

A fast algorithm for simulation of periodic flows using discrete vortex particles

Túlio R. Ricciardi¹ · William R. Wolf¹  · Alex M. Bimbato²

Received: 12 May 2017 / Accepted: 28 August 2017 / Published online: 6 September 2017
© The Brazilian Society of Mechanical Sciences and Engineering 2017

Abstract We present a novel fast algorithm for flow simulations using the discrete vortex method, DVM, for problems with periodic boundary conditions. In the DVM, the solution of the velocity field induced by interactions among N discrete vortex particles is governed by the Biot–Savart law and, therefore, leads to a computational cost proportional to $O(N^2)$. The proposed algorithm combines exponential and power series expansions implemented using a divide and conquer strategy to accelerate the calculation of the cotangent kernel that models periodic boundary conditions. The fast multipole method, FMM, is applied for the solution of singular terms appearing in the power series expansion and also for the exponential series expansion. Error and computational cost analyses are performed for the individual steps of the algorithm for double and quadruple machine precision. The current method presents more accurate solutions when compared to those obtained by periodic domain replication using the free-field FMM kernel. The novel algorithm provides computational savings of nearly 240 times for double-precision simulations with one million particles when compared to the direct calculation of the Biot–Savart law.

Keywords Fast algorithm · Cotangent kernel · Periodic boundary conditions · Discrete vortex method · Fast multipole method

1 Introduction

In fluid mechanics, several problems can be modeled using periodic boundary conditions and a number of applications can be found in mechanical, civil, aeronautical and naval engineering. For example, one can cite numerical simulations involving periodic homogeneous directions such as in isotropic turbulence [25, 27], turbulent channel flows [21] or the turbulent flow past an airfoil [24]. Other problems may present more complex periodic flow arrangements such as in a turbine cascade configuration [5] or in a heat exchanger [28]

In flows with regions of concentrated vorticity, mesh-free numerical methods such as the discrete vortex method, DVM, have become a suitable tool for simulation [4, 6]. Alternatively to the classical Eulerian methods, the DVM applies a Lagrangian technique where vorticity is advected with the local flow velocity. Recently, the DVM has been applied to solve problems including airfoil leading edge separation [12, 19], unsteady motion of a pitching airfoil [18] and cylinder wake instability [8].

Vortex cloud modeling offers great potential for numerical analysis of relevant problems in fluid mechanics. Some of the main advantages of the Lagrangian vortex methods are: (1) the absence of mesh generation which can be cumbersome for complex configurations; (2) computational efforts are directed only to regions with non-zero vorticity, differently from Eulerian formulations which require solutions along the entire discretized flowfield; (3) far-field boundary conditions are taken care automatically

Technical Editor: André Cavalieri.

✉ William R. Wolf
wolf@fem.unicamp.br

Túlio R. Ricciardi
tulioricci@fem.unicamp.br

Alex M. Bimbato
alexbimbato@feg.unesp.br

¹ University of Campinas, Campinas, SP 13083-860, Brazil

² São Paulo State University, Guaratinguetá, SP 12516-410, Brazil

by the method and (4) vortex methods do not suffer from numerical dispersion and dissipation errors.

Despite its advantages, in mesh-free vortex methods, the number of operations at each time step is proportional to the square of the number of vortices present in the flow since it is necessary to compute the interactions among all N discrete vortex particles using the Biot–Savart law. This leads to large CPU requirements, proportional to N^2 . To circumvent this difficulty, faster algorithms are required to accelerate the vortex–vortex interactions computed by the Biot–Savart law. Lindsay and Krasny [15] presented a divide and conquer methodology with an adaptive local refinement to accelerate the solution of vortex sheet motion. Yokota et al. [25–27] employed the fast multipole method [10, 17], FMM, to perform fast simulations of isotropic turbulence and homogeneous shear flows using vortex methods.

Mesh-free methods can be applied to solve problems with periodic boundary conditions. For instance, several authors have employed the DVM to study the development of vortex sheet roll-up [1, 3, 13, 14, 22, 23] which requires the application of a cotangent kernel to represent the periodicity of the vortex particles. These authors studied several aspects of the problem including the regularization of the kernel and the development of instabilities in the solution. Fast algorithms such as the FMM have been applied in combination with vortex methods for the solution of flows with periodic boundary conditions [7, 20, 25–27]. In general, periodicity is implemented in the FMM through replication of the multipole expansions of the computational domain. Ricciardi [20] presents an error analysis of such implementation for a vortex sheet roll-up problem. Other authors have also applied fast algorithms together with Lagrangian methods to solve problems in a different context, such as in Stokes flow simulations with periodic boundary conditions [9, 16].

In this work, we present a novel fast algorithm for flow simulations using the discrete vortex method for problems with periodic boundary conditions. The proposed algorithm combines exponential and power series expansions implemented using a divide and conquer strategy to accelerate the calculation of the cotangent kernel that models periodic boundary conditions. The fast multipole method is applied for the solution of singular terms appearing in the power series expansion and also for the exponential series expansion. Error and computational cost analyses are performed for the individual steps of the algorithm for double and quadruple machine precision. The current method presents more accurate solutions when compared to those obtained by periodic domain replication using the free-field FMM kernel. The novel algorithm provides computational savings of nearly 240 times for double-precision

simulations with one million particles when compared to the direct calculation of the Biot–Savart law.

2 Numerical formulation

The discrete vortex method solves the Navier–Stokes equations in the vorticity form. For an incompressible Newtonian fluid, the 2D vorticity transport equation in non-dimensional variables is given by

$$\frac{\partial \omega}{\partial t} + \mathbf{u} \cdot \nabla \omega = \frac{1}{Re} \nabla^2 \omega, \quad (1)$$

where $\mathbf{u} = u \hat{\mathbf{i}} + v \hat{\mathbf{j}}$ is the velocity vector, $\omega = \nabla \times \mathbf{u}$ is the z-vorticity component and Re is the Reynolds number. Chorin [6] proposed a numerical solution of Eq. 1 solving separately the inviscid and viscous terms in two fractional steps. The first step considers the flow to be inviscid, i.e., $Re \rightarrow \infty$, leading to

$$\frac{\partial \omega}{\partial t} + \mathbf{u} \cdot \nabla \omega = \frac{D \omega}{D t} = 0, \quad (2)$$

while the second step solves the viscous effects

$$\frac{\partial \omega}{\partial t} = \frac{1}{Re} \nabla^2 \omega. \quad (3)$$

The calculation of the inviscid term requires the solution of the Biot–Savart law, which computes the velocity field due to the vortex interactions. For a two-dimensional problem, the Biot–Savart law is written as

$$\mathbf{u}(\mathbf{x}, t) = \frac{1}{2\pi} \int_{S'} \frac{\boldsymbol{\omega}(\mathbf{x}', t) \times (\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|^2} dS', \quad (4)$$

where $\boldsymbol{\omega}(\mathbf{x}', t)$ is the vorticity and S' represents a two-dimensional surface which surrounds the vorticity field.

In complex notation, the velocity induced by a potential vortex j at a point k is given by

$$w_j = u_j - i v_j = i \frac{\Gamma_k}{2\pi z_{jk}}. \quad (5)$$

In this equation, Γ_k is the circulation of vortex k defined as positive in the clockwise direction. The complex distance between vortices j and k is given by $z_{jk} = (x_j - x_k) + i (y_j - y_k)$. Here, u and v are the Cartesian components of the velocity vector in the x and y directions, respectively. If the vortex j is replicated n times, from $-\infty$ to ∞ , with spatial periodicity λ , it is possible to calculate the induced velocity of this array at z_j according to

$$w_j = \sum_{n=-\infty}^{\infty} i \frac{\Gamma_k}{2\pi} \left(\frac{1}{z_{jk} + n\lambda} \right). \quad (6)$$

For $\lambda = 1$, this equation can be simplified to

$$w_j = i \frac{\Gamma_k}{2\pi} \left[\frac{1}{z_{jk}} + \sum_{n=1}^{\infty} \left(\frac{1}{z_{jk} - n} + \frac{1}{z_{jk} + n} \right) \right], \tag{7}$$

where the summation inside brackets has an exact solution given by the cotangent function [2]

$$\pi \cot(\pi z_{jk}) = \frac{1}{z_{jk}} + \sum_{n=1}^{\infty} \left(\frac{1}{z_{jk} - n} + \frac{1}{z_{jk} + n} \right). \tag{8}$$

This leads to an exact solution for the induced complex velocity by an infinite array of periodical vortex particles k at a point z_j . Therefore, it is possible to rewrite Eq. 7 as

$$w_j = i \frac{\Gamma_k}{2\pi} [\pi \cot(\pi z_{jk})]. \tag{9}$$

As previously discussed, the direct calculation of the induced velocity by Eq. 9 leads to an algorithm with computational cost proportional to $O(N^2)$, being N the number of vortex particles in the discretization. To accelerate the calculation of the Biot–Savart law, we propose the following approximations which will be combined in a fast algorithm.

2.1 Exponential series expansion

The first approximation investigated is an exponential series expansion of the cotangent function given by

$$\cot(\pi z) = +i - 2i \sum_{n=0}^{\infty} \exp(2\pi i n z), \tag{10}$$

where the argument z is given by $z = x + iy$. If one truncates the series with p -terms, it is possible to write an approximation as

$$\cot(\pi z) \approx +i - 2i \sum_{n=0}^p [\exp(2\pi i n x) \exp(-2\pi n y)]. \tag{11}$$

Hence, two distinct exponential terms are present. The one dependent on the distance along the x -axis is imaginary and purely oscillatory, so it does not contribute to the convergence. On the other hand, the term based on the separation along the y -axis is real and decays if $y > 0$ and, since $|\exp(-2\pi y)|$ is smaller than unit, the series is convergent. For $y < 0$, the series diverges since the exponential grows.

