Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

On the Influence of the Mask Function in Volume Penalisation in Pseudo-Compressible Flow Simulations

Eden Getachew

eden.getachew@tnt.tu-berlin.de

Technische Universität Berlin

Thomas Engels

Aix-Marseille Université

Mathias Lemke

Technische Universität Berlin

Kai Schneider

Aix-Marseille Université

Angela Busse

Technische Universität Berlin

Research Article

Keywords: Volume penalisation, Mask functions, Order of convergence, Boundary smoothing, Artificial compressibility method

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11404422/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

On the Influence of the Mask Function in Volume Penalisation in Pseudo-Compressible Flow Simulations

Eden Getachew[1,2,3*], Thomas Engels[2], Mathias Lemke[1], Kai Schneider[3], Angela Busse[1]

[1*]Institut für Strömungsmechanik und Technische Akustik (ISTA), Technische Universität Berlin, Berlin,Germany.
[2]Institut des Sciences du Mouvement Etienne-Jules Marey (I2M, UMR 7287), Aix-Marseille Université, CNRS, Marseille, France.
[3]Institut de Mathématique de Marseille (I2M, UMR 7373), Aix-Marseille Université, CNRS, Marseille, France.

*Corresponding author(s). E-mail(s): eden.getachew@tnt.tu-berlin.de;

Abstract

Accurately imposing boundary conditions on complex geometries remains a central challenge in computational fluid dynamics. The volume penalisation method provides a simple and flexible framework for representing arbitrarily shaped solid obstacles by a mask function without requiring body-fitted meshes. The accuracy of the method is influenced by both modelling and discretisation errors, which depend on the choice of mask function and penalisation parameter. In this work, these error sources are systematically analysed for one-, two-, and three-dimensional benchmark problems. Particular emphasis is placed on assessing the optimised smooth mask function proposed by Hester et al. [1]. While its favourable convergence properties were derived analytically and demonstrated for highly resolved grid-aligned configurations, its performance in practical simulations with non-body-conformal Cartesian grids remains insufficiently investigated. The present study therefore examines how these properties transfer to increasingly complex configurations involving non-aligned grids, fully developed turbulence, and moving boundaries. A one-dimensional Poisson problem is employed to investigate fundamental error mechanisms and convergence behaviour. The study is then extended to a two-dimensional flow past a circular cylinder at $\boldsymbol{Re = 40}$, a three-dimensional turbulent channel flow at $\boldsymbol{Re_\tau = 180}$, and a three-dimensional flapping wing with prescribed motion.

The artificial compressibility method is used to relax the incompressibility constraint. The optimised mask function demonstrates improved robustness with respect to parameter variations and non-grid-aligned interfaces. In comparison, both the conventional discontinuous and cosine mask functions exhibit a stronger sensitivity to parameter tuning. For the two-dimensional cylinder flow, convergence rates of up to $\mathcal{O}(\Delta x^{1.6})$ are obtained for the velocity field. In the turbulent channel flow, both the discontinuous and optimised mask functions reproduce mean velocity profiles and turbulence statistics in good agreement with reference data. Overall, our results highlight the importance of mask function design for accuracy and robustness in volume penalisation methods, especially for non-grid-aligned interfaces. Further improvements in multi-dimensional and fluid-structure interactions configurations will primarily depend on continuing to mitigate discretisation errors.

Keywords: Volume penalisation, Mask functions, Order of convergence, Boundary smoothing, Artificial compressibility method

1 Introduction

In computational fluid dynamics (CFD), fluid–structure interactions (FSI) remain a fundamental challenge. Enforcing boundary conditions on geometrically complex and possibly moving interfaces is essential for reliable simulations in many applications. Traditional body-fitted methods depend on mesh generation, which increases computational cost and complicates the treatment of moving geometries [2]. To overcome these limitations, immersed boundary methods (IBM) were developed, originally proposed by Peskin for the simulation of cardiac flows [3]. IBM enforces boundary conditions by adding forcing terms to the governing equations on a fixed Cartesian grid [3]. This approach provides geometric flexibility while avoiding repeated mesh generation. IBM has been successfully applied to a wide range of problems, including biological flows, such as heart valve dynamics, blood flow in complex vascular geometries [4–6], fish-like swimming [7], benchmark problems such as flow past a cylinder [8], as well as complex external flows around road vehicles [2, 8], and simulations of confined plasma turbulence in magnetohydrodynamics [9, 10].

Immersed boundary (IB) formulations are commonly distinguished according to whether the fluid–solid interface is smeared over a finite grid width or retained as a sharp boundary conforming geometry [11]. The former, diffuse interface approaches, rely on forcing methods, like direct forcing [12], discrete delta [3] or mask functions [13]. In contrast, sharp-interface methods impose boundary conditions through local modifications of the flux, such as cut-cell, immersed-interface, or reconstruction-based techniques [11]. This distinction has direct implications for accuracy: classical IB formulations generally exhibit first-order spatial accuracy, whereas sharp-interface formulations can achieve second-order spatial accuracy through an accurate reconstruction of velocity and pressure boundary conditions at the geometrical interface [14]. More recent developments include high-order sharp immersed formulations for stationary and moving boundaries [15], treatments enforcing geometric conservation

for moving boundaries [16], and local grid refinement and wall-modelling strategies for high-Reynolds-number flows [17].

This work focuses on the volume penalisation method (VPM), a diffuse interface approach where the entire solid region is penalised by damping the velocity inside the obstacle [13]. In contrast to surface-based forcing approaches, the penalisation force therefore acts throughout the solid volume. A distinctive property of VPM is the existence of analytical convergence results showing that, for vanishing penalisation parameter, the penalised incompressible Navier–Stokes equations converge to the no-slip boundary conditions [13, 18]. Moreover, integral hydrodynamic forces can be evaluated by volume integration of the penalisation terms, avoiding an explicit reconstruction of pressure and viscous stresses at the immersed surface. From a practical perspective, VPM allows simple integration into existing CFD solvers, since no auxiliary fields or constraint equations are introduced. The method has been applied to incompressible flows, ranging from fundamental convergence studies to large-scale turbulent configurations [13, 18–22]. In addition, VPM has been extended to compressible flows [23–25] and combined with the artificial compressibility method (ACM) [26]. Representative benchmark cases include moving obstacles [27] and wall-bounded turbulence [22]. Biological applications of VPM include fish-like swimming [28] and insect flight. For the latter, studies include fluid-structure interaction with flexible flapping foils as simplified models of insect wings [21], three-dimensional simulations of rigid flapping wings with the parallel Fourier-based solver `FluSI` [29], and, more recently, multiscale simulations of flapping insects within the wavelet-adaptive `WABBIT` framework [26].

A related methodological refinement of VPM is the active penalty method, proposed by Shirokoff and Nave [30]. This approach constructs an extension function inside the solid domain to match the boundary values and derivatives. Although this improves accuracy, it is less straightforward to apply to complex or moving geometries.

In recent years, VPM has increasingly been incorporated into hybrid numerical and optimisation frameworks. Kou et al. [31] combined volume penalisation with a high-order flux reconstruction scheme. They investigated implicit treatments of the penalisation term to alleviate stiffness-related time-step restrictions. More recently, Liu and Hasegawa [32] developed a continuous adjoint-based shape optimisation approach for compressible flows that combines VPM with a level-set representation of the geometry. This approach avoids repeated mesh generation during optimisation. For transonic flows around cylinders and spheres, drag reductions of more than 20% were reported. These developments demonstrate the continued relevance and versatility of VPM.

Despite its broad applicability, VPM is limited in terms of accuracy and convergence. Solid obstacles are represented as porous regions, where the velocity is damped by the penalty term. As the penalisation parameter $C_\eta \to 0$, the solution converges to the desired boundary condition [13]. The penalisation introduces a thin boundary layer inside the solid domain near the fluid–solid interface [18]. Nguyen van yen et al. [33] analysed the approximation properties of VPM from a spectral perspective. For a given discretisation, they showed that an optimal penalisation parameter C_η exists. This optimum corresponds to a balance between penalisation and discretisation errors. Early theoretical analysis by Carbou and Fabrie [18] proved that the penalised

solution converges to the Navier–Stokes solution at a rate of $\mathcal{O}(\sqrt{C_\eta})$ in the L^2-norm. Hester et al. [1] later identified the dominant error source as the displacement length error, i.e. the geometric offset between the penalised and the true interface. By introducing an optimised smooth mask function, they analytically demonstrated second-order convergence in C_η. However, these results were primarily derived from continuous analysis and validated using highly resolved numerical experiments with grid-aligned interfaces. In practical applications of VPM, solid obstacles are generally represented on finite Cartesian grids and are rarely aligned with the underlying discretisation. It therefore remains important to assess how different mask formulations behave once discretisation, finite grid resolution, and lack of grid alignment are taken into account. The objective of the present work is to systematically compare three mask functions established in the literature, the discontinuous [13], the cosine [29] and the optimised mask function [1], and to assess their accuracy, convergence behaviour, and sensitivity to the penalisation parameter across increasingly complex numerical configurations. Addressing this gap between continuous analysis and practical discrete implementations constitutes a central motivation for this comparison.

To this end, the different mask functions and penalisation parameters are evaluated in four test cases: a one-dimensional (1D) Poisson problem, a two-dimensional (2D) approximately incompressible cylinder flow, a three-dimensional (3D) approximately incompressible turbulent channel flow, and a 3D simulation of a flapping wing. The test cases address complementary aspects of the numerical method. The 1D problem is used to analyse penalisation in detail, including discretisation, and grid-alignment effects. The 2D cylinder introduces a curved, non-grid-aligned interface, while the turbulent channel assesses the mask formulations in a fully developed 3D turbulent flow, and the 3D flapping-wing case extends the analysis to moving boundaries.

The remainder of this work is structured as follows. Section 2 introduces the governing equations and numerical methods, sections 3.1 – 3.4 present the numerical test cases with implementation and results. Section 4 concludes with a summary of the main findings.

2 Governing equations and Numerical Methods

2.1 Governing equations

We consider a viscous, approximately incompressible Newtonian fluid in a domain denoted Ω_f with an immersed solid obstacle Ω_s. The fluid motion is governed by the Navier-Stokes equations. The no-slip condition is enforced numerically using VPM [13]. Rather than prescribing a Dirichlet condition directly at the boundary $\partial\Omega_s$, the constraint is imposed within the solid domain Ω_s.

The resulting formulation combines the penalised momentum equation with the artificial compressibility method (ACM) and reads as

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} + \frac{1}{\rho}\frac{\partial p}{\partial x_i} - \nu \frac{\partial^2 u_i}{\partial x_j^2} = -\frac{\chi}{C_\eta}(u_i - u_{s,i}), \tag{1}$$

$$\frac{1}{C_0^2}\frac{\partial p}{\partial t} + \frac{\partial u_i}{\partial x_i} = 0. \tag{2}$$

Here, ρ denotes the constant fluid density, $u_i(\underline{x}, t)$ the velocity, $u_{s,i}(\underline{x}, t)$ the solid velocity, $p(\underline{x}, t)$ the pressure, and ν the kinematic viscosity. The above equation is written in index notation using Einstein's summation convention with $i, j \in \{1, 2, 3\}$. In the following, the index and vector notation is used interchangeably ($x_1 = x, u_1 = u_x, x_2 = y, u_2 = u_y$, etc).

