



HAL
open science

Mesh adaptation and dynamic positioning for the efficient simulation of lifting hydrofoil flows

Jeroen Wackers, Alban Leroyer, Jules Richeux, Pierre Robin, Ganbo Deng,
Hayriye Pehlivan Solak, David de Prémorél

► **To cite this version:**

Jeroen Wackers, Alban Leroyer, Jules Richeux, Pierre Robin, Ganbo Deng, et al.. Mesh adaptation and dynamic positioning for the efficient simulation of lifting hydrofoil flows. *Ocean Engineering*, 2025, 315, pp.119751. hal-04659314v2

HAL Id: hal-04659314

<https://hal.science/hal-04659314v2>

Submitted on 9 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Mesh adaptation and dynamic positioning for the efficient simulation of lifting hydrofoil flows

Jeroen Wackers^{*a}, Alban Leroyer^a, Jules Richeux^a, Pierre Robin^{a,b}, Ganbo Deng^a, Hayriye Pehlivan Solak^a, David de Prémorel^b

^a*LHEEA Lab, CNRS UMR 6598, Centrale Nantes, France*

^b*Finot-Conq Architectes Navals, Vannes, France*

Abstract

Accurate simulations of the flow around lifting hydrofoils are challenging, since they need to capture the flow near the foil surface precisely, represent the free surface, and take into account body motion and deformation. Therefore, these simulations are often computationally expensive.

This paper studies numerical methods to limit the costs of hydrofoil simulation. Mesh adaptation is used to efficiently capture the free-surface flow, to resolve flow details around the foil surface, and to ensure the accuracy of mesh motion techniques, like overset meshing. For maximum precision of the boundary-layer flow, adaptation is started from dedicated body-aligned meshes.

Hydrofoil flexibility is taken into account through a linear eigenmode-based reduced-order model of the structural response. This approach removes the need to couple directly the fluid and structure solvers and reduces the computational overhead for fluid-structure simulation. Equilibrium positions for flexible and rigid motion are determined with a fixed-point iteration based on approximate models for the forces, which eliminate the need for costly time-accurate simulation.

Test cases demonstrate that these methods work together, providing accurate simulation of realistic hydrofoils with reasonable computational costs.

Keywords: Lifting hydrofoil, RANS simulation, mesh adaptation, fluid-structure interaction, quasi-static positioning, sock mesh

1. Introduction

Hydrofoils, which reduce ship resistance by lifting the hull out of the water, are important in modern ship design. Created originally for fast motor vessels, hydrofoils are now extensively used on racing sailboats, for kite and windsurfing, and on pleasure craft; they can also increase the energy efficiency of fast passenger transport. Considering this trend, naval architects require design methodologies for lifting hydrofoils. Thus, efficient, accurate and easy-to-use hydrodynamic simulation is of major importance for the industry.

*Corresponding author

Email address: `jeroen.wackers@ec-nantes.fr` (Jeroen Wackers*)

And yet, not many simulations of realistic hydrofoils exist in the open literature. The research often focuses on ventilation (Harwood et al., 2016; Matveev et al., 2019; Zhi et al., 2022), and cavitation (Brizzolara, 2019; Perali, 2023). Non-ventilating flows are mostly simulated within yacht racing teams, where the work is kept a secret. Open discussions include (Banks et al., 2014; Marimon Giovannetti et al., 2022; Knudsen et al., 2024) who present combined experimental and numerical studies of hydrofoils, (Prasad et al., 2015) who study the effects of free-surface proximity, and (Sacher et al., 2018; Ploé, 2018) who simulate hydrofoils in the context of automatic shape optimisation.

Accurate simulation of hydrofoils is computationally expensive, for three reasons. First, hydrofoils are often slender elongated structures, so the compression and suction peaks are small with respect to the foil’s main dimensions. Thus, very fine meshes are required to capture the pressure. Second, the foil motion and deformation must be taken into account, since a foil’s stable attitude varies strongly with its charge and immersion depth, and foils are flexible structures whose deformation under load influences their performance. Finally, the free-surface position is hard to predict before the computation if the hydrofoil moves or makes waves, so resolving the surface accurately is a challenge for meshing.

This paper investigates four numerical methods to increase the efficiency of hydrofoil simulation: (1) Accurate and reliable resolution of local flow features is obtained through adaptive grid refinement. (2) Fluid-structure interaction (FSI) is simulated efficiently with a structural model based on modal analysis, while (3) quasi-static approaches are used to determine the equilibrium position and shape for moving hydrofoils. Finally, (4) a dedicated initial mesh topology increases the precision of the results and the possibilities of FSI simulations. The four methods function together and reinforce each others’ efficiency.

Adaptive mesh refinement. Mesh adaptation locally and automatically refines the mesh during the simulation, according to the requirements of the flow. Adaptation is reaching maturity for the simulation of complex industrial flows, notably in aeronautics (Balan et al., 2020; Park et al., 2021; Alauzet and Frazza, 2021). For hydrodynamic simulations, adaptation is used with flow solvers such as OpenFOAM (Eskilsson and Bensow, 2012; Wang et al., 2020), StarCCM+ (Perić, 2022) and FINE/Marine (Korkmaz et al., 2023; Wackers et al., 2017).

Recently, Wackers et al. (2022) used adaptive mesh refinement for the automated simulation of ship hull resistance. By combining anisotropic refinement of unstructured hexahedral meshes with standardised user guidelines, they obtained simulations that require little user intervention and produce better accuracy on coarser meshes than traditional meshing methods.

Here, this approach is used for efficient meshing of hydrofoils. Mesh adaptation is applied specifically to capture the free surface for any position of the foil, to accurately resolve the flow around the hydrofoil, and to maintain the correct cell sizes when the mesh is deformed or moved.

FSI modal approach. Since modern hydrofoils deform under load, (at least) the static deformation must be taken into account for accurate force predictions. Full simulation of such fluid-structure interaction problems requires a two-way coupling between the fluid and structure solver, which is complicated and expensive for industrial simulation.

Therefore, we consider an alternative: to replace the structural solver by a reduced-order model based on a modal analysis, that is pre-computed with a finite-element solver. This approximation is cheap to evaluate and thanks to the tight coupling with the fluid solver, the procedure converges fast. A disadvantage is that the structural deformation is linearised, but for typical hydrofoil deformations this is an acceptable approximation.

Quasi-static positioning and deformation. Hydrofoils are often simulated at a specified lift force (corresponding to the weight of the hull carried by the foil) and sideforce. We achieve this by dynamically adjusting the incidence and sideslip angles of the foil during the simulation.

To achieve optimal convergence, quasi-static procedures have been implemented both for the modal structural deformation and for the dynamic positioning: using simplified models, the equilibrium position is estimated regularly and the actual deformations are brought smoothly to this estimated equilibrium. Such procedures are more robust than time-accurate resolution of the motions and, with good estimations of the equilibrium, they converge rapidly.

‘Sock’ initial mesh topology. On unstructured hexahedral meshes, the highest accuracy is obtained when the cells are aligned as much as possible with the geometry. However, for the intricate curved shapes of hydrofoils, this is a challenge. Following Mallol et al. (2019), we show how a nearly body-fitted mesh can be created in a small tube-like domain around the hydrofoil. This domain is then included as an overset mesh into a much larger background domain.

The advantages of this approach are the possibility to handle large deformations of the hydrofoil, as well as high accuracy of the solutions for reasonable mesh sizes, especially when combined with adaptive refinement.

The study uses the Navier-Stokes solver ISIS-CFD (Queutey and Visonneau, 2007; Wackers et al., 2011). This flow solver is described in section 2, with special attention for its mesh refinement method and the simulation of body motion. Sections 3 to 6 present hydrofoil mesh adaptation, the modal deformation, the quasi-static approach and the ‘sock’ meshes. These methods are applied to various hydrofoil test cases in section 7, to test their versatility and accuracy. Comparisons with non-adapted meshes, different solvers, and experiments are provided. To conclude, the implications of the tests are discussed in section 8.

2. The ISIS-CFD flow solver

The flow solver ISIS-CFD is developed by CNRS / Centrale Nantes and distributed by Cadence Design Systems as part of the FINE/Marine flow simulation suite. This introductory section presents the flow solver (section 2.1), mesh deformation (section 2.2) and overset (section 2.3) techniques that allow 6 degree-of-freedom rigid and flexible body motions, and the integrated mesh adaptation method (section 2.4).

2.1. The flow solver

ISIS-CFD is an incompressible unsteady Navier-Stokes solver for multifluid flow, based on the finite-volume method to build the spatial discretisation of the transport equations (Duvigneau and Visonneau, 2003; Queutey and Visonneau,

2007). The velocity field is obtained from the momentum conservation equations and the pressure field is extracted from the mass conservation constraint transformed into a pressure equation, similar to the SIMPLE method (Rhie and Chow, 1983). Among others, classical Reynolds-averaged Navier-Stokes turbulence models like $k - \omega$ SST (Menter, 1994) are available, with wall-resolved and law-of-the-wall boundary conditions. Both steady and unsteady free-surface flows are solved with a time integration technique.

Free-surface flow is simulated by treating the fluid everywhere as a mixture of water and air (Wackers et al., 2011), which are distinguished with a conservation equation for the volume fraction of water that has a discontinuous inflow condition. This discontinuity is convected through the flow, implicitly defining the free-surface position. Specific compressive discretisation schemes keep the interface as sharp as possible. This approach is robust but it requires fine regular meshes at the interface location.