However, as cotangent is an odd function, it is possible to achieve convergence for both $|y| > 0$ based on the identity $\cot(z) = -\cot(-z)$. Therefore, it is possible to write

$$\cot(\pi z) \approx +i - 2i \sum_{n=0}^p [\exp(+2\pi i n x) \exp(-2\pi n y)],$$

for $y > 0$,

$$\tag{12}$$

and

$$\cot(\pi z) \approx -i + 2i \sum_{n=0}^p [\exp(-2\pi i n x) \exp(+2\pi n y)],$$

for $y < 0$.

$$\tag{13}$$

The series convergence is investigated to determine the number of p terms required to achieve an arbitrary precision E . To do so, one must solve

$$\left| \cot(\pi + i \pi y)_{x=1} - \left\{ i - 2i \sum_{n=0}^p [\exp(2\pi i n)_{x=1} \exp(-2\pi n y)] \right\} \right| \leq E, \tag{14}$$

which is valid for $y > 0$. Results are indicated by the different curves in Fig. 1, ranging from $E = 1 \times 10^{-6}$ to 1×10^{-30} . They are obtained for an argument where the imaginary part y ranges from 0 to 1, while the real part is fixed at $x = 1$. It is possible to see faster convergence for larger $|y|$, since fewer terms are required in the series truncation. Recently, Ricciardi et al. [20] have shown that simulations of vortical flows using the DVM may have numerical instabilities in the inviscid limit due to numerical truncation errors that destabilize the solutions. In this sense, they show that increasing the machine precision to quadruple precision may delay the formation of spurious solutions in the DVM. For the current formulation, the role of machine precision is presented in Figs. 1a, b, for double and quadruple precision, respectively.

2.2 Power series expansion

The second approximation is given by a power series expansion that approximates the cotangent function for an argument $|z| < 1$. After a truncation with p terms, this approximation is given by

$$\cot(\pi z) \approx \sum_{n=0}^p \frac{(-1)^n 2^{2n} B_{2n}}{(2n)!} (\pi z)^{2n-1}, \tag{15}$$

where the Bernoulli numbers B_{2n} are defined as

$$B_{2n} = (-1)^{n+1} \frac{(2n)! 2}{(2\pi)^{2n}} \zeta(2n). \tag{16}$$

In the above equation, the Riemann zeta function ζ is computed for $n \geq 1$ as $\zeta(2n) = \sum_{j=1}^{\infty} j^{-2n}$ and, for $n = 0$, it is given by $\zeta(0) = -\frac{1}{2}$.

Combining the equations above, it is possible to obtain a simplified series expansion for the cotangent function as

$$\cot(\pi z) \approx \sum_{n=0}^p \frac{(-2)^n \zeta(2n)}{\pi^{2n}} (\pi z)^{2n-1}, \tag{17}$$

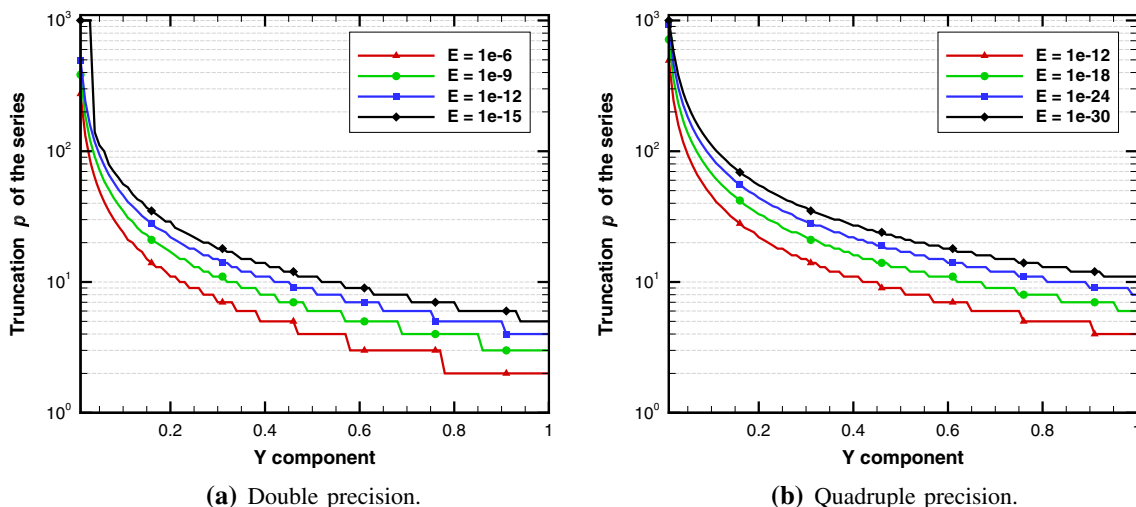


Fig. 1 Convergence properties of the exponential series expansion of the cotangent function

which leads to the cotangent power series equation

$$\cot(\pi z) = \frac{1}{\pi z} - \frac{1}{3}(\pi z) - \frac{1}{45}(\pi z)^3 - \frac{2}{945}(\pi z)^5 + \dots + O[(\pi z)^{2p+1}]. \tag{18}$$

It is possible to simplify Eq. 17, leading to

$$\cot(\pi z) \approx \sum_{n=0}^p C_{2n-1} z^{2n-1}, \tag{19}$$

with coefficients C_{2n-1} that can be precomputed and stored for faster calculations

$$C_{2n-1} = \frac{(-2)^n \zeta(2n)}{\pi}. \tag{20}$$

The convergence of Eq. 19 is investigated for $|z| < 1$, where $z = x + iy$. In these computations, the imaginary part is null, i.e., $y = 0$, but non-zero values for y can be used as well. Hence, only the influence of the separation of two particles along the x -axis is investigated, such that $z = x$. The absolute error from the series expansion to the exact solution can be evaluated as

$$\left| \cot(\pi z) - \sum_{n=0}^p C_{2n-1} z^{2n-1} \right| = E. \tag{21}$$

Figure 2 shows that for small distances the cotangent behaves like $1/z$, and few terms are required for good convergence of the series. When the distance increases, $z \rightarrow 1$, the method lacks precision and the error grows, as

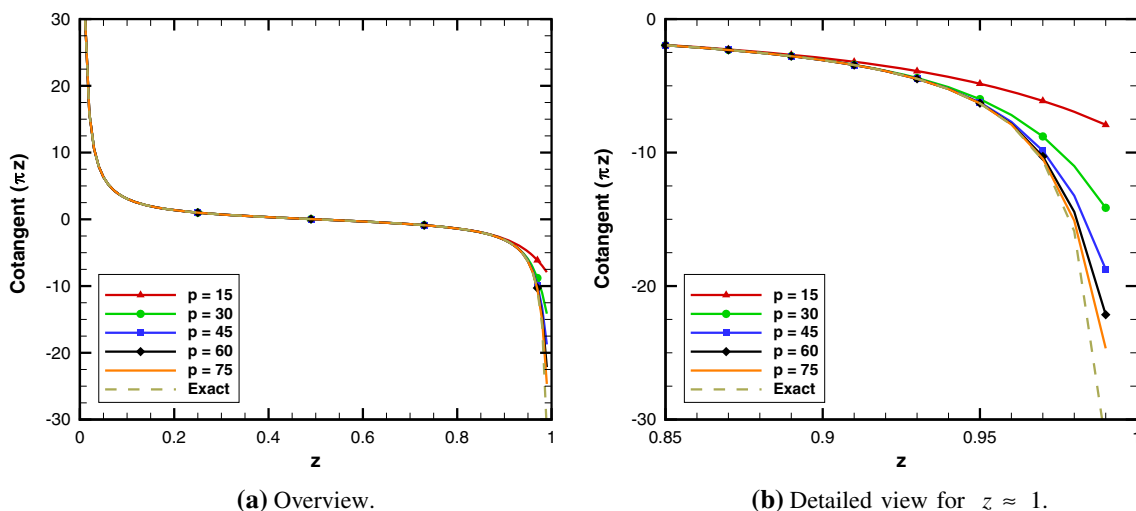


Fig. 2 Absolute value of the cotangent power series expansion based on the argument z for several truncation terms p evaluated with double precision

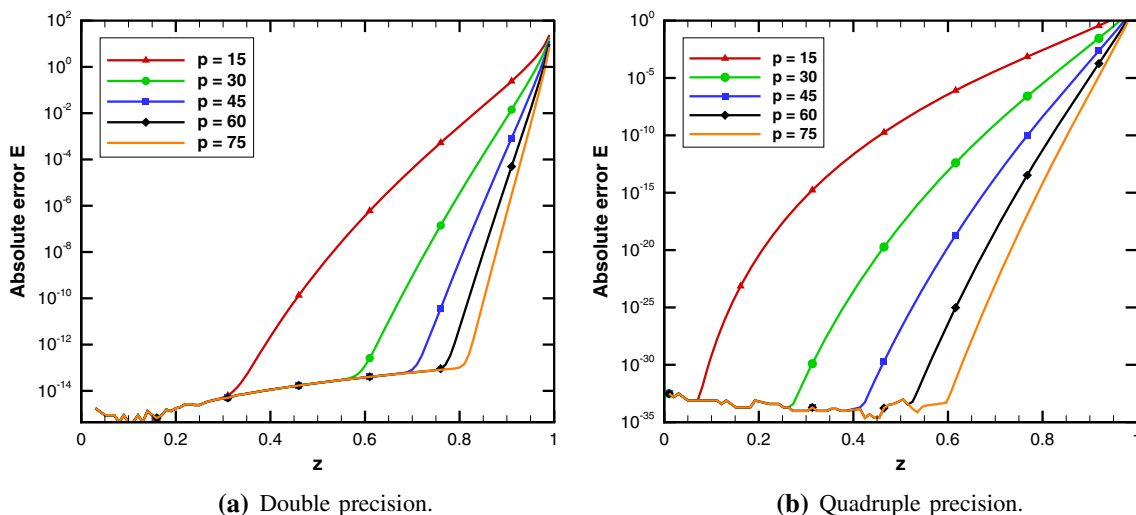


Fig. 3 Error of the cotangent power series expansion based on the argument z for several truncation terms

presented in Fig. 3 for a different number of p terms in the series. Finally, if the argument of the cotangent function is larger than π ($|z| > 1$), the present power series expansion is not applicable, i.e., this formulation should not be used for vortex particles far from each other.