ACM [34, 35] relaxes the incompressibility constraint by introducing a pseudo-time derivative of the pressure. The parameter C_0 denotes the artificial speed of sound. In the steady-state limit, the pressure becomes time-independent and the velocity field satisfies the incompressibility condition.

The discontinuous mask function χ is defined as

$$\chi_d(\underline{x}) = \begin{cases} 1 & \underline{x} \in \Omega_s \\ 0 & \underline{x} \in \Omega_f, \end{cases} \tag{3}$$

and encodes the geometry of the solid body. The penalisation parameter C_η represents the permeability of the solid body. In the limit $C_\eta \to 0$, the no-slip condition is recovered. The parameter C_η has the dimension of time and can be interpreted as a relaxation timescale. For explicit time integration schemes, this implies the stability restriction $\Delta t = \mathcal{O}(C_\eta)$ [29]. The time step must therefore resolve the penalisation dynamics, i.e. the velocity decay inside the solid. Otherwise, the damping becomes under-resolved and the numerical scheme may become unstable.

2.2 Convergence Properties and Parameters of Volume Penalisation

The convergence behaviour of volume penalisation is governed by both modelling and discretisation errors. The modelling error associated with VPM has been analysed in previous studies [13, 33]. In addition, ACM introduces a divergence-related error due to the relaxation of the incompressibility constraint [34]. This error depends on the choice of the modelling parameter C_0 and can be reduced by pressure correction schemes [36]. In the present work, the artificial speed of sound is selected such that its contribution to the overall error remains negligible for the simulations considered. Residual divergence errors are transported through the domain by the pressure evolution term in a wave-like manner with propagation speed C_0 [37, 38]. Depending on the boundary treatment, these errors are diluted and therefore less localised, or transported out of the domain.

Discretisation Error and Parameter Scaling. Optimal convergence of VPM requires the discretisation error to scale with the same order as the dominant modelling error [33]. This relation can be analysed through the penalisation boundary layer that forms inside the solid domain near the fluid–solid interface. In the following, we assume a uniform grid resolution in all spatial directions with $\Delta x_i = \Delta x$. The thickness of the penalisation boundary layer is given by $\epsilon = \sqrt{\nu C_\eta}$ [18]. To quantify its numerical resolution, Engels et al. [29] introduced the penalisation constant $K_\eta = \epsilon / \Delta x$, which represents the number of grid points across the boundary layer. Maintaining a constant value of K_η during grid refinement yields the scaling relation $C_\eta \propto \Delta x^2$.

Since ϵ depends on the viscosity, the penalisation boundary-layer thickness decreases with increasing Reynolds number [18]. This behaviour is particularly advantageous for turbulent flows. Following this argument, the penalisation parameter is defined as

$$C_\eta = \frac{(\Delta x K_\eta)^2}{\nu}. \tag{4}$$

Influence of Mask Regularity. In practice, K_η is typically selected in the interval $0.1 - 0.4$ [29]. Within this range, a balance between penalisation and discretisation errors minimises the total error. However, the discontinuous mask function (3) limits the discretisation error to $\mathcal{O}(\Delta x)$. Smooth mask functions are therefore commonly used to improve regularity. These include cosine-based smoothing [29] and hyperbolic tangent profiles [24, 39]. The present study therefore hence compares mask functions of different regularity to assess their influence onto the overall error.

Mask Functions Considered. Hester et al. [1] proposed an optimised smooth mask function to minimise the modelling error. The method is based on matched asymptotic analysis and constructs a smooth transition profile that compensates for the displacement length. The resulting mask is

$$\chi_h(\xi) = \frac{1}{2}\left(1 - \tanh\left(\frac{2\xi}{\delta_{opt}}\right)\right), \quad \text{with} \quad \delta_{opt} = 2.64822828, \tag{5}$$

hereafter referred to as Hester's mask function. The variable $\xi = \xi(\underline{x})$ denotes the dimensionless signed distance to the fluid-solid interface defined as $\xi = \phi/\epsilon = \phi/\sqrt{\nu C_\eta}$. Here, $\phi = \phi(\underline{x})$ denotes the signed distance from the spatial position $\underline{x}$ to the interface $\partial\Omega_s$. The mask (5) adapts naturally to both the penalisation parameter and the grid resolution.

The present study also considers a second smooth mask function, namely the cosine mask [29]. It serves as a reference to assess the properties of the optimised mask.

$$\chi_c(\phi) = \begin{cases} 1 & \phi \leq -l \\ \frac{1}{2}\left(1 + \cos\left(\pi \frac{\phi + l}{2l}\right)\right) & -l < \phi < +l \\ 0 & \phi > +l, \end{cases} \tag{6}$$

where $l = \Delta x$ defines the half-width of the smoothing region relative to the grid spacing.

Error Measure. In the following, the discontinuous, cosine, and Hester's mask functions are quantitatively compared. To this end, the total error is evaluated in the physically relevant fluid domain and on the solid surface $\mathcal{D} \in \{\Omega_f, \partial\Omega_s\}$. It is measured component-wise using the relative L^2-error,

$$\varepsilon_{q,2} = \left(\frac{\int_\mathcal{D} (q - q_0)^2 \, \mathrm{d}\mathcal{D}}{\int_\mathcal{D} q_0^2 \, \mathrm{d}\mathcal{D}}\right)^{1/2}, \qquad q = \{u_x, u_y, u_z, p\} \tag{7}$$

with q_0 denoting either the exact or a reference solution, depending on the test case.

2.3 Discretisation

The different test cases are discretised on a periodic Cartesian grid. The grid resolution is defined as $\Delta x_i = L_i/N_i$, where L_i is the domain length and N_i the number of grid points in the i-th direction. Spatial derivatives are computed using fourth-order centred finite differences. The same spatial discretisation is employed for all test cases. The non-linear convective terms are formulated in skew-symmetric form, which ensures discrete L^2 stability of the scheme [40].

The implementation varies across the different setups. The 1D problem (3.1) is solved in MATLAB. The higher-dimensional problems (3.2-3.4) are simulated with the open source solver `WABBIT` [1] [26]. The code is designed for block-adaptive grids, which are parallelised using MPI. The cylinder and channel-flow cases are computed on non-adaptive equidistant grids, whereas the flapping-wing case employs the block-adaptive setting. Spatial derivatives are discretised with central finite-difference schemes. Through dynamic wavelet adaptation, the grid automatically refines where flow gradients require it and coarsens in regions where the flow is smooth. Each refinement level increases (or decreases) block-wise by a factor of two in each coordinate direction.

The time step size is constrained by penalisation, the CFL condition (CFL < 1), and viscosity. For stability of the time discretisation, Δt is selected adaptively as the minimum of these constraints:

$$\Delta t \leq \min\left(C_\eta,\ \mathrm{CFL}\frac{\Delta x}{u_{\max}},\ \frac{\Delta x^2}{2\nu}\right), \tag{8}$$

where $u_{\max}$ denotes the maximum velocity magnitude in the computational domain. Time integration is performed with the classical fourth-order Runge–Kutta (RK4) scheme. In cases where the time-step restriction is dominated by penalisation or viscous effects, the Runge–Kutta–Chebyshev (RKC) [41, 42] scheme is employed to allow larger stable time steps.

3 Numerical Experiments

This section presents four test cases: the 1D Poisson equation, 2D flow past a cylinder, 3D turbulent channel flow, and 3D flapping wing simulation. For each case, the numerical implementation and results are discussed.

3.1 Poisson Equation

To investigate the fundamental properties of volume penalisation, a 1D Poisson problem is considered. The penalised form of the governing equation is given by

$$-\frac{\mathrm{d}^2 u}{\mathrm{d}x^2} = f - \frac{\chi}{C_\eta}(u - u_s), \tag{9}$$

[1] https://github.com/adaptive-cfd/WABBIT

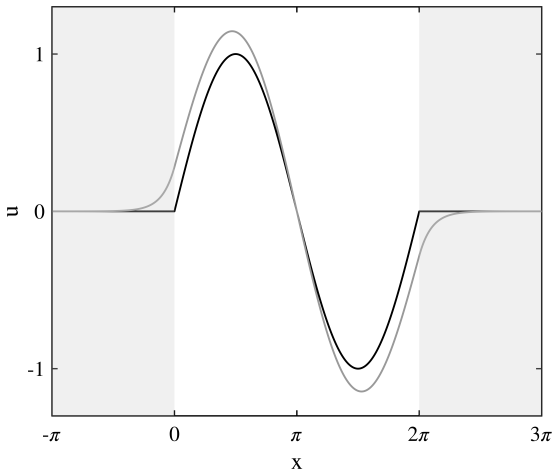


Fig. 1 Poisson equation (1D). Analytical (black) and penalised (grey) solutions; shaded regions indicate solid domains Ω_s, where the discontinuous mask function is applied. Here, the penalisation parameter is chosen large enough to visualise the boundary layer inside the solid.

where the forcing term is defined as $f(x) = \sin(x)$ for $x \in \Omega_f$, which yields the analytic solution $u(x) = \sin(x)$. In the solid domain Ω_s, the forcing function is set to zero. The computational domain and boundary conditions are specified in the numerical setup below.

Figure 1 illustrates the exact and penalised solution for a representative parameter choice. The penalisation term effectively damps the velocity but induces an artificial boundary layer inside the solid domain.

3.1.1 Numerical setup

The computational domain spans $x \in [-\pi, 3\pi]$, consisting of the fluid region $\Omega_f \in]0, 2\pi[$ and solid regions on either side. The no-slip condition is imposed by setting the solid velocity $u_s(x) = 0$ (Eq. (9)).

The spatial convergence rate is studied through successive grid refinement using $N = 100$, 200, 400, 800, 1600 and 3200 grid points. Three different mask functions, Eqs. (3), (6), (5) are considered and evaluated over $K_\eta \in [10^{-2}, 10]$ (Eq. 4), sampled with 6000 logarithmically spaced points.

To analyse the effect of grid alignment, two configurations are considered. In the aligned case, the fluid–solid interface coincides with grid points. In the misaligned case, the penalised and reference solution are evaluated at offset positions, defined by $x^* = x + \Delta x/a$, with $a = 2$, 4, 8. This configuration is used to mimic the staircase effect arising in higher-dimensional Cartesian discretisations.

In their study, Hester et al. [1] proposed shifting the discontinuous mask to compensate for the displacement length error. This approach is also assessed in the present work. Although it reduces the absolute error, no improvement in convergence rate is observed. Furthermore, a discrete mask shift is only feasible when the penalisation boundary layer remains sufficiently resolved by the grid $K_\eta \in \mathbb{N}$. However, for stronger penalisation ($K_\eta < 1$), the resolution of the displacement length becomes smaller than

Table 1 Poisson equation (1D). Optimal penalisation parameters in the grid-aligned boundary study, determined at the highest grid resolution $N = 3200$.

	discontinuous	cosine	Hester's
$K_{\eta,\mathrm{opt}}$	0.2684	0.1898	0.2137

the grid spacing ($\epsilon = K_\eta \Delta x$). The required correction can therefore no longer be represented by a discrete shift of the mask. Consequently, this approach is not considered further.