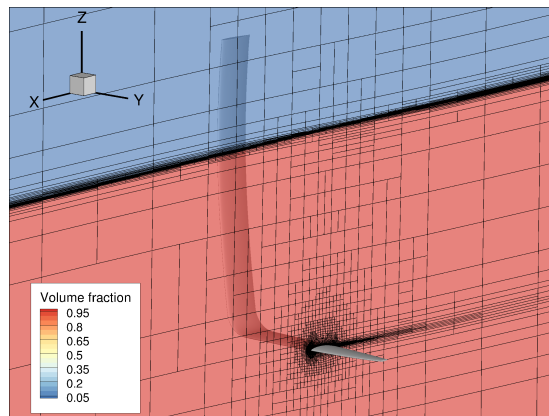


Figure 1: An adapted mesh for a lifting hydrofoil, coloured with the water volume fraction.

The unstructured finite-volume discretisation is face-based: fluxes are computed face by face using only the two neighbour cells, which means that cells with any number of arbitrarily shaped faces are accepted. ISIS-CFD is mostly used with unstructured hexahedral grids from the Hexpress grid generator which is also part of FINE/Marine. These meshes combine semi-structured regions with body-fitted viscous layer grids near the walls (see figures 1 and 3). The grids consist purely of hexahedral cells, with mesh size variations obtained by placing larger cells next to two or four smaller neighbour cells. Thanks to its face-based nature, ISIS-CFD treats these grids the same as any other type of mesh.

2.2. Mesh deformation and motion

Body motion is taken into account with both mesh deformation and rigid transformation of the mesh (Leroyer and Visonneau, 2005; Leroyer et al., 2008), or sliding interfaces and overset meshes which allow a part of the mesh to move within the rest (see section 2.3). To solve the flow on moving meshes, ISIS-CFD uses an Arbitrary Lagrangian Eulerian (ALE) formulation of the flow equations.

Rigid-body motion. For free-surface flows with rigid-body motion, the alignment of the free surface with the mesh must be preserved as much as possible close to

external boundaries. Therefore, for single-domain configurations, motions that do not displace the free surface (x - and y -translation, z -rotation, with z the vertical axis) are treated by rigid-body displacement of the entire mesh. Those that do (z -translation, x - and y -rotation) require mesh deformation where the outer boundaries of the mesh are not deformed, to make sure the mesh at the free surface moves as little as possible. If the free-surface grid is created with adaptive refinement, z -translation can also be treated with rigid-body movement.

Mesh deformation is performed with analytical weighted regriding: the required motion is applied to all the nodes of the mesh, multiplied by a weighting coefficient $w(\mathbf{x})$ which goes to zero on the outer boundaries so that only the inner part of the mesh moves (figure 2). An initial coefficient $w_0(\mathbf{x})$ is computed as a solution of Laplace’s equation with Dirichlet boundary conditions 1 on the body surface and 0 on the outer boundaries. This coefficient is then modified to produce the actual weighting:

$$w(\mathbf{x}) = \max(w_0(\mathbf{x})^q/c_{\max}, 1). \quad (1)$$

The power q reduces the zone where the mesh deforms significantly, to leave the free-surface mesh far away from the body undisturbed. The factor c_{\max} on the other hand, imposes a zone of rigid-body motion $w = 1$ close to the body, to preserve the shape of the boundary layer mesh. For hydrofoil simulations with mesh refinement, the default values are $q = 1.6$ and $c_{\max} = 0.55$.

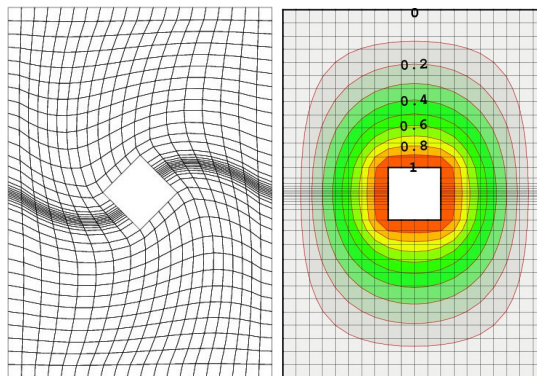


Figure 2: Deformed mesh for rotation of the central body (left) and the weighting coefficient $w(\mathbf{x})$ (right).

Flexible-body deformation. The deformation of flexible bodies cannot be performed with the analytical weighting presented above. For the modal FSI described in section 4, the mesh is deformed with radial basis functions (RBF) scattered over the geometry, which are used to interpolate the displacement of each of the modal bodies throughout the mesh (Mouton et al., 2018).

RBF methods were initially developed as general interpolation techniques (Hardy, 1990; Franke, 1982), but have been widely adopted for mesh deformation in the ALE framework (de Boer et al., 2007; Rendall and Allen, 2009; Bos et al., 2013; Biancolini et al., 2014; Groth et al., 2019). Separate RBF interpolations are used for the displacements in each coordinate direction. To accurately reproduce data which resemble polynomials, the so-called *augmented*

RBF method is commonly used. Given a set of points in space $\{\mathbf{x}_j\}_{j=1}^{n_p}$ for which data $\{d_j\}_{j=1}^{n_p}$ is known, the augmented RBF interpolant $s(\mathbf{x})$ is given by:

$$s(\mathbf{x}) = \sum_{j=1}^{n_p} \lambda_j \phi(|\mathbf{x} - \mathbf{x}_j|) + \pi(\mathbf{x}), \quad (2)$$

$$\text{with } \pi(\mathbf{x}) = \sum_{k=1}^{m_d} \gamma_k \pi_k(\mathbf{x}), \quad (3)$$

$$\{\pi_k(\mathbf{x})\}_{k=1}^{m_d} \text{ basis } \mathbb{P}_m^d \text{ of d-variate polynomials} \quad (4)$$

of degree at most m , $m_d = \dim \mathbb{P}_m^d$,

where ϕ is the RBF kernel (by default, we use CTPS C_b^2 from de Boer et al. (2007) for hydrofoils, with a characteristic radius equal to the chord c). $|\cdot|$ indicates a vector length. The coefficients $\{\lambda_j, j \in \llbracket 1, n_p \rrbracket\}$ and the coefficients $\{\gamma_k, k \in \llbracket 1, m_d \rrbracket\}$ of the polynomial $\pi(\mathbf{x})$ are determined by the interpolation conditions $s(\mathbf{x}_j) = d_j$ for $j \in \llbracket 1, n_p \rrbracket$ augmented with the additional constraints $\sum_{j=1}^{n_p} \lambda_j \pi_k(\mathbf{x}_j) = 0, 1 \leq k \leq m_d$. These additional conditions translate naturally into a constrained optimization problem (Schaback, 1995; Bayona, 2019).
 In this work, polynomials of degree 1 are chosen.

In terms of computational time, obtaining the coefficients requires $\mathcal{O}(n_p^3)$ floating-point operations, while evaluating the solution at n_n nodes is of the order $\mathcal{O}(n_n \times n_p)$. Even if the RBF technique can quickly become prohibitively expensive when the number of structure nodes and/or fluid nodes are large, it will be seen in section 4.3 how its use with the modal approach leads to an effective and robust algorithm.

2.3. Overset meshing

Overset meshing is the dynamic connection of two or more mesh domains which partially or fully cover each other. This allows for unlimited mesh displacement and is therefore the method of choice to simulate large body motions. Figure 20 shows an overset mesh.

There are three different types of cells in an overset computation: active, inactive, and interpolation cells. The treatment of active cells is exactly the same as for computations without overset: the governing equations are discretised and solved to determine the unknowns of the problem in those cells. Inactive cells are cells that are overlapped by active cells in another domain. They are discarded in the computation. Cells located at the interface between the active cells and inactive cells are interpolation cells. Unknowns, as well as all other quantities required for the discretisation such as the gradient of those unknowns, need to be interpolated from a host cell located in another domain.

Domain connection. In each time step, the first task for overset meshing is to determine the status of the different cells. In our implementation, a domain is declared as either a background domain, or an overlapping domain which must contain a solid body. When a cell in a background domain is covered by an overlapping domain, it becomes inactive. When a cell in an overlapping domain is covered by another overlapping domain, it is active if its cell centre is closer to the body in its own domain than the one in the other overlapping domain.

Otherwise, it is inactive. After assigning active or inactive status to all cells, the inactive cells next to an active cell are changed to interpolation cells.

185 The second step in the overset implementation is to search for host cells. A host cell is an active cell in another domain that contains the cell centre point of an interpolation cell. This straightforward search is performed with a tree algorithm whose details will not be given here.

Interpolation. The last step is to interpolate the solution from the host cells to the interpolation cells using the neighbouring cells of each host cell. While the first two steps mainly have an impact on the cost of the computation, this one is crucial both for the accuracy and the stability of the overset algorithm. In the SIMPLE algorithm used in ISIS-CFD, second derivatives of the pressure appear in the continuity equation, so the pressure must be smooth to avoid spurious velocity oscillations. Hence, it is mandatory to ensure a second order accurate interpolation for the pressure. To achieve this, a least-squares approach based on linear polynomials is used. By default, the neighbours of the faces of a host cell are used to create the stencil. If the minimum interpolation coefficient is smaller than a prescribed value (the default value being -0.1), we switch to a stencil based on the neighbours of the nodes. This is usually the case when the interpolation point is closer to a node than to the centre of the host cell. If the minimum interpolation coefficient is still smaller than the prescribed value, a distance-weighted interpolation scheme is used to ensure stability.

200 An exception is made near the free surface. Since the pressure is a C^0 -continuous function at the free surface, least-squares interpolation can create oscillations in this region. Hence, the distance-weighted scheme is always used. And since the volume fraction is discontinuous, second order interpolation makes no sense. Therefore, no interpolation is performed: the volume fraction of the host cell is assigned directly to the interpolation cell.