However, since the cotangent function is a periodic function with period π , it may be possible to rewrite it using integer numbers n ($n \in \mathbb{Z}$) as $\cot(z) = \cot(z + n\pi)$. Hence, when the distance of two particles along the real axis is large, it is possible to account for their interaction indirectly using its image-vortex, leading to the same result. This idea is shown in Fig. 4 where, if $|x| > 0.5$, the interaction is performed not with the original vortex particle, but with its periodic image. Hence, the convergence is faster everywhere inside an unitary domain.

3 Fast summation algorithm

To reduce the computational cost from the direct calculation of the Biot–Savart law from $O(N^2)$ towards $O(N)$, a divide and conquer strategy will be used. Therefore, clusters of vortex particles should interact with each other instead of the direct particles.

As presented in Figs. 2, 3 and 4, particles close to each other are evaluated using the power series expansion approach. As will be later shown, while the singular term $1/z$ present in this series must be handled by the FMM, all the non-singular terms can be evaluated using a Newton binomial expansion. Hence, these two methodologies are coupled to solve the approximation through the power series expansion. The methodology of the FMM used here

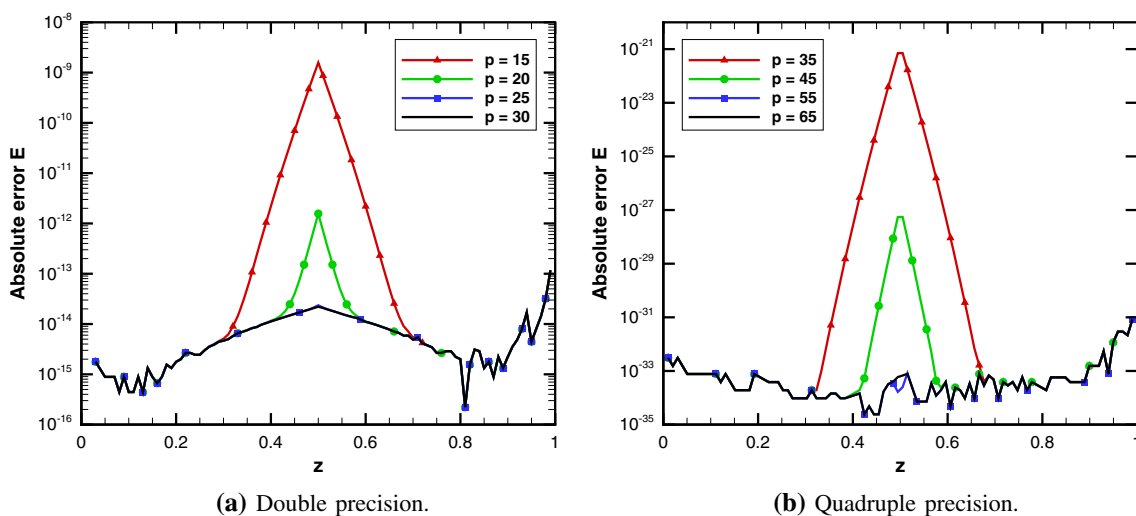


Fig. 4 Convergence of the cotangent power series expansion using periodic images of the vortex particles for $|x| > 0.5$

is the standard one presented in the literature and the authors refer to the following references for details [10, 17, 20]. For particles with large separation along the y -axis, the exponential series expansion is employed.

The present fast algorithm follows the ideas of the fast multipole method where a computational box involves all the discrete vortex particles. This computational domain is recursively divided to create smaller boxes that separate the particles, defining clusters of vortices. To provide fast convergence everywhere, the different series expansions discussed in the previous sections are employed depending on the relative position among the boxes.

3.1 Fast summation of the exponential series

Given a cluster b with N_b vortex particles with intensity Γ_k placed at z_k , if one accounts for the particles and all their periodic images, the velocity field at a point z_j is given by

$$\begin{aligned} & \sum_{k=1}^{N_b} \Gamma_k \cot(\pi z_{jk}) \\ &= \sum_{k=1}^{N_b} \Gamma_k \left\{ +i - 2i \sum_{n=0}^{\infty} \exp[2\pi i n(+z_j - z_k)] \right\}, \\ & \text{for } y_j - y_k > 0, \end{aligned} \quad (22)$$

or

$$\begin{aligned} & \sum_{k=1}^{N_b} \Gamma_k \cot(\pi z_{jk}) \\ &= \sum_{k=1}^{N_b} \Gamma_k \left\{ -i + 2i \sum_{n=0}^{\infty} \exp[2\pi i n(-z_j + z_k)] \right\}, \\ & \text{for } y_j - y_k < 0. \end{aligned} \quad (23)$$

One must recall the requirement for convergence such that the imaginary part y is positive. If it is negative, the series given by Eq. 22 diverges. Hence, one must use the odd-function property of the cotangent, as previously explained, to work with negative values of the imaginary part using Eq. 23. The deduction below is presented only for the case of positive imaginary part, but the modifications to work with negative values are straightforward.

To perform fast summations, it is necessary to split the argument of the cotangent function in a product of two functions. This is straightforward for an exponential function since

$$\exp(z_j - z_k) = \exp(+z_j) \exp(-z_k). \quad (24)$$

Furthermore, the *middleman* strategy of using the box centers to form equivalent clusters of particles can be

employed. Defining z_{c_o} and z_{c_s} as the center of observer and source clusters, respectively, one can evaluate

$$\begin{aligned} \hat{z} &= z_{c_s} - z_k, \\ \tilde{z} &= z_{c_o} - z_{c_s}, \\ \bar{z} &= z_j - z_{c_o}. \end{aligned} \quad (25)$$

The terms above represent the distance from source cluster to the source particles, the separation among clusters and the distance to observer particles from the center of their clusters, respectively. From the definitions provided in Eq. 25, one may recover Eq. 24 as

$$\exp[(z_{c_s} - z_k) + (z_{c_o} - z_{c_s}) + (z_j - z_{c_o})] = \exp[z_j - z_k]. \quad (26)$$

A similar technique of particle-to-multipole (P2M), multipole-to-local expansion (M2L), and local-to-particle expansion (L2P) steps from the free-domain FMM can be employed [10, 17]. There is no need to use M2M and L2L steps since no multi-level operations are required.

From Eqs. 25–26, it is possible to rewrite Eq. 22 as

$$\begin{aligned} & \sum_{k=1}^{N_b} \Gamma_k \cot(\pi z_{jk}) = \sum_{k=1}^{N_b} i \Gamma_k + \\ & \sum_{k=1}^{N_b} \left\{ -2i \Gamma_k \sum_{n=0}^p [\exp(2\pi i n \hat{z}) \exp(2\pi i n \tilde{z}) \exp(2\pi i n \bar{z})] \right\}, \end{aligned} \quad (27)$$

where, after manipulation of the order of the summations one has

$$\begin{aligned} & \sum_{k=1}^{N_b} i \Gamma_k + \\ & \sum_{n=0}^p \left\{ -2i \exp(2\pi i n \tilde{z}) \exp(2\pi i n \bar{z}) \sum_{k=1}^{N_b} [\Gamma_k \exp(2\pi i n \hat{z})] \right\}. \end{aligned}$$

Therefore, a multipole-type expansion dependent only on the source-particles and defined for $n = 0$ to p can be obtained using P2M operations as

$$\mathcal{M}_n^{(b)} = \sum_{k=1}^{N_b} [\Gamma_k \exp(2\pi i n \hat{z})]. \quad (28)$$

Hence, one can write the right-hand side of Eq. 27 as

$$\sum_{k=1}^{N_b} i \Gamma_k + \sum_{n=0}^p \left\{ -2i \exp(2\pi i n \tilde{z}) \exp(2\pi i n \bar{z}) \mathcal{M}_n^{(b)} \right\},$$

where it is possible to evaluate local representations for a box b' of far away clusters from the b boxes using M2L operations given by

$$\mathcal{L}_n^{(b')} = -2 \exp(2\pi i n \tilde{z}) \mathcal{M}_n^{(b)}. \quad (29)$$

Finally, for $n = 0$ to p , the use of L2P operations accounts for the induced velocity from all N_b source particles

