness (the vortex core radius is $a \approx 10^{-10}\,\mathrm{m}$), and the circulation (or strength) of each vortex takes a constant value $\kappa \approx 9.97 \times 10^{-8}\,\mathrm{m}^2/\mathrm{s}$ dictated by quantum mechanical constraints. The velocity induced at a point $\boldsymbol{x}$ away from a vortex core is then given by the Biot–Savart law,

$$\boldsymbol{v}(\boldsymbol{x}) = \frac{\kappa}{4\pi} \oint_{\mathcal{C}} \frac{(\boldsymbol{s}-\boldsymbol{x}) \times \mathrm{d}\boldsymbol{s}}{|\boldsymbol{s}-\boldsymbol{x}|^3}, \tag{1.1}$$

where $\mathcal{C}$ denotes one or more oriented curves representing the vortex filament geometry, and $\boldsymbol{s} \in \mathcal{C}$ denotes a vortex location.

In classical fluid dynamics, the Biot–Savart law is at the core of vortex methods [24, 52] used to describe incompressible viscous flows. These are commonly used in aerodynamics applications [78], but have also been applied to problems as varied as

*Univ. Grenoble Alpes, CNRS, Grenoble INP, LEGI, 38000 Grenoble, France (juan-ignacio.polanco@cnrs.fr).

the simulation of self-propelled swimmers in viscous fluids [37]. Note that, in viscous fluids, vortex methods need to account for effects including vorticity diffusion and energy dissipation due to viscosity. To achieve this, these methods typically deal with vortex *particles*, which should not be interpreted as physical objects but as an ensemble of point charges generating a fluid flow. In particular, the connectivity of vortex particles is not relevant to vortex particle methods [52].

Here we focus on the conceptually simpler application of vortex methods to superfluid helium-4, where viscous effects are absent near the absolute zero and vortex filaments are in fact the main physical object of interest. In this context, the Biot–Savart law is the basic ingredient of the vortex filament model (VFM), which is one of the most common approaches for describing superfluid flows [70, 42, 14]. This model is valid at scales much larger than the atomic vortex thickness, and is therefore well adapted for describing macroscopic vortex motion. Numerically, the standard approach for representing vortex filaments consists in discretising them as a series of connected vortex points. The connectivity is required in order to obtain derived geometrical information such as local tangents and curvatures. Each vortex point evolves in time according to the Biot–Savart law. In fact, the integral (1.1) diverges when evaluated on a vortex point $\boldsymbol{x} = \boldsymbol{s}_0$, but the singularity can be avoided by taking into account the finite (but small) vortex core radius a [70] as detailed in section 2.

This work is motivated by the study of quantum turbulence [14], which is a state of superfluid flows characterised by a wide range of energetically active length scales. In this state, collective behaviour – in the form of *bundles* of polarised quantum vortices – has been shown to be responsible for scaling laws comparable to those observed in classical turbulent flows [9, 63]. Numerically, to investigate such a turbulent state, one needs to compute the non-local interactions between large numbers of vortices, requiring in particular a large number N of discrete vortex points. If one explicitly accounts for all pair interactions, obtaining the velocities at the N points has a $O(N^2)$ cost, which quickly limits the size of the systems which can be numerically studied. In the context of particle simulations, this problem has been solved for a long time, using techniques such as Barnes–Hut (BH) trees [15] or the fast multipole method (FMM) [39], which reduce the complexity to $O(N \log N)$ and $O(N)$ respectively. Such techniques have also been applied to vortex methods for classical fluids [24]. In the case of quantum vortex flows, the BH approach adopted by Baggaley and Barenghi in 2012 [7] is to our knowledge the only attempt to accelerate VFM simulations, and remains the state of the art to this day.

The above mentioned methods are mostly adapted to open systems such that the fields induced by the particles (or vortices) decay at infinity. Here we are rather interested in *periodic* infinite systems, in which a finite set of N_f vortices is replicated an infinite number of times in a spatially periodic fashion. Periodic boundary conditions are commonly used to model a variety of physical systems when one wants to describe phenomena far from boundaries. Furthermore, compared to spatially decaying systems, periodicity allows to study spatially homogeneous configurations, in which all regions in space are equally "active" in a statistical sense. The BH and FMM techniques mentioned above can be adapted to simulate periodic boundary conditions, but this is generally achieved by explicitly accounting for periodic images over one or more layers of periodic cells adjacent to the main computational domain [54, 4]. In three dimensions, accounting for just a single periodicity layer means including the effect of the 26 periodic cells in contact with the central domain [42]. This not only increases the computational cost, but also incurs in a truncation error due to neglecting long-range interactions beyond a few domain sizes.

A natural way of dealing with periodicity in Cartesian domains is via a Fourier series representation, as done for example in Fourier pseudo-spectral methods [22, 19]. These take advantage of the fast Fourier transform (FFT) to efficiently evaluate non-local operations. This suggests the idea of using FFTs to evaluate costly far-field interactions between particles or vortices in periodic systems. However, since particles and vortices are respectively represented as 0D and 1D singularities, one cannot directly describe the associated source fields (e.g. electric charge density or vorticity) using a truncated Fourier series, which in practice precludes the use of FFTs. In particle simulations, one way around this issue is provided by *fast Ewald summation* methods, which are commonly used in molecular dynamics simulations to speed-up the evaluation of electrostatic interactions between charged particles in periodic systems [32, 47, 26, 29, 5]. There, the basic strategy is to additively split the singular source field onto a smooth field responsible for far-field interactions and a correction field accounting for interactions between nearby particles. Fast Ewald summation methods are characterised by a $O(N \log N)$ complexity in the number of particles N. In molecular dynamics simulation benchmarks [4], it has been observed that FFT-based Ewald methods can perform slightly better than the periodic FMM at the same accuracy, actually displaying near linear complexity over a wide range of problem sizes N. The same benchmarks also show that, when running at low accuracy levels, the latter can display a slight energy drift over time which is not observed in the former.

The aim of this paper is to adapt Ewald methods to the evaluation of the Biot–Savart integral (1.1) in three-dimensional periodic systems and evaluate the relevance of this approach. The paper begins in section 2 with an introduction to the VFM in the context of quantum vortex dynamics. We also propose an approach for accurately estimating the kinetic energy in periodic systems. Section 3 describes the Ewald-based method used to evaluate Biot–Savart integrals. Our approach takes advantage of the nonuniform fast Fourier transform (NUFFT) algorithm to speed-up computations. In section 4, we provide analytical estimates of the approximation errors incurred by the proposed approach in terms of tunable parameters. The relevance of these estimates is then verified in section 5 using numerical experiments of different test cases. That section also showcases the accuracy of energy and impulse conservation by the method, and finishes with numerical evidence of near linear $O(N \log N)$ complexity over a wide range of problem sizes N. Finally, section 6 is devoted to conclusions.

2. The vortex filament model. The VFM introduced by Schwarz [70] is one of the main approaches used to describe theoretically and numerically the three-dimensional hydrodynamics of quantum fluids such as low-temperature liquid helium [42, 14]. Unlike classical fluids, quantum fluids near the absolute zero are characterised by having zero viscosity and being irrotational almost everywhere. In fact, rotational motions are confined to very thin vortex filaments carrying a quantized circulation $\kappa = h/m$, where h is Planck's constant and m the mass of one atom in the case of helium-4. In other words, a straight vortex filament induces a circular motion of the fluid around it with velocity $v = \kappa/(2\pi r)$, where r is the distance to the vortex.

We restrict our attention to the zero temperature limit. Indeed, in finite temperature superfluid helium, quantum vortices can be seen as coexisting and interacting with a viscous normal fluid [13]. This is described by more complex models (some of them based on the VFM) which are still the subject of active research [14]. A second very important aspect which is not discussed here is the reconnection of vortices when they are close to collision. While this phenomenon is crucial for describing energy

dissipation and quantum turbulence, its modelling is orthogonal to the subject of this work, and can be disregarded as long as vortex elements stay sufficiently far from each other.

2.1. Filaments as oriented curves. In the VFM, one is interested in describing the motion induced by a collection of vortex filaments on themselves. Each filament is represented as an oriented curve $\{\boldsymbol{s}(\xi), 0 \leq \xi \leq \mathcal{L}\}$ where $\boldsymbol{s}(\xi)$ is a location on the curve, ξ is the arc length, and $\mathcal{L}$ is the total length of the curve. Here the curve is represented using the natural arc length parametrisation, such that the unit tangent to the curve is $\boldsymbol{s}'(\xi) = \mathrm{d}\boldsymbol{s}/\mathrm{d}\xi$ with $|\boldsymbol{s}'(\xi)| = 1$ for all ξ. In practice, it is often more convenient to deal with arbitrary parametrisations which will be denoted $\boldsymbol{s}(\tau)$ for $0 \leq \tau \leq \mathcal{T}$, such that $|\mathrm{d}\boldsymbol{s}/\mathrm{d}\tau| \neq 1$ in general. Here $[0, \mathcal{T}]$ is the *main* domain of validity of the parameter τ, meaning that in principle the curve could be evaluated outside of it. This makes sense for closed or periodic curves as detailed below.

We will consider the vortex filaments to be embedded in a triply periodic domain, so that every filament is replicated an infinite number of times along each Cartesian direction. For simplicity, throughout this paper, the domain period is set to L in all directions, but all definitions and results can be readily generalized to different periodicities in each direction. In fluid dynamics, Helmholtz' theorems [46] state that vortex lines cannot end in the fluid. Therefore, in the absence of solid boundaries, they must either be closed curves or extend to infinity. Hence we only consider these two cases, with the additional restriction that infinite curves must be described by a periodic function matching the domain periodicity. Both cases are defined by the property $\boldsymbol{s}(\mathcal{T} + \tau) = \boldsymbol{s}(\tau) + \boldsymbol{n}L$ with $\boldsymbol{n} \in \mathbb{Z}^3$. In particular, closed curves satisfy this property with $\boldsymbol{n} = \boldsymbol{0}$. An example of an infinite periodic curve is $\boldsymbol{s}(\tau) = [\cos(2\pi\tau/L), 0, \tau]$ with $\mathcal{T} = L$. This curve satisfies the above property with $\boldsymbol{n} = [0, 0, 1]$, extending infinitely along the third Cartesian direction.

2.2. The Biot–Savart law. We now consider a set of N_f vortex filaments and their periodic images, and denote $\mathcal{C} = \cup_{j=1}^{N_\mathrm{f}} \mathcal{C}_j$ the set of spatial curves $\mathcal{C}_j$ representing their locations. The vorticity field $\boldsymbol{\omega} = \boldsymbol{\nabla} \times \boldsymbol{v}$ associated to these filaments is then

$$\boldsymbol{\omega}(\boldsymbol{x}) = \kappa \sum_{\boldsymbol{n} \in \mathbb{Z}^3} \oint_{\mathcal{C}} \delta(\boldsymbol{s} - \boldsymbol{x} + \boldsymbol{n}L)\, \mathrm{d}\boldsymbol{s} \tag{2.1}$$

where $\delta(\boldsymbol{x})$ is the Dirac delta function, and the infinite sum over $\boldsymbol{n}$ accounts for the periodic vortex images. Here κ is the circulation of the vortex filaments, related to the magnitude of the velocity induced by them. Throughout this work κ will be taken to be constant (as is actually the case of quantum vortices), but in principle one could also consider a variable circulation $\kappa(\boldsymbol{s})$.

Inverting the curl operator leads to the Biot–Savart law (1.1) describing the velocity induced by the set of vortex filaments on a point $\boldsymbol{x}$. Including periodicity effects, this law writes

$$\boldsymbol{v}(\boldsymbol{x}) = \frac{\kappa}{4\pi} \sum_{\boldsymbol{n} \in \mathbb{Z}^3} \oint_{\mathcal{C}} \frac{(\boldsymbol{s} - \boldsymbol{x} + \boldsymbol{n}L) \times \mathrm{d}\boldsymbol{s}}{|\boldsymbol{s} - \boldsymbol{x} + \boldsymbol{n}L|^3}. \tag{2.2}$$

Note that periodicity requires the total vorticity to be zero within a periodic cell, or otherwise the curl operator cannot be inverted (the velocity diverges). This condition corresponds to $\oint_{\mathcal{C}} \mathrm{d}\boldsymbol{s} = \boldsymbol{0}$. This is trivially satisfied for closed filaments, while it requires special care when dealing with infinite unclosed filaments.

2.3. Desingularisation of the Biot–Savart integral. In the VFM, one is generally interested in the velocity induced by the vortex filaments *on the filaments themselves*. In other words, one wants to evaluate (2.2) on positions $\boldsymbol{x} = \boldsymbol{s}_0 \in \mathcal{C}$, where the Biot–Savart integral clearly diverges [21]. More precisely, performing a Taylor expansion of $\boldsymbol{s}(\xi)$, one can show that close to $\boldsymbol{s}_0 = \boldsymbol{s}(\xi_0)$ the integrand behaves as [3]

$$\frac{[\boldsymbol{s}(\xi) - \boldsymbol{s}_0] \times \boldsymbol{s}'(\xi)}{|\boldsymbol{s}(\xi) - \boldsymbol{s}_0|^3} \,\mathrm{d}\xi = \frac{\boldsymbol{s}_0' \times \boldsymbol{s}_0''}{2|\xi - \xi_0|} \,\mathrm{d}\xi + O(1), \tag{2.3}$$

where primes denote derivatives with respect to the arc length ξ. In particular, $\boldsymbol{s}_0' \equiv \boldsymbol{s}'(\xi_0)$ and $\boldsymbol{s}_0'' \equiv \boldsymbol{s}''(\xi_0)$ are respectively the unit tangent and the curvature vectors at $\boldsymbol{s}_0$.

The divergence of the Biot–Savart integral is unphysical since (2.1) does not account for the finite (but very small) radius a of the vortex core. This issue can be circumvented in a physically consistent manner by introducing a small cut-off coefficient ϵ_v such that locations $|\boldsymbol{s}(\xi) - \boldsymbol{s}_0| < \epsilon_v a$ close to the singularity at $\boldsymbol{s}_0$ are omitted from the integral [66, Ch. 11]. It is convenient to write $\epsilon_v = e^{\Delta}/2$, where Δ is a constant which depends on the actual vorticity profile within the vortex core. This constant can be analytically derived for commonly used vorticity profiles, under the assumption that the local curvature radius $1/|\boldsymbol{s}''|$ stays much larger than the core size a. Some usual values for axisymmetric circular cores are [66, Ch. 10]:

- $\Delta = 1/2$ for a hollow vortex core, $\omega(r) = \kappa\,\delta(r-a)/(2\pi a)$;
- $\Delta = 1/4$ for a uniform vortex core, $\omega(r) = \kappa/(\pi a^2)$ for $r \le a$;
- $\Delta \approx 0.558 + \frac{1}{2}\ln(2)$ for a Gaussian core, $\omega(r) = \kappa/(2\pi a^2)\, e^{-r^2/2a^2}$.

In practice, the VFM is generally used to describe length scales which are several orders of magnitude larger than a. In numerical simulations, vortex filaments are discretised with a typical line resolution $\ell \gg a$. Therefore, directly computing the Biot–Savart integral with a cut-off of order a as described above can lead to large numerical error, especially when quadrature rules are used to estimate the integrals. For this reason, the usual approach is to split the evaluation of the Biot–Savart integral on a filament location $\boldsymbol{s}_0$ onto *local* and *non-local* contributions [70], as represented in Figure 1 for a single filament. The local contribution is estimated by analytically integrating the leading-order term of the Taylor expansion (2.3) on the local portion $\mathcal{C}_0$ of the curve and excluding the cut-off region of length $2\epsilon_v a$. This leads to

$$\boldsymbol{v}_{\mathrm{local}}(\boldsymbol{s}_0) = \frac{\kappa}{4\pi} \boldsymbol{s}_0' \times \boldsymbol{s}_0'' \left[\ln\left(\frac{2\sqrt{\ell^- \ell^+}}{a} \right) - \Delta \right], \tag{2.4}$$

where ℓ^- and ℓ^+ are the lengths of the adjacent segments composing $\mathcal{C}_0$. The local velocity is thus oriented along the local binormal direction and is proportional to the local curvature $|\boldsymbol{s}_0''|$. As illustrated in Figure 1, when filaments are discretised by connecting vortex points and $\boldsymbol{s}_0$ is one of these points, the adjacent segments are usually defined as those linking $\boldsymbol{s}_0$ to its neighbours.