210 2.4. Adaptive refinement

Mesh adaptation in ISIS-CFD (Wackers et al., 2012, 2017) uses local cell division (figure 1), starting from a coarse initial mesh (figure 3). To efficiently capture flat or stretched flow features such as vortical wakes or the water surface, the adaptation is anisotropic: cells can either be split in one or several directions. For unsteady flows or to accommodate the convergence of a steady flow, existing refinement can be undone and the procedure is fully parallel, including adaptive load balancing. Since a cell can only be refined or derefined once during each adaptation step, the procedure is called every 10 – 25 time steps, until the flow has converged and the mesh is no longer changed by the adaptation.

Metric-based refinement. Anisotropic grid refinement is based on metric tensors (George et al., 1991; Loseille et al., 2010; Alauzet and Frazza, 2021). For metric-based adaptation, the refinement criterion is a symmetric tensor field $\mathcal{C}(\mathbf{x})$ which indicates the ideal size of the cells everywhere in the domain. Starting from this continuous field computed from the flow, the grid is adapted until the three dimensions $\mathbf{d}_{i,j}$ of each hexahedral cell i , which are the vectors between the face centres for the three pairs ($j = 1, 2, 3$) of opposing faces, satisfy:

$$|\mathcal{C}_i \mathbf{d}_{i,j}| = T_r, \quad \forall i, \forall j \in [1, 2, 3], \quad (5)$$

220 as closely as possible. Cells are refined in a direction j whenever the modified distance $|\mathbf{C}_i \mathbf{d}_{i,j}|$ exceeds T_r and derefined whenever it is below $T_r/2.5$, to make
 225 sure the derefined cells are not immediately refined again. T_r is a global specification of the mesh fineness: all the cell sizes are proportional to this parameter.

Free surface – Hessian criterion. The criteria \mathbf{C} come from the Hessian matrix
 225 of second spatial derivatives of the solution, which is a measure of linear interpolation errors (Loseille et al., 2010). The refinement criterion must react to all the flow features which are relevant for hydrofoil resistance: pressure fields, boundary layers, wakes, and shear layers. Therefore, it is based on the Hessians of both the pressure and the velocity (Wackers et al., 2017).

To get the same order of magnitude for all Hessians, a common weight ρV is assigned to the the velocity Hessians. In ALE, the criterion is made independent of the (potentially moving) reference frame by choosing $V = |\mathbf{v} - \mathbf{v}_\infty|$, where $\mathbf{v} = [u, v, w]$ and \mathbf{v}_∞ is the inflow velocity¹. Our tests show that this weighting, which emphasizes the velocity Hessians close to the geometry, is more effective for resistance computation than other formulations. The criterion then becomes:

$$\mathbf{C}_H = \left(\max(\|\mathcal{H}(p)\|, \rho V \|\mathcal{H}(u)\|, \rho V \|\mathcal{H}(v)\|, \rho V \|\mathcal{H}(w)\|) \right)^\varphi. \quad (6)$$

230 \mathcal{H} is the Hessian operator. $\|\cdot\|$ indicates a matrix having the same eigenvectors as the original one and the absolute values of its eigenvalues; the power φ is again applied to the eigenvalues, $\varphi = 0.5$ is used here. Finally, the maximum of two tensors is computed following the procedure in Wackers et al. (2012), which is based on George and Borouchaki (1998).

A smooth anisotropic mesh at the free surface is required for the volume fraction equation, but the pressure gradient is discontinuous at the free surface so the pressure Hessian is undefined there. Therefore, the Hessian criterion is extrapolated from below the surface, and the refined mesh at the free surface is created with a second criterion (Wackers et al., 2012), based on the water volume fraction α . Directional refinement normal to the surface is obtained from normal vectors $\mathbf{n}_S = \nabla \alpha_A / |\nabla \alpha_A|$, where α_A is a smeared volume fraction field created with Laplacian smoothing. The criterion is:

$$\mathbf{C}_S = \begin{cases} \mathbf{n}_S \otimes \mathbf{n}_S & \text{if } 0.1 \leq \alpha_A \leq 0.9, \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (7)$$

235 The choice based on α_A creates a buffer layer of a few cells around the surface position for safety.

The two tensor criteria are combined into one by taking a weighted maximum of the tensors. Since the free-surface criterion always has a unit eigenvalue, a weighting constant a is applied only to the Hessian criterion:

$$\mathbf{C}_C = \max(\mathbf{C}_S, a \mathbf{C}_H). \quad (8)$$

In practice, we specify separate thresholds T_{rS} and T_{rH} for the free-surface and Hessian criterion. These are then used to set $T_r = T_{rS}$ and $a = \frac{T_{rS}}{T_{rH}}$.

¹This criterion is the same as the one from Wackers et al. (2022) which uses $V = |\mathbf{v}|$, since $\mathbf{v}_\infty = \mathbf{0}$ in the moving reference frames used for free-surface flows.

In the adapted mesh of figure 1, refinement in the pressure peaks at the leading and trailing edge can be seen, as well a anisotropic refinement in the wake. The free surface is refined and the Hessian criteria capture the region where the pressure disturbance from the foil interacts with the free surface.

Grid quality measures. Post-treatment of the refinement criterion is used to improve the regularity of the adaptively refined meshes. First, a minimum cell size is imposed: cells smaller than this size are no longer refined. This option prevents spurious refinement in case of locally large errors in the refinement criterion, which may appear in the high aspect-ratio cells of the near-wall boundary layer grid. Also, it prevents infinite refinement around flow singularities.

Furthermore, adaptive refinement can be forbidden outside of a specified limiting box. This option is used to maintain a coarse grid near the outflow boundary in order to damp out perturbations of the flow. For the grid refinement in the x -, y - and z -directions, different limiting boxes can be chosen.

3. Adaptive mesh refinement for hydrofoil simulation

Using the strategy for standardising computational setup presented in section 3.1, we discuss how mesh adaptation can address the hydrofoil simulation challenges outlined in section 1: creating the mesh around the unknown free-surface position (section 3.2), providing fine cells to capture the flow details around the foil geometry (section 3.3), and ensuring the continuity of the grid size during mesh deformation and overset motion (section 3.4).

3.1. Standardised computational setup

Systematic use of mesh adaptation requires that the simulations are easy to configure, with reliable guidelines for the simulation parameters which are valid over a large range of cases and which do not require trial and error to produce correct simulations. This concerns the choice of the refinement criterion and numerical parameters such as the threshold T_r .

To get the same type of mesh for similar objects, independent of their size or speed, Wackers et al. (2022) use dimension analysis. The guidelines for the refinement criterion and all other parameters are non-dimensionalised in terms of the object velocity U , reference length L , and density ρ . The remaining dependence on Fr , Re and the geometry is weak for a class of geometries; for example, simulation parameters for resistance evaluation of model-scale displacement ships can be chosen independent of the hull shape or the flow regime.

Guidelines for the non-dimensional parameters are obtained through systematic testing. A full definition of the hydrofoil simulation protocol is given by Richeux (2022). Without repeating all the details of the tests, the following subsections present some elements of the protocol and this subsection introduces its basis: the reference length and the initial mesh.

Hydrofoil length scale. For dimension-independent guidelines, all length parameters (refinement thresholds, minimum cell size etc.) are expressed in terms of a reference length L . This length can be the largest dimension of the object, or it can be a characteristic dimension of the flow. For a ship, these lengths coincide: the flow around a hull develops over the distance from bow to stern, which is

also the ship’s largest dimension. Therefore, Wackers et al. (2022) choose this length scale for their ship refinement protocol.

285 For hydrofoils, the largest dimension is typically the span. However, the most important flow phenomena for a correct prediction of the foil performance are the boundary layer and the leading-edge and trailing-edge peaks in the pressure field. These are all linked with the streamwise flow, so it makes sense to standardise the meshes in the wing sections, independent of the overall shape
290 of the foil. To achieve this, the chord c , the main dimension in the streamwise direction, is taken as the reference length. We use the maximum chord; an average value can be chosen if the wing is highly tapered.

Initial mesh. Since the refinement criterion $\mathcal{C}(\mathbf{x})$ directly specifies the target size and aspect ratio of the cells, the adapted mesh should be independent
295 of the initial Hexpress mesh from which the computation is started. But for refinement by local division, this is only partially true, which means that the initial mesh should be created with care.

Division-based refinement can only divide the initial cell sizes by a power of 2. Thus, if a specific cell size is desired for the refined grid, this must be taken
300 into account for the initial grid. Hexpress creates meshes by octree division of a cartesian base mesh, followed by projection on the geometry; for hydrofoils, we use a base mesh with cell sizes of $4c$. Since any further refinement, either in Hexpress or in the adaptation, is accomplished by cutting these cells in half, the cell sizes $c/2^n$ are available for the mesh adaptation.

305 Furthermore, the adaptation cannot change the orientation of the cells. For accurate solutions, the initial mesh must therefore be aligned with the main flow features. In the boundary layer, Hexpress inserts a body-aligned grid whose height is proportional to the cell sizes outside the viscous layer, so the coarser the initial mesh, the thicker (i.e. better) the boundary layer grid will be. Unfortunately, very coarse grids risk bad quality cells around the trailing and leading
310 edges. For hydrofoils, a good compromise is an initial cell size of $c/8$ on the foil surfaces, with $c/32$ around the leading and trailing edge (Richeux, 2022).