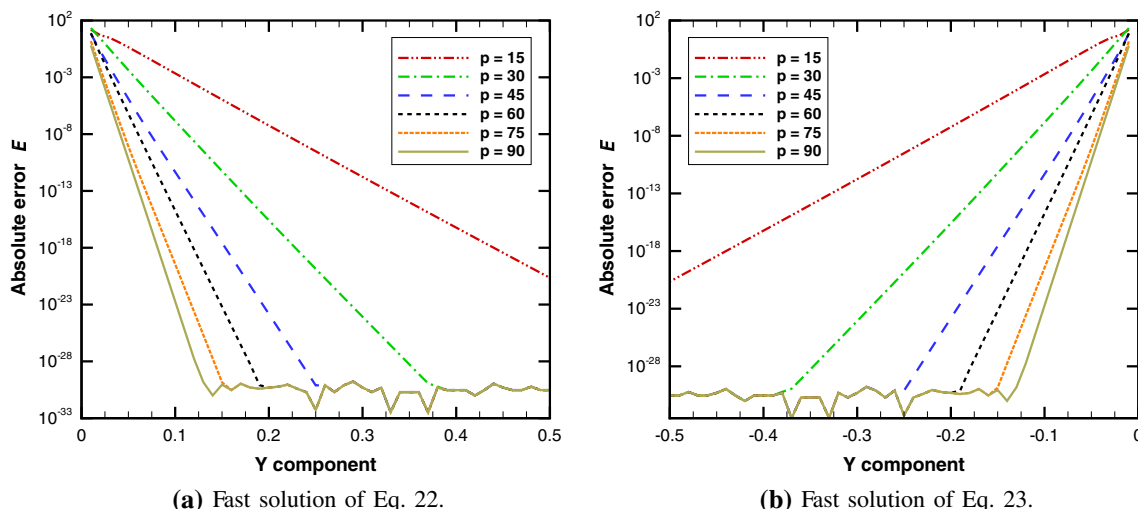


Fig. 5 Error of the fast exponential series summation for different truncation terms

located in the far-field of certain box b' at an observer placed at z_j inside this box. This is written as

$$\sum_{k=1}^{N_b} \Gamma_k \cot(\pi z_{jk}) = i \sum_{k=1}^{N_b} \Gamma_k + i \sum_{n=0}^p \left\{ \mathcal{L}_n^{(b')} \exp(2\pi i n \bar{z}) \right\}. \tag{30}$$

Given the truncation of the series in p terms, the computational cost of P2M and L2P operations, Eqs. 28 and 30, respectively, is $O(pN)$. Hence, they are linearly dependent on the number of particles. Furthermore, the method has a fixed cost which depends only on the box-to-box operations given by Eq. 29. It is proportional to the desired precision as well as the number of non-empty boxes at a level ℓ , given by $O(4^\ell)$. Hence, the exponential series summation has an overall cost of $O(2pN + 4^\ell p)$.

A quick test to validate this algorithm is performed with the computation of the absolute error E , based on the vertical distance for a far-field observer and series truncation using p terms. To do so, one thousand potential vortex particles are placed along the x -axis, while the observer is positioned along the y -axis from $y = -0.5$ to 0.5 . The solution obtained with the algorithm is then compared to the exact solution evaluated by the Biot–Savart law and the deviation is shown in Fig. 5. In the figure, each line represents a different number of truncation terms in the series, showing a convergence for quadruple machine precision.

3.2 Fast summation of the power series

The fast summation of the series given by Eq. 18 was first applied by Gumerov and Duraiswami [11] to solve Fourier transforms. Here, it is applied to solve problems involving vortex dynamics using the discrete vortex method. The

local velocity field at z_j induced by the N_b particles at z_k inside a box b is written as

$$w_j = i \sum_{k=1}^{N_b} \frac{\Gamma_k}{2\pi} \pi \cot(\pi z_{jk}), \tag{31}$$

where $z_{jk} = z_j - z_k$. Two vortex particles can be far from each other in the central domain, but they can be close to the other’s image. Hence, to have better convergence, the nearest distance among two vortex particles, following the fact that $\cot(z) = \cot(z + n\pi)$, is to compute z_{jk} as

$$z_{jk} = \begin{cases} z_{jk} - 1, & \text{if } |z_{jk} - 1| = \min(|z_{jk} - 1|, |z_{jk}|, |z_{jk} + 1|) \\ z_{jk}, & \text{if } |z_{jk}| = \min(|z_{jk} - 1|, |z_{jk}|, |z_{jk} + 1|) \\ z_{jk} + 1, & \text{if } |z_{jk} + 1| = \min(|z_{jk} - 1|, |z_{jk}|, |z_{jk} + 1|). \end{cases} \tag{32}$$

Using Eq. 18 as an approximation for the velocity field leads to

$$w_j = \sum_{k=1}^{N_b} \left\{ \left(\frac{\Gamma_k}{2\pi} \right) \pi \left[\frac{1}{\pi z_{jk}} - \frac{1}{3} (\pi z_{jk}) - \frac{1}{45} (\pi z_{jk})^3 - \frac{2}{945} (\pi z_{jk})^5 + \dots \right] \right\}, \tag{33}$$

where it is possible to obtain two main terms in the induced velocity as

$$w_j = \sum_{k=1}^{N_b} \left\{ \frac{\Gamma_k}{2\pi z_{jk}} \right\} + \sum_{k=1}^{N_b} \left\{ \frac{\Gamma_k}{2\pi} \pi \left[-\frac{1}{3} (\pi z_{jk}) - \frac{1}{45} (\pi z_{jk})^3 - \frac{2}{945} (\pi z_{jk})^5 + \dots \right] \right\}. \tag{34}$$

The first series involves the singular terms and they have exactly the same kernel used in the free-domain FMM,

$$\sum_{k=1}^{N_b} i \left\{ \frac{\Gamma_k}{2\pi z_{jk}} \right\}.$$

One should remind that it is possible to avoid the singular values of the cotangent function ($z_{jk} = n\pi$, for $n \in \mathbb{Z}$) using the Lamb–Oseen vortex with a viscous core σ [20] as

$$w_j = i \frac{\Gamma_k}{2\pi} \left\{ \pi \cot(\pi z) - \sum_{n=1}^1 \left(\frac{1}{z_{jk} + n} \right) \exp \left[-\frac{(x_{jk} + n)^2 + (y_{jk})^2}{\sigma^2} \right] \right\}. \tag{35}$$

The second series contains the regular terms written as

$$\sum_{k=1}^{N_b} \left\{ \frac{\Gamma_k}{2\pi} \pi \left[-\frac{1}{3} (\pi z_{jk}) - \frac{1}{45} (\pi z_{jk})^3 - \frac{2}{945} (\pi z_{jk})^5 + \dots \right] \right\},$$

with the coefficients shown in Eq. 20. After truncating the series using p terms, it is possible to use a simplified notation for this series given by

$$\sum_{k=1}^{N_b} \left\{ \frac{\Gamma_k}{2} \sum_{n=1}^p \left(C_{2n-1} z_{jk}^{2n-1} \right) \right\}.$$

The term z_{jk}^{2n-1} , for $z_{jk} = z_j - z_k$, can be solved exactly using a Newton binomial expansion

$$z_{jk}^{2n-1} = \sum_{m=0}^{2n-1} \binom{2n-1}{m} z_j^{2n-1-m} (-z_k)^m. \tag{36}$$

Considering a complex distance z_{jk} and swapping the summation order, it is possible to use a binomial expansion and the series becomes

$$\frac{1}{2} \sum_{n=1}^p \left\{ C_{2n-1} \sum_{m=0}^{2n-1} \left[\binom{2n-1}{m} z_j^{2n-1-m} \sum_{k=1}^{N_b} (\Gamma_k (-z_k)^m) \right] \right\}.$$

Since the inner summation depends only on the source particle k , a multipole cluster $\mathcal{B}_m^{(b)}$ can be defined for a box b as

$$\mathcal{B}_m^{(b)} = \sum_{k=1}^{N_b} (\Gamma_k (-z_k)^m). \tag{37}$$

Furthermore, $\mathcal{B}_m^{(b)}$ can be precomputed to avoid a convolution in the vortex particles. For m from 0 to $2p - 1$, its computational cost is $O(2pN)$. From this new variable, it is possible to write the power series as

$$\frac{1}{2} \sum_{n=1}^p \left\{ C_{2n-1} \sum_{m=0}^{2n-1} \left[\binom{2n-1}{m} z_j^{2n-1-m} \mathcal{B}_m^{(b)} \right] \right\}.$$

The cotangent series, Eq. 31, is then written in its final form as

$$w_j = \sum_{k=1}^{N_b} \left\{ \frac{\Gamma_k}{2\pi z_{jk}} \right\} + \frac{1}{2} \sum_{n=1}^p \left\{ C_{2n-1} \sum_{m=0}^{2n-1} \left[\binom{2n-1}{m} z_j^{2n-1-m} \mathcal{B}_m^{(b)} \right] \right\}, \tag{38}$$

where the free-domain FMM is employed to solve the singular term.

The most important time reduction in fast summation methods is due to the cluster–cluster operations if the number of clusters is smaller than the number of particles. The function which governs the interaction among particles can be manipulated without loss of precision using the centroids of the clusters as *middleman* for the interactions. This is shown in Eqs. 24–26 for the exponential series expansion.

For the power-series expansion, this strategy is not applicable since one does not have explicit functions for the centroid of the source and observer clusters. Hence, in this binomial series, the local expansions are performed via cluster–particle, where the number of particles is now relevant. In other words, at level ℓ , there are $O(4^\ell)$ non-empty boxes which interact with N particles. Moreover, due to the desired precision in the cotangent series evaluation, the inner and outer summation costs are, respectively, $O(2p)$ and $O(p)$. Finally, adding the cost to create the clusters \mathcal{B} , the total cost of the binomial series expansion is $O(4^\ell 2p^2 N + 2pN)$.

To validate the algorithm and to define parameters for the computations, a measurement of the error is performed. In addition, an assessment of the number of discrete potential vortex particles in the computations is performed. A shear-layer is placed along $y = 0$ To evaluate its self-induced velocity. The singular term of the cotangent series is solved using the FMM series expansion truncated with 40 and 75 terms for double and quadruple precision, respectively. These values guarantee machine precision [20]. In Fig. 6, the RMS deviation from the exact solution is shown. There, the dashed lines indicate computations in double precision (8 bytes), while solid lines represent results with quadruple precision (16 bytes). It is possible to see that the series diverges so the error increases when more terms are used in the series.