Accounting for the local correction, the total induced velocity of a point $\boldsymbol{s}_0$ on a vortex can be written as

$$\boldsymbol{v}(\boldsymbol{s}_0) = \boldsymbol{v}_{\mathrm{local}}(\boldsymbol{s}_0) + \frac{\kappa}{4\pi} \sum_{\boldsymbol{n} \in \mathbb{Z}^3} \int_{\mathcal{C}}' \frac{(\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L) \times \mathrm{d}\boldsymbol{s}}{|\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L|^3}, \tag{2.5}$$

where the prime over the integral indicates that integration is omitted on the local portion $\mathcal{C}_0$ of the filament containing $\boldsymbol{s}_0$.

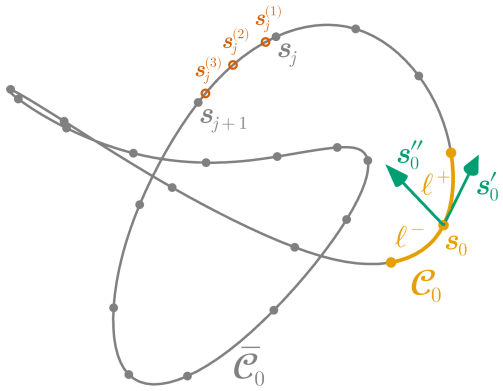

FIG. 1. *Discretisation of a closed vortex filament, and distinction between local and non-local contributions to the Biot–Savart law. Filled grey circles represent vortex discretisation points. The velocity induced on the vortex point* $\boldsymbol{s}_0$ *includes the* local *contribution* (2.4) *of its adjacent segments* ($\mathcal{C}_0$, *in orange), and the* non-local *contribution of the rest* ($\overline{\mathcal{C}}_0$, *in grey). Periodic images and other filaments (not depicted) also contribute to the non-local velocity. Near the evaluation point* $\boldsymbol{s}_0$, *the lengths* ℓ^- *and* ℓ^+ *of the adjacent segments are represented, as well as the local unit tangent* $\boldsymbol{s}_0'$ *and curvature vector* $\boldsymbol{s}_0''$. *Numerically, non-local contributions are computed by approximating line integrals by quadrature sums over each non-local segment. As an example, the open red circles between points* $\boldsymbol{s}_j$ *and* $\boldsymbol{s}_{j+1}$ *are the quadrature points associated to that segment (here, using* $Q = 3$ *point Gauss–Legendre quadratures on each segment). These are the points where the integrand is actually evaluated when computing non-local contributions.*

Note that neglecting the non-local contribution in (2.5) – thus keeping only the local part (2.4) – and replacing the segment lengths $\ell^\pm$ with a constant length scale leads to the *local induction approximation* (LIA), which is a widely used model of vortex motion [25, 3, 66]. This equation conserves vortex length as it disallows line stretching. Interestingly, the LIA can be mapped onto a non-linear Schrödinger equation [44], a connection which is still today the subject of active research [11]. Numerically, the LIA is very convenient since it only requires local information. Filament velocities can be efficiently computed in $O(N)$ time where N is the number of vortex discretisation points. However, since it neglects all interactions between vortices, the LIA is simply not suitable for the simulation of complex configurations and in particular turbulent vortex flows [1, 42], where non-local effects and collective behaviour are crucial ingredients.

2.4. Energy estimation in periodic systems. The VFM described above includes no dissipation mechanisms, and kinetic energy is expected to be a conserved quantity. It is crucial to be able to accurately estimate the kinetic energy of the system, as this may allow (i) to verify the accuracy of energy conservation at the numerical level or (ii) to estimate energy decay rates when the model is extended with dissipative mechanisms.

In a periodic domain, the kinetic energy per unit mass is defined as

$$E = \frac{1}{2V} \int_\Omega |\boldsymbol{v}(\boldsymbol{x})|^2 \, \mathrm{d}^3\boldsymbol{x} \tag{2.6}$$

where $\Omega = [0, L]^3$ is the main periodic cell and $V = L^3$ is its volume. In principle, this expression requires knowing the velocity field $\boldsymbol{v}(\boldsymbol{x})$ at every point $\boldsymbol{x}$ in space. In the VFM this is not only impractical – as it would require many evaluations of the Biot–Savart law (2.2) – but is also delicate since the velocity field presents strong gradients near vortex filaments. For these reasons, in non-periodic VFM simulations [68, 6, 41]

the above expression is commonly replaced by[1]

$$E_{\text{non-periodic}} = \kappa \oint \boldsymbol{v}(\boldsymbol{s}) \cdot (\boldsymbol{s} \times \mathrm{d}\boldsymbol{s}), \tag{2.7}$$

which only requires knowing the vortex geometry and the velocity of the vortex filaments. However, this expression assumes that the velocity and vorticity fields decay to zero at infinity [66, Ch. 3], which is not the case in periodic domains.

An alternative expression for the kinetic energy per unit mass can be obtained from (2.6) using integration by parts [66],

$$E = \frac{1}{2V} \int_\Omega \boldsymbol{\psi}(\boldsymbol{x}) \cdot \boldsymbol{\omega}(\boldsymbol{x}) \, \mathrm{d}^3\boldsymbol{x} = \frac{\kappa}{2V} \oint_{\mathcal{C}} \boldsymbol{\psi}(\boldsymbol{s}) \cdot \mathrm{d}\boldsymbol{s}, \tag{2.8}$$

where the last equality is obtained using (2.1). Note that the boundary terms of the integration by parts vanish in periodic domains. In (2.8), we have introduced the streamfunction vector $\boldsymbol{\psi}(\boldsymbol{x})$ (or *vector potential* [33, 66]), which is related to the velocity by $\boldsymbol{\nabla} \times \boldsymbol{\psi} = \boldsymbol{v}$, and thus to the vorticity by Poisson's equation $-\nabla^2\boldsymbol{\psi} = \boldsymbol{\omega}$. In other words, the kinetic energy can be obtained from a line integral requiring knowledge of streamfunction values on vortex filaments. Interestingly, (2.8) also allows to interpret the tangential streamfunction, $\psi_t = \boldsymbol{\psi} \cdot \boldsymbol{s}'$, as the linear energy density of a vortex point (up to a multiplicative constant $\kappa/2V$), and thus as the contribution of a vortex element to the total kinetic energy. Finally, note that (2.8) is also valid in non-periodic unbounded domains under the same assumptions leading to (2.7). To our knowledge, this expression has never been used before in the context of the VFM. Perhaps one of the reasons is that $\boldsymbol{\psi}(\boldsymbol{s})$ is a priori not available in VFM simulations, and computing it comes at an additional cost.

2.5. Obtaining the streamfunction vector. As mentioned above, the streamfunction is the solution of the Poisson equation $-\nabla^2\boldsymbol{\psi} = \boldsymbol{\omega}$. In three dimensions, the solution can be explicitly written as a convolution of $\boldsymbol{\omega}$ with the Green's function $G(\boldsymbol{x}, \boldsymbol{y}) = G(\boldsymbol{x} - \boldsymbol{y}) = 1/(4\pi|\boldsymbol{x} - \boldsymbol{y}|)$,

$$\boldsymbol{\psi}(\boldsymbol{x}) = (G * \boldsymbol{\omega})(\boldsymbol{x}) = \frac{\kappa}{4\pi} \sum_{\boldsymbol{n} \in \mathbb{Z}^3} \oint_{\mathcal{C}} \frac{\mathrm{d}\boldsymbol{s}}{|\boldsymbol{x} - \boldsymbol{s} + \boldsymbol{n}L|}, \tag{2.9}$$

where once again we have used (2.1) to express the vorticity field. Noting that $\boldsymbol{\nabla} G(\boldsymbol{r}) = -\boldsymbol{r}/(4\pi r^3)$, where $r = |\boldsymbol{r}|$, one can readily show that taking the curl of (2.9) leads to the Biot–Savart law (2.2).

Similarly to the velocity (subsection 2.3), the streamfunction obtained via (2.9) diverges when evaluated on a filament location $\boldsymbol{s}_0 \in \mathcal{C}$. Thus, the streamfunction integral must be desingularised by accounting for the finite radius a of the vortex core. As for the velocity, we introduce a cut-off coefficient ϵ_ψ so that locations $|\boldsymbol{s} - \boldsymbol{s}_0| < \epsilon_\psi a$ are omitted from the line integral in (2.9). Crucially, the cut-off coefficient ϵ_ψ need not be equal to the coefficient $\epsilon_v = e^\Delta/2$ used for the velocity. In fact, in order for kinetic energy to be properly conserved, one must rather set $\epsilon_\psi = e^{\Delta-1}/2$, as argued analytically in subsection 2.6 and verified numerically in subsection 5.2.

Using the same notation as in (2.3), the integrand of (2.9) behaves close to $\boldsymbol{s}_0 = \boldsymbol{s}(\xi_0)$ as

$$\frac{\boldsymbol{s}'(\xi)}{|\boldsymbol{s}(\xi) - \boldsymbol{s}_0|} \, \mathrm{d}\xi = \frac{\boldsymbol{s}_0'}{|\xi - \xi_0|} \, \mathrm{d}\xi + O(1). \tag{2.10}$$

[1] Here we express the energy is per unit *density* and not *mass*, since the domain volume V (and thus the fluid mass) is infinite.

Then, in analogy with the local velocity (2.4), integrating over $\xi \in [\xi_0 - \ell^-, \xi_0 - \epsilon_\psi a] \cup [\xi_0 + \epsilon_\psi a, \xi_0 + \ell^+]$ ($\mathcal{C}_0$ segments in Figure 1) leads to the local contribution

$$\boldsymbol{\psi}_{\text{local}}(\boldsymbol{s}_0) = \frac{\kappa \boldsymbol{s}_0'}{2\pi}\left[\ln\left(\frac{2\sqrt{\ell^-\ell^+}}{a}\right) - (\Delta - 1)\right]. \tag{2.11}$$

Interestingly, this term is tangent to the filament, and therefore it fully contributes to the kinetic energy (2.8). In the context of vortex dynamics, this local contribution is commonly referred to as *vortex tension* [59, 14].

Finally, in analogy with (2.5), the streamfunction vector on a filament location $\boldsymbol{s}_0$ can be written as

$$\boldsymbol{\psi}(\boldsymbol{s}_0) = \boldsymbol{\psi}_{\text{local}}(\boldsymbol{s}_0) + \frac{\kappa}{4\pi}\sum_{\boldsymbol{n}\in\mathbb{Z}^3}\int_{\mathcal{C}}' \frac{\mathrm{d}\boldsymbol{s}}{|\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L|}, \tag{2.12}$$

where once again the prime over the integral denotes the omission of local vortex elements around $\boldsymbol{s}_0$ ($\mathcal{C}_0$ segments in Figure 1).

2.6. An application: vortex ring dynamics. As an illustration of the VFM and in order to justify the chosen streamfunction integration cut-off ϵ_ψ, we apply the above definitions to the classical example of an isolated circular vortex ring. Vortex rings are ubiquitous in classical and quantum fluid dynamics [72]. In the absence of viscosity or the influence of other vortices or obstacles, they are characterised by propagating with a constant self-induced speed U without changing their shape.

We consider a circular vortex ring of radius R and circulation κ, which can be parametrised by its arc length ξ as $\boldsymbol{s}(\xi) = [R\cos(\xi/R), R\sin(\xi/R), 0]$ for $\xi \in [0, 2\pi R]$. We assume periodicity effects to be negligible, i.e. $R \ll L$, so that the influence of periodic images can be discarded. Then, under the assumption that $\ell^\pm \ll R$, one can analytically compute the non-local integrals in (2.5) and (2.12) (see e.g. [70] for the case of the velocity). Including the local contributions (2.4) and (2.11), the total velocity and streamfunction on a location $\boldsymbol{s}(\xi)$ on the ring are

$$\boldsymbol{v}(\boldsymbol{s}) = \frac{\kappa \hat{\boldsymbol{z}}}{4\pi R}\left[\ln\left(\frac{8R}{a}\right) - \Delta\right] = U\hat{\boldsymbol{z}}, \tag{2.13}$$

$$\boldsymbol{\psi}(\boldsymbol{s}) = \frac{\kappa \boldsymbol{s}'}{2\pi}\left[\ln\left(\frac{8R}{a}\right) - (\Delta + 1)\right]. \tag{2.14}$$

We remark that the $\Delta + 1$ term in (2.14) directly comes from the choice of cut-off coefficient $\epsilon_\psi = e^{\Delta-1}/2$ from the previous section, which by (2.8) also influences the estimated kinetic energy of the velocity field induced by the ring. The latter is explicitly given by

$$E = \frac{\kappa^2 R}{2V}\left[\ln\left(\frac{8R}{a}\right) - (\Delta + 1)\right]. \tag{2.15}$$

One can show that the above choice of ϵ_ψ is the only possible choice guaranteeing Hamilton's equation $U = \partial E/\partial p$ [65] to be satisfied, under the hypothesis that the vortex core size a stays constant (which is imposed in the VFM). Here, $p = |\boldsymbol{p}|$ is the magnitude of the hydrodynamic impulse (per unit mass) defined as

$$\boldsymbol{p} = \frac{\kappa}{2V}\oint_{\mathcal{C}} \boldsymbol{s} \times \mathrm{d}\boldsymbol{s}. \tag{2.16}$$

In the absence of non-conservative external body forces, the total impulse of a set of vortex filaments is also a conserved quantity [66, Ch. 3]. For a circular vortex ring, the impulse is aligned with the ring velocity (2.13), and its magnitude $p = (\kappa/V)\pi R^2$ is proportional to the area πR^2 enclosed by the ring. One can readily differentiate both conserved quantities with respect to the ring radius R,

$$\frac{\partial E}{\partial R} = \frac{\kappa^2}{2V}\left[\ln\left(\frac{8R}{a}\right) - \Delta\right] \quad \text{and} \quad \frac{\partial p}{\partial R} = \frac{2\kappa\pi R}{V}, \tag{2.17}$$

in order to verify that $\partial E/\partial p = U$ is indeed satisfied, using the analytical vortex ring speed U in (2.13).

Ensuring that the adopted energy definition obeys Hamilton's equation is critical for accurate diagnostics of energy conservation in the absence of dissipative mechanisms in the underlying model. Later in subsection 5.2, we verify that using definition (2.11) for the local contribution to the streamfunction leads to proper energy conservation up to numerical accuracy. This appears to be the case not only for circular vortex rings, but also for complex configurations containing multiple closed vortices of arbitrary shape, and even unclosed vortices extending to infinity (as in Figure 2, centre), in which case energy computations have been deemed impractical in the past [6]. Finally, note that the alternative energy expression (2.7) commonly used in open non-periodic systems is not consistent with Hamilton's equation for a circular ring, as it corresponds to replacing the $\Delta + 1$ term in the ring energy (2.15) with Δ. This suggests that (2.7) does not properly account for the structure of the vortex core.

2.7. Discretisation of spatial curves and line integrals. We finish this section by discussing the approach we adopt to represent vortex filaments and estimate line integrals in numerical simulations. In the context of vortex filament simulations for describing quantum fluids, filaments have been almost invariably described by a set of points in space connected by straight segments [70, 67, 7, 41, 82]. Integration of the Biot–Savart law over lines is then straightforward and can even be done analytically over each straight segment. Curve derivatives (local tangents, curvatures) are estimated on the discrete points using e.g. finite difference approximations [7]. While this approach can be implemented with relative ease, it is a low-order discretisation, often requiring very small distances between discretisation points to achieve accurate results. It also presents an inconsistency between the C^0 continuity of the lines and the C^2 continuity required to estimate curvatures.

Here we take a different route and consider the filaments as smooth curves passing through a set of points, as illustrated in Figure 1. As introduced in subsection 2.1, each curve is parametrised as $\boldsymbol{s}(\tau)$ for some scalar parameter $\tau \in [0, \mathcal{T}]$. The numerical degrees of freedom are the positions of discrete vortex points or *nodes*. Some interpolation method is then applied to evaluate vortex positions in-between nodes. The interpolation also gives direct access to curve derivatives along vortex lines, allowing to estimate local tangent and curvature vectors required for Biot–Savart computations. In the numerical experiments of section 5 we use quintic spline interpolations which have global continuity C^4.