Figure 3: Typical chordwise cut through an initial mesh, repeated in spanwise direction to mesh the entire foil.

Figure 3 shows an example of such a mesh. Since the protocol is defined in terms of the reference length c , the initial mesh in profile cross-sections is
315 similar for all foils (c.f. figure 13).

3.2. Free-surface refinement

The volume fraction equation which determines the position of the free surface (section 2.1) requires a fine, regular mesh around the discontinuity. However, due to the mesh motion and deformation, as well as the large waves which some hydrofoils generate, the free-surface position within the mesh is impossible to predict before a simulation. Thus, apart from filling most of the domain with fine cells, the only way to create the free-surface mesh is with adaptation.

Guidelines for ISIS-CFD non-adapted meshes suggest a constant cell size at the free surface, independent of the mesh size below the surface. This strategy is conserved for the adapted meshes, so the threshold T_{rS} (section 2.4) is fixed, independent of the other parameters. The free-surface criterion (7) has unit eigenvalues, so T_{rS} directly specifies the size of the refined cells at the free surface.

To capture the free surface while keeping the number of cells as low as possible, we search the highest value for T_{rS} which does not perturb the solution. The ISIS-CFD recommendation for hydrofoils is a free-surface mesh size around $c/100$ vertically, so $T_{rS} = c/128$ and $c/64$ are tested here on the L-foil of section 7.1. Table 1 shows four cases where the meshes around the hydrofoil become progressively finer, while the foil is gradually brought closer to the free surface. Both T_{rS} produce similar forces, while the number of cells is significantly reduced for $T_{rS} = c/64$. In the cases where the hydrofoil tip is close to the water surface, the simulations with $T_{rS} = c/128$ crashed due to the large mesh deformation needed to reach the equilibrium position (see section 5.3). Thus, the higher free-surface threshold $T_{rS} = c/64$ contributes both to reducing the number of cells and to stabilising the mesh deformation. $T_{rS} = c/64$ is therefore adopted.

Table 1: Influence of the free-surface threshold on the drag value and the number of cells for different submersions and Hessian thresholds (section 3.3).

T_{rH}	Tip angle	Immersion	T_{rS}	Nb. cells	F_x (N)
$2c$	90°	100%	$c/128$	1.235M	1066.3
			$c/64$	0.588M (-52%)	1066.8 (-0.04%)
$c/4$	60°	83%	$c/128$	4.041M	1046.4
			$c/64$	3.051M (-24%)	1043.4 (-0.28%)
$c/4$	60°	54%	$c/128$		crash
			$c/64$	2.997M	820.3
$c/11.3$	60°	54%	$c/128$		crash
			$c/64$	12.480M	795.1

This free-surface threshold seems high compared to the threshold for ships $T_{rS} = L/1000$ (Wackers et al., 2022). However, for a typical sailing yacht, the hydrofoil chord is about 30 times smaller than the ship length, so $c/64$ corresponds to $L/1920$. Thus, the choice of $T_{rS} = c/64$ is coherent.

3.3. Hessian refinement

Solving correctly the flow around the geometry is essential to compute hydrofoil forces. Since most of these forces are pressure-based, an accurate resolution of the pressure field is needed. Furthermore, the boundary layer has a major influence on the viscous drag and the displacement effect of the flow in the boundary layer changes the pressure field. High-accuracy solutions therefore

require the precise representation of the velocity in the boundary layer. This motivates the use of the pressure–velocity Hessian criterion of section 2.4.

Threshold. Since the size of all the cells below the surface is proportional to T_{rH} , this is a convenient parameter to control the overall fineness of the mesh. The relation between T_{rH} and the numerical uncertainty is similar over a class of test cases (Wackers et al., 2022), so T_{rH} can be seen as a direct specification of the numerical uncertainty.

For the values of T_{rH} , Wackers (2019) finds a suitable range of $[0.4L, 0.04L]$ for a diverse set of simulations including ships, wing profiles, and 3D hydrofoils. Wackers et al. (2022) propose $T_{rH} \in [0.2L, 0.025L]$ for displacement ships. For hydrofoils, Serani et al. (2019) use $T_{rH} \in [0.4c, 0.1c]$.

The tests in this paper (section 7) confirm that suitable threshold values for hydrofoils are somewhat higher than for ships. The main reason for this is practical: hydrofoils are slender, elongated shapes with highly concentrated pressure peaks on the leading and trailing edges. Such flows require large numbers of cells for accurate solutions, even with adapted meshes. Thus, to keep the cell sizes reasonable, a slightly lower accuracy than for ships must be accepted; hence the higher thresholds. The range used in this paper is $T_{rH} \in [2.0c, 0.0625c]$.

Refinement limiting. Unmodified application of the pressure-velocity Hessian criterion leads to the accurate resolution of the entire flow field, which may not be necessary to compute the forces on a hydrofoil. Limiting the refinement criterion can reduce the number of cells in the refined meshes, while still producing good force predictions.

The refinement far behind the foil is removed since, for most hydrofoils, it has little influence on the forces. A limiting box is used to forbid refinement in the x - and y -directions from a certain distance behind the foil. The refinement along z is allowed everywhere in the domain to correctly capture the free surface and thus prevent it from diffusing. A series of tests on the L-foil in (Richeux, 2022) shows that the limiting box does not have any influence on the drag as long as it is more than $0.25c$ from the trailing edge. To increase the safety margin, the current protocol uses a box at $1c$ behind the trailing edge.




Figure 4: L-foil drag as a function of T_{rH} for three minimum cell sizes.

Furthermore, the mesh is not refined below a minimum cell size. This acts as a cutoff filter for refinement in small flow details, which do not contribute to

the forces. Figure 4 shows the convergence of the drag with T_{rH} for the L-foil, for three minimum cell sizes: $h_{\min} = c/128$, $c/256$ and $c/512$. For low numbers of cells, these three behave similarly. Afterwards, the $c/128$ series converges faster, but it levels off at a value that is too high. The two other series have a similar convergence behaviour, although the $c/256$ series probably converges faster for the higher numbers of cells. Thus, $h_{\min} = c/256$ is the most interesting value, that provides faster convergence than finer minimum cell sizes but does not significantly modify the converged forces. While $c/512$ may be the safest choice for very fine meshes, $c/256$ is retained in the protocol.

3.4. Deformation and overset

When the mesh deformation and overset techniques are used to accompany body motion (section 2.2), adaptive refinement helps to preserve the simulation accuracy. These approaches interact as follows.

For analytical mesh deformation, the weighting $w(\mathbf{x})$ in (1) is obtained on the refined meshes by interpolation from the existing nodes to the new ones. For division-based refinement, new nodes are always surrounded by existing ones, so this interpolation is straightforward and fast. The distances $d_{i,j}$ for the metric criterion (5) are computed on the deformed mesh. Thus, if a part of the mesh gets stretched by the deformation, this is compensated by extra refinement.

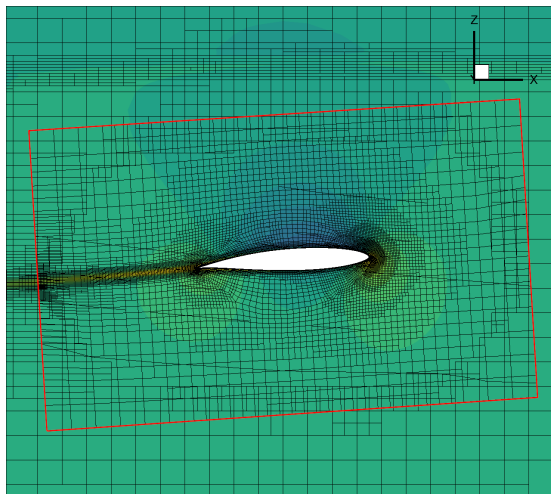


Figure 5: Cut through an overset domain around a hydrofoil, with matched refinement in the background and overset domains. The axial velocity is shown in colour.

For overset meshing, there are reasons to place the domain boundaries close to the hydrofoil (section 6), but this requires a high accuracy for the interpolations between the overset and background domains. The best interpolation is achieved when the cell sizes in both domains are locally the same. Since the cell sizes for metric-based refinement are linked to the refinement criterion value, this condition is satisfied naturally since the Hessian-based refinement criterion is continuous over the domain interface. Figure 5 shows this continuity of the cell sizes; the interpolation is good enough to preserve the strength of the wake.

Without adaptive refinement, the background domain needs fine cells in all the possible positions of the overset interface. The resulting increase in

the number of cells cancels the benefits of the overset approach. Thus, mesh adaptation is essential to make overset efficient.

415 4. Modal approach for fluid-structure interaction

Modern composite hydrofoils deform during their operation, which has an influence on their performance (Marimon Giovannetti et al., 2022). It is becoming clear in the industry that flow-induced deformation should be taken into account for hydrodynamic simulation.

420 Ideally, this fluid-structure interaction (FSI) should be solved with computational costs that are similar to stationary rigid-body simulations. For fully coupled FSI, this is difficult since a finite-element (FEM) solver, which computes the deformation, has to run in parallel with the fluid solver. The coupling between the solvers can be loosened by alternately solving rigid-body flow and
425 applying the fluid forces in the FEM solver to get a new deformed geometry. However, this iterative process may converge slowly.

Our strategy for reducing the computational overhead of FSI is to maintain the strong coupling, but to replace the structure solver by a reduced model of the structural response, based on eigenvalue analysis. For small to medium
430 deformations, this provides a good approximation of the structural deformation with negligible overhead. This modal approach is presented in section 4.1. Section 4.2 discusses the iterative solution of FSI coupling using added-mass stabilisation. Finally, section 4.3 presents the mesh deformation and shows the natural link between modal FSI and RBF interpolation.