Although the series is convergent in Figs. 3 and 4, there, its calculation is performed using a straightforward evaluation of Eq. 18. Hence, each term in the power series goes to zero as the exponent becomes larger since there are no source errors. However, the fast summation algorithm given by Eq. 38 employs a Newton binomial expansion which may have truncation in several operations, for instance, the grouping of source-terms in \mathcal{B}_m , as well as the evaluation of the binomial coefficients. This last term is the main source of error in the method since it quickly

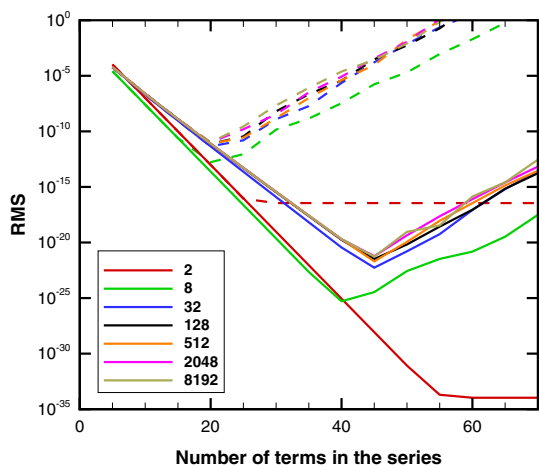


Fig. 6 Error in velocity magnitude for a shear layer discretized by a different number of vortex particles. *Dashed and solid lines* present results for double and quadruple precision, respectively. *Legend* indicates the number of vortex particles used in the simulation

becomes large for the computer to handle all its digits (the largest factorials without truncation for double and quadruple precision are, respectively, 21 and 37). Therefore, the calculation of the Newton binomial expansion may not achieve machine precision due to truncation errors. Consequently, the cotangent series has the largest errors in the algorithm and, to reduce these errors, the number of terms used in the series expansion is that which is closer to the inflection point in Fig. 6. For simulations involving more than 32 particles, we do not observe much difference between the inflection point. Summarizing, to obtain the minimum error, we choose the number of terms in the series to be around 20, for double precision, or 40, for quadruple precision.

3.3 Refinement of the initial domain

As explained, the separation among vortex particles plays an important role in the present fast algorithm. Based on Figs. 3 and 5, it is necessary that the minimum separation of the imaginary part to use the exponential series, as well as the argument z of the power series, should obey

$$|y| \gtrsim 0.125 \text{ to use exponential series or}$$

$$|z| \lesssim 0.70 \text{ to use power series.}$$

An efficient algorithm for fast summation of randomly scattered vortex particles should be based on clustering. Hence, the domain refinement idea from the FMM is employed here as previously discussed. One must not forget the convergence criterion for Eq. 15, $|z| < 1$, which imposes the largest domain to have unitary size, i.e., $S_0 = 1$.

Using the ideas of divide and conquer, the domain is refined into smaller boxes which interact among each other.

The separations of the boxes depend exclusively on the division of the domain. Using the same box refinement from the free-domain FMM to simultaneously satisfy the requirements of box separation for both series expansions discussed above, the refinement level used should be $\ell = 3$. In this refinement level, the size of a box is $S_3 = 0.125S_0$ since $S_\ell = S_0/2^\ell$.

A methodology can be imposed to regulate the series expansion to be solved. The interactions among clusters with minimum vertical separation $|y| = 0.125$ is performed via exponential series. In other words, the separation is, at least, one box from the source to the observer cluster, as presented in Fig. 7 for a generic observer box. In this figure, one can see an observer particle indicated in green inside a dark gray box. This particle interacts with a source particle indicated in cyan within a light gray box and their separation is given by $|y| = 0.125$. Hence, all boxes with $y_j - y_k > 0$ must solve Eq. 22, while the ones where $y_j - y_k < 0$ are solved by Eq. 23.

In the region where the vertical separation $|y|$ among two particles is smaller than 0.125, one must solve the power series expansion, Eq. 38. This is shown in Fig. 8, where a generic observer box indicated by dark gray interacts with all source boxes, colored by light gray. Furthermore, the maximum separation z among two particles is 0.673 and it guarantees good convergence based on Fig. 3.

If the distance among two particles is greater than the convergence range, $|z| < 1$, this series expansion cannot be applied. However, due to the periodic properties of the cotangent function, one can use the replica of a source particle instead, so the argument z is within the convergence range, as presented in Eq. 32. To implement this logic, if an observer box is close to the border of the domain, it should not interact with the original source

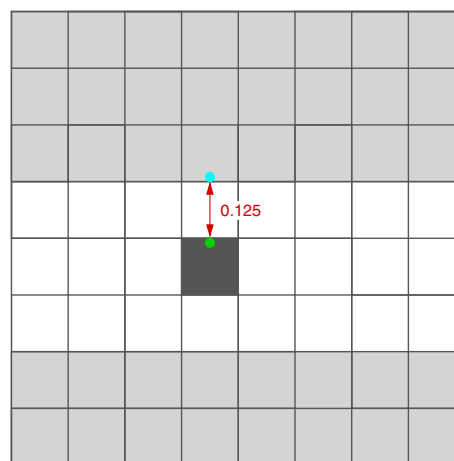


Fig. 7 Light gray boxes in the exponential interaction list of a dark gray box




Fig. 8 Light gray boxes in the power series interaction list of that in dark gray

cluster, but instead, it interacts with the replicated source box within the nearest level 0 neighbor. One should remind that this logic is only valid for the calculation using the power series. The exponential series expansion does not have any limitation with respect to the separation along the x -axis and, therefore, for this case, all interactions are performed within the original domain.

In Fig 9a below, the light gray boxes are descendants of the left level 0 neighbor of the central domain, *a.k.a.* domain $n = -1$, and they interact with the dark gray box in Fig. 9b. This artifice forces that the interactions are always performed for the nearest clusters to guarantee convergence.

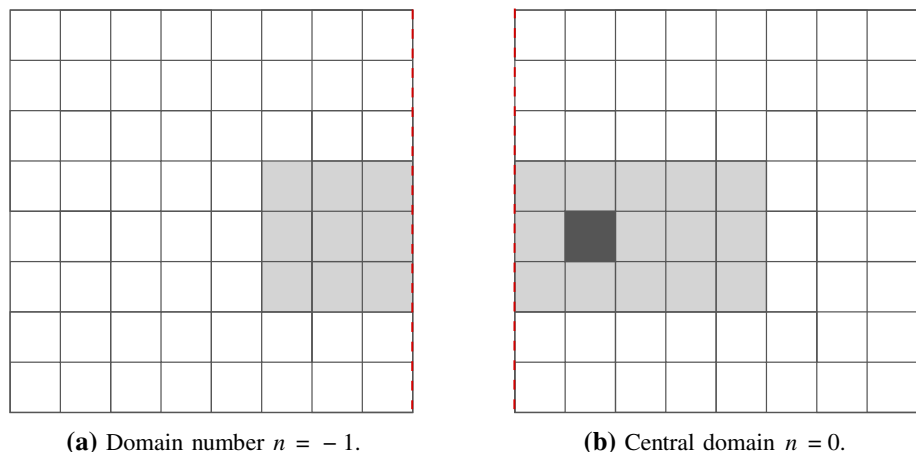
One important remark is that level $\ell = 3$ guarantees the fastest evaluation of both exponential series and Newton binomial expansion. The only drawback of such coarse level is the treatment of the singular term by the FMM, where direct evaluation of the cotangent function is still necessary to solve all terms of the power series expansion.

These direct evaluations are performed along near-field clusters which do not satisfy the separation criteria of “well-separated” boxes in the FMM [10, 17].

To solve the FMM, one must solve both near and far-field calculations. The fast summation of the far-field interactions is performed using box-to-box operations in a hierarchical pattern, where a multi-level approach with its “interaction list” regulates the interactions among clusters. In the current implementation for the cotangent kernel, part of the domain is solved by the exponential series expansion, which does not require the evaluation of the $1/z$ term. This singular term is present only in the power series and, therefore, not all boxes interact via FMM inside the domain. In addition, the interactions of both exponential and power series expansions are performed at level $\ell = 3$. For the two reasons presented, modifications in the interaction list are necessary to account for all the 24 boxes, at level 3, marked to perform the interaction using the power series expansion. The list of boxes for a generic observer box, indicated by dark gray in Fig. 8, is illustrated in light gray in the same figure. In addition, to simplify the implementation, no operations at level $\ell = 2$ are performed.

The near-field evaluations in the FMM at level $\ell = 3$ are expensive if a large number of particles have to be accounted for. Therefore, further refinement levels are necessary to reduce this cost. Due to the criterion of separation to apply Eqs. 22 and 23, the eight neighbors of a box at level 3 are always accounted for using power series, where the $1/z$ term is present. Hence, these nine boxes (the eight neighbors and the observer box itself) can be recursively refined to level $\ell = 4$ and so forth, without any restriction. This is shown in Fig. 10, where the boxes around the dark gray one have been further refined in the multi-level FMM. All the light gray boxes interact, either directly via M2L operations or indirectly with the combination of M2L and L2L operations from boxes at level 3,

Fig. 9 Interaction list with nearest clusters (light gray) of a box (dark gray)



(a) Domain number $n = -1$.

(b) Central domain $n = 0$.




Fig. 10 Multi-level FMM refinement around the observer box for faster power series solution of the singular term

with the dark gray box. The white boxes are accounted for via exponential series and, hence, they are ignored in this step.