In the context of the Biot–Savart problem, a second difficulty is the need to estimate integrals on vortex lines. Here we estimate line integrals using Gauss–Legendre quadratures on each segment connecting two neighbouring nodes. Concretely, if one considers a closed curve $\mathcal{C}$ parametrised as $\boldsymbol{s}(\tau)$ with nodes $\{\boldsymbol{s}_j = \boldsymbol{s}(\tau_j); j = 1, \dots, N\}$,

then an integral over such curve is approximated as

$$\int_0^{\mathcal{T}} F[\boldsymbol{s}(\tau)]\,\mathrm{d}\tau \approx \sum_{j=1}^{N}\sum_{i=1}^{Q} w_i F\big[\boldsymbol{s}_j^{(i)}\big]\Delta\tau_j, \quad \boldsymbol{s}_j^{(i)} \equiv \boldsymbol{s}(\tau_j + h_i\Delta\tau_j), \tag{2.18}$$

where Q is the number of quadrature points per segment (typically one chooses $Q \leq 4$ in simulations). Here, $F[\boldsymbol{s}]$ is a functional which may depend on curve locations $\boldsymbol{s}(\tau)$ as well as its derivatives. For convenience, the quadrature rule above is defined in the domain $[0,1]$, such that $h_i \in [0,1]$ are the quadrature locations, and w_i the associated weights satisfying $\sum_{i=1}^{Q} w_i = 1$. Moreover, $\Delta\tau_j = \tau_{j+1} - \tau_j$ is the parameter increment associated to a single segment, with the convention that $\tau_{N+1} = \tau_1 + \mathcal{T}$ for a closed or an infinite periodic curve. As an example, the open red circles in Figure 1 are the quadrature points associated to a single filament segment (using $Q = 3$).

Remark 2.1. The use of the truncated Taylor expansions (2.3) and (2.10) can in principle induce an error on the computation of the local velocity (2.4) and streamfunction (2.11). The error can be expected to be negligible when the vortex line discretisation is sufficiently fine. However, for computational efficiency, it may be desirable to use a coarser discretisation, in which case this source of error can become dominant. To remediate this issue, one possibility is to subdivide the local segments $\mathcal{C}_0$ (Figure 1) onto (i) a smaller central segment containing $\boldsymbol{s}_0$, and (ii) two *outer* segments respectively ending at the two nodes adjacent to $\boldsymbol{s}_0$. The Taylor expansions are then applied in the central segment only. In the outer segments, some quadrature rule can be used to directly evaluate the Biot–Savart and streamfunction integrals. Since these integrations are performed near the singularity at $\boldsymbol{s}_0$, using Gauss–Legendre quadratures with a small number of points Q as in (2.18) can result in accuracy loss, defeating the purpose of subdividing the local segments. We have found that using adaptive tanh-sinh (double exponential) quadratures [74, 76] instead can provide an important gain of precision at a relatively small additional cost.

3. Fast Ewald summation for the Biot–Savart problem. In this section we adapt fast Ewald summation methods to the evaluation of velocity and streamfunction fields induced by a set of vortex filaments. This requires some adjustments since these methods generally deal with a scalar-valued source term (e.g. electrostatic charge) supported on discrete points, while here the source term (vorticity) is vector-valued and defined on spatial curves. In the context of the VFM, the basic idea of these methods is to split the line integrals in (2.5) and (2.12) onto short- and long-range components. Concretely, including the local terms appearing in those expressions, the streamfunction and the velocity at a vortex location $\boldsymbol{s}_0$ are decomposed as

$$\boldsymbol{\psi}(\boldsymbol{s}_0) = \boldsymbol{\psi}_{\mathrm{local}}(\boldsymbol{s}_0) + \boldsymbol{\psi}^{<}(\boldsymbol{s}_0) + \boldsymbol{\psi}^{>}(\boldsymbol{s}_0) - \boldsymbol{\psi}^{>}_{\mathrm{self}}(\boldsymbol{s}_0), \tag{3.1}$$

$$\boldsymbol{v}(\boldsymbol{s}_0) = \boldsymbol{v}_{\mathrm{local}}(\boldsymbol{s}_0) + \boldsymbol{v}^{<}(\boldsymbol{s}_0) + \boldsymbol{v}^{>}(\boldsymbol{s}_0) - \boldsymbol{v}^{>}_{\mathrm{self}}(\boldsymbol{s}_0), \tag{3.2}$$

where the $<$ and $>$ superscripts denote short- and long-range components, and $\boldsymbol{\psi}^{>}_{\mathrm{self}}$ and $\boldsymbol{v}^{>}_{\mathrm{self}}$ are corrections to the latter which are discussed further below.

3.1. Ewald splitting. Ewald summation methods split the singular Green's function $G(\boldsymbol{r})$ associated to Poisson's equation into (i) a fast decaying part $G^{<}(\boldsymbol{r})$ accounting for short-range interactions and (ii) a slowly decaying part $G^{>}(\boldsymbol{r})$, which is well-behaved at $\boldsymbol{r} = \boldsymbol{0}$ and describes long-range interactions. In three dimensions, this is usually achieved via the identity $\operatorname{erf}(x) + \operatorname{erfc}(x) = 1$,

$$G(\boldsymbol{r}) = \frac{1}{4\pi r} = \frac{\operatorname{erfc}(\alpha r)}{4\pi r} + \frac{\operatorname{erf}(\alpha r)}{4\pi r} = G^{<}(\boldsymbol{r}) + G^{>}(\boldsymbol{r}), \tag{3.3}$$

where $\operatorname{erf}(x) = (2/\sqrt{\pi}) \int_0^x e^{-u^2}\,\mathrm{d}u$ and $\operatorname{erfc}(x) = (2/\sqrt{\pi}) \int_x^\infty e^{-u^2}\,\mathrm{d}u$ are respectively the error function and the complementary error function. Due to linearity, the solution $\boldsymbol{\psi} = G * \boldsymbol{\omega}$ to the Poisson equation $-\nabla^2 \boldsymbol{\psi} = \boldsymbol{\omega}$ is accordingly split into $\boldsymbol{\psi} = \boldsymbol{\psi}^< + \boldsymbol{\psi}^>$. Similarly, the velocity field can be decomposed as $\boldsymbol{v} = \boldsymbol{\nabla} \times \boldsymbol{\psi} = \boldsymbol{v}^< + \boldsymbol{v}^>$. Importantly, α in (3.3) is the Ewald splitting parameter, which defines an inverse length scale setting the transition between short- and long-range interactions. This is a purely numerical parameter, as the physical fields $\boldsymbol{\psi}$ and $\boldsymbol{v}$ obtained by adding both contributions are in theory independent of α. To facilitate the interpretation of the splitting (3.3), it is helpful to consider the "long-range" vorticity field defined by $\boldsymbol{\omega}^> = -\nabla^2 \boldsymbol{\psi}^>$, which can be written as the convolution of $\boldsymbol{\omega}$ with $-\nabla^2 G^>$. In fact this is a Gaussian-filtering operation, since $-\nabla^2 G^>(\boldsymbol{r}) = (\alpha/\sqrt{\pi})^3 e^{-(\alpha r)^2}$. In other words, $\boldsymbol{\psi}^>$ is the streamfunction associated to a coarse-grained (smoothed) version of the original (singular) vorticity $\boldsymbol{\omega}$. Similarly, $\boldsymbol{v}^> = \boldsymbol{\nabla} \times \boldsymbol{\psi}^>$ is the corresponding coarse-grained velocity field.

Taking the gradient of (3.3), one arrives to the corresponding splitting for the Biot–Savart kernel $\boldsymbol{\nabla} G$,

$$\boldsymbol{\nabla} G(\boldsymbol{r}) = -\frac{\boldsymbol{r}}{4\pi r^3} = -\left[g^<(\boldsymbol{r}) + g^>(\boldsymbol{r})\right] \frac{\boldsymbol{r}}{4\pi r^3} = \boldsymbol{\nabla} G^<(\boldsymbol{r}) + \boldsymbol{\nabla} G^>(\boldsymbol{r}) \tag{3.4}$$

with the weight functions

$$g^<(\boldsymbol{r}) = \operatorname{erfc}(\alpha r) + \frac{2\alpha r}{\sqrt{\pi}}\, e^{-(\alpha r)^2}, \quad g^>(\boldsymbol{r}) = \operatorname{erf}(\alpha r) - \frac{2\alpha r}{\sqrt{\pi}}\, e^{-(\alpha r)^2}. \tag{3.5}$$

It can be verified that the long-range kernels are non-singular, behaving as $G^>(\boldsymbol{r}) = \alpha/(2\pi^{3/2}) + O(r^2)$ and $\boldsymbol{\nabla} G^>(\boldsymbol{r}) = -\alpha^3 \boldsymbol{r}/(3\pi^{3/2}) + O(r^3)$ near $r = 0$. On the other hand, the short-range kernels asymptotically decay as $G^<(\boldsymbol{r}) \sim e^{-(\alpha r)^2}/\left[4\pi r(\alpha r)^{3/2}\right]$ and $\boldsymbol{\nabla} G^<(\boldsymbol{r}) \sim \alpha \boldsymbol{r}\, e^{-(\alpha r)^2}/\left(2\pi^{3/2} r^2\right)$ for large r.

3.2. Estimation of short-range interactions. The short-range streamfunction $\boldsymbol{\psi}^<$ is obtained by replacing $G(\boldsymbol{r})$ in (2.9) with $G^<(\boldsymbol{r})$ defined in (3.3), resulting in

$$\boldsymbol{\psi}^<(\boldsymbol{s}_0) = \frac{\kappa}{4\pi} \sum_{\boldsymbol{n} \in \mathbb{Z}^3} \int_{\mathcal{C}}' f^<(\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L)\, \frac{\mathrm{d}\boldsymbol{s}}{|\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L|}, \tag{3.6}$$

where $f^<(\boldsymbol{r}) = \operatorname{erfc}(\alpha r)$. We recall that the prime over the integral symbol denotes the omission of local vortex elements around $\boldsymbol{s}_0$ ($\mathcal{C}_0$ segments in Figure 1). Similarly, following the decomposition (3.4), the short-range velocity field is explicitly given by the modified Biot–Savart integral

$$\boldsymbol{v}^<(\boldsymbol{s}_0) = \frac{\kappa}{4\pi} \sum_{\boldsymbol{n} \in \mathbb{Z}^3} \int_{\mathcal{C}}' g^<(\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L)\, \frac{(\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L) \times \mathrm{d}\boldsymbol{s}}{|\boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L|^3}. \tag{3.7}$$

The original integrals in (2.5) and (2.12) are recovered by setting $\alpha = 0$ in the above expressions. In practice, we approximate the above line integrals using quadrature sums over discrete vortex line locations $\boldsymbol{s}_j^{(i)}$ according to (2.18). Moreover, since the short-range kernels $G^<$ and $\boldsymbol{\nabla} G^<$ decay exponentially with $(\alpha r)^2$, one can safely define a cut-off distance r_c beyond which short-range interactions can be neglected. The truncation errors associated to the choice of r_c are discussed in section 4.

3.3. Estimation of long-range interactions. Since the long-range Green's function $G^>$ is non-singular, one can use a truncated Fourier series representation to indirectly solve $-\nabla^2 \boldsymbol{\psi}^> = \boldsymbol{\omega}^>$. We start by writing the periodic vorticity field as

$$\boldsymbol{\omega}(\boldsymbol{x}) \approx \sum_{\boldsymbol{k}\in\mathbb{K}^3} \hat{\boldsymbol{\omega}}_{\boldsymbol{k}}\, e^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}} \quad \text{where} \quad \hat{\boldsymbol{\omega}}_{\boldsymbol{k}} = \frac{1}{V}\int_\Omega \boldsymbol{\omega}(\boldsymbol{x})\, e^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}}\, \mathrm{d}^3\boldsymbol{x}. \tag{3.8}$$

Here $\mathbb{K} = \left\{ \frac{2\pi m}{L}; m = -\left\lfloor \frac{M}{2} \right\rfloor, \dots, \left\lfloor \frac{M-1}{2} \right\rfloor \right\}$ is the set of M resolved Fourier wavenumbers in each Cartesian direction and $\lfloor\cdot\rfloor$ denotes the floor operation (for simplicity, we take M to be the same in all directions). We recall that $\Omega = [0, L]^3$ represents the main periodic cell and $V = L^3$ is its volume. Since the vorticity (2.1) is singular and supported on spatial curves, its Fourier coefficients are then

$$\hat{\boldsymbol{\omega}}_{\boldsymbol{k}} = \frac{\kappa}{V}\oint_{\mathcal{C}} e^{-i\boldsymbol{k}\cdot\boldsymbol{s}}\, \mathrm{d}\boldsymbol{s} \approx \frac{\kappa}{V}\sum_{l=1}^{N_\mathrm{f}}\sum_{j=1}^{N_l}\sum_{i=1}^{Q} w_i\, \Delta\tau_{jl}\, \boldsymbol{s}'_{ijl}\, e^{-\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{s}_{ijl}}, \tag{3.9}$$

where the last expression is obtained from the quadrature approximation (2.18). The outermost sum is over the N_f vortex filaments of the system, each being discretised by a possibly different number of nodes N_l. Moreover, $\boldsymbol{s}_{ijl}$ is a shorthand for $\boldsymbol{s}_l(\tau_j + h_i\Delta\tau_j)$, where $\boldsymbol{s}_l(\tau)$ is the parametrisation of the l-th filament, while $\boldsymbol{s}'_{ijl}$ is the derivative with respect to τ at that location (aligned with the local tangent vector). The triple sum in (3.9) can be interpreted as a sum of vector charges $\boldsymbol{q}_{ijl} = w_i\,\Delta\tau_{jl}\,\boldsymbol{s}'_{ijl}$ on locations $\boldsymbol{s}_{ijl}$. Note that using the same quadrature nodes for short- and long-range computations means that interpolated values of $\boldsymbol{s}$ and $\boldsymbol{s}'$ can be shared among both components, reducing the computational cost associated to interpolations.

Now, since $\boldsymbol{\omega}^>$ is the convolution of $\boldsymbol{\omega}$ with a Gaussian kernel, its Fourier coefficients are $\hat{\boldsymbol{\omega}}^>_{\boldsymbol{k}} = \hat{\boldsymbol{\omega}}_{\boldsymbol{k}}\, e^{-k^2/4\alpha^2}$ where $k = |\boldsymbol{k}|$. Unlike $|\hat{\boldsymbol{\omega}}_{\boldsymbol{k}}|$, the amplitudes $|\hat{\boldsymbol{\omega}}^>_{\boldsymbol{k}}|$ can be expected to decay quickly for $k \gg \alpha$, which justifies the truncation of the Fourier series. Moreover, solving Poisson's equation amounts to division by k^2 in Fourier space. This ultimately allows us to express the long-range velocity in physical space as

$$\boldsymbol{v}^>(\boldsymbol{x}) = \sum_{\substack{\boldsymbol{k}\in\mathbb{K}^3 \\ |\boldsymbol{k}|\neq 0}} \hat{\boldsymbol{v}}^>_{\boldsymbol{k}}\, e^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}} \quad \text{with} \quad \hat{\boldsymbol{v}}^>_{\boldsymbol{k}} = \mathrm{i}\boldsymbol{k}\times\hat{\boldsymbol{\psi}}^>_{\boldsymbol{k}} = \mathrm{i}\boldsymbol{k}\times\frac{\hat{\boldsymbol{\omega}}^>_{\boldsymbol{k}}}{k^2}. \tag{3.10}$$

Note that this requires $\hat{\boldsymbol{\omega}}_{\boldsymbol{0}} = \boldsymbol{0}$, i.e. the mean vorticity within a periodic cell must be zero. The long-range streamfunction field $\boldsymbol{\psi}^>(\boldsymbol{x})$ can be written similarly to (3.10) from its Fourier coefficients $\hat{\boldsymbol{\psi}}^>_{\boldsymbol{k}} = \hat{\boldsymbol{\omega}}^>_{\boldsymbol{k}}/k^2$.

When evaluated on a vortex location $\boldsymbol{s}_0$, the obtained long-range velocity and streamfunction vectors include contributions of the local segments adjacent to $\boldsymbol{s}_0$. However, integration over these segments must be excluded according to (2.5) and (2.12). This issue is analogous to the spurious self-interaction electrostatic potential appearing in standard Ewald methods [4]. In the present case, the spurious local integrals are explicitly given by

$$\boldsymbol{\psi}^>_{\mathrm{self}}(\boldsymbol{s}_0) = \frac{\kappa}{4\pi}\int_{\mathcal{C}_0} f^>(\boldsymbol{s}-\boldsymbol{s}_0)\,\frac{\mathrm{d}\boldsymbol{s}}{|\boldsymbol{s}-\boldsymbol{s}_0|}, \tag{3.11}$$

$$\boldsymbol{v}^>_{\mathrm{self}}(\boldsymbol{s}_0) = \frac{\kappa}{4\pi}\int_{\mathcal{C}_0} g^>(\boldsymbol{s}-\boldsymbol{s}_0)\frac{(\boldsymbol{s}-\boldsymbol{s}_0)\times\mathrm{d}\boldsymbol{s}}{|\boldsymbol{s}-\boldsymbol{s}_0|^3}, \tag{3.12}$$

where $f^{>}(\boldsymbol{r}) = \operatorname{erf}(\alpha r)$, $g^{>}(\boldsymbol{r})$ is defined in (3.5), and $\mathcal{C}_0$ consists of the two segments adjacent to $\boldsymbol{s}_0$ (see Figure 1). They can be estimated using a variant of the quadrature procedure (2.18), thus requiring $2Q$ evaluations of the integrand for each vortex location $\boldsymbol{s}_0$. As indicated by (3.1) and (3.2), these two "self-interaction" terms must be subtracted from the long-range estimations to avoid an unphysical dependence of the results on the Ewald parameter α.