3.1.2 Convergence Study: Influence of grid-alignment

Figure 2(a-b) shows the relative L^2-error as a function of the penalisation constant K_η. In all cases, the error decreases towards a distinct minimum at $K_{\eta,\mathrm{opt}}$ and increases again for both smaller and larger values of K_η. For $K_\eta > K_{\eta,\mathrm{opt}}$, the displacement length error dominates. In this regime, stronger penalisation reduces the total error, and clear differences between the mask functions become visible. The effect is most pronounced for Hester's mask, yielding a substantial reduction in the overall error. For $K_\eta < K_{\eta,\mathrm{opt}}$, the discretisation error becomes dominant, leading to an error saturation that is similar for all mask functions. Figure 2(b) further shows that the optimal penalisation parameter converges towards a resolution-independent value.

The convergence rate is evaluated at the value of K_η corresponding to the minimal error (Fig. 2(c-e)). This analysis is performed for both grid-aligned and misaligned configurations. For the aligned case (coloured lines), the grey shaded areas indicate the sensitivity of the error to small variations of the optimal parameter value, $K_{\eta,\mathrm{opt}} \pm 0.05$ (Tab. 1). Under optimal penalisation, all mask functions achieve a convergence rate of $\mathcal{O}(\Delta x^{1.5})$ and comparable minimum error levels. However, the sensitivity to deviations from $K_{\eta,\mathrm{opt}}$ clearly differs notably. For the discontinuous and cosine masks, even small deviations of the penalisation parameter lead to a pronounced reduction in convergence order. In contrast, Hester's mask exhibits a flatter error curve near to the minimum, indicating reduced sensitivity to imperfect parameter tuning. A slight loss of convergence is nevertheless observed for strongly sub-optimal choices of K_η.

For misaligned boundaries (black lines, dotted, dashed), clear differences between the tested mask functions emerge. As expected, the discontinuous mask drops to first-order convergence, with increasing error as the distance to the interface grows. The cosine mask exhibits non-monotonic behaviour. For a half-cell offset $x + \Delta x/2$, the error improves marginally (Fig. 2d), whereas for intermediate offsets the error increases and the convergence rate deteriorates. Hester's mask (Fig. 2e) remains the most robust across all tested offsets. It consistently maintains a convergence rate of $\mathcal{O}(\Delta x^{1.5})$, although a moderate increase in the error magnitude is observed for misaligned interfaces.

The 1D Poisson problem shows that the convergence behaviour is ultimately limited by the spatial discretisation once the penalisation error becomes sufficiently small.

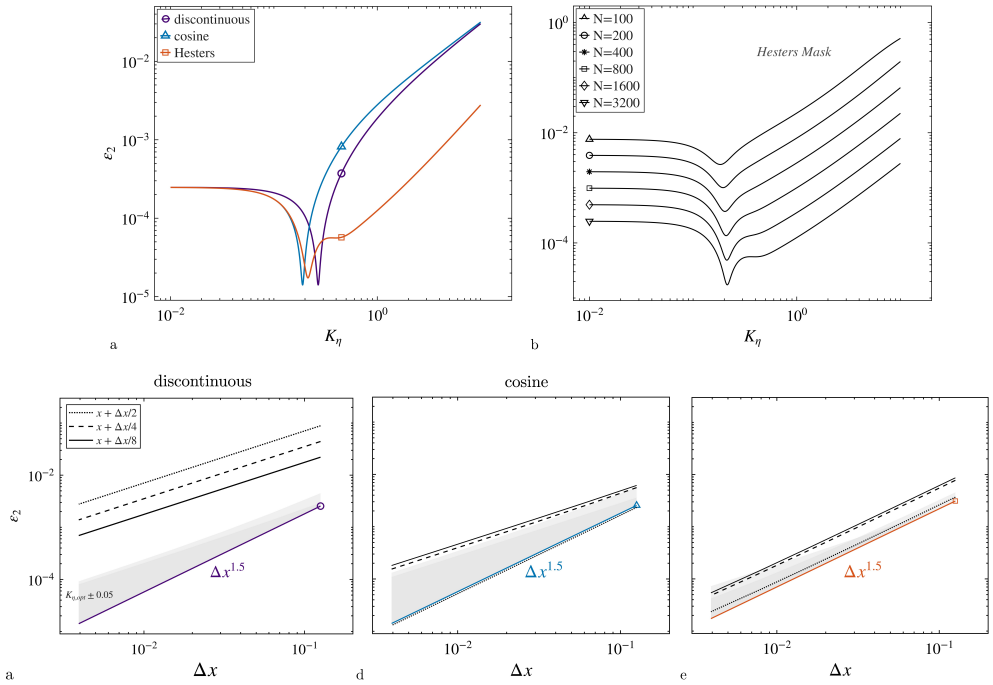


Fig. 2 Poisson equation (1D). Relative L^2-error ε_2 over K_η, Δx; (a) mask function, grid-aligned at $N = 3200$, (b) Hester's mask, grid-aligned for all grid resolutions, (c)-(e) grid-aligned boundaries: $K_{\eta,opt}$ with deviations ± 0.05 (grey shaded areas), not grid-aligned boundaries: grid shifts of $+\Delta x/2$ (dotted, black), $+\Delta x/4$ (dashed, black), $+\Delta x/8$ (solid, black)

Consequently, reducing the penalisation parameter alone does not lead to further improvements. The sensitivity to the penalisation parameter K_η differs significantly between the mask functions. While the discontinuous and cosine masks exhibit a substantial loss of accuracy for small deviations from the optimal parameter value, Hester's mask remains considerably more robust. This behaviour indicates that the optimised smooth mask function effectively reduces the displacement length error, even for not grid-aligned boundaries.

3.2 Flow past a cylinder

A 2D approximately incompressible flow past a cylinder at $Re = 40$ is considered, representing a canonical benchmark problem. Reference data is provided by the pseudo-spectral solution of Gautier et al. [43]. The configuration combines a simple geometry with a steady separated wake and is well suited to assess the accuracy of VPM. The governing equations are approximated using volume penalisation and artificial compressibility, as introduced in Eqs. (1) and (2). The equations are solved in non-dimensional form using the cylinder diameter and streamwise velocity as reference length and velocity scales. A constant density of $\rho = 1$ is assumed, and the kinematic viscosity is set to $\nu = 1/Re = 0.025$. For the artificial speed of sound, we

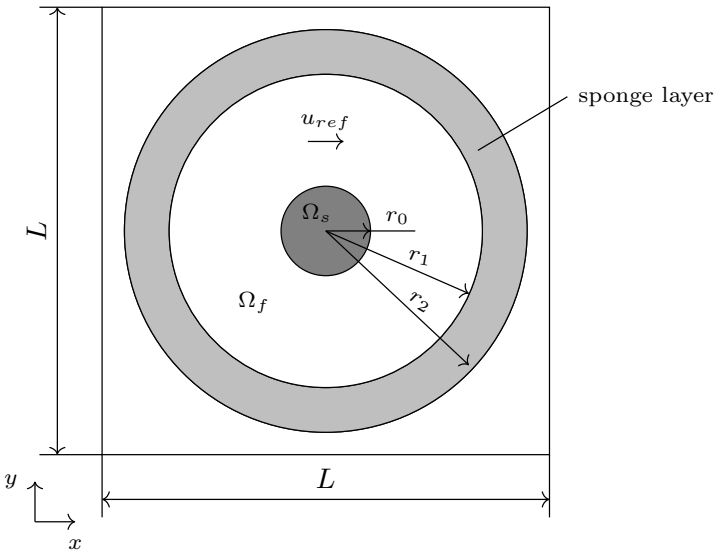


Fig. 3 Flow past a cylinder (2D). Computational domain with $L = 5$. The cylinder radius is $r_0 = 0.5$, and the sponge layer extends from $r_1 = 1.75$ to $r_2 = 2.25$. Inflow conditions are applied within the sponge layer, while the outer region enables periodic boundary conditions.

chose $C_0 = 10$. Higher values were also tested but did not produce any measurable improvement, indicating that artificial compressibility effects are already negligible at this value.

3.2.1 Numerical setup

Figure 3 shows the simulation setup. The computational domain is defined as a square with $L_x = L_y = L = 5$ with a cylinder of radius $r_0 = 0.5$ in the centre. Two concentric radii, $r_1 = 1.75$ and $r_2 = 2.25$, define the sponge layer that imposes inflow conditions from the reference solution [43], i.e. $u^{\mathrm{sp}}_{s,i} = u^{\mathrm{sp}}_{\mathrm{ref},i}$. A no-slip condition is enforced at the cylinder surface with $u^{\mathrm{cyl}}_{s,i} = 0$. Periodic conditions are applied at the outer domain boundaries.

Uniform grid resolutions with $N_x = N_y = N = 128$, 256, 512, 1024 and 2048 points are used in both spatial directions. The penalisation constant is sampled over $K_\eta \in [0.1, 2.0]$ with 20 uniformly spaced values. Parameter variations, including the three mask functions, are applied locally at the cylinder surface. In contrast, the inflow conditions are penalised in the sponge layer using fixed parameters, Hester's mask and $K_\eta = 0.2$. These were determined in a preliminary study and are kept constant throughout the present analysis to isolate the effect of penalisation at the cylinder surface.

Time integration is performed with Runge-Kutta schemes. For coarse grids ($N < 512$), the classical explicit RK4 method is used with CFL = 1. At the intermediate resolution $N = 512$, however, occasional instabilities are observed. These are likely

caused by transient peaks in the maximum velocity, which can lead to an overprediction of the allowable time step according to Eq. (8). To maintain a stable time integration, the CFL number is therefore reduced to CFL = 0.5.

For finer grids ($N = 1024, 2048$), the RKC method [41, 42] is employed with CFL = 0.5. At these resolutions, the penalisation and viscous terms become the dominant constraints on the time step size (Eq. (8)). RKC has an extended stability region along the real axis, making it well suited for such stiff problems. In order to ensure a stable time step, the eigenvalues of the system have to be estimated. Although RKC is formally second-order accurate, the selected time step remains sufficiently small to maintain the overall accuracy. The method permits significantly larger stable time steps. For the smallest K_η, a speed-up factor of up to 25 is obtained compared to RK4 at the same resolution. For coarser grids, however, a reliable estimate of the system's eigenvalues could not be obtained. In addition, the time-step restriction is primarily determined by the CFL condition, so that RKC provides little advantage over RK4 in this regime. Therefore, RK4 is used for the coarser resolutions.

3.2.2 Results

For a qualitative assessment, Fig. 4 shows the pointwise relative error distribution of the velocity component u_x within the fluid domain ($r_0 \leq \sqrt{x_i x_i} < r_1$), in the steady state. For each mask function, two representative cases are considered: strong penalisation ($K_\eta = 0.1$) and weak penalisation ($K_\eta = 2$). The colour map represents the decadic logarithm of the error. This enhances contrast across several orders of magnitude and reveals localised error structures. For readability, the colour bar is labelled by the corresponding error magnitudes. For comparison, the limits are fixed to $[10^{-7}, 10^{-1}]$. Values outside this interval are truncated in the colour scale to avoid distortion of the dominant error structures. The variations of the penalisation parameters at the cylinder surface have a pronounced effect on the error distribution. For weak penalisation, elevated error levels are observed in the near-wall region, which persist into the wake. Across all cases shown, Hester's mask exhibits lower error levels and a more confined near-wall error region. The discontinuous and cosine masks display broader regions of increased error magnitude. With fixed penalisation parameters in the sponge layer, the inflow conditions are accurately enforced for all cases.

The relative L^2-error is evaluated for the velocity components u_x, u_y and the pressure p (Fig. 5). Panel (a) shows the error as a function of K_η for the streamwise velocity component u_x at the highest resolution ($N = 2048$). For the cosine and Hester's masks, a minimum of the error is observed. In contrast, no distinct minimum is detected for the discontinuous mask within the tested parameter range.