It is worth noticing that the maximum level of the FMM algorithm to solve the near-field interactions has no limitations in terms of separation criterion due to the series convergence. However, the box cannot be further refined if it conflicts with the particularities of each problem, e.g., the viscous core of a Lamb–Oseen vortex becomes larger than the box size. For more details about such issues, the authors refer to Ref. [20].

Summarizing the fast algorithm: the fast interactions among particles are given either by Eq. 30 (and its counterpart that can be obtained from Eq. 23) if the particles are distant, or by the computation of Eq. 38. Although the exponential series is valid for positive arguments of $\text{Im}(z)$, one can benefit from the even–odd property of the cotangent function and work with a similar equation for negative arguments of $\text{Im}(z)$. Hence, one must compute first the boxes with positive separation and then evaluate the boxes with negative separation.

The power series is solved using two different methodologies: one is the FMM for the singular term $1/z$ and the other is the Newton binomial expansion which evaluates all the non-singular terms of the power series. The FMM still requires direct evaluations of the cotangent function, while the Newton binomial expansion can be solved without restriction among neighbor boxes. In addition, due to the convergence criterion, boxes that are far-away from each other may interact via one of their image replicas, to guarantee convergence using the smallest argument.

In addition, the FMM can be accounted for using refinement levels $\ell = 3$ or higher; the Newton binomial and the exponential series expansion are accounted for at $\ell = 3$. This level guarantees good convergence for all series and leads to the smallest computational time. Using this combination of free-domain FMM, Newton binomial expansion

and fast summation of exponential series, three different formulations are employed simultaneously to solve the convolution among vortex particles with periodic boundary conditions for which the convolution kernel is the cotangent function.

4 Results

In this section, the fast algorithm is employed to solve a two-dimensional velocity field with periodic boundary conditions and a vortex sheet roll-up problem using the DVM. Furthermore, we present an assessment of the error in velocity magnitude as well as an analysis of the computational cost of the algorithm.

4.1 Evaluation of a random periodic velocity field

To evaluate a periodic velocity field using the fast algorithm, 256 vortex particles are randomly placed inside a square with unitary side, centered at the origin of the Cartesian coordinate system. The circulation of each vortex particle ranges from $-0.1 \leq \Gamma \leq 0.1$.

To illustrate the problem, Figs. 11a, b shows the velocity magnitude for the direct calculation solving the Biot–Savart law and the fast algorithm, respectively. The velocity field is evaluated along a grid with 251 points in each direction. The deviation of the fast numerical algorithm is presented in Fig. 12, where the maximum error in velocity magnitude is about 1×10^{-12} .

4.2 Vortex sheet roll-up problem

Here, the fast algorithm is employed to study the evolution of a vortex sheet roll-up where the formation of the Kelvin–Helmholtz instability is observed. Such flow problem has been investigated in the literature by several authors [1, 3, 13, 14, 22, 23]. A periodical shear layer is discretized using N vortex particles and its temporal evolution is simulated. The vorticity γ of a continuous shear layer with length λ , with a jump in velocity ΔU , is given by

$$\gamma = \Delta U \lambda. \tag{39}$$

Discretizing the shear layer using N vortex particles with constant length λ/N leads to

$$\Gamma_k = \frac{\Delta U \lambda}{N}, \tag{40}$$

and setting $\Delta U = 1$ and $\lambda = 1$, the circulation of each vortex is given by

$$\Gamma_k = \frac{1}{N}. \tag{41}$$

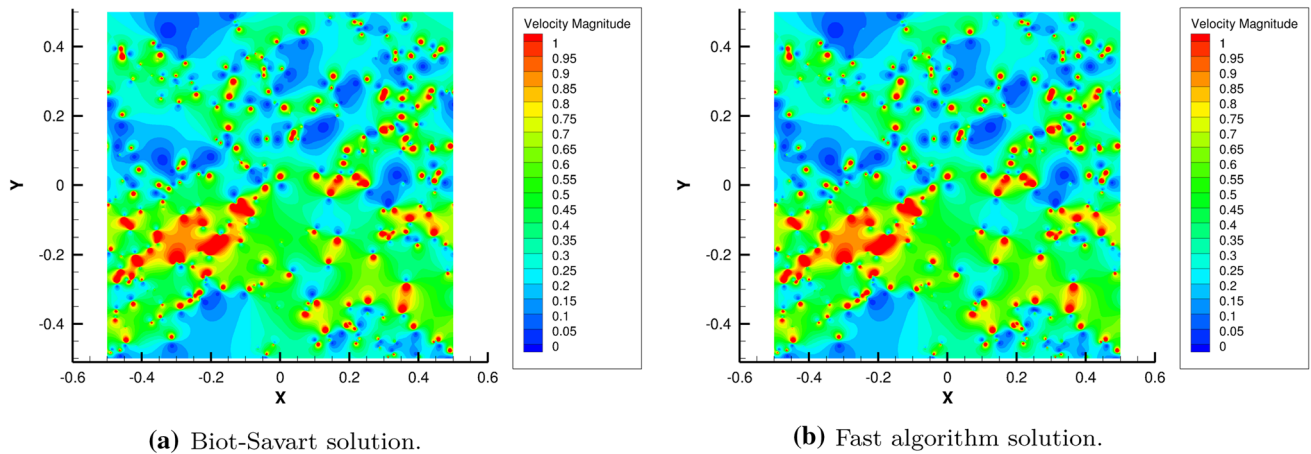


Fig. 11 Magnitude of the velocity field




Fig. 12 Error in velocity magnitude of the proposed fast algorithm compared to the Biot–Savart law computed using machine quadruple precision




Fig. 13 Sinusoidal initial displacement, with amplitude $A = 0.01$, of a shear layer

The initial solution includes a sinusoidal displacement with length λ and amplitude A , as shown in Fig. 13 and it is given by

$$y_k = A \sin\left(\frac{2\pi}{\lambda}x_k\right), \tag{42}$$

with

$$A < < \lambda, \tag{43}$$

and

$$-0.5 < x_k < 0.5. \tag{44}$$

Figure 14a presents the solutions of the shear layer discretized by $N = 5120$ vortex particles. The simulation is performed using quadruple precision and a Lamb–Oseen vortex with viscous core $\sigma = 0.02$. The time step is set as $\Delta t = 0.01$ and the calculation is run for 200 iterations. One can see that the solution obtained by the fast algorithm has an excellent agreement with the Biot–Savart calculation. If periodic boundary conditions are not enforced, the shear layer develops a spurious precession and its tips roll towards the core [20]. In Fig. 14b, one can observe the RMS error assessment as a function of the time evolution for a simulation performed by the Biot–Savart law in double precision and another performed using the fast algorithm, in quadruple precision. The error is computed with respect to the quadruple precision Biot–Savart law.

4.3 Error analysis

We perform an error analysis for the evaluation of a random periodic velocity field, as in Sect. 4.1. To measure the error based on the number of vortex particles, an investigation is performed from two up to one million vortex particles, randomly distributed inside the domain. The velocity magnitude is evaluated at each vortex location using the Biot–Savart law, the coupled fast method as well as a “brute-force” approach which replicates the FMM domains to simulate the periodic boundary conditions. Results from the brute-force approach are obtained using 2^{30} FMM replicated boxes. This last method is still faster than the direct solution of the Biot–Savart law as shown in Ref. [20]. The results of RMS deviation from a comparison

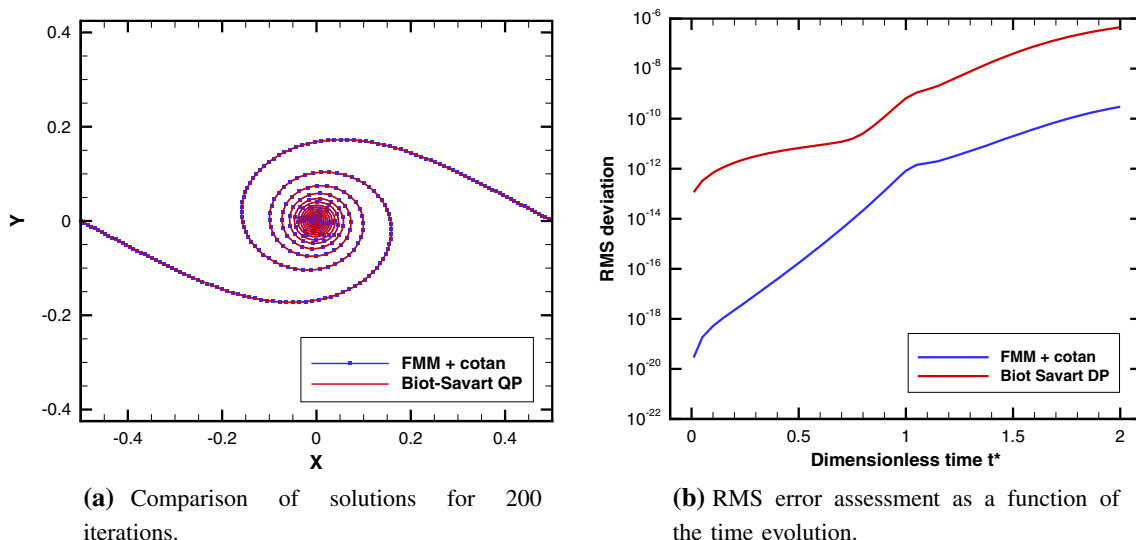


Fig. 14 Simulation of vortex sheet roll-up problem using $N = 5120$ particles, a viscous core $\sigma = 0.02$ and $\Delta t = 0.01$