In the present context, the main difficulty of fast Ewald summation methods is the estimation of the Fourier coefficients of the vorticity field (3.9) and the evaluation of the resulting coarse-grained fields on arbitrary physical locations (3.10). While these two operations can in principle be directly evaluated according to their definitions, this is extremely inefficient, and fast approximations are required in practical applications. In fact, there are many variants of fast Ewald summation (so-called $\mathrm{P^3M}$, PME, SPME, $\mathrm{P^2NFFT}$, ...) which mainly differ on how these operations are performed [29, 4]. What they all have in common is that they introduce a regular 3D grid, which enables the acceleration of these operations using FFTs. First, point charges are spread onto the grid by some smoothing (spreading) operation, and (3.9) is then estimated using FFTs. Then, the Laplacian is generally inverted in Fourier space using some kind of modified Green's function [47, 29, 10] which accounts for errors introduced by the spreading and interpolation operations. Finally, (3.10) is estimated using backwards FFTs followed by an interpolation from the grid to the target locations. An important realisation [45, 4, 62] is that all of these operations can be conveniently and efficiently performed using the non-uniform fast Fourier transform (NUFFT) algorithm [31]. In fact, the sums in (3.9) and (3.10) respectively correspond to type-1 (non-uniform to uniform) and type-2 (uniform to non-uniform) NUFFTs [38]. From a practical standpoint, this eases the development of Ewald methods and allows one to rely on recent theoretical developments and on existent fast and accurate NUFFT implementations. When relying on the NUFFT, the accuracy of long-range computations is then controlled by the NUFFT accuracy and the truncation wavenumber $k_{\max} = (2\pi/L)\lfloor (M-1)/2 \rfloor$. These two sources of error are respectively discussed in Appendix B and section 4.

4. Truncation error estimates. We now provide estimates of the root-mean-square errors associated to the short- and long-range cut-offs r_{c} and $k_{\max}$. Similar estimates have already been provided for the electrostatic problem [51, 30], but these do not directly apply to the VFM since (i) the quantities of interest are not the same, and (ii) the singular sources in the present case are spatial curves and not points. We also show that the effect of the parameters α, r_{c} and $k_{\max}$ can be reduced to a unique non-dimensional coefficient β which can be tuned to control the method accuracy, thus greatly simplifying the parameter selection procedure.

The derivation of the error estimates is detailed in Appendix C. In summary, the estimates make the simplifying assumptions that (i) vortex filaments are homogeneously distributed in the spatial domain, (ii) the relative orientation of two vortex elements at a distance $r > r_{\mathrm{c}}$ is completely decorrelated, and (iii) $k_{\max}$ is much larger than the typical curvature $|\boldsymbol{s}''|$ of the vortices. Introducing the non-dimensional cut-off coefficients $\beta^{<} = \alpha r_{\mathrm{c}}$ and $\beta^{>} = k_{\max}/2\alpha$, we show that both short- and long-range errors decay exponentially with the square of these respective coefficients, which justifies reducing them into a unique coefficient β. This finally leads to the truncation

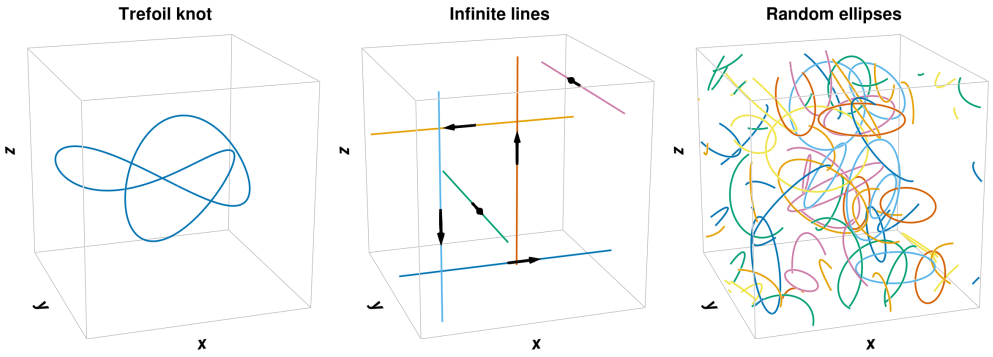

FIG. 2. *Visualisation of considered test cases within a periodic cell of size $L^3 = (2\pi)^3$. Left: trefoil knot with characteristic size $R = \pi/3$. Centre: infinitely extended straight lines. Arrows indicate the line orientation. Right: 40 ellipses with random locations, orientations, sizes and aspect ratios. For visualisation purposes, curves crossing the domain boundaries are broken onto multiple parts, which are periodically wrapped to fit within the periodic cell. Colours allow to distinguish between individual vortex lines.*

error estimates

$$\varepsilon_v \approx \kappa \left[\frac{\mathcal{L}}{\sqrt{\pi}\alpha V} + \frac{1}{\beta} \left(\frac{\mathcal{L}}{8\pi V} \right)^{1/2} \right] e^{-\beta^2}, \tag{4.1}$$

$$\varepsilon_\psi \approx \frac{\kappa}{2\alpha} \left[\frac{\mathcal{L}}{\alpha V} + \frac{1}{\beta^{1/2}} \left(\frac{\mathcal{L}}{8\pi V} \right)^{1/2} \right] \frac{e^{-\beta^2}}{\beta^{3/2}}, \tag{4.2}$$

where the short- and long-range contributions to the errors correspond to the first and second terms within each square bracket. The accuracy of the method is thus mainly driven by a unique non-dimensional parameter β. This sets both physical- and Fourier-space cut-offs $r_c = \beta/\alpha$ and $k_{\max} = 2\beta\alpha$ for a given value of the inverse splitting distance α, which is left as a free parameter that can be adjusted to optimise performance (subsection 5.3).

5. Numerical experiments. We now perform different numerical experiments to evaluate the accuracy, conservation properties and computational complexity of the proposed FFT-based method. All computations are performed in double precision arithmetic and the methods are implemented in the Julia programming language [17]. In all numerical tests below, the domain is periodic with period $L = 2\pi$ in all directions. The vortex core size and the vorticity profile parameter – both appearing in the locally induced velocity (2.4) – are respectively set to $a = 10^{-8}$ and $\Delta = 1/4$, and the vortex circulation is set to $\kappa = 1$. Vortex lines are represented as smooth curves using quintic splines (see Appendix A for details), which allows to estimate derivatives and evaluate quantities in-between discretisation points. Moreover, line integrals are estimated using $Q = 3$ quadrature points per vortex segment (see subsection 2.7).

5.1. Accuracy. To numerically assess the accuracy of the proposed method, we consider three different test cases of varying complexity. The test cases are detailed below and illustrated in Figure 2.

1. Trefoil knot. The trefoil knot consists in a single knotted curve (Figure 2, left), interacting with itself and its periodic images. The trefoil knot curve is parametrically

defined by

$$\boldsymbol{X}(\theta) = R \begin{bmatrix} \sin(\theta) + 2\sin(2\theta) \\ \cos(\theta) - 2\cos(2\theta) \\ -\sin(3\theta) \end{bmatrix} \quad \text{for } \theta \in [0, 2\pi], \tag{5.1}$$

where R determines the size of the trefoil knot. We take $R = L/6$ so that periodicity effects are non-negligible. The curve is discretised by evaluating (5.1) on $N = 512$ equally spaced values of θ.

2. Infinite lines. This test case consists of a set of straight infinitely extended lines inducing a three-dimensional velocity field. Each periodic cell is crossed by three pairs of lines aligned positively and negatively with each Cartesian direction (Figure 2, centre). The filament orientations are such that the velocity induced on each line is non-zero and positively aligned with the line ($\boldsymbol{v} \cdot \boldsymbol{s}' > 0$). Since the lines have zero curvature, the velocity induced by a line on itself is zero, and therefore the velocity of a vortex point is completely due to non-local interactions. Therefore, this test case emphasises the accuracy of long-range computations. We discretise each filament with $N_l = 128$ independent nodes (for $l = 1, \dots, 6$).

3. Random ellipses. The third test case is closer to a disordered configuration relevant to turbulent vortex flows. It consists of 40 ellipses which are randomly positioned and oriented within the spatial domain (Figure 2, right). The size and aspect ratio of each ellipse is also random. Concretely, the minor and major radii are random values uniformly and independently distributed in $[L/16, L/4]$. We discretise each ellipse with $N_l = 128$ nodes (for $l = 1, \dots, 40$).

Empirical relative errors. To numerically quantify the truncation errors in each test case, we start by computing a reference solution with very high accuracy. Concretely, we set the cut-off coefficient to $\beta = 8$ and compute NUFFTs with a nominal relative tolerance $\sim 10^{-14}$. Here, the reference solution corresponds to a set of velocity and streamfunction values, $\boldsymbol{v}_i^{\text{ref}}$ and $\boldsymbol{\psi}_i^{\text{ref}}$, on each discrete filament node $\boldsymbol{s}_i$. We then compare this reference solution with the results $\boldsymbol{v}_i$ and $\boldsymbol{\psi}_i$ obtained by varying β. The splitting parameter is set to $\alpha = 24/L$ in all numerical experiments. Accuracy estimates are obtained by evaluating the relative ℓ_2 errors

$$\varepsilon_v^{\text{rel}} = \frac{\|\boldsymbol{v} - \boldsymbol{v}^{\text{ref}}\|}{\|\boldsymbol{v}^{\text{ref}}\|} \quad \text{and} \quad \varepsilon_\psi^{\text{rel}} = \frac{\|\boldsymbol{\psi} - \boldsymbol{\psi}^{\text{ref}}\|}{\|\boldsymbol{\psi}^{\text{ref}}\|}, \tag{5.2}$$

where the ℓ_2 norm of a vector field $\boldsymbol{u}$ evaluated at N discrete locations is defined by $\|\boldsymbol{u}\|^2 = \frac{1}{N}\sum_{i=1}^{N} |\boldsymbol{u}_i|^2$. Here $|\cdot|$ denotes the vector magnitude.

The obtained relative errors are plotted in Figure 3 (lines with symbols). In almost all cases, the relative error decays exponentially with the squared cut-off parameter β^2, until roundoff errors become dominant. The decays are in remarkably good agreement with the rms error estimates (4.1) and (4.2) (dashed lines), which in almost all cases provide upper bounds to the actual errors. This confirms the relevance of the truncation estimates in a variety of configurations.

5.2. Timestepping and conservation properties. The VFM as described in section 2 is expected to conserve the total kinetic energy (2.6) and the hydrodynamic impulse (2.16). In this section, this is verified by evolving in time the positions of a set of vortex filaments according to the Biot–Savart velocity (2.5). The test case we consider here is inspired by the classical problem of two coaxial circular vortex rings which travel in the same direction. Due to their mutual interaction, the rings are

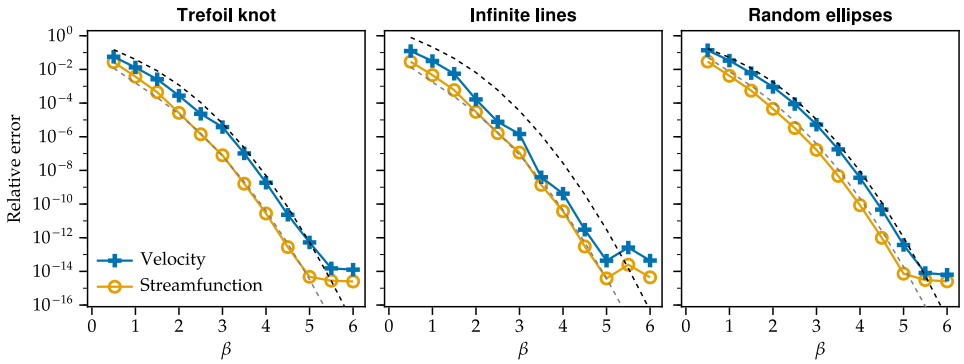

FIG. 3. *Accuracy of full Biot–Savart computations for three different test cases. Relative root-mean-square error of filament velocity* (+) *and streamfunction values* (×) *at varying non-dimensional cut-off parameter* β. *The reference solution is obtained using* $\beta = 8$. *The Ewald splitting parameter is kept constant,* $\alpha = 24/L$. *The black dashed lines represent the total velocity error estimate* (4.1). *The grey dashed lines represent the total streamfunction error estimate* (4.2).

expected to continuously change their sizes and therefore their translational velocities. As a result, they pass through one another in a cyclic fashion in "leapfrogging" motion. This phenomenon has received much attention in classical viscous fluids [72] and has also been investigated in VFM simulations [79].

Here we consider a slightly more complex variant of the classical problem: instead of perfectly circular vortex rings we simulate elliptical ones, so that not only their size changes in time but also their shape. The two elliptical vortex rings evolve in a cubic periodic domain of size L. The vortices are initially located on two parallel planes – both orthogonal to the x axis – at a distance $L/4$ from each other. Their minor and major radii are respectively set to $R_a = L/4$ and $R_b = R_a\sqrt{2}$. Their size comparable to L means that the effect of periodic images cannot be fully neglected. Besides, the ellipses are rotated relative to each other so that their major axes are respectively aligned with the y and z axes. This initial configuration is visible in Figure 4(a).

Aiming at a relative accuracy of $\sim 10^{-6}$, each vortex is discretised with a relatively small number of points $N = 32$ (see Appendix A for the impact of spatial resolution on accuracy). Such a coarse discretisation can lead to errors in the local terms (2.4) and (2.11) due to the use of Taylor expansions. As detailed in Remark 2.1, we subdivide filament segments so that Taylor expansions are applied on only about 10% of the local segments, while adaptive tanh-sinh quadratures are used in the remaining local segment portions. Each discrete vortex point $\boldsymbol{s}_i$ is evolved in time according to the equation $\frac{\mathrm{d}\boldsymbol{s}_i}{\mathrm{d}t} = \boldsymbol{v}(\boldsymbol{s}_i)$, where $\boldsymbol{v}(\boldsymbol{s}_i)$ is the Biot–Savart velocity (2.5). For the temporal discretisation we adopt the standard fourth-order explicit Runge–Kutta (RK4) scheme. To guarantee near 6-digit accuracy, we set $\beta = 3.5$ (Figure 3), while the splitting parameter is $\alpha = 7/L$ so that the short-range cut-off distance is $r_c = \beta/\alpha = L/2$.[2]

To preserve stability, the maximum allowed timestep Δt in VFM simulations is roughly proportional to $\ell_{\min}^2$ [70, 42] where $\ell_{\min}$ is the smallest spacing between vortex points – in this test case $\ell_{\min} \approx 2\pi R_a/N$. Physically, this CFL-type condition can be understood by considering a small-amplitude sinusoidal perturbation of wavelength λ on an otherwise straight vortex. The perturbation is physically expected to rotate

[2]The condition $r_c \leq L/2$ ensures that, in short-range computations, a pair of vortex points sees each other at most once. This simplifies the implementation as one does not need to explicitly deal with periodic images. In molecular dynamics this is called the *minimum image convention* [5].