435 4.1. Linearised modal approach for fluid-structure interaction

Let a full 3D model of the structure (here the foil) be defined through any FEM code, prior to running the fluid solver. In this model, the linearised form of the equations governing the motion of the structure is given by:

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{f}(t), \quad (9)$$

where \mathbf{M} , \mathbf{C} , and \mathbf{K} mean respectively the mass, damping and stiffness matrices, and \mathbf{u} the displacement of the nodes in the FEM model ($\mathbf{u} = \mathbf{0}$ represents the equilibrium position around which linearisation has been performed). $\mathbf{f}(t)$ refers to the external load, i.e. the fluid load.

The linear system (9) can be treated by the FEM code with a so-called modal analysis by computing its natural vibration eigenmodes and describing the solution as:

$$\mathbf{u}(t) = \mathbf{\Psi}\mathbf{q}(t), \quad (10)$$

where the i^{th} column $\tilde{\boldsymbol{\psi}}_i$ of the eigenmodal matrix $\mathbf{\Psi}$ represents the i^{th} eigenmode, whose size is equal to the number of DOF n of (9), and $\mathbf{q}(t)$ the vector of amplitudes for all the modes. Given the \mathbf{M} and \mathbf{K} orthogonal property of the natural vibration modes and assuming the hypothesis of Rayleigh damping, i.e. taking \mathbf{C} as a linear combination of \mathbf{M} and \mathbf{K} , the problem is now simplified to a resolution of n single-DOF equations. Each single equation (corresponding to the i^{th} row of the system) has the form:

$$m_i\ddot{q}_i + c_i\dot{q}_i + k_iq_i = \tilde{\boldsymbol{\psi}}_i^T \mathbf{f}(t).$$

The eigenmodal matrix Ψ can be mass-normalized to have $\Psi^T M \Psi = I$ ($\tilde{\psi}_i \rightsquigarrow \psi_i = \tilde{\psi}_i/m_i$) leading to the modal equations:

$$\ddot{q}_i + 2\varepsilon_i \omega_i \dot{q}_i + \omega_i^2 q_i = \psi_i^T \mathbf{f}(t), \quad i \in [1, n], \quad (11)$$

with $\omega_i = 2\pi f_i$, where f_i are the eigenfrequencies. The total deformation is:

$$\mathbf{u} = \sum_{i=1}^n q_i(t) \psi_i. \quad (12)$$

440 Model reduction is achieved by excluding from the resolution the modes which are without major interest for the studied configuration. In practice, only the few first modes n_m corresponding to the lower natural frequencies are retained as input to the modal FSI module integrated in the fluid solver. Figure 23 presents an example of these first modes.

445 The modal shapes for the nodes located on the surface of the structure, denoted ψ_{s_i} , with the frequencies f_i and damping coefficients ε_i , are retained from this off-line resolution and provided to the fluid solver as an input file. During the FSI coupling, the deformed shape at the structure nodes of the body surface is mapped continuously onto the fluid domain with the RBF technique
450 described in section 2.2.

4.2. Flow coupling and stabilisation

The n_m decoupled modal equations selected from the first modes are solved with a specific module integrated in the ISIS-CFD code and coupled with the flow resolution, in the same manner as the resolution of the Newton's laws for
455 rigid bodies when solving unsteady configurations (Yvin et al., 2018). Within each time step, the structure solution and the implicit flow discretisation are updated alternately in a series of nonlinear iterations.

Just like with rigid bodies, the added mass effects are naturally included in the source term through the fluid forces $\mathbf{f}(t)$ and failing to account for them has the same destabilising effect on the nonlinear iterations. To make the coupled FSI solution process stable, an added mass coefficient is computed at the beginning of the simulation, by solving the pressure field $\tilde{\rho}_i$ while imposing a unit acceleration for the i^{th} mode. The coefficients a_i are deduced by projecting the pressure field $\tilde{\rho}_i$ on the i^{th} mode shape: $a_i = \psi_i^T \tilde{\rho}_i$. Equations (11) are then modified to equations (13), which are solved at each non-linear iteration k .

$$(1+a_i) \ddot{q}_i|_{t+dt}^{k+1} + 2\varepsilon_i \omega_i \dot{q}_i|_{t+dt}^{k+1} + \omega_i^2 q_i|_{t+dt}^{k+1} = \psi_i^T \mathbf{f}|_{t+dt}^k + a_i \ddot{q}_i|_{t+dt}^k, \quad i \in [1, n_m]. \quad (13)$$

The added-mass coefficient can be updated during the simulation if it is required, typically when the immersion of the body drastically changes. Alternatively, a
460 quasi-static approach can be used for configurations where only the dynamic equilibrium is of interest, see section 5.1.

4.3. Mesh deformation and data transfer at the interface

Coupling a fluid and a structure model implies the interpolation of (a) forces from the fluid to the structure and (b) the structure deformation to the fluid
465 domain. With the modal approach, the structure is viewed by the fluid solver as

a set of points without connectivity or mesh; the RBF interpolation (section 2.2) is chosen, on the one hand, because it is suitable for this type of data. On the other hand, it allows both problems (mesh deformation and force interpolation) to be treated in a unified, efficient, and precise manner.

470 In general, RBF interpolation is not the most efficient method in terms of computation time, especially when the number of structural points increases and the entire interpolation process (obtaining coefficients and evaluating the solution at desired nodes) needs to be repeated each time. However, in the present case, the first selected modes can be accurately described by a limited
475 set of points, and, most importantly, this data remains unchanged throughout the simulation, even if the fluid mesh changes through adaptive refinement. As a result, the evaluation step of the RBF coefficients λ_i and γ_k in equation (2) is done only once at the beginning of the simulation and stored.

Furthermore, if the modal shape extraction defined during the off-line FEM resolution leads to an excessive density of points compared to the requirement, a greedy algorithm can be used to select only a subset. The implemented method, inspired by Rendall and Allen (2009) iteratively selects a subset of structure points such that the reduced RBF interpolation generates a relative error in displacement lower than a tolerance T_{RBF} over the entire set of points available. For each mode, initially a very small set of points is chosen: the nodes with null displacement or the four nodes with the smallest displacement if no such nodes exist, and the node associated with the maximal displacement. Then, the algorithm iteratively adds the nodes for which the error between the interpolated reduced mode $\psi_{s_i}^{red}$ and ψ_{s_i} is the highest. The process stops when the relative error as described by equation (14) is lower than the tolerance T_{RBF} .

$$E_{RBF,i} = \frac{\max_{j \in [1, n_s]} |\psi_{s_i}^{red}(\mathbf{x}_j) - \psi_{s_i}(j)|}{\max_{j \in [1, n_s]} |\psi_{s_i}(j)|}. \quad (14)$$

By default, $T_{RBF} = 0.001$. The subset of points obtained reduces the size of the
480 matrices to be solved, without greatly compromising the quality of the interpolation or altering the structural response. The computation time is drastically reduced for obtaining the RBF coefficients as well as for the evaluation of fluid node positions, which is performed at least at each time step.

The combination modal/RBF becomes even more attractive as it seamlessly integrates with the interpolation of forces, naturally ensuring the conservation of energy transmitted at the fluid-structure interface (Rendall and Allen, 2009). If \mathbf{H}_S defines the interpolation operator obtained with RBF from the structure surface nodes to the fluid surface nodes, then fluid-domain deformation modes are defined as:

$$\psi_{f_i} = \mathbf{H}_S \psi_{s_i}, \quad \text{for } i \in [1, n_m]. \quad (15)$$

To ensure energy conservation at the interface, the fluid forces are interpolated using $\mathbf{f}_s = \mathbf{H}_S^T \mathbf{f}_f$. With the source term of the modal equation, this naturally gives:

$$\psi_{s_i}^T \mathbf{f}_s = \psi_{s_i}^T \mathbf{H}_S^T \mathbf{f}_f = \underbrace{(\mathbf{H}_S \psi_{s_i})^T}_{\psi_{f_i}} \mathbf{f}_f = \psi_{f_i}^T \mathbf{f}_f. \quad (16)$$

485 The specificity of modal deformations is also exploited for multi-domain configurations (overset) in which one of the domains contains a modal body. Since the deformations are linear combinations of modal shapes invariant during the simulation, each process of a parallel computation can compute the positions of the domain boundary nodes directly by calling the RBF procedure. Thus, overset receives the deformed domain boundary without requiring any inter-process
490 communication.

5. Quasi-static computation of equilibrium

An issue with dynamic motion, both rigid-body and flexible, is that time-accurate resolution of the coupled fluid and body motion equations often requires smaller time steps than the equivalent fixed-body simulations, to keep
495 the simulations stable. This adds to the cost of the simulations and is undesired when time integration is used only to find a steady equilibrium position.

Therefore, we introduce ‘quasi-static’ resolution of the body motion, where the flow at a given instant is used to predict the equilibrium position and shape of the body, which is then gradually brought to this position over several time steps.
500 Subsequently, a new prediction is computed, etc. This procedure rapidly finds the equilibrium positions and, since the motion is not solved time-accurately, the strict time step limits do not apply.

The key element of this method is the prediction of the equilibrium position. While the principle of the quasi-static approach is the same for FSI deformation
505 and for rigid-body motion, this prediction is different for the two. The solution of the modal FSI equations (section 4) is shown in section 5.1 and rigid-body equilibrium is solved in section 5.2. Finally, mesh-related limits for the quasi-static motions are discussed in section 5.3.