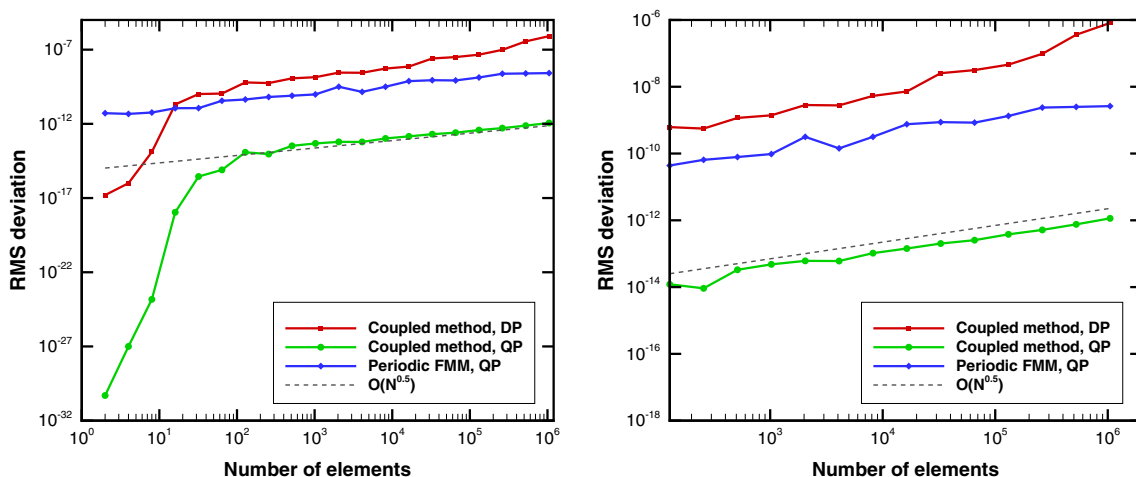


Fig. 15 Error for a single evaluation of local velocity field for different methods

of both fast methods to the Biot–Savart law are presented in Fig. 15. One can see that the cotangent approximation using the coupled power and exponential series expansions is more accurate than the replicated FMM solution, called periodic FMM in the figure, and which has algebraic decay. Furthermore, this slow rate of decay implies that the usage of quadruple precision does not lead to better results. On the other hand, the coupled method benefits from increased machine precision, since the error drops six orders of magnitude. Finally, the RMS deviation of both methods grows proportional to $O(N^{0.5})$.

4.4 Computational time

The overall computational cost of the algorithm results from the combination of the solutions of the free-domain

FMM and the Newton binomial expansions or the exponential expansions. This section presents results of computational time for simulations performed in serial using a single thread in a 2.5 GHz Intel®Xeon™ E5-2670v2 CPU.

The results of overall computational cost are shown in Fig. 16a, b for double and quadruple precision, respectively. The cost behavior is similar as that for the free-domain FMM, where the increase in the number of particles requires finer levels of refinement to reduce the quadratic dependence of near-field Biot–Savart calculations. This is shown by different curves in the legend, where the near-field is tested for different refinement levels in the FMM, which handles the singular term. One must not forget that the Newton binomial series and the exponential series are solved only at level 3. In the figure, one can see that for double precision (8 bytes), time savings of

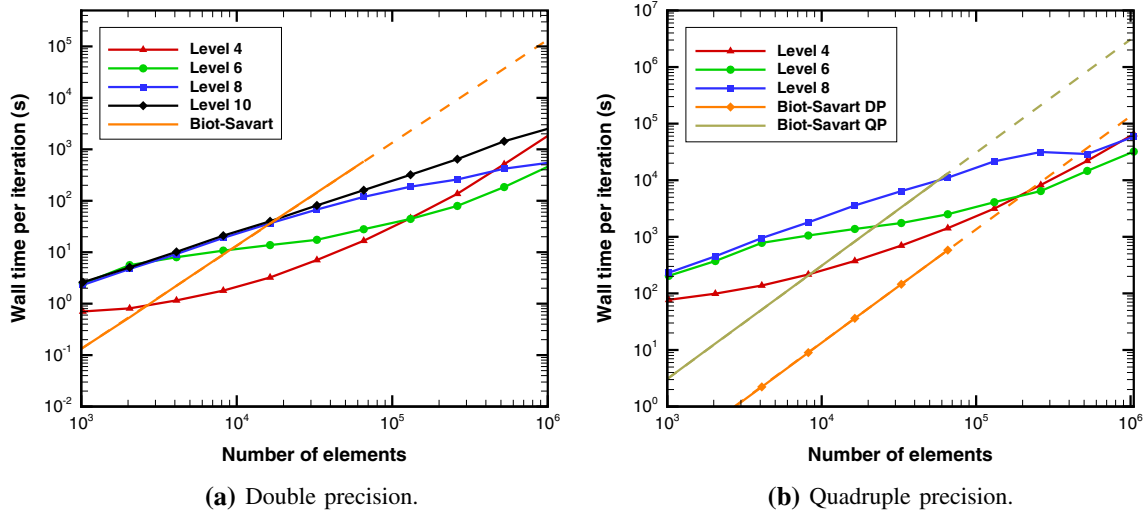


Fig. 16 Overall computational cost of the combined fast method compared to the Biot–Savart solution

nearly 240 times are obtained for 1 million vortex particles, while in quadruple precision (16 bytes) the time savings for the same number of particles is about 100 times. Furthermore, the fast algorithm using 16 bytes can be even faster than the direct summation in double precision for more than 200,000 vortex particles.

As discussed, the computational cost from the new algorithm is a combination of the costs from the free-domain FMM, the binomial series as well as the exponential series calculations. The FMM has a mixed computational cost depending linearly and quadratically with the number of particles [20], depending on near and far-field balancing of the FMM, while the other approaches present a linear cost, depending on the number of particles. The method also has a constant fixed cost. Hence, the total cost is given by

$$\mathcal{A}_2 N^2 + \mathcal{A}_1 N + \mathcal{A}_0,$$

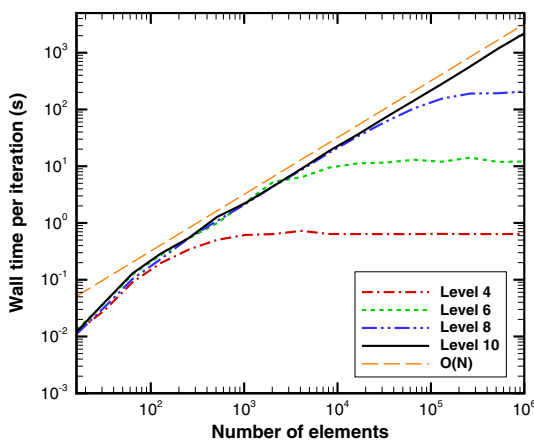


Fig. 17 Computational time for box-to-box operations

where \mathcal{A}_0 only depends on the number of non-empty boxes. For large refinement levels, there are several non-empty boxes for a small number of particles, which delays the characterization of a fixed cost.

On the other hand, as shown in Fig. 17, for a low number of particles in this random cloud case, the coupled M2L and L2L operations have a delayed plateau and the cost \mathcal{A}_0 is no longer independent of the number of particles. It is now based on the ratio of particles per number of boxes. If it is smaller than one, there are empty boxes indicating that fewer cluster operations are performed, so there is an initial slope on the curves. When there is at least one particle per box, without empty boxes, the domain becomes *saturated*. This way, the cost \mathcal{A}_0 of box interactions is dumped for a low number of particles. However, when the domain is saturated, the plateau appears and \mathcal{A}_0 becomes constant.

The additional operations to solve both exponential and power-series increase mainly the linear cost \mathcal{A}_1 . These series are solved only at level $\ell = 3$, independently of the maximum refinement level L of the singular term solved by the free-domain FMM. In addition, the P2M and L2P operations in the FMM do not depend on the maximum refinement level and they are linearly proportional to the number of vortex particles $O(pN)$. For the exponential series, the additional cost is $O(pN)$ as well. However, using the Newton binomial expansion at level $\ell = 3$, the additional cost of $O(4^3 2p^2 N)$ is not only present but also dominant for a small number of vortices.

It is possible to see in Fig. 18a the most expensive operations dependent on the number of particles, i.e., both P2M and M2L operations, as well as the cost for the exponential and Newton binomial expansions. There, the computations are performed for double precision. In

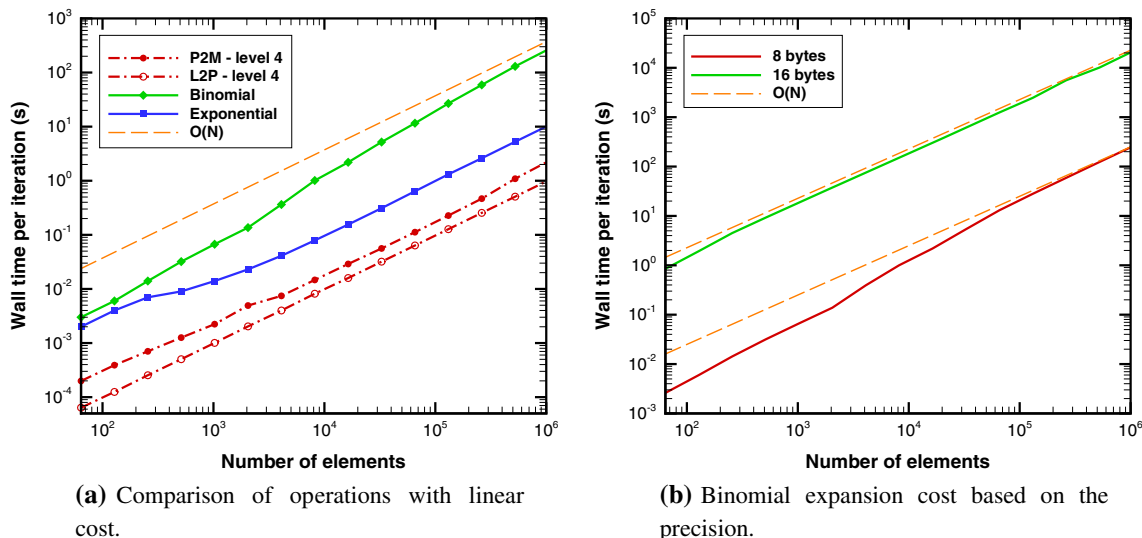


Fig. 18 Computational cost of steps linearly dependent on the number of vortex particles