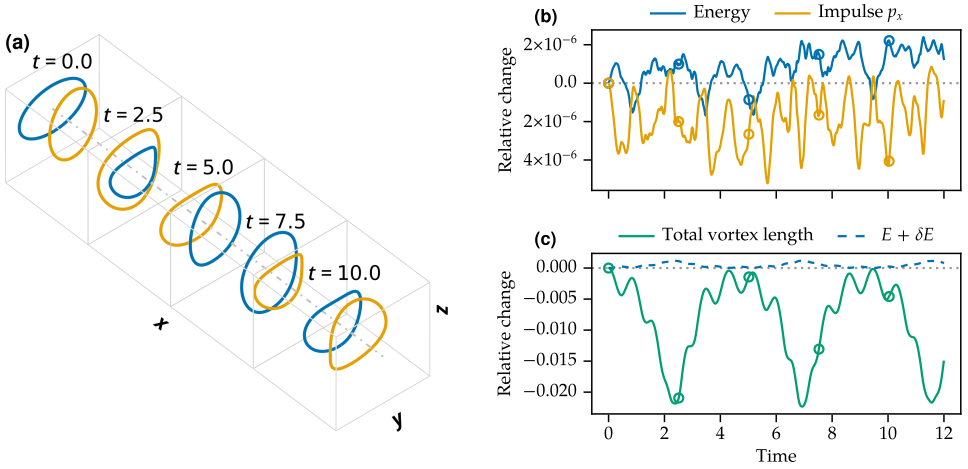

FIG. 4. *Energy and impulse conservation of two leapfrogging elliptical rings. (a) Vortex geometry at different time instants. Vortices mainly advance in the x direction, visiting 4 periodic boxes (light grey cubes) in the shown simulation time. Also shown is the symmetry axis aligned with the x direction (dash-dotted line). (b) Temporal variation of kinetic energy E and hydrodynamic impulse $\boldsymbol{p}$ in the x direction relative to their initial values. (c) Similar for the total vortex length $\mathcal{L}$. The dashed line corresponds to an incorrect estimation of the energy $E' = E + \delta E$ discussed in the main text. Circles in (b) and (c) correspond to the snapshots in (a). Times are normalised by R_a^2/κ.*

about the vortex axis with a period $T_{\mathrm{kw}}(\lambda) \approx \frac{2\lambda^2}{\kappa}\left[\ln\left(\frac{\lambda}{\pi a}\right) + \frac{1}{2} - (\Delta + \gamma)\right]^{-1}$, where $\gamma \approx 0.57721$ is the Euler–Mascheroni constant [75, 8]. Such periodic motion is known as a Kelvin wave. Numerically, the timestep Δt must therefore be small enough to capture the Kelvin wave oscillations associated to the smallest resolved scale $\lambda \sim \ell_{\mathrm{min}}$. In practice, we find that setting $\Delta t = C\, T_{\mathrm{kw}}(\ell_{\mathrm{min}})$ with $C = 1$ is enough for the RK4 scheme to remain stable. We finish by mentioning that the above restriction is actually imposed by the local self-interaction term (2.4) [20, 28]. This suggests the use of splitting methods [57, 18] – perhaps in combination with the Hasimoto transform [44, 11] to deal with the local term – to relax such strong restriction in future works.

Different temporal snapshots of the simulation are visualised in Figure 4(a). The two vortices travel in the x direction while passing through one another as their size and shape change. In particular, at time $t = 2.5$ the first vortex is going through the second one, while the opposite has just occurred at time $t = 7.5$. This cyclic motion is also visible in Figure 4(c), which shows the temporal variation of the total vortex length $\mathcal{L}$. As expected, this quantity is not conserved, displaying relatively slow variations of up to 2% which are clearly correlated with the physical leapfrogging motion. These slow variations coexist with higher-frequency fluctuations which are explained by the non-circular shape of each vortex. On the other hand, as seen in Figure 4(b), the total kinetic energy (2.8) and the total hydrodynamic impulse (2.16) are approximately conserved. Indeed, they display temporal variations of about $10^{-4}\%$ which are not correlated with the physical phenomenon and which are consistent with the $\sim 10^{-6}$ accuracy of the calculations.

We stress that, as discussed in subsection 2.5, observing a high-accuracy energy conservation rests on the proper choice of near-singularity cut-off $\epsilon_\psi = e^{\Delta - 1}/2$ for the streamfunction integral (2.12). Remarkably, while this choice was justified in subsection 2.6 for circular vortex rings, it appears to stay valid for more complex

TABLE 1

Parameters used in performance tests and summary of benchmark results. ε_{tol}, nominal relative tolerance; β, cut-off parameter; C_β, constant prefactor in (5.3); σ and w, NUFFT parameters (see Appendix B). The last two columns summarise the obtained runtimes: (1) prefactor C_{fit} of the least-squares fit $T_{fit}(N) = C_{fit} N \log_{10} N$ and (2) percentage of relative error $\varepsilon_{fit} = \|T_{fit} - T\|/\|T\|$.

ε_{tol}	β	C_β	σ	w	C_{fit} (µs)	ε_{fit} (%)
10^{-3}	2.0	2.06	1.5	4	1.41	6.07
10^{-4}	2.5	1.92	1.5	6	2.07	2.73
10^{-6}	3.5	1.71	1.5	8	4.52	2.25
10^{-8}	4.0	1.64	1.5	10	6.56	4.01
10^{-10}	4.5	1.57	1.5	12	8.95	1.81
10^{-12}	5.0	1.52	1.5	14	12.17	0.70
10^{-14}	5.5	1.47	1.5	16	15.63	0.99

vortex geometries. If instead one used the same cut-off $e^\Delta/2$ as for the velocity integral (exchanging $\Delta - 1$ with Δ in (2.11)), this would result according to (2.8) in an extra (negative) energy contribution $\delta E = -\frac{\kappa^2}{4\pi V}\mathcal{L}$. The resulting energy $E' = E + \delta E$ would not be conserved since the vortex length $\mathcal{L}$ is allowed to fluctuate in time. This is verified in our simulations, where E' displays fluctuations of up to 0.1% (dashed line in Figure 4(c)), which are correlated with the fluctuations of $\mathcal{L}$ and are about 3 orders of magnitude larger than the fluctuations of E.

5.3. Performance. We finally investigate the computational complexity of the proposed method in terms of the total number of vortex points N. We start by discussing how to choose the splitting parameter α to achieve optimal complexity for a fixed value of the accuracy parameter β. Intuitively, α controls the relative amount of work done by short- and long-range computations. From the point of view of short-range computations, one may attempt to choose α so that it reduces their expected complexity from $O(N^2)$ to $O(N)$. Concretely, assuming a spatially homogeneous vortex point distribution, the runtime associated to the short-range part can be estimated as $T^< \propto QN^2(r_c/L)^3$. In terms of α and β, this becomes $T^< \propto QN^2\beta^3/(\alpha L)^3$. Therefore, choosing

$$\alpha(N) = C_\beta N^{1/3}/L \tag{5.3}$$

can be expected to lead to linear complexity, namely $T^< \propto QN(\beta/C_\beta)^3$. While C_β should be regarded as a constant when N is varied, its optimal value may vary with the accuracy parameter β, hence the notation. The next question is how does the choice (5.3) impact the complexity of long-range computations. These are dominated by the cost of the NUFFTs, which basically consist of $\propto QN$ spreading and interpolation operations and $O(M^3 \log M)$ 3D FFTs (see Appendix B). The former are already linear in N and do not depend on α. As for the latter, the size of each 1D FFT is $M \approx k_{\max} L/\pi = 2\beta\alpha L/\pi$, so that choosing α according to (5.3) leads to $M \propto \beta C_\beta N^{1/3}$ and therefore $O(N \log N)$ complexity. In conclusion, choosing α as in (5.3) can be expected to lead, at worst, to $O(N \log N)$ complexity of the full method.

To numerically assess the performance of the method, we adapt the random ellipses test case introduced in subsection 5.1, illustrated in Figure 2 (right). Each ellipse is now discretised with $N_l = 32$ vortex points, and the Biot–Savart velocity is evaluated on all vortex points for systems with up to 8192 vortices ($N = 2^{18}$ points). The benchmarks are executed on a single core of an Intel Core i7-12700H laptop

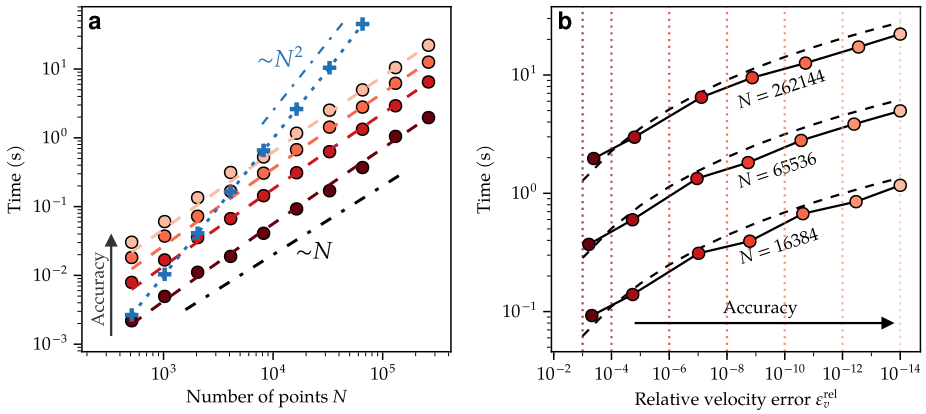

FIG. 5. *Performance of Biot–Savart velocity computations. (a) Runtime as a function of the number of vortex points N. Circles: proposed method for relative tolerances 10^{-3}, 10^{-6}, 10^{-10} and 10^{-14} (from darker to lighter). Dashed lines represent least-squares fits $T_{fit}(N) = C_{fit} N \log_{10} N$; see Table 1 for values of C_{fit}. Blue crosses: naïve Biot–Savart computation in non-periodic domain. (b) Influence of requested accuracy on runtimes, for different values of N. Except for the rightmost circles, errors are relative to the highest accuracy case (10^{-14} relative tolerance). Dotted vertical lines represent the* nominal *relative tolerances (Table 1). Dashed black lines correspond to the empirical scaling $T \propto p_\varepsilon^2 N \log_{10} N$ where $p_\varepsilon = -\log_{10} \varepsilon_{tol}$ is the estimated number of precision digits.*

processor using Julia 1.10.3. The standard cell lists algorithm [47] is used to speed-up short-range computations. We consider 7 different accuracy levels, determined by β and the NUFFT parameters as detailed in Table 1. In each case, we choose C_β so that the time spent in short- and long-range computations is approximately equal. Empirically, we find that $C_\beta \approx 2.6\beta^{-1/3}$ roughly fulfills this criterion, but the proper choice this prefactor is likely dependent on the actual implementation and on the machine where the tests are run.

In Figure 5 we present the results of the numerical experiments. Each timing is obtained as the average over 10 runs. Panel (a) shows the runtime associated to evaluating the Biot–Savart velocity on all vortex points for a wide range of problem sizes N. The proposed method displays near linear $O(N \log N)$ complexity for all different accuracy levels (least-square fits are detailed in Table 1). For comparison, we also show the result of a naïve computation of the Biot–Savart velocity (blue crosses), which displays the expected $O(N^2)$ scaling. Note that the naïve computation does not account for periodic boundary conditions, and thus the results are not directly comparable to those obtained by the proposed method.

From Figure 5(a), it is clear that increasing the accuracy level can have an important impact on runtimes. This is quantified in more detail in panel (b) for a few different problem sizes N and for all different tolerances listed in Table 1. Empirically, runtimes are seen to roughly increase as $T \propto p_\varepsilon^2$ (black dashed lines), where $p_\varepsilon = -\log_{10} \varepsilon_{\text{tol}}$ is the nominal number of accuracy digits. The figure also confirms that the nominal tolerances ε_{tol} (vertical dotted lines) are reasonable (and in fact pessimist) estimates of the actual errors $\varepsilon_v^{\text{rel}}$ (markers). To be clear, except for the $\varepsilon_{\text{tol}} = 10^{-14}$ case, the plotted errors are computed via (5.2) by taking the $\varepsilon_{\text{tol}} = 10^{-14}$ case as reference.

6. Conclusions. We have introduced an efficient technique for the numerical evaluation of the Biot–Savart law in periodic systems when the inducing field (e.g.

current or vorticity) is supported on three-dimensional spatial curves. The approach is adapted from fast Ewald summation methods commonly used to accelerate particle simulations. The present work is relevant to applications in electromagnetism and fluid dynamics, where such fields respectively induce magnetic or velocity fields around the curves. As an application, we have considered the vortex filament model (VFM) describing superfluid helium flows near the absolute zero. In this context, the Fourier-space representation of the induced fields is a very attractive feature of the method, as it allows the efficient evaluation of physically relevant quantities such as energy spectra [12], coarse-grained velocity fields [9, 55] and velocity circulation [60, 63].

Several tunable parameters are introduced by the method, including an inverse "splitting" length scale α and two cut-off parameters r_c and k_max in physical and Fourier spaces. We have shown that the choice of the latter two can be reduced to that of a single non-dimensional coefficient β setting the accuracy of the method, while α can be tuned to maximise performance. For a fixed accuracy level, choosing $\alpha \propto N^{1/3}$ leads to near-linear $O(N \log N)$ complexity of the method, where N is the number of vortex discretisation points. To our knowledge, this parameter selection procedure contrasts with the seemingly common approach in molecular dynamics which consists in first specifying reasonable cut-offs r_c and k_max in dimensional units, and then choosing a value of α which maximises accuracy given these cut-offs [30, 10, 4].

In terms of performance, our implementation is already capable of simulating turbulent systems with a reasonably large number of vortices in reasonable simulation time using thread-based parallelism. We believe the methods presented here are also very amenable to distributed-memory (MPI) or GPU computing, since short-range computations may be treated using standard domain decomposition, while the long-range part could be computed using a distributed [62] or GPU [73] implementation of the NUFFT. A second challenge concerns the timestep stability requirement $\Delta t \propto \ell_\mathrm{min}^2$, which becomes very restrictive as the smallest resolved scale ℓ_min is decreased. Within this work, this issue is partially remediated by representing spatial curves using highly continuous splines, which allows to use relatively large values of ℓ_min. Potential strategies for further improving on these limitations can include the use of splitting [18] or multirate [69] timestepping methods, in which the evolution equations would be split onto a fast term responsible for oscillations (e.g. the local self-induced velocity) and a slower term (e.g. all non-local contributions).

The proposed method can be readily generalised to cases where the vortex circulation κ varies along the filaments. In particular, it may be used in the implementation of vortex filament methods for classical flows [24]. Another possible extension concerns the simpler two-dimensional (2D) case, in which vortex filaments are replaced by point vortices. Periodic point vortex systems have received much attention [80, 61, 77, 53, 40], in part due to their link with statistical mechanics and with 2D turbulence. In that case, Ewald summation can be achieved using a function splitting similar to (3.3) [43, 23]. Finally, and more directly relevant to the motivation of this work, the proposed method may be applied to the study of *finite-temperature* liquid helium-4, in which the superfluid flow coexists and interacts with a normal viscous fluid. A very relevant model for describing this regime consists in coupling the VFM for the superfluid with the incompressible Navier–Stokes equations governing the normal fluid [49, 83, 35]. Note that, when the latter are solved using standard pseudo-spectral methods [22], the Fourier-space representation included in the present method can enable an accurate and efficient two-way transfer of interaction forces between the vortices and the normal fluid.

Acknowledgments. The author acknowledges support from the French Agence Nationale de la Recherche through the QuantumVIW project (Grant No. ANR-23-CE30-0024-04).

Appendix A. Discretisation of spatial curves.

The evaluation of the Biot–Savart law (2.2) and its local regularisation (2.4) on vortex filaments require the estimation of local tangent and curvature vectors from a set of discrete points $\{\boldsymbol{s}_j\}_{j=1}^N$ representing a spatial curve. This assumes that curves can be parametrised by some sufficiently regular function $\boldsymbol{s}(\tau)$ whose derivatives on discrete locations τ_j (such that $\boldsymbol{s}(\tau_j) = \boldsymbol{s}_j$) are well defined. Furthermore, the use of quadratures to estimate line integrals (see (2.18)) requires being able to evaluate (interpolate) $\boldsymbol{s}(\tau)$ and its derivatives at any arbitrary location τ.

Here we compare different discretisation strategies allowing to achieve the above requirements. First, note that in general the nodes $\boldsymbol{s}_j$ are not equispaced, and that the choice of the parametrisation (the τ_j values) is arbitrary. A natural and common choice when describing arbitrary curves is the so-called chordal parametrisation. This consists of setting $\tau_j = \tau_{j-1} + |\boldsymbol{s}_j - \boldsymbol{s}_{j-1}|$ for $j \in \{2, \dots, N\}$, $\tau_1 = 0$, where $|\cdot|$ represents the Euclidean distance. In this case, τ is a rough approximation of the curve arc length ξ, tending to ξ as the number of points increases. Here we adopt this choice as well, except in the case of Fourier series representation for reasons that will be clear below.

A.1. Curve discretisation schemes.

A.1.1. Finite difference approximations. One possible strategy is to start by estimating derivatives on the nodes $\boldsymbol{s}_j$ using some finite difference scheme on a non-uniform grid (since, in general, the τ_j values are not equispaced). Here we consider an explicit fourth-order scheme [36] which is commonly used in VFM calculations [6]. This scheme provides an estimation of the derivatives $\boldsymbol{s}_i'$ and $\boldsymbol{s}_i''$ on node i using a five-point stencil $\{\boldsymbol{s}_{i-2}, \dots, \boldsymbol{s}_{i+2}\}$ centred at $\boldsymbol{s}_i$.

As a second step, once the derivatives on filament nodes are known, one can use Hermite interpolations, which rely on the positions and derivatives on nodes $\boldsymbol{s}_i$ and $\boldsymbol{s}_{i+1}$, to reconstruct the segment connecting these nodes. Here we use quintic Hermite interpolations, which require up to second derivatives and have global C^2 continuity.