Panels (b–d) present the spatial convergence at mask-specific optimal values of K_η. For the velocity component u_x, Hester's mask consistently yields the lowest error levels across all resolutions. At $K_\eta = 0.2$, a convergence rate of approximately $\mathcal{O}(\Delta x^{1.6})$ is obtained. The cosine and discontinuous masks exhibit convergence rates closer to first order. Similar trends are observed for the velocity component u_y, however, with slightly larger error values. The pressure field p exhibits higher error magnitudes overall. Nevertheless, Hester's mask again achieves the highest convergence rate, reaching approximately $\mathcal{O}(\Delta x^{1.8})$.

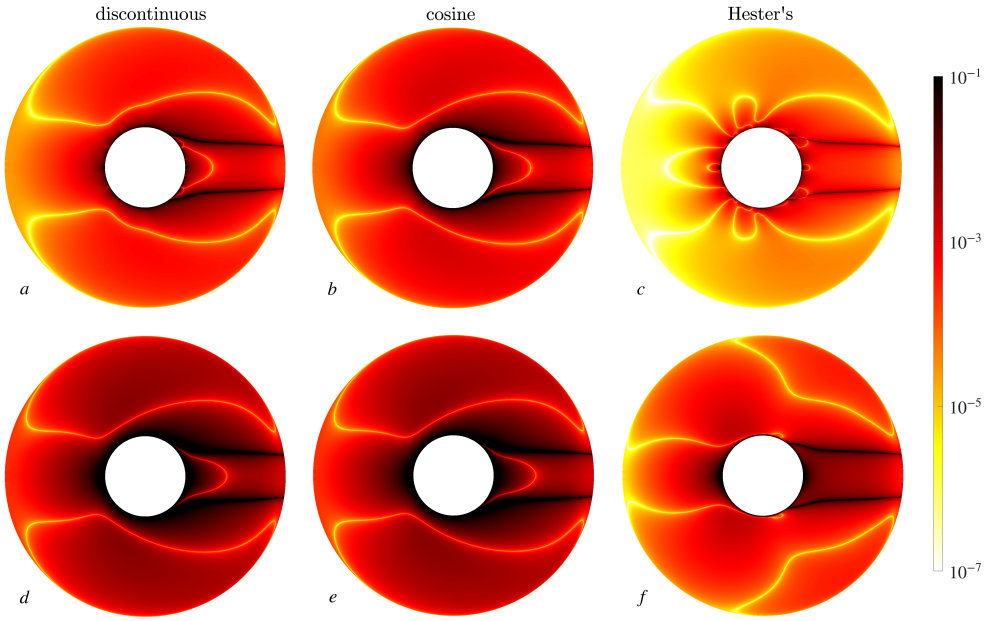


Fig. 4 Flow past a cylinder (2D). Decadic logarithm of the pointwise relative error of the velocity component u_x in the fluid domain surrounding the cylinder ($r_0 \leq \sqrt{x_i x_i} < r_1$). The colourbar labels indicate the corresponding error magnitudes in the range $[10^{-7}, 10^{-1}]$. (a)-(c) $K_\eta = 0.1$, (d)-(f) $K_\eta = 2.0$.

The 2D cylinder flow confirms the strong influence of the mask function on the accuracy and convergence behaviour of volume penalisation. Among the tested mask functions, Hester's mask provides the best overall performance and is less sensitive to the penalisation parameter. In contrast, the discontinuous mask does not exhibit a pronounced minimum or a clear saturation of the total error. This indicates that the optimal parameter value may lie between the sampled points or at even smaller values of K_η. Exploring the range $K_\eta \leq 0.1$ was computationally too expensive in the present study, due to the resulting small Δt. In realistic higher-dimensional simulations, the computational cost becomes increasingly significant, making such parameter values rarely practical.

Compared to the 1D Poisson problem, the measured convergence rates are slightly higher. This can be attributed to the spatial averaging within the fluid domain. In the 1D study, the grid misalignment directly affects the global error. In contrast, the curved boundary in the 2D cylinder flow introduces a distribution of local boundary offsets along the surface. The resulting error is therefore partially averaged out in the global L^2-measure. Nevertheless, the consistently superior performance of Hester's mask supports the concept of displacement-length compensation through optimised smooth mask transitions. The results demonstrate that the advantages reported in the continuous analysis by Hester et al. [1] remain observable for non-grid-aligned obstacles.

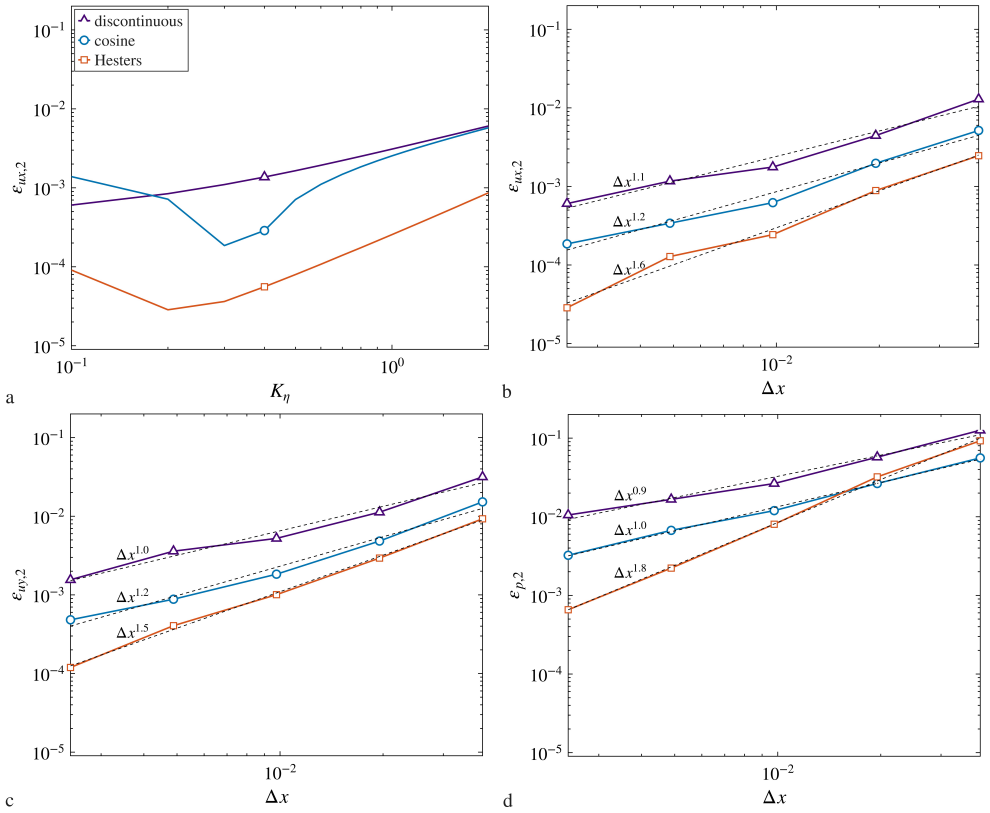


Fig. 5 Flow past a cylinder (2D). relative L^2-error $\varepsilon_{i,2}$ over K_η, Δx; (a) mask functions for u_x at $N = 2048$, (b)-(c) u_x, u_y: discontinuous: $K_\eta = 0.1$, cosine: $K_\eta = 0.3$, Hester's: $K_\eta = 0.2$; (d) p: discontinuous: $K_\eta = 0.7$, cosine: $K_\eta = 0.5$, Hester's: $K_\eta = 0.7$.

3.3 Turbulent Channel Flow

In the following, the analysis is extended to a 3D turbulent channel flow at a friction Reynolds number of $Re_\tau = u_\tau h/\nu = 180$, with friction velocity u_τ and channel half height $h = 1$. The reference data for this classical benchmark case were originally reported by Kim et al. [44]. More recently, Vreman and Kuerten [45, 46] presented an updated, higher-accuracy data set. The grid-aligned geometry is deliberately retained to separate the influence of the mask formulation from geometric misalignment effects and to assess the ACM-VPM framework in a fully developed turbulent flow. In contrast to the preceding laminar cases, the channel flow additionally provides strong near-wall gradients and statistically fluctuating quantities. The governing equations are given by the artificial compressibility formulation of the Navier–Stokes equations with volume penalisation, see Eqs. (1) and (2). A non-dimensional formulation is used based on the channel height and bulk velocity. The artificial speed of sound is fixed at $C_0 = 10$ ($Ma \approx 0.1$). At this Mach number, ACM provides good agreement with the reference data for the incompressible turbulent channel flow [36].

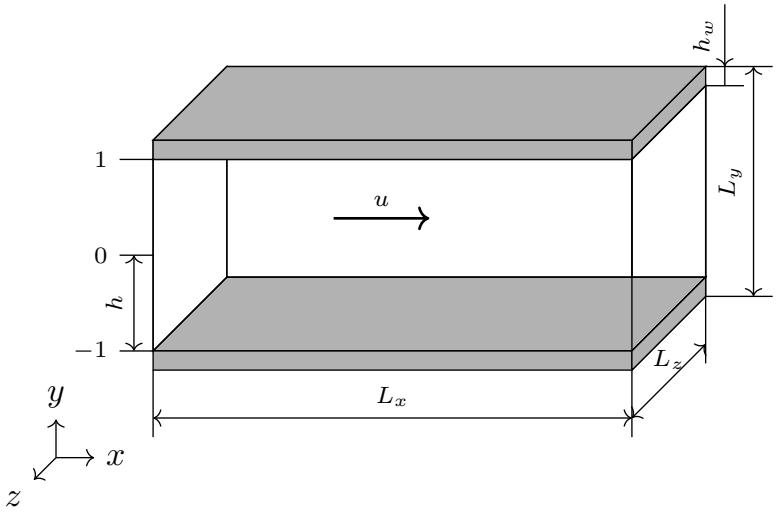


Fig. 6 Turbulent channel flow (3D). Computational domain of the turbulent channel flow with dimensions $L_x = 2\pi$, $L_y = 2.56$ and $L_z = 4\pi/3$. The two solid walls have thickness $h_w = 0.28$.

3.3.1 Numerical setup

The computational domain (Fig. 6) follows the configuration of Sawamura et al. [22] with dimensions $L_x = 2\pi$, $L_y = 2.56$, and $L_z = 4\pi/3$. The fluid region extends over $-1 < y < 1$. At the channel walls, no-slip boundary conditions are imposed using volume penalisation. The domain is extended in the wall-normal direction by solid layers of thickness $h_w = 0.28$. This choice ensures a grid-aligned discretisation of the boundary with $N_y = 512$ grid points. The streamwise and spanwise directions are resolved with $N_x = N_z = 192$ grid points. Periodic boundary conditions are applied in all spatial directions. Time integration is performed using the classical RK4 scheme. The time step size is determined according to Eq. (8) with a CFL number of 1.

Due to the high computational cost of the turbulent channel-flow simulations, $K_\eta = 0.2$ is selected as a representative case for the statistical flow analysis for all considered mask functions. Based on the previous test cases, this value was found to correspond to a favourable penalisation strength. Since there is no deterministic predictability in the turbulent flow regime, only statistics can be compared and a systematic grid-convergence study is not pursued for this case. Therefore, the comparison focuses on the statistical accuracy of the different mask formulations at DNS-level resolution. In addition, the spread in the reference data sets [45] precludes a reliable convergence analysis.