5.1. Quasi-static modal resolution

510 For modal FSI, when a steady deformation is expected, the FSI coupling is not solved using Newton’s law (11) at each time step. Instead, the equilibrium amplitude for each mode is estimated at given time instants k . This estimation is based on an approximation of the Jacobian of the forces projected on each mode (F_i) with respect to the amplitudes from the previous steps.

The final steady condition of the system (11) which should be achieved reads:

$$\omega_i^2 q_i = \boldsymbol{\psi}_i^T \mathbf{f} = F_i. \quad (17)$$

A Newton-Raphson-like method is used to estimate a new stable amplitude q_i^{k+1} from previous states (q_i^k, F_i^k) and (q_i^{k-1}, F_i^{k-1}) . Assuming the new state has to verify $\omega_i^2 q_i^{k+1} = F_i^{k+1}$, we eliminate the unknown F_i^{k+1} by linearisation (18) and then deduce (19).

$$\omega_i^2 (q_i^k + dq_i) = F_i^k + \frac{\partial F_i}{\partial q_i} dq_i, \quad \text{with } dq_i = q_i^{k+1} - q_i^k, \quad (18)$$

$$dq_i = \frac{F_i^k / \omega_i^2 - q_i^k}{1 - \frac{1}{\omega_i^2} \frac{\partial F_i}{\partial q_i}}. \quad (19)$$

515 The Jacobian $\frac{\partial F_i}{\partial q_i}$ is approximated by $(F_i^k - F_i^{k-1}) / (q_i^k - q_i^{k-1})$ to compute dq_i which then provides the new amplitude q_i^{k+1} .

The new set of modal amplitudes $\{q_i^{k+1}, i \in [1, n_m]\}$ are applied progressively using an imposed law of deformation, which prevents classical instabilities due to added mass effects (section 4.2). After reaching the new shape and an additional time period to stabilize the fluid load, a new estimation of the dynamic equilibrium is computed again until convergence. Since hydrofoil forces converge rapidly, the corrections are applied often: for the current simulations, the first correction comes after 30 time steps and each subsequent one after 20 time steps, with 10 to 15 time steps for the shape modification (for comparison, the foil traverses one chord length c in 20 time steps). An under-relaxation with a coefficient of 0.5 is applied to the newly computed amplitude, to avoid divergence of the process due to a too rough approximation of the Jacobian.

5.2. Quasi-static positioning approach

Besides deformation, accurate simulation of hydrofoils requires dynamic positioning of the foil. Since the drag of a foil depends strongly on its load (lift and sideforce), simulations are often performed at imposed loads and the foil attitude is adjusted to obtain these loads. The equilibrium position is hard to predict before a simulation, so large changes of the foil attitude may be needed.

Simulating a lifting hydrofoil at a fixed load requires adjusting the orientation of the foil with respect to the flow. To control both the vertical and the lateral force, the rake (pitch) and leeway (yaw) angles are adjusted. Their effect is usually coupled: a change in rake or leeway modifies both the lateral force and the lift force. For this rigid-body motion (section 2.2) the mesh is deformed to change the rake, while the leeway is modified by rotating the entire mesh around the z -axis. The quasi-static positioning was introduced by Ploé (2018).




Figure 6: Surface element for the determination of the vertical and lateral lift slope.

Contrary to the solution-derived Jacobians of section 5.1, the relation between the forces and the angles is estimated analytically here. Let F_y and F_z be the lateral and vertical forces, R_y the rake, and R_z the leeway. For a small section of the hydrofoil surface dS with normal vector \mathbf{n} (figure 6), the surface parallel to the x -axis is:

$$d\tilde{S} = dS \frac{\sqrt{n_y^2 + n_z^2}}{\sqrt{n_x^2 + n_y^2 + n_z^2}}. \quad (20)$$

Let θ be the angle of $d\tilde{S}$ with respect to the horizontal plane: $\theta = \arctan(n_y/n_z)$. Then the change in angle of attack for $d\tilde{S}$ due to changes in the main angles

dR_y and dR_z is:

$$d\alpha = dR_y \cos \theta + dR_z \sin \theta. \quad (21)$$

Using the theoretical 2D lift slope $dC_l = 2\pi d\alpha$ and the dynamic pressure $\frac{1}{2}\rho U^2$, the force change on the surface is $df = \frac{1}{2}\rho U^2 2\pi d\alpha d\tilde{S}$. Finally, the force change is decomposed as $df_y = df \sin \theta$ and $df_z = df \cos \theta$. Integrating these incremental forces over the entire wetted surface S leads to the total change in the forces:

$$\begin{bmatrix} dF_y \\ dF_z \end{bmatrix} = \frac{1}{2}\rho U^2 2\pi \begin{bmatrix} \frac{1}{2} \int_S \sin^2 \theta d\tilde{S} & \frac{1}{2} \int_S \sin \theta \cos \theta d\tilde{S} \\ \frac{1}{2} \int_S \sin \theta \cos \theta d\tilde{S} & \frac{1}{2} \int_S \cos^2 \theta d\tilde{S} \end{bmatrix} \begin{bmatrix} dR_z \\ dR_y \end{bmatrix}. \quad (22)$$

The coefficients $\frac{1}{2}$ before the integrals are added because the hydrofoil has an upper and a lower side, so the projected surface of the foil is integrated twice. Numerically, the integrals are approximated by summing over the wall faces in the fluid mesh.

545 For each quasi-static correction, the new angles are computed from (22) by substituting $[dF_y, dF_z]^T = [F_y^{\text{target}} - F_y^k, F_z^{\text{target}} - F_z^k]$, solving for the angle change, and deducing R_y^{k+1} , R_z^{k+1} like in (18). The same movement procedure and time intervals as for the quasi-static modal resolution are used here, except that no explicit under-relaxation of the computed angle corrections is applied:
550 the 2D lift slope used in (22) is always an overestimation in 3D so the angle corrections are systematically too small, which stabilizes the process.

Depending on the geometry of the hydrofoil, it may not be possible to solve equation (22). For example, a vertical foil (i.e. a rudder) has $\cos \theta = 0$ everywhere so dR_y cannot be computed; a rudder cannot generate vertical forces. A
555 fortiori, for any straight foil, $\cos \theta$ and $\sin \theta$ are constant so the determinant of the matrix in (22) is zero. Since the lift force on a straight foil is normal to its projected surface, such a foil cannot generate arbitrary forces in vertical and lateral directions independently.

5.3. Mesh-related motion limits

560 Since the quasi-static approach is stable, large attitude changes can be applied quickly which is good for the convergence speed. However, rapid ALE mesh motion can perturb the solution, especially the free surface which has to pass through the moving mesh. If the mesh moves more than 0.2 – 0.5 cells per time step, the discretisation of the volume fraction equation smears out the volume fraction at the surface. Besides losing accuracy, this smearing causes mesh
565 refinement over a larger zone at the surface, so the number of cells increases.

Figure 7 shows how solid-body rotation with analytical weighting (1) deforms a layer of free-surface cells. The rotation dR_y over a period dT provokes a vertical motion of the mesh, which has a maximum dz_{max} because the weighting
570 $w(\mathbf{x})$ goes to zero on the outer boundaries. To prevent smearing, dz_{max} cannot exceed $0.5h_{min,FS} dT/dt$ with dt the time step, which puts a limit on dR_y .

This reasoning leads us to impose limits for dR_y , dR_z and the modal coefficients q_i . For weighted deformations, dz_{max} can be analytically linked with the angle changes but for FSI, this is not possible. Therefore, based on several test
575 cases, standard limits of 0.25° per quasi-static cycle are adopted for dR_y and dR_z , and of 0.15 per cycle for dq_i .

Still, this artificial limiting reduces the convergence of the quasi-static procedures. For overset meshing, where only the small overset domain moves, no such limits apply. This is one of several advantages of the overset approach.





Figure 7: Analytical weighted deformation: the effect of a quasi-static pitch rotation on the free-surface mesh.

580 **6. Increased accuracy and flexibility with ‘sock’ meshes**

585 Proper alignment of the mesh with the flow and the foil geometry is important for good accuracy, which is the motivation for using thick boundary layer meshes (section 3.1). However, another alignment issue exists. Hexpress, which generates the initial meshes for the adaptive refinement, starts from a coarse uniform cartesian grid which is adjusted to the geometry by local octree refinement and projection of the grid to the geometry. This volume-to-surface approach gives the best meshes when the original cartesian grid is aligned with the surface. For highly curved hydrofoils however, this is not always possible.



(a) The foil-shaped domain

(b) Initial mesh created in this domain

Figure 8: Example of a sock domain for the Nacra 17 foil (section 7.4).

590 A solution is to use a close-fitting curved overset domain around the foil (figure 8a), in which a curved starter mesh generated using Cadence’s IGG

structured mesher (figure 8b) replaces the cartesian grid. This mesh is then imported in Hexpress, where local refinements are performed as usual to capture the foil geometry. This approach, dubbed ‘sock’ meshing for obvious reasons, was invented by Cadence and Centrale Nantes (Mallol et al., 2019).

595 Since the sock mesh follows the shape of the foil, the total number of cells is reduced because the geometry of the foil is captured more efficiently, especially at the leading and trailing edges (figure 9). In addition, the mesh is of better quality, as it contains less ‘diamond’ cells which degrade the accuracy of the discretised flow equations. And as seen above, the overset domain allows large
600 motions and deformations without imposing limits on the convergence speed.