This scheme presents the advantage that all operations are local and therefore efficient: derivatives can be estimated from neighbouring points, and interpolations in-between two nodes only require values at those nodes.

A.1.2. Spline interpolation. Secondly, we consider parametric splines of order k [27]. Here by convention the polynomial order is $k - 1$, so that $k = 4$ and 6 respectively correspond to cubic and quintic splines. The basic idea is to parametrise a spatial curve in terms of a B-spline basis,

$$\boldsymbol{s}(\tau) = \sum_{j=1}^{N} \boldsymbol{c}_j b_j(\tau), \quad \text{for } \tau \in [0, \mathcal{T}], \tag{A.1}$$

where $b_j(\tau)$ is a B-spline basis function. The basis functions are fully defined by the choice of spline order k and spline *knots* t_j, which are a set of locations in $[0, \mathcal{T}]$. Because the curves considered here are periodic, we can simply set $t_j = \tau_j$ for all j (free-ended curves would require some extra care). An important and convenient property of the B-splines is their compact support. Namely, for even k, the B-spline $b_j(\tau)$ is zero outside of the interval $[t_{j-k/2}, t_{j+k/2}]$.

Spline interpolation consists in determining the coefficients $\boldsymbol{c}_j$ given the equalities $\boldsymbol{s}(\tau_i) = \boldsymbol{s}_i$ for $i \in \{1, \dots, N\}$. This leads to a linear system of the form $A_{ij}\boldsymbol{c}_j = \boldsymbol{s}_i$ where $A_{ij} = b_j(\tau_i)$. Thanks to the compact support of the B-splines, $\mathbf{A}$ is a cyclic banded matrix, which enables the use of specialised algorithms to efficiently solve the system. More precisely, with our choice of knots, $\mathbf{A}$ has $k-1$ bands, as well as a few out-of-bands entries in the top right and bottom left corners due to periodicity. In particular, for cubic splines ($k = 4$), $\mathbf{A}$ is almost tridiagonal with two out-of-bands entries (one on each corner). A standard solution is to convert the system to a fully tridiagonal one using the Sherman–Morrison formula and then use the Thomas algorithm to solve the system in $O(N)$ time [81]. A similar procedure [56] can be applied to efficiently solve the cyclic pentadiagonal system arising from periodic quintic spline interpolation ($k = 6$).

The derivative of a spline $\boldsymbol{s}(\tau)$ of order k is itself a spline of order $k-1$, whose coefficients $\boldsymbol{c}'_j$ can be directly obtained from the coefficients $\boldsymbol{c}_j$ using a local differentiation formula. Moreover, splines can be efficiently evaluated at locations τ using de Boor's algorithm [27]. Therefore, knowing the coefficients $\boldsymbol{c}_j$ one can evaluate curve positions and derivatives at any location τ.

Splines have global continuity C^{k-2}, and therefore they are very interesting for describing smooth curves. Moreover, they can be readily used when the spatial distribution of discretisation points is inhomogeneous, which can be convenient for accurately representing curves with non-uniform features – such as localised large-curvature regions where a dense point distribution may be desired.

A.1.3. Fourier series representation. Since we consider periodic curves, it may also be natural to represent them using a truncated Fourier series,

$$\boldsymbol{s}(\tau) = \sum_{k=-N/2}^{N/2-1} \hat{\boldsymbol{s}}_k \, e^{\mathrm{i}k\tau} \quad \text{for } \tau \in [0, 2\pi]. \tag{A.2}$$

Unlike the previous cases, here we consider the locations τ_i to be equispaced, $\tau_i = 2\pi(i-1)/N$, which allows the use of FFTs to efficiently obtain the complex coefficients $\hat{\boldsymbol{s}}_k$. Moreover, since the locations $\boldsymbol{s}(\tau)$ are real-valued, the coefficients satisfy the Hermitian symmetry $\hat{\boldsymbol{s}}_{-k} = \hat{\boldsymbol{s}}_k^*$, and one can use real-to-complex transforms to avoid computing negative wavenumbers k. Curve derivatives $\boldsymbol{s}'_i$ on nodes are obtained by backwards FFT of $\hat{\boldsymbol{s}}'_k = \mathrm{i}k\hat{\boldsymbol{s}}_k$ (and similarly for $\boldsymbol{s}''_i$). In principle, it is possible to directly evaluate $\boldsymbol{s}(\tau)$ and its derivatives in-between filament nodes using (A.2) (or via NUFFTs). However, for convenience and performance we prefer to approximate these evaluations using quintic Hermite interpolations, as already described in Appendix A.1.1.

One can expect exponential convergence with N when describing simple smooth spatial curves using Fourier series. However, when the discretisation points $\boldsymbol{s}_i$ are strongly non-uniformly spaced, the equidistant parametrisation required for the FFTs can lead to unwanted artefacts (such as artificial cusps or loops) in the resulting curve $\boldsymbol{s}(\tau)$.

A.2. Accuracy of different discretisation methods. To test the accuracy and convergence rate of the different discretisation methods introduced above, we consider two simple settings where numerical results can be compared with analytical predictions. The three approaches introduced above are compared, and both cubic and quintic splines are considered.

The first setting is a trefoil knot vortex (5.1), as represented in Figure 2 (left). The hydrodynamic impulse (2.16) of this curve can be analytically obtained, $\boldsymbol{p}_{\text{trefoil}} =$

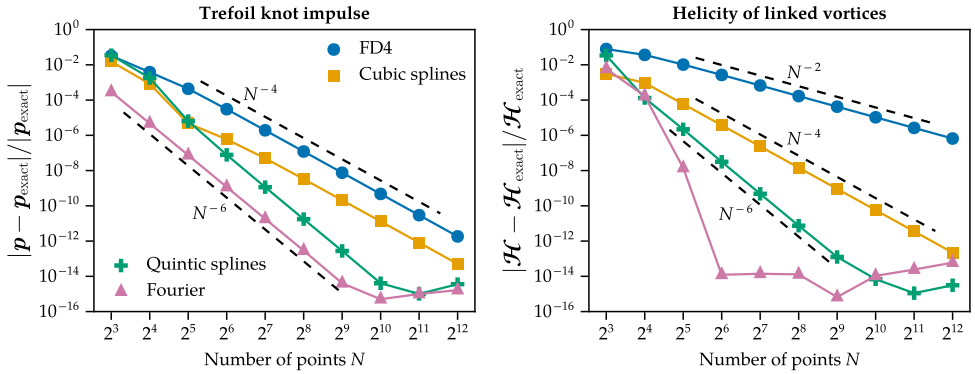

FIG. 6. *Accuracy of different filament discretisation methods for a varying number of discretisation points N. Left: impulse estimation $\boldsymbol{p}$ for a trefoil knot filament. Right: helicity $\mathcal{H}$ of two linked vortices. Different markers correspond to different discretisation schemes (circles, 4th order finite differences; squares, cubic splines; crosses, quintic splines; triangles, Fourier series). In all cases, integrations are estimated using 3-point Gauss–Legendre quadratures on each filament segment (4-point quadratures yield the same results).*

$7\pi\kappa R^2/V\hat{z}$. Note that the impulse is a purely geometrical quantity, and therefore the numerical results presented in the following do not depend on the accuracy of Biot–Savart computations. For each considered method, we first discretise the curve by evaluating (5.1) on N locations $\boldsymbol{s}_i = \boldsymbol{X}(\theta_i)$ with $\theta_i = 2\pi(i-1)/N$. Then, the curve is re-parametrised as $\boldsymbol{s}(\tau)$ by setting the discrete locations τ_i from the points $\boldsymbol{s}_i$ as discussed in the previous section, and interpolation coefficients are then computed. The numerical results for this setting are shown in Figure 6 (left). All methods display polynomial convergence with the number of discretisation points N. Finite differences and cubic splines display fourth-order convergence, while quintic splines and Fourier series converge as N^{-6}. The polynomial (as opposed to exponential) convergence of Fourier series is likely explained by the use of quintic Hermite interpolations instead of exact evaluation to obtain values on filaments.

The second case corresponds to two linked elliptical vortices, topologically equivalent to a Hopf link. In this case, independently of the shape of the curves (provided they are unknotted), the expected total helicity is simply $\mathcal{H} = 2\kappa^2\mathfrak{L}$ [58] where $\mathfrak{L} = 1$ is the linking number of a Hopf link. Here the helicity is defined as

$$\mathcal{H} = \oint_\Omega \boldsymbol{v}(\boldsymbol{x}) \cdot \boldsymbol{\omega}(\boldsymbol{x}) \, \mathrm{d}^3\boldsymbol{x} = \kappa \oint_{\mathcal{C}} \boldsymbol{v}(\boldsymbol{s}) \cdot \mathrm{d}\boldsymbol{s}. \tag{A.3}$$

We choose the ellipses to have semi-major and semi-minor axes $a = 2$ and $b = 1$, and their centres are separated by $\delta = 2/3$. Their semi-major axes are aligned, while their semi-minor ones are relatively rotated by an angle $7\pi/20$ on the plane orthogonal to their semi-major axes. Note that estimating the total helicity of the system via (A.3) requires knowing the filament velocities, which are obtained using the Biot–Savart law. Moreover, the vortices interact with their periodic images (the system is triply periodic with $L = 2\pi$), but this in theory does not affect the total helicity (unless vortices are linked to their periodic images, which is not the case here). To ensure that accuracy is unaffected by errors associated to fast Ewald summation, we set cut-off parameters to $\beta^< = \beta^> = 6$ (see Figure 3), and NUFFTs are computed to very high accuracy as well. The number of resolved Fourier modes for long-range interactions is $M^3 = 48^3$. Convergence results are shown in Figure 6 (right). Fourth-order finite

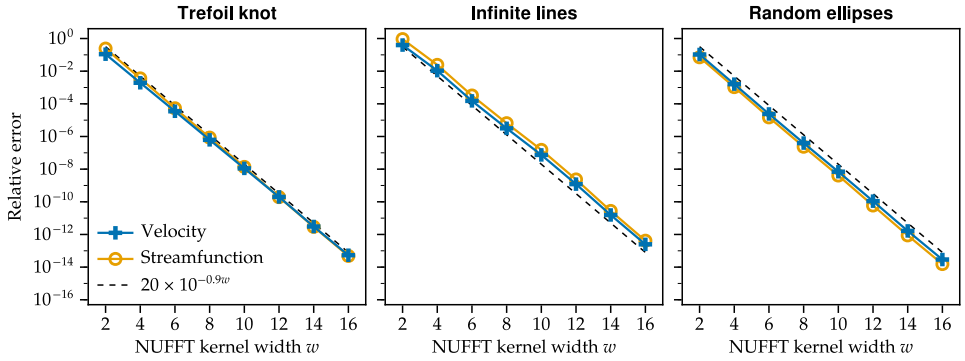

FIG. 7. *Velocity and streamfunction errors associated to the NUFFT precision, for the test cases illustrated in Figure 2. Relative root-mean-square error of long-range filament velocity* (+) *and streamfunction values* (×) *at varying NUFFT spreading widths* w*. NUFFTs are computed using the backwards Kaiser–Bessel spreading kernel* (B.1) *with an oversampling factor* $\sigma = 1.5$*. Dashed lines represent an exponential decay* $\sim 10^{-0.9w}$*. In all cases, the long-range cut-off parameter is set to* $\beta^> = 8$*, ensuring subdominant cut-off errors. The Ewald splitting parameter is kept constant,* $\alpha = 24/L$*.*

differences actually display second-order convergence, while splines of order k converge as N^{-k} as in the trefoil knot case. Fourier series display a much faster (exponential) convergence, suggesting that these results are less affected by Hermite interpolations than in the trefoil knot case.

In both cases, to estimate line integrals according to (2.18), we have used 3-point Gauss–Legendre quadratures on each filament segment (i.e. a total of $3N$ points on the filament), which is enough for results to converge. However, using as few as 2 points degrades the convergence rates for finite differences and splines (not shown). More generally, our numerical experiments empirically suggest that a minimum of $k/2$ Gauss–Legendre quadrature points per segment are required for proper convergence of k-th order splines.

Appendix B. NUFFT details and influence on Biot–Savart accuracy.

Here we consider the error induced by the NUFFT tolerance on the long-range component of Biot–Savart computations. Briefly, a type-1 NUFFT can be summarised in three steps: (i) spreading values from a set of non-uniform locations onto a uniform grid using some smoothing kernel, (ii) performing a regular FFT on that grid, and (iii) undoing the spreading operation by deconvolution in Fourier space [38]. The type-2 NUFFT is the adjoint operation, roughly consisting of the same steps in opposite order. To reduce aliasing errors, the first two steps are done in a grid of σM points in each direction, where M is the number of desired Fourier modes and $\sigma > 1$ is an oversampling factor. The accuracy of the NUFFT is then determined by (i) the choice of spreading kernel, (ii) the width of the kernel (usually in number of grid points w), and (iii) the oversampling factor σ [31, 64, 16]. The computational complexity of a 3D NUFFT on N non-uniform points can be estimated to be $O(w^3 N + (\sigma M)^3 \log M)$.

To evaluate NUFFTs, here we use the (backwards) Kaiser–Bessel function as spreading kernel [48, 64]. For a support in $x \in [-1, 1]$, the kernel and its Fourier transform are

$$\varphi(x) = \frac{\sinh\left(\zeta\sqrt{1-x^2}\right)}{\pi\sqrt{1-x^2}} \quad \leftrightarrow \quad \hat{\varphi}(k) = I_0\left(\sqrt{\zeta^2 - k^2}\right), \tag{B.1}$$

where I_0 is the zero-th order modified Bessel function of the first kind. The shape parameter is chosen as $\zeta = \gamma w\pi(1 - 1/2\sigma)$ to minimise NUFFT errors [64]. Here $\gamma = 0.995$ is an empirical "safety factor" [16] which slightly improves accuracy. For performance reasons, the kernel $\varphi(x)$ is evaluated via a fast and accurate piecewise polynomial approximation [16, 71] which takes advantage of SIMD vectorisation on modern CPUs. Our NUFFT implementation is written in Julia and is openly available on https://github.com/jipolanco/NonuniformFFTs.jl. Compared to other existent libraries available in Julia [16, 50], our implementation allows real-valued non-uniform data, in which case it takes advantage of real-to-complex FFT routines of the FFTW library [34].

In Figure 7, we evaluate the effect of varying the NUFFT accuracy on the relative error of long-range Biot–Savart computations, for the cases illustrated in Figure 2. Concretely, we set the oversampling factor to $\sigma = 1.5$ and vary the spreading width w. For comparison, we compute a reference solution with $\sigma = 2$ and $w = 16$, which ensures that its accuracy is dominated by roundoff error. As seen in Figure 7, the relative error associated to the long-range streamfunction and velocity decay exponentially as $\sim 10^{-0.9w}$ when using an oversampling factor $\sigma = 1.5$. For $\sigma = 2$, the same errors decay as $\sim 10^{-w}$ (not shown), but at the expense of costlier FFTs.

Appendix C. Derivation of truncation error estimates.