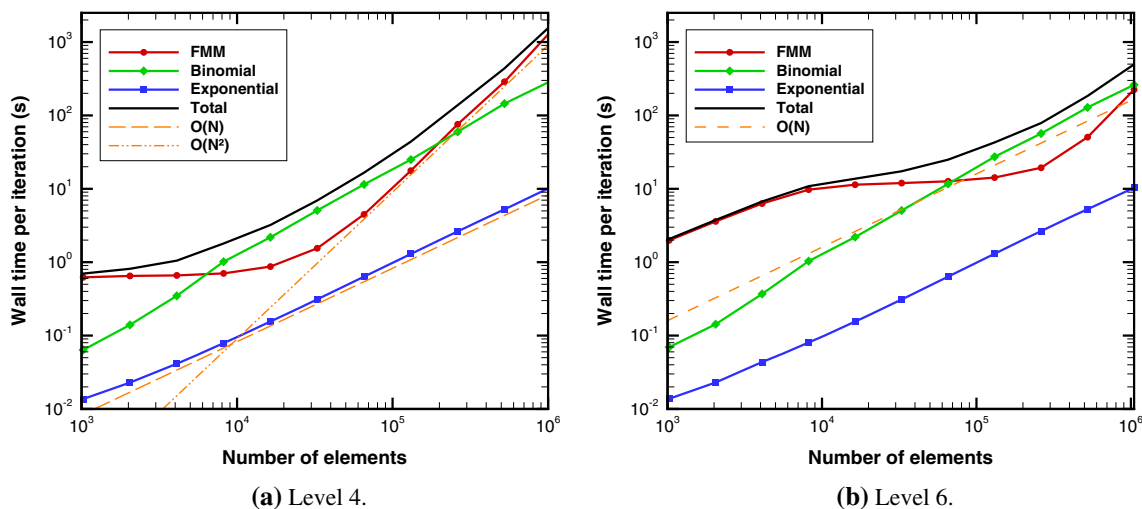


Fig. 19 Computational cost of the three approaches employed in the fast summation algorithm

Fig. 18b, one can see the cost of the binomial expansion for both double and quadruple precision. There, this cost converges to $O(N)$ when more particles are present in the domain.

Finally, since neither the exponential nor the power series have singular terms, particles close to each other interact exclusively via fast algorithms. This is not the case for the FMM, which still depends on convolutions among particles close to each other. Hence, the quadratic dependency \mathcal{A}_2 is that typical of fast multipole methods [20]. This means that the quadratic cost increases for lower levels of the FMM when more direct interactions need to be evaluated.

In Fig. 19 it is possible to see a comparison of the cost of each individual methodology used by the algorithm. For a small number of particles, the far-field operations in the

FMM are the most expensive. However, when the domain becomes *saturated*, the cost for the binomial series is the highest. Finally, for a large number of particles, the near-field cost in the FMM is dominant.

5 Conclusions

A novel fast algorithm is proposed to perform flow simulations using the discrete vortex method including periodic boundary conditions. The algorithm combines exponential and power series expansions to accelerate the calculation of the cotangent kernel used in the Biot–Savart law. A divide and conquer strategy is employed so clusters of vortex particles interact with each other instead of the particle–particle interactions. The exponential series expansion is

implemented using a *middleman* context similar to that from the fast multipole method. On the other hand, the power series expansion is implemented using a Newton binomial expansion for computing regular terms and using a fast multipole method for the calculation of singular terms.

An error analysis shows that the current method provides more accurate solutions than a “brute-force” implementation of the fast multipole method using image particles to represent the periodic boundary conditions. An analysis of the computational cost of the individual steps of the algorithm is provided and savings of order 240 are observed for double-precision simulations with one million particles.

Acknowledgements The authors acknowledge the financial support received from Fundação de Amparo à Pesquisa do Estado de São Paulo, FAPESP, under Grants No. 2013/03413-4 and 2013/07375-0 and from Conselho Nacional de Desenvolvimento Científico e Tecnológico CNPq, under Grants No. 470695/2013-7 and 305277/2015-4. The authors also acknowledge CAPES for providing a M.Sc. scholarship to the first author. The computational resources provided by CENAPAD-SP under Project No. 551 are also acknowledged.

References

1. Abid M, Verga A (2002) Stability of a vortex sheet roll-up. *Phys Fluids* 14:3829–3834
2. Abramowitz M, Stegun IA (1964) Handbook of mathematical functions with formulas, graphs and mathematical tables. National Bureau of Standards, Gaithersburg
3. Baker GR, Pham LD (2006) A comparison of blob methods for vortex sheet roll up. *J Fluid Mech* 547:297–316
4. Barba LA, Leonard A, Allen CB (2004) Advances in viscous vortex methods—meshless spatial adaption based on radial basis function interpolation. *Int J Numer Methods Fluids* 46:001–032
5. Bhaskaran R, Lele SK (2010) Large eddy simulation of free-stream turbulence effects on heat transfer to a high-pressure turbine cascade. *J Turbul* 11:1–15
6. Chorin AJ (1973) Numerical study of slightly viscous flow. *J Fluid Mech* 57:785–796
7. Cocle R, Winckelmans G, Daeninck G (2008) Combining the vortex-in-cell and parallel fast multipole methods for efficient domain decomposition simulations. *J Comput Phys* 227:9091–9120
8. Dymnikova GY, Dymnikov YA, Guvernuyuk SV (2016) Mechanism underlying Kármán vortex street breakdown preceding secondary vortex street formation. *Phys Fluids* 28:054101
9. Greengard L, Kropinski MC (2004) Integral equation methods for stokes flow in doubly-periodic domains. *J Eng Math* 48:157–170
10. Greengard L, Rokhlin V (1987) A fast algorithm for particle simulations. *J Comput Phys* 73:325–348
11. Gumerov NA, Duraiswami R (2004) Fast multipole method based filtering of non-uniformly sampled data. Tech. rep, University of Maryland
12. Hammer P, Altman A, Eastep F (2014) Validation of a discrete vortex method for low Reynolds number unsteady flows. *AIAA J* 52:643–649
13. Krasny R (1986) Desingularization of periodic vortex sheet roll-up. *J Comput Phys* 65:292–313
14. Krasny R (1986) A study of singularity formation in a vortex sheet by the point-vortex approximation. *J Fluid Mech* 167:65–93
15. Lindsay K, Krasny R (2001) A particle method and adaptive treecode for vortex sheet motion in three-dimensional flow. *J Comput Phys* 172:879–907
16. Marple G, Barnett A, Gillman A, Veerapaneni S (2015) A fast algorithm for simulating multiphase flows through periodic geometries of arbitrary shape. arXiv preprint [arXiv:1510.05616](https://arxiv.org/abs/1510.05616)
17. Nishimura N (2002) Fast multipole accelerated boundary integral equation methods. *Appl Mech Rev* 55:299–324
18. Quinn DB, Moored KW, Dewey PA, Smits AJ (2014) Unsteady propulsion near a solid boundary. *J Fluid Mech* 742:152–170
19. Ramesh K, Gopalarathnam A, Granlund K, Ol MV, Edwards JR (2014) Discrete-vortex method with novel shedding criterion for unsteady aerofoil flows with intermittent leading-edge vortex shedding. *J Fluid Mech* 751:500–538
20. Ricciardi TR, Wolf WR, Bimbato AM (2017) Fast multipole method applied to Lagrangian simulations of vortical flows. *Commun Nonlinear Sci Numer Simul* 51:180–197
21. Schiavo LACA, Jesus AB, Azevedo JLF, Wolf WR (2015) Large eddy simulations of convergent divergent channel flows at moderate Reynolds numbers. *Int J Heat Fluid Flow* 56:137–151
22. Shelley MJ (1992) A study of singularity formation in vortex-sheet motion by a spectrally accurate vortex method. *J Fluid Mech* 244:493–526
23. Tryggvason G (1989) Simulation of vortex sheet roll-up by vortex methods. *J Comput Phys* 80:1–16
24. Wolf WR, Lele SK (2012) Trailing-edge noise predictions using compressible large-eddy simulation and acoustic analogy. *AIAA J* 50:2423–2434
25. Yokota R, Barba LA (2013) FMM-based vortex method for simulation of isotropic turbulence on GPUs, compared with a spectral method. *Comput Fluids* 80:17–27
26. Yokota R, Obi S (2010) Comparing vortex methods and finite difference methods in a homogeneous shear flow. *Int J Numer Methods Fluids* 63:828–846
27. Yokota R, Sheel TK, Obi S (2007) Calculation of isotropic turbulence using a pure Lagrangian vortex method. *J Comput Phys* 226:1589–1606
28. Zhang J, He Y, Tao W (2009) 3D numerical simulation on shell-and-tube heat exchangers with middle-overlapped helical baffles and continuous baffles. Part I: numerical model and results of whole heat exchanger with middle-overlapped helical baffles. *Int J Heat Mass Transf* 52:5371–5380