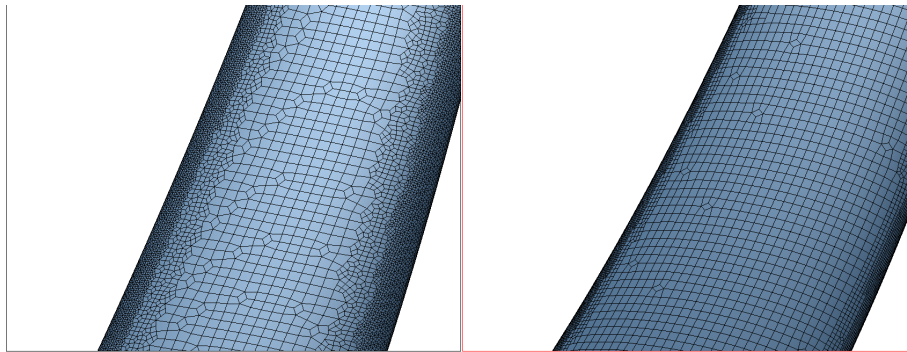


Figure 9: Detail of the surface mesh for an IMOCA foil obtained without (left) and with a sock domain (right).

Finally, the better alignment is an advantage for mesh adaptation, since directional refinement in the flow direction can be used to capture pressure peaks, while diamond cells have to be refined isotropically. The sock meshes are fully hexahedral like standard Hexpress meshes, so they function naturally with
605 adaptive refinement and the adaptation protocol of section 3 is valid. And since the sock domains are narrow (our standard size is $3c$ in axial and $2c$ in lateral direction), the overset boundary is close to the hydrofoil. Therefore, adaptation is essential for accurate overset interpolation (section 3.4).

7. Test cases

610 The developments described in the previous sections are tested on four different test cases with varying shapes, velocities, and motions: a catamaran foil (section 7.1) where the test focuses on the convergence of the quasi-static positioning and the accuracy of adaptive refinement, a small slow windsurfer hydrofoil (section 7.2) that tests the suitability of the adaptation protocol for
615 varying conditions, a monohull foil test (section 7.3) which focuses on the sock mesh, and a Nacra 17 foil (section 7.4) for modal FSI. Except for the last case, no experimental data are available so comparisons are made with non-adapted mesh solutions and potential-flow solutions.

7.1. L-foil grid convergence and accuracy

620 The convergence of the quasi-static positioning is studied on a rigid L-hydrofoil, followed by a mesh dependence study on adapted grids. Since no

experimental results are available, the solutions are compared with simulations on non-adapted meshes.

The case (figure 1) is an enlarged version of a foil for a C-class two-person catamaran. The geometry, produced with the hydrofoil modeller of Ploé (2018), has a chord $c = 0.35\text{m}$, a span of 1.5m , and is simulated at $U = 10\text{m/s}$ with imposed forces $F_y^{\text{target}} = 6000\text{N}$ and $F_z^{\text{target}} = 8000\text{N}$. The density is $\rho = 1026.02\text{kg/m}^3$, the gravity 9.81m/s^2 and the dynamic viscosity is 0.00122kg/(ms) , giving a Froude number $Fr = 5.40$ and a Reynolds number $Re = 2.94 \cdot 10^6$. Turbulence is modelled with $k - \omega$ SST and a wall-modelled boundary condition with $y^+ = 50$ is used. The time step is chosen such that the foil advances one chord length in 20 steps; 800 time steps are run for convergence, which is enough for a high-velocity hydrofoil. Sign conventions are as in figure 10.




Figure 10: Sign conventions for forces and rotations of the L-foil. The hydrofoil advances in the positive X -direction. Figure not to scale.

The simulation setup follows the mesh adaptation protocol of section 3, except that the free-surface threshold (section 3.2) is chosen as $T_{r,S} = c/128$ to better agree with the non-adapted meshes described below. A series of meshes is produced by varying the Hessian threshold $T_{r,H}$ between $4c$ and $c/16$.

For comparison, a series of non-adapted meshes is created with Hexpress, featuring uniform grids of varying sizes on the foil surface with two supplementary levels of refinement on the leading and trailing edge. Around the free surface, a refinement box is used with cell size $c/128$ on the second finest grid (the free-surface cell sizes vary for this series); the box has a height of $0.7c$ to ensure the good resolution of the volume fraction even in the case of mesh deformation. This setup represents the existing state of the art for FINE/Marine hydrofoil simulations, except that the free surface is normally captured with adaptive refinement instead of a fixed refinement box (see section 7.3).

Convergence of the quasi-static positioning. The rake (R_y) and leeway (R_z) are adjusted during the computation, to attain the requested forces in y - and z -direction (see section 5.2). The convergence of the forces and rotations during a typical simulation is shown in figure 11.

This figure shows the alternation of linear angle modifications over 10 time steps, followed by steady intervals of 10 time steps to let the forces stabilise. During the rotations, the forces change notably and the acceleration at the




Figure 11: Time-convergence of the adapted-grid L-foil simulation for $T_{rH} = c/8$. Top: rake R_y and leeway R_z , middle: targeted forces F_y and F_z , bottom: resistance F_x .

beginning and the end of each rotation provokes spikes in the force due to added mass effects. This confirms the importance of having steady periods before estimating a new equilibrium position.

The rake increases to about 1.6° nose-up. Until 0.25s, this increase is constrained by the limit of $0.25^\circ/\text{step}$ (section 5.3) and afterwards, the natural under-relaxation of the procedure is visible: 6 more steps are needed before R_y is close to convergence. The leeway changes never reach the rotation limit and the target F_y , which is most affected by R_z , is quickly reached. However, since the shaft of the foil is not vertical, nose-up rake also increases F_y . The quasi-static procedure compensates for the increasing rake with negative leeway, which eventually reaches -0.8° . Due to this coupling, the convergence speed for R_z is also determined by the limits on R_y .

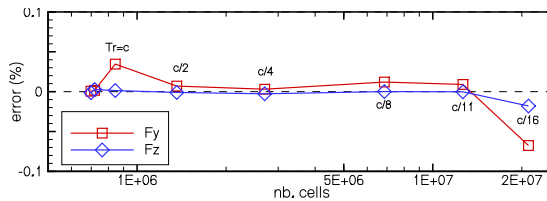


Figure 12: Relative errors $(F - F_{\text{target}})/F_{\text{target}}$ in the controlled forces F_y and F_z for the L-foil series of adapted meshes.

Figure 12 shows that the position control is accurate: in all but one case

the vertical and sideways forces are within 0.05% of the target values, much less than the numerical uncertainty in the non-controlled quantities (see below). The angles converge after about 400 time steps which is acceptable for hydrofoils, whose forces typically need 600-800 time steps to converge, even for fixed positions. Therefore, the quasi-static positioning maintains both the accuracy and the efficiency of fixed-position simulations.




Figure 13: Cuts at $y = 0.08$ for the L-foil, showing (top to bottom) the mesh for $T_{rH} = c/2$, $c/4$, $c/8$ and $c/16$.

Mesh dependence and numerical accuracy. A part of the series of adapted meshes created by varying T_{rH} is shown in figure 13. As the cells of the original grid can only be divided in powers of two, these meshes are not strictly similar, but they are close enough to allow mesh convergence studies.

Figure 14 shows the convergence of F_x with the mesh size. For the adapted meshes, the convergence is excellent: fast, smooth, and monotonic. This is difficult to obtain on unstructured meshes, as shown by the non-adapted series which also exhibit monotonic convergence, but with small irregularities in the slope. Furthermore, the adapted meshes are more accurate: the non-adapted meshes require about 4 times more cells to converge to the same value for F_x . Figure 13 explains this good accuracy: fine cells in the pressure peaks on the leading and trailing edges play a major role in obtaining accurate drag predictions for hydrofoils. Still, the two series converge to the same value which confirms the reliability of both.




Figure 14: Convergence of the resistance for the L-hydrofoil, with uncertainty following Eça and Hoekstra (2014). Comparison of adapted and non-adapted mesh solutions.




Figure 15: Convergence of the incidence angles for the L-foil, with uncertainties (R_z uncertainty on the finest non-adapted grid: 67%).

Like in previous works (Wackers et al., 2017, 2022; Korkmaz et al., 2023), the procedure of Eça and Hoekstra (2014) is used to estimate the numerical uncertainty of the resistance. This method computes uncertainties based on a range of fits for the force data as a function of the grid size. For the adapted-mesh solutions, the threshold T_{rH} can be used as a convenient measure of this mesh size; for the non-adapted meshes, the number of cells to the power $-\frac{1}{3}$ is used. To determine the uncertainty on a given grid, the force from this grid and all the coarser ones are used in the extrapolation.

In figure 14, the errorbars for coarser meshes overlap the ones for the finer meshes, which shows that the estimations are reliable. The adapted-mesh solutions produce the same uncertainty with about four times less cells than the non-adapted meshes. However, the uncertainties are relatively high: the adapted meshes give 5% uncertainty on the finest 21M cells grid, whereas ship simulations lead to about 1% uncertainty at 6M cells (Wackers et al., 2022). This is one of our motivations to consider sock meshes (see section 7.4).

The forces F_y and F_z are kept constant by the quasi-static procedure, so the incidence angles R_y and R_z are converging quantities (figure 15). The adapted-grid results show good but not strictly monotonic convergence, with lower uncertainties than for the resistance (which was also observed by Wackers et al. (2017)). For the non-adapted meshes, the convergence is oscillatory so the uncertainties are high; it is unclear if the non-adapted series will converge to the same angles as the adapted grids. Overall, the adapted grids produce accurate and reliable results on much coarser meshes than non-adapted grids.