The kinematic viscosity was chosen such that $Re_m = U_m 2h/\nu = 5600$ following the reference configurations [22, 44]. Here, the bulk velocity is defined as $U_m = 1/2h \int_{-1}^{1} \overline{u}(y)\mathrm{d}y$. This yields a viscosity of $\nu = 2/5600$.

Initial condition

All simulations are initialised with the same fully developed turbulent flow field. This field is generated in a precursor simulation by evolving a perturbed laminar Poiseuille flow until a statistically stationary turbulent state is reached. The final snapshot of the turbulent field is then used as the initial condition for all subsequent tests and simulations.

In the precursor simulation, the laminar profile is perturbed using a divergence-free velocity field $\alpha_i = \alpha_i(x, y, z)$. It is constructed from the curl of a vector potential A_k such that

$$\alpha_i = \varepsilon_{ijk} \partial_j A_k, \tag{10}$$

which guarantees $\partial_i \alpha_i = 0$. The vector potential is defined component-wise by choosing $A_x = 0, A_y = f$, and $A_z = g$, where $f(x, z)$ and $g(x, y)$ denote spatially correlated random fields obtained from Gaussian-filtered noise within the fluid region. This yields the divergence-free components

$$\begin{aligned} \alpha_x &= \partial_y g - \partial_z f, \\ \alpha_y &= -\partial_x g, \\ \alpha_z &= \partial_x f. \end{aligned} \tag{11}$$

The streamwise velocity component is given by

$$u_x(x, y, z, t = 0) = u_c(1 - y^2) + \alpha_x, \quad y \in \Omega_f, \tag{12}$$

where $u_c = 1.5$ denotes the centreline velocity. The added perturbations have an amplitude of order 10^{-2} relative to u_c. To remain consistent with the test case [44], the profile is normalised such that the bulk velocity satisfies $U_m = 1$. The transverse velocity components are initialised using α_y and α_z, while the pressure field is set to zero. The flow was evolved until the total shear-stress profile matched the characteristic linear total shear-stress profile across the channel [47, p. 269].

Definition of Flow Statistics

Flow statistics are evaluated using the reference data set S3 by Vreman and Kuerten [46]. After restarting with the fully developed turbulent field, the simulation is advanced for approximately 48 wash-through times. This allows transient effects to decay before statistical averaging is started. A wash-through time is defined as $T_w = L_x/U_m$. Subsequently, statistics are accumulated over approximately 280 wash-through times. Mean values and fluctuations are collected using an online average during the simulation, resulting in about 2.5×10^6 samples.

The mean streamwise velocity profile is defined as

$$\overline{u}_x(y) = \frac{1}{T L_x L_z} \int_0^T \int_0^{L_x} \int_0^{L_z} u_x(x, y, z, t) \, \mathrm{d}z \, \mathrm{d}x \, \mathrm{d}t, \tag{13}$$

where T denotes the averaging interval.

Root-mean-square (RMS) fluctuations are defined relative to the mean velocity, e.g. $u_{x,\mathrm{RMS}}(y) = \sqrt{\overline{(u_x - \overline{u}_x)^2}}$, with analogous expressions for u_y, u_z and p. Flow

statistics are expressed in wall units using the friction velocity u_τ. It is defined with the wall shear stress [47, p. 275], for example for the lower wall:

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}}, \quad \text{with} \quad \tau_w = \nu\rho \left.\frac{\partial \overline{u}_x}{\partial y}\right|_{y=-1}, \tag{14}$$

where the density $\rho = 1$. Quantities normalised by the viscous scales are indicated with superscript plus, e.g. $\overline{u}_x^+ = \overline{u}_x/u_\tau$, $u^+_{\mathrm{RMS},x} = u_{\mathrm{RMS},x}/u_\tau$, $p^+_{\mathrm{RMS}} = p_{\mathrm{RMS}}/u_\tau^2$ and $y^+ = u_\tau(y+1)/\nu$. The spatial resolution in wall units is given by $\Delta x^+ = 5.4$, $\Delta y^+ = 0.8$, and $\Delta z^+ = 3.6$, which is consistent with a typical resolution of a direct numerical simulation (DNS) of turbulent channel flow [45].

Channel Forcing

At each time step, a constant mass flux rate (CFR) condition is enforced. This is achieved by removing the temporal change of the bulk velocity from the streamwise momentum equation, thereby maintaining a constant bulk velocity.

$$\left(\frac{\partial u_x}{\partial t}\right)^{n,\mathrm{corr}} = \left(\frac{\partial u_x}{\partial t}\right)^{n} - \frac{1}{V_{\mathrm{fluid}}}\int_{\Omega_f}\left(\frac{\partial u_x}{\partial t}\right)^{n} \mathrm{d}V. \tag{15}$$

As a result, the bulk velocity remains constant throughout the simulation. For the numerical evaluation of the volume integral, a second-order trapezoidal rule is employed.

3.3.2 Results

Figure 7a-b shows the mean streamwise velocity and pressure profile in wall units, $\overline{u}_x^+(y^+)$. The discontinuous and Hester's masks reproduce the reference profile well, particularly in the near-wall region, with only minor differences towards the channel centre. In contrast, the cosine mask exhibits substantial deviations throughout the domain.

Figure 7c-f presents the RMS velocity and pressure fluctuations. The cosine mask overpredicts the fluctuation amplitudes throughout the channel and shows the largest discrepancies. The deviations are most pronounced in the near-wall peak regions and persist towards the channel centre. In contrast, the discontinuous and Hester's masks remain close to the reference data, with only minor differences near the fluctuation peaks. Hester's mask shows the closest agreement for the pressure fluctuations and the wall-normal velocity component.

At the selected value of $K_\eta = 0.2$, the discontinuous and the Hester's mask reproduce both the mean flow and RMS fluctuations in good agreement with the reference data, whereas the cosine mask shows larger deviations. Since the channel walls are grid-aligned, this configuration provides a controlled setting in which geometric misalignment effects are absent, allowing the mask formulations to be assessed in fully developed turbulent flow. The results support the applicability of the ACM–VPM framework to 3D wall-bounded turbulence. The subsequent flapping-wing case complements this configuration by introducing a moving geometry.

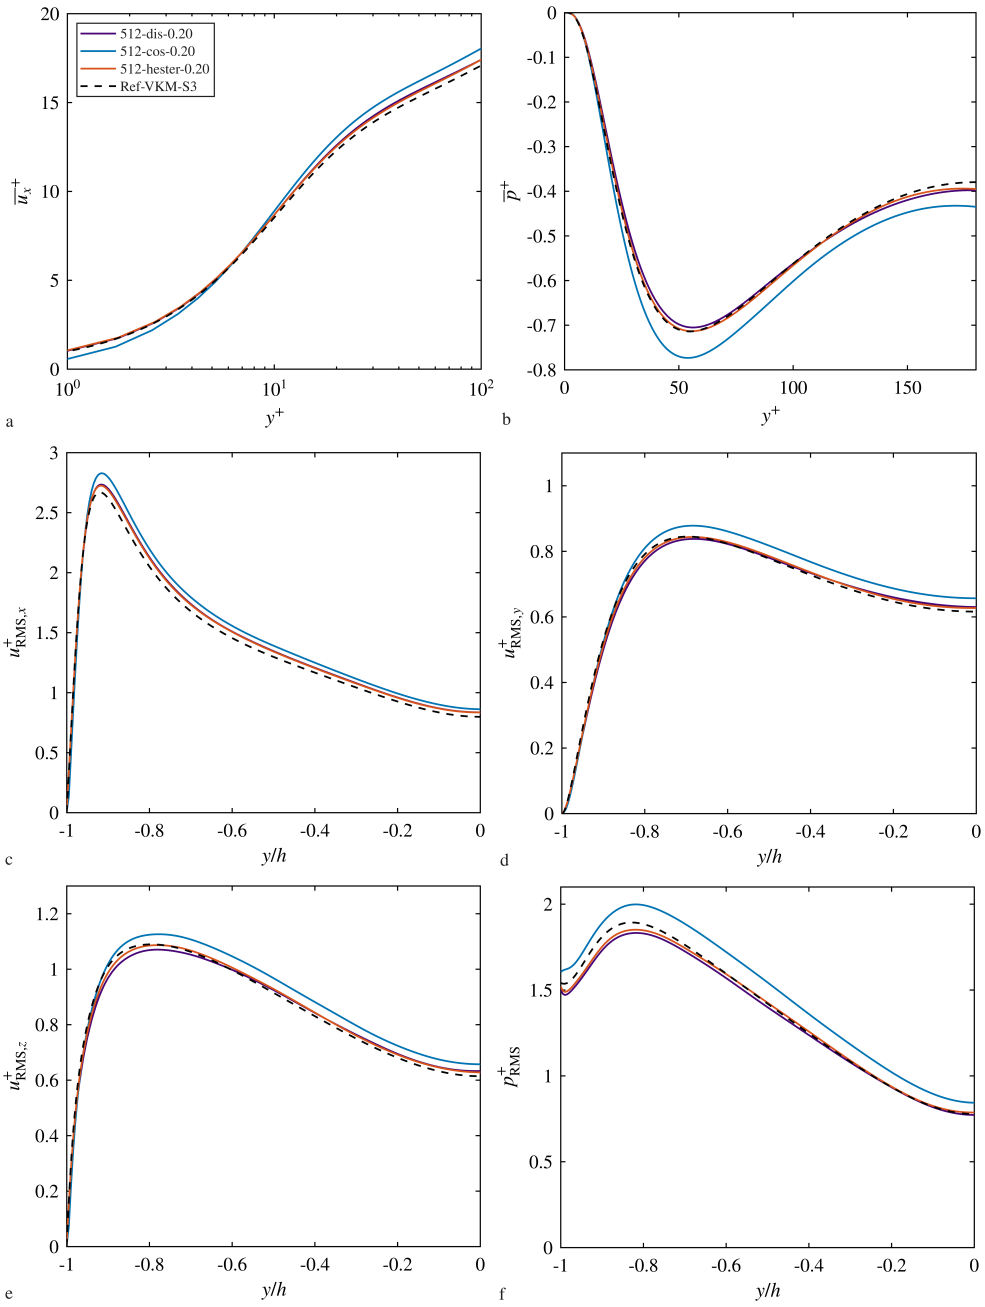


Fig. 7 Turbulent channel flow (3D). (a)-(b) Dimensionless mean flow values $\overline{u_x}^+$ and $\overline{p}^+$ over y^+ for the three mask functions: discontinuous, cosine, Hester's, $K_\eta = 0.2$. (c)-(f) Dimensionless fluctuation RMS values $u^+_{\mathrm{RMS},x}$, $u^+_{\mathrm{RMS},y}$, $u^+_{\mathrm{RMS},z}$ and p^+_{RMS} over dimensionless channel hight y/h for the three mask functions: discontinuous, cosine, Hester's, $K_\eta = 0.2$.