C.1. Short-range errors. The absolute error associated to the truncation of the short-range velocity integral (3.7) at a single location $\boldsymbol{s}_0$ is given by

$$\boldsymbol{v}^{<}_{\mathrm{err}}(\boldsymbol{s}_0) = \frac{\kappa}{4\pi} \sum_{\boldsymbol{n}\in\mathbb{Z}^3} \int_{r>r_\mathrm{c}} g^{<}(\boldsymbol{r})\, \frac{\boldsymbol{r}\times \mathrm{d}\boldsymbol{s}}{|\boldsymbol{r}|^3} \quad \text{with } \boldsymbol{r} = \boldsymbol{s} - \boldsymbol{s}_0 + \boldsymbol{n}L, \tag{C.1}$$

where the integral is over all vortex elements at a distance $|\boldsymbol{r}| = r > r_\mathrm{c}$ from the point $\boldsymbol{s}_0$. We first note that, for large r, the weight function $g^{<}(\boldsymbol{r})$ defined in (3.5) behaves as $g^{<}(\boldsymbol{r}) = \frac{2\alpha r}{\sqrt{\pi}} e^{-(\alpha r)^2} \left[1 + O(r^{-5/2})\right]$. Secondly, we assume that the vortices are homogeneously distributed in the spatial domain, so that the total vortex length within a spherical shell of radii $[r, r + \mathrm{d}r]$ with infinitesimal thickness $\mathrm{d}r$ is the product between the mean vortex line density and the volume of the shell, namely $\mathrm{d}\mathcal{L} = (\mathcal{L}/V)(4\pi r^2 \mathrm{d}r)$. Here $\mathcal{L}$ is the total vortex length within a periodic cell of volume $V = L^3$. Thirdly, we assume that the orientations of the vortex elements $\mathrm{d}\boldsymbol{s}$ within the shell are random and independent of the separation vector $\boldsymbol{r}$. Under these assumptions, we can formally estimate the root mean square (rms) short-range velocity error to be

$$\varepsilon^{<}_v \approx \frac{\kappa}{4\pi} \int_{r_\mathrm{c}}^{\infty} \frac{2\alpha r}{\sqrt{\pi}} e^{-(\alpha r)^2} \frac{r\,\mathrm{d}\mathcal{L}}{r^3} = \frac{2\kappa\mathcal{L}}{\sqrt{\pi}\alpha V} \int_{\beta^<}^{\infty} x\, e^{-x^2}\, \mathrm{d}x, \tag{C.2}$$

where $\beta^< = \alpha r_\mathrm{c}$. This integrates exactly to

$$\varepsilon^{<}_v \approx \frac{\kappa\mathcal{L}}{\sqrt{\pi}\alpha V}\, e^{-(\beta^<)^2}. \tag{C.3}$$

The same procedure can be applied to the short-range streamfunction (3.6). Considering that $\operatorname{erfc}(x) = \frac{e^{-x}}{x^{3/2}} \left[1 + O(x^{-2})\right]$ for large x, one obtains the associated rms error

$$\varepsilon^{<}_\psi \approx \frac{\kappa}{4\pi} \int_{r_\mathrm{c}}^{\infty} \frac{e^{-(\alpha r)^2}}{(\alpha r)^{3/2}} \frac{\mathrm{d}\mathcal{L}}{r} = \frac{\kappa\mathcal{L}}{\alpha^2 V} \int_{\beta^<}^{\infty} \frac{e^{-x^2}}{\sqrt{x}}\, \mathrm{d}x = \frac{\kappa\mathcal{L}}{2\alpha^2 V}\, \Gamma\left(\frac{1}{4}, (\beta^<)^2\right). \tag{C.4}$$

Here $\Gamma(s,x) = \int_x^\infty u^{s-1}e^{-u}\,\mathrm{d}u$ is the upper incomplete gamma function, which behaves asymptotically as $\Gamma(s,x) = x^{s-1}e^{-x}\left[1+O(x^{-1})\right]$. Replacing this in (C.4) finally leads to the estimate

$$\varepsilon_\psi^< \approx \frac{\kappa\mathcal{L}}{2\alpha^2 V}\,\frac{e^{-(\beta^<)^2}}{(\beta^<)^{3/2}}. \tag{C.5}$$

In both cases, the truncation error is dominated by the exponential decay with the square of the non-dimensional cut-off parameter $\beta^<$.

C.2. Long-range errors. The rms error associated to the long-range velocity field is defined by

$$\left(\varepsilon_v^>\right)^2 = \frac{1}{V}\int_\Omega \left|\boldsymbol{v}^>(\boldsymbol{x}) - \tilde{\boldsymbol{v}}^>(\boldsymbol{x})\right|^2 \mathrm{d}^3\boldsymbol{x}, \tag{C.6}$$

where $\boldsymbol{v}^>(\boldsymbol{x})$ is the Fourier-truncated velocity field defined in (3.10), and $\tilde{\boldsymbol{v}}^>(\boldsymbol{x})$ is the non-truncated velocity field – obtained by replacing $\mathbb{K}$ with $\mathbb{K}_\infty \equiv \left\{\frac{2\pi m}{L}; m\in\mathbb{Z}\right\}$ in (3.10). Using Parseval's theorem, the error can be equivalently written as

$$\left(\varepsilon_v^>\right)^2 = \sum_{\boldsymbol{k}\in(\mathbb{K}_\infty\setminus\mathbb{K})^3} |\hat{\boldsymbol{v}}_{\boldsymbol{k}}^>|^2 \;\le\; \sum_{|\boldsymbol{k}|>k_{\max}} |\hat{\boldsymbol{v}}_{\boldsymbol{k}}^>|^2, \tag{C.7}$$

where $\hat{\boldsymbol{v}}_{\boldsymbol{k}}^>$ are the Fourier coefficients of the long-range velocity. The last inequality results from converting from a "cubic" truncation ($k_i > k_{\max}$ for each Cartesian component i) to a spherical truncation $|\boldsymbol{k}| > k_{\max}$.

In physical terms, the velocity error (C.7) can be directly related to the kinetic energy spectrum $E(k)$, defined in a periodic domain as

$$E(k) = \frac{1}{2\Delta k}\sum_{|\boldsymbol{k}|\in\mathcal{I}_k} |\hat{\boldsymbol{v}}_{\boldsymbol{k}}|^2 \quad \text{for } k = \frac{0\pi}{L}, \frac{2\pi}{L}, \frac{4\pi}{L}, \ldots, \tag{C.8}$$

where $\hat{\boldsymbol{v}}_{\boldsymbol{k}}$ is a Fourier coefficient of the (unsmoothed) velocity field, $\Delta k = 2\pi/L$ is the distance between two successive discrete wavenumbers, and the sum is over all 3D wavenumbers $\boldsymbol{k}$ within the spherical shell $\mathcal{I}_k \equiv \left[k-\frac{\Delta k}{2}, k+\frac{\Delta k}{2}\right)$. Formally, the total kinetic energy (per unit mass) is then $E = \int_0^\infty E(k)\,\mathrm{d}k$. In the present context, (C.8) will allow us to quantify the effect of truncating Fourier-space computations at a certain wavenumber $k_{\max}$.

At sufficiently small scales – smaller than the typical distance between vortices and the typical curvature radius of the vortices – vortex filaments can be considered as isolated straight lines. Under that assumption, it can be shown that the energy spectrum is $E(k) = \frac{\kappa^2\mathcal{L}}{4\pi V}\,k^{-1}$ for sufficiently large k [2]. This prediction is accurately verified in Figure 8 (blue curves) for the three different vortex configurations illustrated in Figure 2. Note that the relatively slow decay as k^{-1} suggests that the integral of this spectrum (and thus the total energy) presents a logarithmic divergence. In fact, physically this is not the case, since the spectrum is regularised near the wavenumber $k_a = 2\pi/a$ associated to the vortex core size a, which is orders of magnitude smaller than the scales described by the VFM.

To provide an estimate of the error in the computation of the coarse-grained velocity $\boldsymbol{v}^>$, we now consider its associated kinetic energy spectrum $E^>(k) = E(k)\,e^{-k^2/2\alpha^2} = \frac{\kappa^2\mathcal{L}}{4\pi V}\,k^{-1}e^{-k^2/2\alpha^2}$ for large k. Note that the integral of $E^>(k)$ is the contribution of the

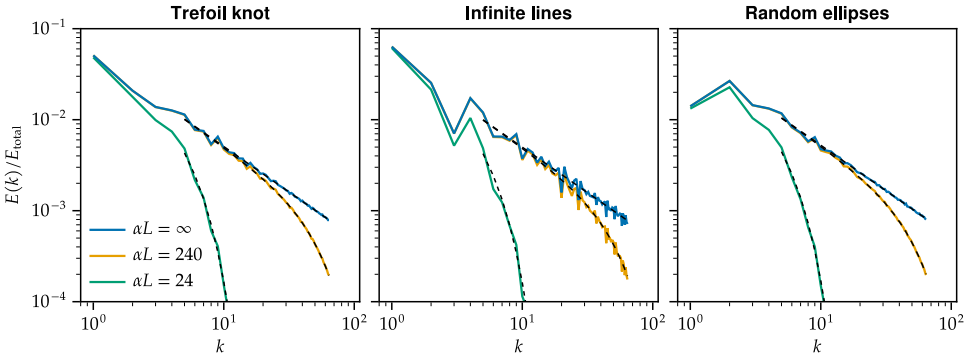

FIG. 8. *Kinetic energy spectra associated to three different test cases. The domain period is* $L = 2\pi$ *and the spectra are truncated at* $k_{max} = 64$ *(the number of Fourier modes in each direction is* $M = 128$*). Each panel shows the truncated energy spectrum associated to the unsmoothed velocity field* ($\alpha L = \infty$)*, as well as the energy spectra associated to two Gaussian-smoothed velocity fields* ($\alpha L = 240$ *and* 24)*. The spectra are normalised by the total kinetic energy* E_{total} *obtained via (2.8). At large wavenumbers* k*, the spectra are well approximated by* $E(k) = \frac{\kappa^2 \mathcal{L}}{4\pi V} k^{-1} e^{-k^2/2\alpha^2}$ *(black dashed lines), where* $\mathcal{L}$ *is the total vortex length in each test case.*

long-range velocity field to the total kinetic energy E. Examples of $E^>(k)$ are shown in Figure 8 for two different values of the smoothing parameter α. In the figure, the long-range truncation is at $k_{\max} = 64$ and the domain period is $L = 2\pi$. It is clear that, if α is too large (e.g. $\alpha L = 240$ in the figure, corresponding to $\beta^> = k_{\max}/2\alpha \approx 0.84$), then this level of truncation is insufficient to capture the whole long-range kinetic energy. A smaller value of α ($\alpha L = 24$ in the figure) means that the full coarse-grained energy can be captured when truncating at the same $k_{\max}$.

Assuming the truncation wavenumber $k_{\max}$ to be sufficiently large, we can now estimate an upper bound for the absolute truncation error (C.7) associated to the long-range velocity,

$$\frac{(\varepsilon_v^>)^2}{2} \approx \int_{k_{\max}}^{\infty} E^>(k)\,\mathrm{d}k = \frac{\kappa^2 \mathcal{L}}{4\pi V} \int_{k_{\max}}^{\infty} \frac{e^{-k^2/2\alpha^2}}{k}\,\mathrm{d}k = \frac{\kappa^2 \mathcal{L}}{8\pi V} E_1\Big[2(\beta^>)^2\Big], \tag{C.9}$$

where $E_1(x) = \int_x^\infty \frac{e^{-u}}{u}\,\mathrm{d}u$ is the exponential integral function. For large x, $E_1(x) = \frac{e^{-x}}{x}\left[1 - \frac{1}{x} + O(x^{-2})\right]$, and therefore $(\varepsilon_v^>)^2 \approx \frac{\kappa^2 \mathcal{L}}{8\pi V} e^{-2(\beta^>)^2}/(\beta^>)^2$. Finally, the estimated error associated to the long-range velocity field is

$$\varepsilon_v^> \approx \kappa \left(\frac{\mathcal{L}}{8\pi V}\right)^{1/2} \frac{e^{-(\beta^>)^2}}{\beta^>}. \tag{C.10}$$

A similar estimation can be done for the long-range truncation error $\varepsilon_\psi^>$ associated to the streamfunction. This is obtained by integrating

$$\frac{(\varepsilon_\psi^>)^2}{2} \approx \int_{k_{\max}}^{\infty} \frac{E^>(k)}{k^2}\,\mathrm{d}k \approx \frac{\kappa^2}{4\alpha^2} \frac{\mathcal{L}}{4\pi V} \frac{e^{-2(\beta^>)^2}}{\left[2(\beta^>)^2\right]^2}, \tag{C.11}$$

where we have used $\int_x^\infty \frac{e^{-u}}{u^2}\,\mathrm{d}u = \frac{e^{-x}}{x} - E_1(x) = \frac{e^{-x}}{x^2}\left[1 + O(x^{-1})\right]$. This results in

$$\varepsilon_\psi^> \approx \frac{\kappa}{2\alpha} \left(\frac{\mathcal{L}}{8\pi V}\right)^{1/2} \frac{e^{-(\beta^>)^2}}{(\beta^>)^2}. \tag{C.12}$$

Similarly to short-range errors, the long-range truncation errors are dominated by the exponential decay with the square of the cut-off parameter $\beta^>$. We stress that the estimates (C.10) and (C.12) are expected to be valid for sufficiently large values of $k_{\max}$ (or $\beta^>$). Furthermore, they can be expected to overestimate the actual errors due to the inequality in (C.7), which discards the contributions of some large-wavenumber modes which are actually resolved in the simulations (more precisely, a factor of $\sim 6/\pi$ Fourier modes is discarded by the above estimates).

C.3. Combined truncation errors. To avoid unnecessary computations and achieve optimal accuracy for a given computational cost, one would like the errors associated to the short- and long-range components to be equivalent. In light of the above estimations, it seems natural to set their respective non-dimensional cut-off parameters to the same value, $\beta^< = \beta^> = \beta$, which also helps reducing the size of the parameter space. This finally leads to the following total truncation error estimates:

$$\varepsilon_v \approx \kappa \left[\frac{\mathcal{L}}{\sqrt{\pi}\alpha V} + \frac{1}{\beta} \left(\frac{\mathcal{L}}{8\pi V} \right)^{1/2} \right] e^{-\beta^2}, \tag{C.13}$$

$$\varepsilon_\psi \approx \frac{\kappa}{2\alpha} \left[\frac{\mathcal{L}}{\alpha V} + \frac{1}{\beta^{1/2}} \left(\frac{\mathcal{L}}{8\pi V} \right)^{1/2} \right] \frac{e^{-\beta^2}}{\beta^{3/2}}. \tag{C.14}$$

In summary, the accuracy of the method is mainly controlled by a unique non-dimensional parameter β, from which both physical- and Fourier-space cut-offs r_c and $k_{\max}$ can be obtained. This still leaves the inverse splitting distance α as a free parameter that can be adjusted to optimise performance (subsection 5.3).

REFERENCES

[1] H. Adachi, S. Fujiyama, and M. Tsubota, *Steady-state counterflow quantum turbulence: Simulation of vortex filaments using the full Biot-Savart law*, Phys. Rev. B, 81 (2010), p. 104511, https://doi.org/10.1103/PhysRevB.81.104511.

[2] T. Araki, M. Tsubota, and S. K. Nemirovskii, *Energy Spectrum of Superfluid Turbulence with No Normal-Fluid Component*, Phys. Rev. Lett., 89 (2002), p. 145301, https://doi.org/10.1103/PhysRevLett.89.145301.

[3] R. J. Arms and F. R. Hama, *Localized-Induction Concept on a Curved Vortex and Motion of an Elliptic Vortex Ring*, Phys. Fluids, 8 (1965), pp. 553–559, https://doi.org/10.1063/1.1761268.

[4] A. Arnold, F. Fahrenberger, C. Holm, O. Lenz, M. Bolten, H. Dachsel, R. Halver, I. Kabadshow, F. Gähler, F. Heber, J. Iseringhausen, M. Hofmann, M. Pippig, D. Potts, and G. Sutmann, *Comparison of scalable fast methods for long-range interactions*, Phys. Rev. E, 88 (2013), p. 063308, https://doi.org/10.1103/PhysRevE.88.063308.

[5] A. Arnold and C. Holm, *Efficient Methods to Compute Long-Range Interactions for Soft Matter Systems*, in Advanced Computer Simulation Approaches for Soft Matter Sciences II, C. Holm and K. Kremer, eds., Advances in Polymer Science, Springer, Berlin, Heidelberg, 2005, pp. 59–109, https://doi.org/10.1007/b136793.

[6] A. W. Baggaley and C. F. Barenghi, *Spectrum of turbulent Kelvin-waves cascade in superfluid helium*, Phys. Rev. B, 83 (2011), p. 134509, https://doi.org/10.1103/PhysRevB.83.134509.

[7] A. W. Baggaley and C. F. Barenghi, *Tree Method for Quantum Vortex Dynamics*, J. Low Temp. Phys., 166 (2012), pp. 3–20, https://doi.org/10.1007/s10909-011-0405-6.

[8] A. W. Baggaley and J. Laurie, *Kelvin-wave cascade in the vortex filament model*, Phys. Rev. B, 89 (2014), p. 014504, https://doi.org/10.1103/PhysRevB.89.014504.

[9] A. W. Baggaley, J. Laurie, and C. F. Barenghi, *Vortex-Density Fluctuations, Energy Spectra, and Vortical Regions in Superfluid Turbulence*, Phys. Rev. Lett., 109 (2012), p. 205304, https://doi.org/10.1103/PhysRevLett.109.205304.

[10] V. Ballenegger, J. J. Cerda, O. Lenz, and Ch. Holm, *The optimal P3M algorithm for computing electrostatic energies in periodic systems*, J. Chem. Phys., 128 (2008), p. 034109, https://doi.org/10.1063/1.2816570.

[11] V. Banica, G. Maierhofer, and K. Schratz, *Numerical Integration of Schrödinger Maps via the Hasimoto Transform*, SIAM J. Numer. Anal., 62 (2024), pp. 322–352, https://doi.org/10.1137/22M1531555.

[12] C. F. Barenghi, V. S. L'vov, and P.-E. Roche, *Experimental, numerical, and analytical velocity spectra in turbulent quantum fluid*, Proc. Natl. Acad. Sci. USA, 111 (2014), pp. 4683–4690, https://doi.org/10.1073/pnas.1312548111.