3.4 Flapping wing

As a final flow configuration we consider the flapping-wing test case reported by Suzuki et al. [48][Appendix B2]. It consists of a single rigid rectangular wing of finite thickness undergoing prescribed flapping and pitching motion. Wing movement is performed in an initially quiescent fluid to mimic hovering flight. The Reynolds number $Re = u_{\mathrm{tip}} c_m/\nu = 100$ is defined using the mean wing tip velocity u_{tip} and the chord length c_m. This test case has been used for validation in several previous studies by Engels et al. [26, 29], Dilek et al. [49], Medina et al. [50] and Huang et al. [51]. We therefore limit the description of the test case to a minimum, further details of geometry and wingbeat kinematics can be found for instance in [29]. To our knowledge, the present study extends these previous investigations by assessing finer spatial resolutions than those reported so far.

3.4.1 Numerical setup

The single flapping wing is set to unit span length $R = 1$ with finite thickness $h_{\mathrm{wing}} = 0.04171$. The wing moves in a horizontal stroke plane (see Fig. 8a). For this test case, the geometry handling in WABBIT is extended to planar polygonal wing shapes of finite thickness using the signed-distance representation described in B. We have a computational domain of size $L_x = L_y = L_z = 6$ with homogeneous Dirichlet conditions $u_{\infty,i} = 0$ and $p_\infty = 0$ at the outer boundary imposed by a sponge with thickness $h_{\mathrm{sponge}} = 1/2$.

Previous studies of VPM with moving obstacles have shown that discontinuous mask functions can lead to spurious oscillations in the hydrodynamic forces as the obstacle moves relative to the Cartesian grid, whereas smoothing the mask substantially reduces these oscillations [20]. Therefore, the discontinuous mask is not considered for this moving-wing configuration. The penalisation constant is set to $K_\eta = 0.3$. In the adaptive setting, the solver ensures that the highest resolution is maintained at the mask function. We select the biorthogonal lifted Cohen–Daubechies–Feauveau (CDF) wavelet [52], CDF 4/2. It ensures a sufficient balance between computational costs and accuracy. The wavelet threshold C_ε, as well as the artificial compressibility parameter C_0 are scaled based on convergence properties to guarantee at least first order convergence of all terms, $C_\varepsilon \sim \Delta x$ and $C_0 \sim \Delta x^{-1/2}$. This ensures a constant numerical damping over the varying grid resolution. Engels et al. [26] derived the more conservative scaling $C_\varepsilon \sim \Delta x^2$ based on an upper bound for the accumulated thresholding error, assuming an error contribution of order $\mathcal{O}(C_\varepsilon)$ at each time step. In practical adaptive simulations, however, only a fraction of the grid is typically affected by coarsening at a given time step, such that the actual accumulation of the thresholding error is expected to be smaller. Additional tests indicate that C_ε does not influence the solution. Table 2 summarises the parameters used for the different grid resolutions, note that a constant block size of $Bs = 24$ is used.

3.4.2 Results

To assess grid sensitivity, a self-convergence study is performed by comparing the lift and drag coefficients C_L and C_D over $3 \leq t \leq 4$. They are defined from the vertical

Table 2 Flapping wing (3D). Numerical parameters for the different grid resolutions.

	$J_{\max}$	Δx	$h_{\text{wing}}/\Delta x$	C_ε	C_0
Coarse	5	1/128	5.34	$1.6 \cdot 10^{-3}$	25
Medium	6	1/256	10.68	$8 \cdot 10^{-4}$	35.36
Fine	7	1/512	21.36	$4 \cdot 10^{-4}$	50

and horizontal force components, respectively. The solution with fine resolution (Tab. 2) serves as a reference for both mask functions.

For the cosine mask, the relative L^2-difference decreases from 3.09% to 0.77% for C_L and from 3.48% to 0.98% for C_D, when increasing the resolution from coarse to medium. The Hester mask exhibits a similar reduction from 3.19% to 0.78% for C_L and from 3.15% to 0.94% for C_D. The relative differences of less than 1% between the two finest resolutions indicate that $\Delta x = 1/512$ provides sufficient spatial resolution for the subsequent benchmark comparison. The results indicate a nearly identical grid sensitivity for both mask functions. A more detailed analysis of the self-convergence is provided in A.

Figure 8 compares the lift and drag coefficients obtained with the Hester mask at the fine resolution with previously published results. Overall, the present results fall within the range of the published solutions and reproduce well the characteristic force evolution throughout the wingbeat. In particular, the lift coefficient agrees closely with Suzuki et al. [48], Medina [50], and Dilek et al. [49] over most of the cycle. This agreement provides further confidence in the present simulations. At the same time, a noticeable spread exists among the published data themselves, most prominently in the lift coefficient during and shortly after stroke reversal and in the magnitude of the lift maxima. Huang et al. [51], for example, predict systematically higher lift peaks. The drag coefficients show a smaller spread around stroke reversal, while larger differences occur near the positive and negative drag extrema. During stroke reversal, the wing also pitches rapidly and changes its direction of translation, making this phase particularly sensitive to the numerical treatment of the moving boundary and the evaluation of aerodynamic forces. The minor changes with increasing grid resolution and the nearly identical convergence behaviour of both mask functions support the numerical consistency of the present results within the spread of the published solutions.

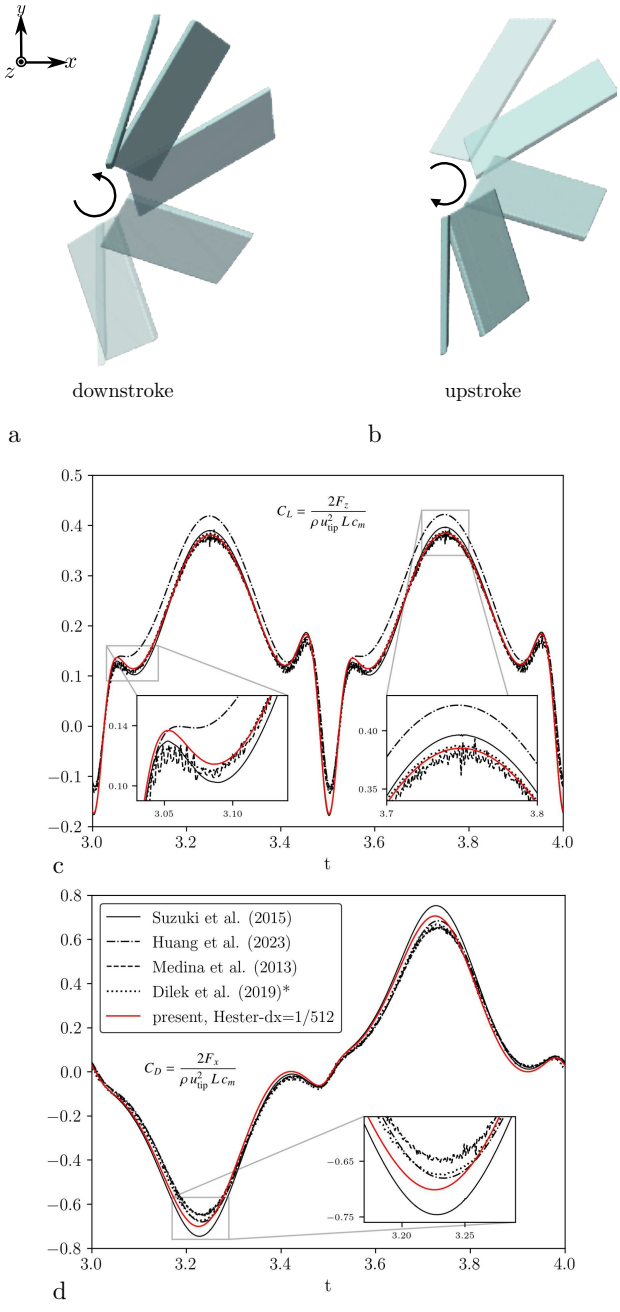


Fig. 8 Flapping wing (3D). (a-b) Wing movement, (c-d) comparison of different published solutions for the flapping-wing test case. *The data from Dilek et al. [49] were digitised from the published figure and therefore include a digitisation error.

4 Conclusion

We assessed the volume penalisation method (VPM) in combination with the artificial compressibility method (ACM) for 1D, 2D, and 3D benchmark problems. The study examines modelling errors associated with penalisation from discretisation errors introduced by the underlying Cartesian grid.

The 1D Poisson problem and the 2D cylinder flow show that the influence of the mask formulation is particularly pronounced when penalisation and geometric discretisation errors dominate. In these configurations, Hester's optimised mask exhibits reduced sensitivity to the penalisation parameter and grid alignment. For the non-grid-aligned cylinder, convergence rates of up to $O(\Delta x^{1.6})$ are obtained for the velocity field, demonstrating that the favourable properties identified in continuous analysis remain observable in a discrete multidimensional configuration.

The 3D turbulent channel flow provides a complementary assessment in a grid-aligned, fully developed turbulent regime. At the selected penalisation parameter, mean velocity profiles and turbulence statistics are reproduced in good agreement with reference data. This demonstrates that the ACM–VPM framework can reproduce both mean and fluctuating quantities in wall-bounded turbulence without introducing geometric misalignment effects. The flapping-wing test case extends the analysis to a moving 3D geometry. The two smooth mask functions exhibit nearly identical grid sensitivity, and the aerodynamic forces obtained at the finest common resolution lie within the range of previously published solutions. In this more complex configuration, the differences between the smooth mask formulations are substantially less pronounced than in the controlled 1D and 2D tests.

Overall, the results show that the influence of the mask formulation is strongest when penalisation and grid-alignment errors dominate. Future work should focus on bridging the remaining gap between continuous penalisation theory and discrete numerical implementations. Possible approaches include active penalisation, a velocity extension technique, and continued investigation of boundary interpolation techniques designed to reduce the geometrical staircase effect.

Acknowledgements. The authors acknowledge the computational and storage resources provided by GENCI at IDRIS and TGCC (allocations 2025-A0182A14152 and A0202A14152; www.genci.fr). Computational resources were allocated on the Jean Zay supercomputer at IDRIS (www.idris.fr) and on the Rome partition of the IRENE supercomputer at TGCC (www-hpc.cea.fr).

The authors also acknowledge access to the HPC resources of the Centre de Calcul Intensif d'Aix-Marseille Université.

The authors gratefully acknowledge the computing time made available on the high-performance computer "Lise" at the NHR Centre NHR@ZIB. This centre is jointly supported by the Federal Ministry of Education and Research and the state governments participating in the NHR (www.nhr-verein.de/unsere-partner).

The authors gratefully acknowledge K. Suzuki, R. Huang, and A. Medina for providing reference data used in this study.

Declarations

Funding

This work was supported by the French-German project "CVSFIT", funded by the ANR (grant ANR-23-CE46-0014) and DFG (grant 529979103).

KS acknowledges funding from the French Federation for Magnetic Fusion Studies (FR-FCM) and the EUROfusion Consortium, funded by the Euratom Research and Training Programme under Grant Agreement No. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

Appendix A Flapping flight: auto-comparison

Figure A1 provides a more detailed comparison of the lift and drag coefficients at the investigated spatial resolutions. The upper panels show that the solutions largely overlap over the entire wingbeat, while the lower panels show the pointwise differences with respect to the Hester mask with $\Delta x = 1/512$. The largest deviations occur for the coarse resolution. For C_L, these are particularly pronounced near the stroke reversals, where the lift changes rapidly, whereas the largest deviations in C_D occur around the drag extrema. These deviations are substantially reduced at the medium resolution. In contrast, the solutions at the fine resolution obtained with the cosine and Hester masks are nearly indistinguishable. Moreover, the cosine and the Hester masks exhibit very similar deviations, suggesting that the observed differences are primarily associated with spatial resolution rather than the choice of the mask function.

Appendix B Signed-distance representation of polygonal wing shapes

The wing geometry is represented in the local wing coordinate system by a closed polygonal contour with vertices $\mathbf{p}_i = (p_{i,x}, p_{i,y})$, $i = 1, \ldots, N$, where the final edge connects $\mathbf{p}_N$ back to $\mathbf{p}_1$. The transformation between the global, body, and wing coordinate systems (Fig. B2) are described in detail by Engels et al. [29]. The implementation is restricted to planar polygonal wings of constant thickness, i.e. extrusion of a 2D polygon.

To reduce computational cost, the signed distance function is only computed within a narrow band around the wing contour. A bounding box enclosing the complete wing is first used to identify the computational blocks that can be affected by the wing. The block structure of the computational grid is used to exclude blocks outside the wing region, while additional bounding boxes around the individual polygon segments restrict the subsequent point-to-segment distance calculations.

For the remaining grid points, the coordinates are transformed into the local wing frame. The two-dimensional distance in the wing plane is obtained from the minimum distance between the projected point $\mathbf{x}_{w,xy} = (x_w, y_w)$ and the edges of the polygon. For the i-th segment connecting $\mathbf{p}_i$ and $\mathbf{p}_{i+1}$, with $\mathbf{p}_{N+1} = \mathbf{p}_1$, we define $\mathbf{v}_i = \mathbf{p}_{i+1} - \mathbf{p}_i$ and $\mathbf{w}_i = \mathbf{x}_{w,xy} - \mathbf{p}_i$.

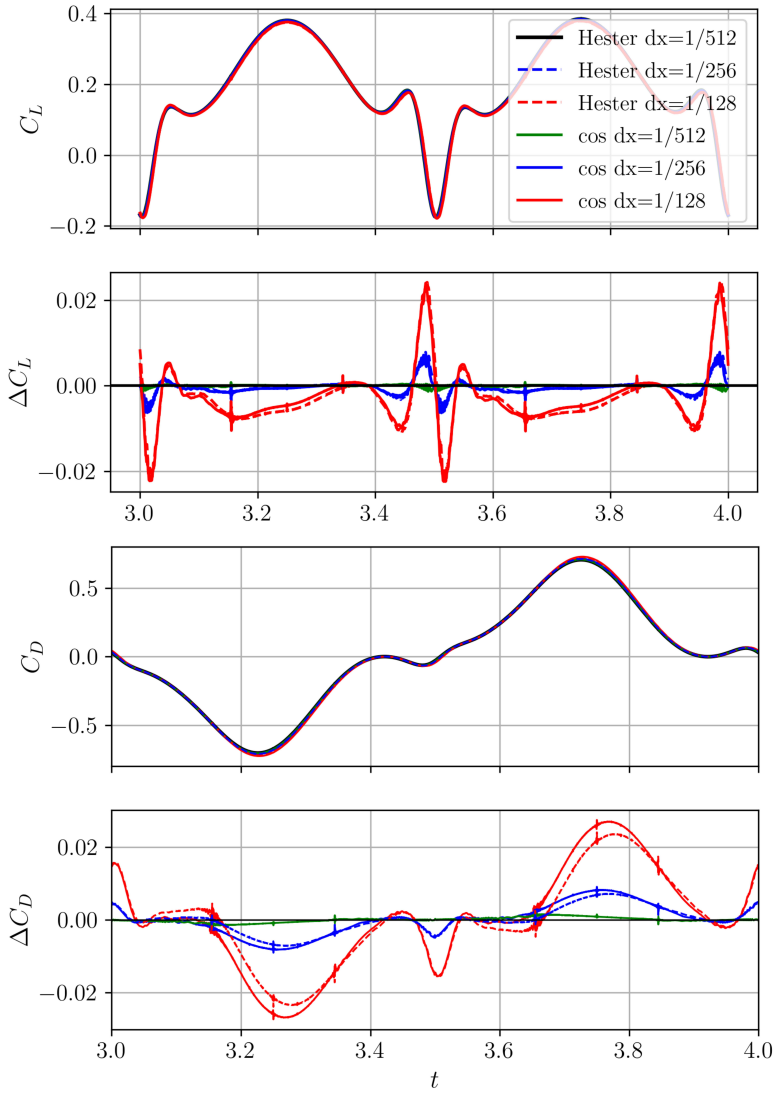


Fig. A1 Flapping wing (3D). Self-comparison of lift and drag coefficients at different spatial resolutions for the cosine and Hester masks over one wingbeat. The corresponding pointwise differences ΔC_L and ΔC_D relative to Hester at $\Delta x = 1/512$.

The closest point on the finite segment is obtained by the clamped projection

$$t_i = \min\left(1, \max\left(0, \frac{\mathbf{w}_i \cdot \mathbf{v}_i}{\mathbf{v}_i \cdot \mathbf{v}_i}\right)\right), \tag{B1}$$

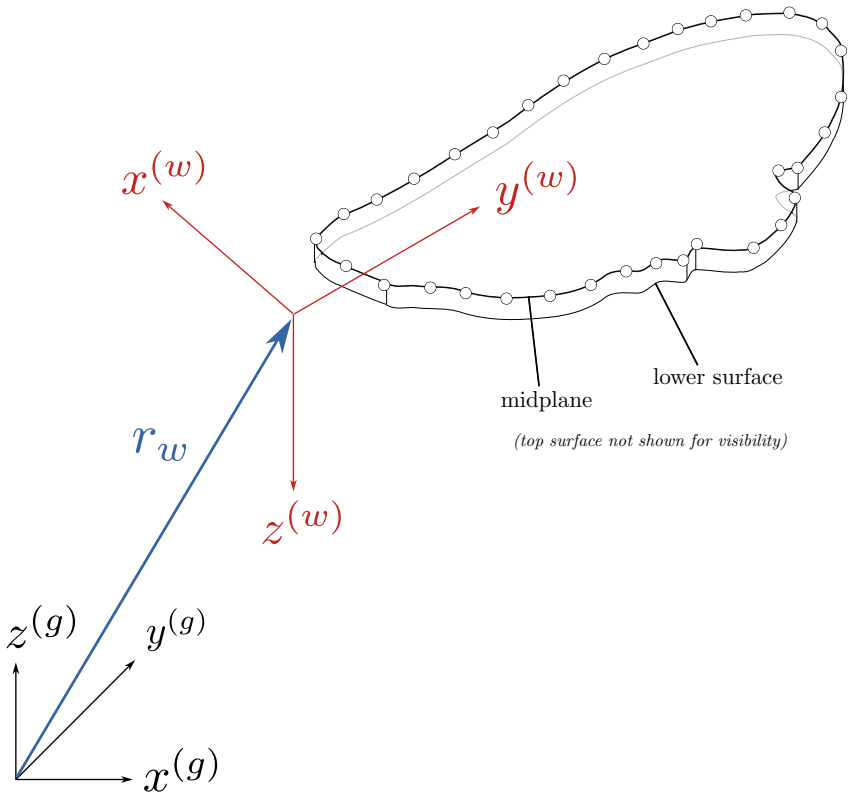


Fig. B2 Flapping wing (3D). Schematic representation of the transformation from the global coordinate system (x, y, z) to wing coordinate system (x_w, y_w, z_w).

such that $\mathbf{x}_{\text{closest},i} = \mathbf{p}_i + t_i \mathbf{v}_i$. The corresponding point-to-segment distance is

$$d_{\text{seg},i} = \|\mathbf{x}_{w,xy} - \mathbf{x}_{\text{closest},i}\|_2 \,. \tag{B2}$$

The two-dimensional unsigned distance to the polygon boundary is then

$$d_{xy} = \min_{i=1,\ldots,N} d_{\text{seg},i}. \tag{B3}$$

Degenerate segments with coincident endpoints are handled by falling back to the distance to the segment endpoint.

To assign the sign of the two-dimensional distance, a ray-crossing point-in-polygon test is used. A horizontal ray is cast in the positive x-direction from the query point. Each crossing of a polygon edge toggles the inside state; an odd number of crossings classifies the point as inside the polygon, and an even number as outside.

The signed two-dimensional distance is therefore defined as

$$\phi_{xy} = \begin{cases} -d_{xy}, & \mathbf{x}_{w,xy} \in \Omega_s, \\ 0, & \mathbf{x}_{w,xy} \in \partial\Omega_s, \\ d_{xy}, & \mathbf{x}_{w,xy} \in \Omega_f, \end{cases} \tag{B4}$$

where Ω_s and Ω_f denote the solid and fluid regions in the local wing plane, respectively, with $\partial\Omega_s$ representing the polygon boundary. Distances exceeding the prescribed narrow-band width are truncated accordingly. Points outside this region are excluded from further evaluation.

The finite wing thickness is incorporated by considering the interval $z_{\min} \leq z_w \leq z_{\max}$ in the local wing coordinate system. The signed distance to this slab is

$$\phi_z = \max\left(z_{\min} - z_w,\, z_w - z_{\max}\right). \tag{B5}$$

Thus, $\phi_z \leq 0$ inside the wing and $\phi_z > 0$ outside it.

The full 3D wing is the intersection of the extruded polygon in the xy-plane and the finite thickness interval in z. Combining these two gives

$$\phi = \begin{cases} \max(\phi_{xy}, \phi_z), & \phi_{xy} \leq 0,\ \phi_z \leq 0, \\ \phi_{xy}, & \phi_{xy} > 0,\ \phi_z \leq 0, \\ \phi_z, & \phi_{xy} \leq 0,\ \phi_z > 0, \\ \sqrt{\phi_{xy}^2 + \phi_z^2}, & \phi_{xy} > 0,\ \phi_z > 0. \end{cases} \tag{B6}$$

The two distance components describe the position relative to the wing in different directions: ϕ_{xy} gives the signed distance to the polygon boundary in the wing plane, while ϕ_z gives the signed distance in the thickness direction. For points inside the wing, the distance to the nearest surface is used. For points outside the wing in both directions, the two distances are combined using the Euclidean norm. The signed distance field is then converted into a smoothed mask function.