[13] C. F. Barenghi, L. Skrbek, and K. R. Sreenivasan, *Introduction to quantum turbulence*, Proc. Natl. Acad. Sci. USA, 111 (2014), pp. 4647–4652, https://doi.org/10.1073/pnas.1400033111.

[14] C. F. Barenghi, L. Skrbek, and K. R. Sreenivasan, *Quantum Turbulence*, Cambridge University Press, Cambridge, 2023, https://doi.org/10.1017/9781009345651.

[15] J. Barnes and P. Hut, *A hierarchical O(N log N) force-calculation algorithm*, Nature, 324 (1986), pp. 446–449, https://doi.org/10.1038/324446a0.

[16] A. H. Barnett, J. Magland, and L. af Klinteberg, *A Parallel Nonuniform Fast Fourier Transform Library Based on an "Exponential of Semicircle" Kernel*, SIAM J. Sci. Comput., 41 (2019), pp. C479–C504, https://doi.org/10.1137/18M120885X.

[17] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, *Julia: A fresh approach to numerical computing*, SIAM Rev., 59 (2017), pp. 65–98, https://doi.org/10.1137/141000671.

[18] S. Blanes, F. Casas, and A. Murua, *Splitting Methods for differential equations*, May 2024, https://arxiv.org/abs/2401.01722. To appear in Acta Numer. (2024).

[19] J. P. Boyd, *Chebyshev and Fourier Spectral Methods*, Dover Publications, Mineola, N.Y, second edition ed., Dec. 2001.

[20] T. F. Buttke, *A numerical study of superfluid turbulence in the self-induction approximation*, Journal of Computational Physics, 76 (1988), pp. 301–326, https://doi.org/10.1016/0021-9991(88)90145-3.

[21] A. J. Callegari and L. Ting, *Motion of a Curved Vortex Filament with Decaying Vortical Core and Axial Velocity*, SIAM J. Appl. Math., 35 (1978), pp. 148–175, https://doi.org/10.1137/0135013.

[22] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang, *Spectral Methods in Fluid Dynamics*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1988, https://doi.org/10.1007/978-3-642-84108-8.

[23] B. Cichocki and B. U. Felderhof, *Electrostatic interactions in two-dimensional Coulomb systems with periodic boundary conditions*, Physica A, 158 (1989), pp. 706–722, https://doi.org/10.1016/0378-4371(89)90487-1.

[24] G.-H. Cottet and P. D. Koumoutsakos, *Vortex Methods: Theory and Practice*, Cambridge University Press, Cambridge, 2000, https://doi.org/10.1017/CBO9780511526442.

[25] L. S. Da Rios, *Sul moto d'un liquido indefinito con un filetto vorticoso di forma qualunque*, Rend. Circ. Matem. Palermo, 22 (1906), pp. 117–135, https://doi.org/10.1007/BF03018608.

[26] T. Darden, D. York, and L. Pedersen, *Particle mesh Ewald: An N·log(N) method for Ewald sums in large systems*, J. Chem. Phys., 98 (1993), pp. 10089–10092, https://doi.org/10.1063/1.464397.

[27] C. de Boor, *A Practical Guide to Splines*, Applied Mathematical Sciences, Springer-Verlag, New York, 1978.

[28] F. De La Hoz, C. J. García-Cervera, and L. Vega, *A Numerical Study of the Self-Similar Solutions of the Schrödinger Map*, SIAM J. Appl. Math., 70 (2009), pp. 1047–1077, https://arxiv.org/abs/27862547.

[29] M. Deserno and C. Holm, *How to mesh up Ewald sums. I. A theoretical and numerical comparison of various particle mesh routines*, J. Chem. Phys., 109 (1998), pp. 7678–7693, https://doi.org/10.1063/1.477414.

[30] M. Deserno and C. Holm, *How to mesh up Ewald sums. II. An accurate error estimate for the particle–particle–particle-mesh algorithm*, J. Chem. Phys., 109 (1998), pp. 7694–7701, https://doi.org/10.1063/1.477415.

[31] A. Dutt and V. Rokhlin, *Fast Fourier Transforms for Nonequispaced Data*, SIAM J. Sci. Comput., 14 (1993), pp. 1368–1393, https://doi.org/10.1137/0914081.

[32] P. P. Ewald, *Die Berechnung optischer und elektrostatischer Gitterpotentiale*, Ann. Phys., 369 (1921), pp. 253–287, https://doi.org/10.1002/andp.19213690304.

[33] R. P. Feynman, R. B. Leighton, and M. L. Sands, *The Feynman Lectures on Physics. Volume 2: Mainly Electromagnetism and Matter*, Basic Books, New York, 2011.

[34] M. Frigo and S. G. Johnson, *The design and implementation of FFTW3*, Proc. IEEE, 93 (2005), pp. 216–231, https://doi.org/10.1109/JPROC.2004.840301.

[35] L. Galantucci, A. W. Baggaley, C. F. Barenghi, and G. Krstulovic, *A new self-consistent approach of quantum turbulence in superfluid helium*, Eur. Phys. J. Plus, 135 (2020), p. 547, https://doi.org/10.1140/epjp/s13360-020-00543-0.

[36] L. Gamet, F. Ducros, F. Nicoud, and T. Poinsot, *Compact finite difference schemes on non-uniform meshes. Application to direct numerical simulations of compressible flows*, Int. J. Numer. Methods Fluids, 29 (1999), pp. 159–191, https://doi.org/10.1002/(SICI)1097-0363(19990130)29:2⟨159::AID-FLD781⟩3.0.CO;2-9.

[37] M. Gazzola, B. Hejazialhosseini, and P. Koumoutsakos, *Reinforcement Learning and Wavelet Adapted Vortex Methods for Simulations of Self-propelled Swimmers*, SIAM J. Sci. Comput., 36 (2014), pp. B622–B639, https://doi.org/10.1137/130943078.

[38] L. Greengard and J.-Y. Lee, *Accelerating the Nonuniform Fast Fourier Transform*, SIAM Rev., 46 (2004), pp. 443–454, https://doi.org/10.1137/S003614450343200X.

[39] L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348, https://doi.org/10.1016/0021-9991(87)90140-9.

[40] F. Grotto and S. Morlacchi, *Decay of time correlations in point vortex systems*, Physica D, (2024), p. 134169, https://doi.org/10.1016/j.physd.2024.134169.

[41] R. Hänninen, *Dissipation enhancement from a single vortex reconnection in superfluid helium*, Phys. Rev. B, 88 (2013), p. 054511, https://doi.org/10.1103/PhysRevB.88.054511.

[42] R. Hänninen and A. W. Baggaley, *Vortex filament method as a tool for computational visualization of quantum turbulence*, Proc. Natl. Acad. Sci. USA, 111 (2014), pp. 4667–4674, https://doi.org/10.1073/pnas.1312535111.

[43] H. Hasimoto, *On the periodic fundamental solutions of the Stokes equations and their application to viscous flow past a cubic array of spheres*, J. Fluid Mech., 5 (1959), pp. 317–328, https://doi.org/10.1017/S0022112059000222.

[44] H. Hasimoto, *A soliton on a vortex filament*, J. Fluid Mech., 51 (1972), pp. 477–485, https://doi.org/10.1017/S0022112072002307.

[45] F. Hedman and A. Laaksonen, *Ewald summation based on nonuniform fast Fourier transform*, Chem. Phys. Lett., 425 (2006), pp. 142–147, https://doi.org/10.1016/j.cplett.2006.04.106.

[46] H. Helmholtz, *Über Integrale der hydrodynamischen Gleichungen, welche den Wirbelbewegungen entsprechen.*, J. Reine Angew. Math., 55 (1858), pp. 25–55.

[47] R. Hockney and J. Eastwood, *Computer Simulation Using Particles*, CRC Press, 1988, https://doi.org/10.1201/9780367806934.

[48] J. Jackson, C. Meyer, D. Nishimura, and A. Macovski, *Selection of a convolution function for Fourier inversion using gridding (computerised tomography application)*, IEEE Trans. Med. Imaging, 10 (Sept./1991), pp. 473–478, https://doi.org/10.1109/42.97598.

[49] D. Kivotides, C. F. Barenghi, and D. C. Samuels, *Triple Vortex Ring Structure in Superfluid Helium II*, Science, 290 (2000), pp. 777–779, https://doi.org/10.1126/science.290.5492.777.

[50] T. Knopp, M. Boberg, and M. Grosser, *NFFT.jl: Generic and Fast Julia Implementation of the Nonequidistant Fast Fourier Transform*, SIAM J. Sci. Comput., 45 (2023), pp. C179–C205, https://doi.org/10.1137/22M1510935.

[51] J. Kolafa and J. W. Perram, *Cutoff Errors in the Ewald Summation Formulae for Point Charge Systems*, Mol. Simul., 9 (1992), pp. 351–368, https://doi.org/10.1080/08927029208049126.

[52] P. Koumoutsakos, *Multiscale Flow Simulations Using Particles*, Annu. Rev. Fluid Mech., 37 (2005), pp. 457–487, https://doi.org/10.1146/annurev.fluid.37.061903.175753.

[53] V. S. Krishnamurthy and T. Sakajo, *The N-vortex problem in a doubly periodic rectangular domain with constant background vorticity*, Physica D, 448 (2023), p. 133728, https://doi.org/10.1016/j.physd.2023.133728.

[54] K. N. Kudin and G. E. Scuseria, *Revisiting infinite lattice sums with the periodic fast multipole method*, J. Chem. Phys., 121 (2004), pp. 2886–2890, https://doi.org/10.1063/1.1771634.

[55] J. Laurie and A. W. Baggaley, *Vorticity locking and pressure dynamics in finite-temperature superfluid turbulence*, Phys. Rev. Fluids, 8 (2023), p. 054604, https://doi.org/10.1103/PhysRevFluids.8.054604.

[56] X.-G. Lv and J. Le, *A note on solving nearly penta-diagonal linear systems*, Appl. Math. Comput., 204 (2008), pp. 707–712, https://doi.org/10.1016/j.amc.2008.07.012.

[57] R. I. McLachlan and G. R. W. Quispel, *Splitting methods*, Acta Numer., 11 (2002), pp. 341–434, https://doi.org/10.1017/S0962492902000053.

[58] H. K. Moffatt and R. L. Ricca, *Helicity and the Călugăreanu invariant*, Proc. R. Soc. Lond. Ser. Math. Phys. Sci., 439 (1992), pp. 411–429, https://doi.org/10.1098/rspa.1992.0159.

[59] D. W. Moore and P. G. Saffman, *The motion of a vortex filament with axial flow*, Phil. Trans. R. Soc. Lond. A, 272 (1972), pp. 403–429, https://doi.org/10.1098/rsta.1972.0055.

[60] N. P. Müller, J. I. Polanco, and G. Krstulovic, *Intermittency of Velocity Circulation in Quantum Turbulence*, Phys. Rev. X, 11 (2021), p. 011053, https://doi.org/10.1103/PhysRevX.11.011053.

[61] P. K. Newton and G. Chamoun, *Vortex Lattice Theory: A Particle Interaction Perspective*,

SIAM Rev., 51 (2009), pp. 501–542, https://doi.org/10.1137/07068597x.
[62] M. Pippig and D. Potts, *Parallel Three-Dimensional Nonequispaced Fast Fourier Transforms and Their Application to Particle Simulation*, SIAM J. Sci. Comput., 35 (2013), pp. C411–C437, https://doi.org/10.1137/120888478.
[63] J. I. Polanco, N. P. Müller, and G. Krstulovic, *Vortex clustering, polarisation and circulation intermittency in classical and quantum turbulence*, Nat. Commun., 12 (2021), p. 7090, https://doi.org/10.1038/s41467-021-27382-6.
[64] D. Potts and G. Steidl, *Fast Summation at Nonequispaced Knots by NFFTs*, SIAM J. Sci. Comput., 24 (2003), pp. 2013–2037, https://doi.org/10.1137/S1064827502400984.
[65] P. H. Roberts and R. J. Donnelly, *Dynamics of vortex rings*, Phys. Lett. A, 31 (1970), pp. 137–138, https://doi.org/10.1016/0375-9601(70)90193-3.
[66] P. G. Saffman, *Vortex Dynamics*, Cambridge University Press, Jan. 1993, https://doi.org/10.1017/cbo9780511624063.
[67] D. C. Samuels, *Velocity matching and Poiseuille pipe flow of superfluid helium*, Phys. Rev. B, 46 (1992), pp. 11714–11724, https://doi.org/10.1103/PhysRevB.46.11714.
[68] D. C. Samuels, *Vortex Filament Methods for Superfluids*, in Quantized Vortex Dynamics and Superfluid Turbulence, C. F. Barenghi, R. J. Donnelly, and W. F. Vinen, eds., Lecture Notes in Physics, Springer, Berlin, Heidelberg, 2001, pp. 97–113, https://doi.org/10.1007/3-540-45542-6_9.
[69] A. Sandu, *A Class of Multirate Infinitesimal GARK Methods*, SIAM J. Numer. Anal., 57 (2019), pp. 2300–2327, https://doi.org/10.1137/18m1205492.
[70] K. W. Schwarz, *Three-dimensional vortex dynamics in superfluid ^{4}He: Line-line and line-boundary interactions*, Phys. Rev. B, 31 (1985), pp. 5782–5804, https://doi.org/10.1103/PhysRevB.31.5782.
[71] D. S. Shamshirgar, J. Bagge, and A.-K. Tornberg, *Fast Ewald summation for electrostatic potentials with arbitrary periodicity*, J. Chem. Phys., 154 (2021), p. 164109, https://doi.org/10.1063/5.0044895.
[72] K. Shariff and A. Leonard, *Vortex Rings*, Annu. Rev. Fluid Mech., 24 (1992), pp. 235–279, https://doi.org/10.1146/annurev.fl.24.010192.001315.
[73] Y.-h. Shih, G. Wright, J. Andén, J. Blaschke, and A. H. Barnett, *cuFINUFFT: A load-balanced GPU library for general-purpose nonuniform FFTs*, in 2021 IEEE Int. Parallel Distrib. Process. Symp. Workshop IPDPSW, June 2021, pp. 688–697, https://doi.org/10.1109/IPDPSW52791.2021.00105.
[74] H. Takahasi and M. Mori, *Double Exponential Formulas for Numerical Integration*, Publ. Res. Inst. Math. Sci., 9 (1973), pp. 721–741, https://doi.org/10.2977/prims/1195192451.
[75] W. Thomson, *3. Vibrations of a Columnar Vortex*, Proc. R. Soc. Edinb., 10 (1880/ed), pp. 443–456, https://doi.org/10.1017/S0370164600044151.
[76] L. N. Trefethen and J. A. C. Weideman, *The Exponentially Convergent Trapezoidal Rule*, SIAM Rev., 56 (2014), pp. 385–458, https://doi.org/10.1137/130932132.
[77] A. van Kan, A. Alexakis, and M.-E. Brachet, *Intermittency of three-dimensional perturbations in a point-vortex model*, Phys. Rev. E, 103 (2021), p. 053102, https://doi.org/10.1103/PhysRevE.103.053102.
[78] L. J. Vermeer, J. N. Sørensen, and A. Crespo, *Wind turbine wake aerodynamics*, Prog. Aerosp. Sci., 39 (2003), pp. 467–510, https://doi.org/10.1016/S0376-0421(03)00078-2.
[79] D. H. Wacks, A. W. Baggaley, and C. F. Barenghi, *Coherent laminar and turbulent motion of toroidal vortex bundles*, Phys. Fluids, 26 (2014), p. 027102, https://doi.org/10.1063/1.4864659.
[80] J. B. Weiss and J. C. McWilliams, *Nonergodicity of point vortices*, Phys. Fluids, 3 (1991), pp. 835–844, https://doi.org/10.1063/1.858014.
[81] M. Yarrow, *Solving periodic block tridiagonal systems using the Sherman-Morrison-Woodbury formula*, in 9th Comput. Fluid Dyn. Conf., Buffalo, NY, U.S.A., June 1989, American Institute of Aeronautics and Astronautics, https://doi.org/10.2514/6.1989-1946.
[82] S. Yui, H. Kobayashi, M. Tsubota, and R. Yokota, *Quantum turbulence coupled with externally driven normal-fluid turbulence in superfluid ^{4}He*, May 2021, https://doi.org/10.48550/arXiv.2105.09499, https://arxiv.org/abs/2105.09499.
[83] S. Yui, M. Tsubota, and H. Kobayashi, *Three-Dimensional Coupled Dynamics of the Two-Fluid Model in Superfluid ^{4}He: Deformed Velocity Profile of Normal Fluid in Thermal Counterflow*, Phys. Rev. Lett., 120 (2018), p. 155301, https://doi.org/10.1103/PhysRevLett.120.155301.