References

[1] Hester, E., Vasil, G., Burns, K.: Improving accuracy of volume penalised fluid-solid interactions. Journal of Computational Physics (2020) https://doi.org/10.1016/j.jcp.2020.110043

[2] Mittal, R., Iaccarino, G.: Immersed boundary methods. Annual Review of Fluid Mechanics **37**, 239–261 (2005) https://doi.org/10.1146/annurev.fluid.37.061903.175743

[3] Peskin, C.: Flow patterns around heart valves: A numerical method. Journal of Computational Physics **10**(2), 252–271 (1972) https://doi.org/10.1016/0021-9991(72)90065-4

[4] McQueen, D.M., Peskin, C.S.: Shared-memory parallel vector implementation of the immersed boundary method for the computation of blood flow in the beating mammalian heart. The Journal of Supercomputing **11**, 213–236 (1997) https://doi.org/10.1023/A:1007951707260

[5] Peskin, C.S.: The immersed boundary method. Acta Numerica **11**, 479–517 (2002) https://doi.org/10.1017/S0962492902000077

[6] Lee, S., Jung, E.: A two-chamber model of valveless pumping using the immersed boundary method. Applied Mathematics and Computation **206**(2), 876–884 (2008) https://doi.org/10.1016/j.amc.2008.09.047

[7] Mittal, R., Dong, H., Bozkurttas, M., Lauder, G.V., Madden, P.: Locomotion with flexible propulsors: II. computational modeling of pectoral fin swimming in sunfish. Bioinspiration & Biomimetics **1**(4), 35–41 (2006) https://doi.org/10.1088/1748-3182/1/4/S05

[8] Iaccarino, G., Verzicco, R.: Immersed boundary technique for turbulent flow simulations. Applied Mechanics Reviews **56**(3), 331–347 (2003) https://doi.org/10.1115/1.1563627

[9] Schneider, K.: Immersed boundary methods for numerical simulation of confined fluid and plasma turbulence in complex geometries: A review. Journal of Plasma Physics **81**(6), 435810601 (2015) https://doi.org/10.1017/S0022377815000598

[10] Kim, W., Choi, H.: Immersed boundary methods for fluid-structure interaction: A review. International Journal of Heat and Fluid Flow **75**, 301–309 (2019) https://doi.org/10.1016/j.ijheatfluidflow.2019.01.010

[11] Sotiropoulos, F., Yang, X.: Immersed boundary methods for simulating fluid–structure interaction. Progress in Aerospace Sciences **65**, 1–21 (2014) https://doi.org/10.1016/j.paerosci.2013.09.003

[12] Uhlmann, M.: An immersed boundary method with direct forcing for the simulation of particulate flows. Journal of Computational Physics **209**(2), 448–476 (2005) https://doi.org/10.1016/j.jcp.2005.03.017

[13] Angot, P., Bruneau, C.-H., Fabrie, P.: A penalization method to take into account obstacles in incompressible viscous flows. Numerische Mathematik **81**, 497–520 (1999) https://doi.org/10.1007/s002110050401

[14] Mittal, R., Dong, H., Bozkurttas, M., Najjar, F.M., Vargas, A., Loebbecke, A.: A versatile sharp interface immersed boundary method for incompressible flows with complex boundaries. Journal of Computational Physics **227**, 4825–4852 (2008) https://doi.org/10.1016/j.jcp.2008.01.028

[15] Ji, X., Shen, C.N., Rees, W.M.: A Fourth Order Sharp Immersed Method for the

Incompressible Navier–Stokes Equations with Stationary and Moving Boundaries and Interfaces (2025). https://doi.org/10.48550/arXiv.2508.15083

[16] Seo, J.H., Mittal, R.: A sharp-interface immersed boundary method with improved mass conservation and reduced spurious pressure oscillations. Journal of Computational Physics **230**(19), 7347–7363 (2011) https://doi.org/10.1016/j.jcp.2011.06.003

[17] Verzicco, R.: Immersed boundary methods: Historical perspective and future outlook. Annual Review of Fluid Mechanics **55**, 129–155 (2023) https://doi.org/10.1146/annurev-fluid-120720-022129

[18] Carbou, G., Fabrie, P.: Boundary layer for a penalization method for viscous incompressible flow. Advances in Differential Equations **8**(12), 1453–1480 (2003)

[19] Kevlahan, N.K.R., Ghidaglia, J.M.: Computation of turbulent flow past an array of cylinders using a spectral method with Brinkman penalization. European Journal of Mechanics - B/Fluids **20**(3), 333–350 (2001) https://doi.org/10.1016/S0997-7546(00)01121-3

[20] Kolomenskiy, D., Schneider, K.: A Fourier spectral method for the Navier–Stokes equations with volume penalization for moving solid obstacles. Journal of Computational Physics **228**(16), 5687–5709 (2009) https://doi.org/10.1016/j.jcp.2009.04.027

[21] Engels, T., Kolomenskiy, D., Schneider, K., Sesterhenn, J.: Numerical simulation of fluid–structure interaction with the volume penalization method. Journal of Computational Physics **281**, 96–115 (2015) https://doi.org/10.1016/j.jcp.2014.10.005

[22] Sawamura, Y., Yoshimatsu, K., Schneider, K.: Numerical simulation of wall-bounded flows using a spectral method with volume penalization. ESAIM: Proceedings and Surveys **63**, 280–289 (2018) https://doi.org/10.1051/proc/201863280

[23] Liu, Q., Vasilyev, O.V.: A brinkman penalization method for compressible flows in complex geometries. Journal of Computational Physics **227**(2), 946–966 (2007) https://doi.org/10.1016/j.jcp.2007.07.037

[24] Reiss, J.: Pressure-tight and non-stiff volume penalization for compressible flows: An immersed boundary method with good conservation properties. Journal of Scientific Computing **90**(3), 86 (2022) https://doi.org/10.1007/s10915-021-01747-x

[25] Schubert, Y., Sarradj, E., Lemke, M.: A volume penalization method for acoustic and aeroacoustic simulation of solid and porous bluff bodies. Computers & Fluids **278**, 106560 (2025) https://doi.org/10.1016/j.compfluid.2025.106560

[26] Engels, T., Schneider, K., Reiss, J., Farge, M.: A wavelet-adaptive method for multiscale simulation of turbulent flows in flying insects. Communications in Computational Physics **30(4)**, 1118–1149 (2021) https://doi.org/10.4208/cicp.OA-2020-0246

[27] Kadoch, B., Kolomenskiy, D., Angot, P., Schneider, K.: A volume penalization method for incompressible flows and scalar advection–diffusion with moving obstacles. Journal of Computational Physics **231**(12), 4365–4383 (2012) https://doi.org/10.1016/j.jcp.2012.01.036

[28] Bergmann, M., Iollo, A.: Modeling and simulation of fish-like swimming. Journal of Computational Physics **230**(2), 329–348 (2011) https://doi.org/10.1016/j.jcp.2010.09.017

[29] Engels, T., Kolomenskiy, D., Schneider, K., Sesterhenn, J.: Flusi: A novel parallel simulation tool for flapping insect flight using a Fourier method with volume penalization. SIAM Journal on Scientific Computing **38**, 3–24 (2016) https://doi.org/10.1137/15M1026006

[30] Shirokoff, D., Nave, J.-C.: A sharp-interface active penalty method for the incompressible Navier–Stokes equations. Journal of Computational Physics **62**, 53–77 (2015) https://doi.org/10.1007/s10915-014-9849-6

[31] Kou, J., Joshi, S., Hurtado-de-Mendoza, A., Puri, K., Hirsch, C., Ferrer, E.: Immersed boundary method for high-order flux reconstruction based on volume penalization. Journal of Computational Physics **448**, 110721 (2022) https://doi.org/10.1016/j.jcp.2021.110721

[32] Liu, M., Hasegawa, Y.: Adjoint-based shape optimization for compressible flow based on volume penalization method. Engineering with Computers, 1–21 (2024) https://doi.org/10.1007/s00366-024-02058-z . Ahead of print, September 30, 2024

[33] Yen, R., Kolomensky, D., Schneider, K.: Approximation of the Laplace and Stokes operators with Dirichlet boundary conditions through volume penalization: a spectral viewpoint. Numerische Mathematik **128**, 301–338 (2014) https://doi.org/10.1007/s00211-014-0610-8

[34] Chorin, A.J.: A numerical method for solving incompressible viscous flow problems. Journal of Computational Physics **2**(1), 12–26 (1967) https://doi.org/10.1016/0021-9991(67)90037-X

[35] Ohwada, T., Asinari, P.: Artificial compressibility method revisited: Asymptotic numerical method for incompressible Navier–Stokes equations. Journal of Computational Physics **229**(5), 1698–1723 (2010) https://doi.org/10.1016/j.jcp.2009.11.003

[36] Dupuy, D., Toutant, A., Bataille, F.: Analysis of artificial pressure equations in numerical simulations of a turbulent channel flow. Journal of Computational Physics **411**, 109407 (2020) https://doi.org/10.1016/j.jcp.2020.109407

[37] Rogers, S.E., Kwak, D., Kiris, C.: Steady and unsteady solutions of the incompressible Navier–Stokes equations. AIAA Journal **29**(4), 603–610 (1991) https://doi.org/10.2514/3.10683

[38] Drikakis, D., Rider, W.: The artificial compressibility method. In: High-Resolution Methods for Incompressible and Low-Speed Flows, pp. 173–208. Springer, Berlin, Heidelberg (2005)

[39] Lemke, M., Reiss, J.: Approximate acoustic boundary conditions in the time-domain using volume penalization. The Journal of the Acoustical Society of America **153**(3), 1219–1228 (2023) https://doi.org/10.1121/10.0017347

[40] Reiss, J., Sesterhenn, J.: A conservative, skew-symmetric finite difference scheme for the compressible Navier–Stokes equations. Computers & Fluids **101**, 208–219 (2014) https://doi.org/10.1016/j.compfluid.2014.06.004

[41] Verwer, J.G., Hundsdorfer, W.H., Sommeijer, B.P.: Convergence properties of the Runge-Kutta-Chebyshev method. Numerische Mathematik **57**(1), 157–178 (1990) https://doi.org/10.1007/BF01386405

[42] Engels, T., Kolomenskiy, D., Lehmann, F.-O.: Flight efficiency is a key to diverse wing morphologies in small insects. Journal of The Royal Society Interface **18**(183), 20210518 (2021) https://doi.org/10.1098/rsif.2021.0518

[43] Gautier, R., Biau, D., Lamballais, E.: A reference solution of the flow over a circular cylinder at Re=40. Computers & Fluids **75**, 103–111 (2013) https://doi.org/10.1016/j.compfluid.2012.12.017

[44] Kim, J., Moin, P., Moser, R.: Turbulence statistics in fully developed channel flow at low Reynolds number. Journal of Fluid Mechanics **177**, 133–166 (1987) https://doi.org/10.1017/S0022112087000892

[45] Vreman, A.W., Kuerten, J.G.M.: Comparison of direct numerical simulation databases of turbulent channel flow at $Re_\tau = 180$. Physics of Fluids **26**(1), 015102 (2014) https://doi.org/10.1063/1.4861064

[46] Vreman, A.W., Kuerten, J.G.M.: A third-order multistep time discretization for a Chebyshev tau spectral method. Journal of Computational Physics **304**, 162–169 (2016) https://doi.org/10.1016/j.jcp.2015.10.022

[47] Pope, S.: Turbulent Flows, 1st edn. Cambrige University Press, ??? (2000)

[48] Suzuki, K., Minami, K., Inamuro, T.: Lift and thrust generation by a butterfly-like

flapping wing–body model: immersed boundary–lattice boltzmann simulations. Journal of Fluid Mechanics **767**, 659–695 (2015) https://doi.org/10.1017/jfm.2015.57

[49] Dilek, E., Erzincanli, B., Sahin, M.: The numerical investigation of lagrangian and eulerian coherent structures for the near wake structure of a hovering drosophila. Theoretical and Computational Fluid Dynamics **33**(3–4), 255–279 (2019) https://doi.org/10.1007/s00162-019-00492-0

[50] Medina, A.: The aerodynamics of deforming wings at low reynolds number. PhD thesis, University of California, Los Angeles (2013)

[51] Huang, Q., Bhat, S.S., Yeo, E.C., Young, J., Lai, J.C.S., Tian, F.-B., Ravi, S.: Power synchronisations determine the hovering flight efficiency of passively pitching flapping wings. Journal of Fluid Mechanics **974**, 41 (2023) https://doi.org/10.1017/jfm.2023.821

[52] Cohen, A., Daubechies, I., Feauveau, J.C.: Biorthogonal bases of compactly supported wavelets. Communications on Pure and Applied Mathematics **45**, 485–560 (1992)