7.2. Kitefoil: low velocity, fixed position

To test the versatility of the hydrofoil refinement protocol, Richeux (2022) applied it blindly to a range of test cases. The kitesurf foil shown here is as different as possible from the L-foil: a small symmetric hydrofoil in a fixed position, very close to the water surface, at a velocity that is too low to permit flight. Again, no experiments exist for this geometry, so the solution is compared with non-adapted mesh results and with a potential-flow solution.

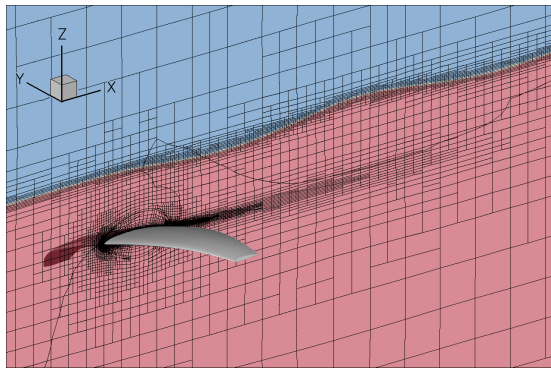


Figure 16: Kitefoil with adapted mesh ($T_{rH} = c/8$) and volume fraction in the $y = 0$ plane.

The geometry (Perali et al., 2024) resembles the main wing of the iQFOIL windsurfer selected for the 2024 Olympics, with a maximum chord $c = 0.12\text{m}$ and a span of 0.474m . The attitude is fixed at zero pitch and yaw, 20° right roll, and an immersion of $0.83c$. The velocity is $U = 0.922\text{m/s}$, the density

$\rho = 1027 \text{ kg/m}^3$, the Froude number $Fr = 0.85$ and the Reynolds number $Re = 8.134 \cdot 10^4$. Turbulence is modelled with $k - \omega$ SST; a wall-resolved boundary condition with $y^+ = 1$ is used for this low Reynolds number. No transition model is used, even though the actual flow is probably partially laminar. The foil advances a chord length in 20 time steps and 3000 steps are run for convergence.

The geometry and a typical mesh are shown in figure 16. At this low speed close to the surface, the hydrofoil generates a wave train which influences the forces on the foil. To resolve this wave train, the limiting box (section 3.3) is placed at $10c$ behind the trailing edge. Otherwise, the standard hydrofoil refinement protocol is used, with free-surface threshold $T_{rS} = c/64$. Non-adapted meshes for comparison are made with the same approach as in section 7.1.




Figure 17: Convergence of the resistance and lift for the kitefoil (F_z uncertainty on the finest non-adapted grid: 26%). F_y is omitted as it is similar to F_z .

For the resistance, despite the difference with the L-foil case, the mesh-convergence on adapted meshes is similar (figure 17, c.f. figure 14). The uncertainty as a function of T_{rH}/c has the same order of magnitude, which shows that T_{rH} is a good choice to control the desired accuracy. The non-adapted series is not obviously converging to the same value, although the errorbars overlap.

The lift F_z has a non-monotone convergence, caused probably by the wave train which reflects on the outflow boundary, as well as on any region where the mesh gets coarser. The marked difference with the non-adapted meshes however is due to a lack of resolution on the leading and trailing edges: when the middle non-adapted mesh was remade with two more levels of refinement on the foil

edges, the lift dropped to about 6.2N (not shown in figure 17). This again shows the importance of adaptive refinement: without adaptation, the solutions seem to converge well but may actually be far from grid-independent.

745 Figure 18 compares the adapted-grid RANS wave pattern with a potential flow simulation using the PUFFin solver (Perali et al., 2024). The potential code, which lacks viscous damping, predicts larger amplitudes as expected. However, the shapes of the wave patterns, the wave lengths, and the asymmetry due to the roll angle, agree closely. This indicates that both results are reliable.




Figure 18: Kitefoil wave pattern, comparison of RANS $T_{rH} = c/8$ (left) and potential flow (right) from Perali et al. (2024).

750 7.3. IMOCA foil and sock mesh

To evaluate the sock mesh approach, the refinement protocol is applied to a hydrofoil of an offshore racing sailboat. The foil pierces the surface at two different locations (Figure 19) and the geometry is notably wider (5.5m) and more slender, but also more complex than the previous geometries. This makes
755 the sock mesh technique presented in section 6 particularly relevant.

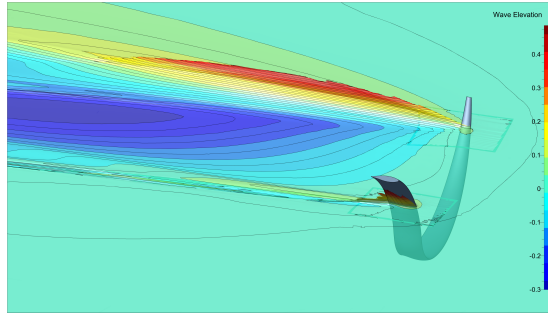


Figure 19: Wave field around the IMOCA foil.

As this test is based on an industrial case, the solutions are compared with the results of the current state-of-the-art industrial approach for hydrofoil simulations. This approach uses a non-adapted mesh around the hydrofoil, which is placed in a box-shaped overset domain. Mesh adaptation is used, but
760 only to capture the free surface. The reference simulations are set up by the FINE/Marine integrated tool C-Wizard, which is commonly used by yacht design teams and therefore relevant for comparison in the scope of industrial use.

The case concerns an IMOCA Open60 monohull foil, placed in conditions close to its position when sailing at beam reach: low immersion and a velocity of $U = 14\text{m/s}$. The attitude is fixed at zero pitch and roll, the density is $\rho = 1025.07\text{kg/m}^3$, the Froude number $Fr = 5.82$ and the Reynolds number is $Re = 8.105 \cdot 10^6$. Turbulence is modelled with $k - \omega$ SST and a wall-modelled boundary condition with $y^+ \approx 51$ is used. To accurately capture the waves, the limiting box for the mesh refinement was extended to $10c$ behind the foil, like for the kitefoil. The Hessian threshold T_{rH} is varied from $4c$ to $c/8$.




Figure 20: Cuts in an x -plane through the IMOCA foil, showing the sock mesh with $T_{rH} = c/4$ (top) and the overset box mesh (bottom). The inserts show the full overset domains.

Mesher. The meshes of the two methods are compared in figures 20 and 21. For the sock-mesh approach, $T_{rH} = c/4$ is shown, while the overset box mesh uses the C-Wizard *coarse* preset. Both lead to refined meshes containing 6.6M elements at the end of the computation, which makes these settings relevant in a foil design process. As shown in figure 20, the distribution of cells is fundamentally different: for the overset box mesh, the geometry crosses the grid directions diagonally while the sock-mesh cells are well-aligned with the foil, which results in a better-quality mesh (c.f. figure 9). At the leading and trailing edge, the sock mesh is at least twice finer than the overset box mesh, so pressure peaks can be captured with accuracy. Since this is similar to previous test cases, it is not shown.

The detail in figure 21 shows the close proximity of the foil and the overset interface for the sock mesh. As explained in section 3.4, the continuity of the refinement criterion over the interface ensures locally equal cell sizes in both domains, which is essential for the interpolation accuracy. At the free surface, the number of cells is higher in the (diagonal) sock mesh than in the (free-surface aligned) overset box mesh. However, since the free-surface zone is small, this is an acceptable price to pay for the better alignment on the foil surface.




Figure 21: Zoom of figure 20, focusing on the foil tip.

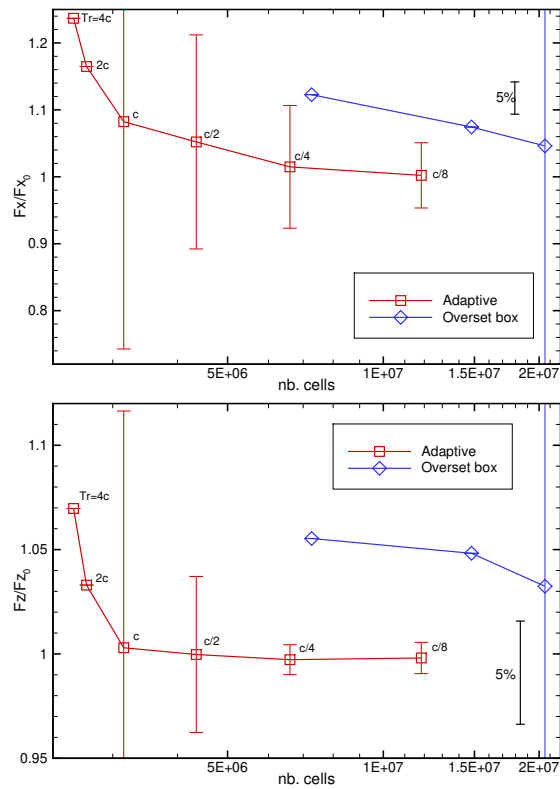


Figure 22: Convergence of the resistance and lift for the IMOCA foil. F_y is omitted as it is similar to F_z .

790 *Convergence of the forces.* Figure 22 shows the forces, scaled with the finest sock-mesh results for confidentiality reasons. The convergence behaviour is similar to the previous sections, while the sock mesh leads to smaller uncertainties: at $c/8$, the uncertainty is below 5% for the resistance and below 1% for the lift, which is about two times less than the L-foil of section 7.1.

795 The overset box series is far from mesh convergence, although all sock-mesh results lie within the uncertainty range for the *fine* mesh preset. And while the $c/4$ sock mesh has three times less cells than this grid, its results are more accurate. These low numerical uncertainties obtained with affordable CPU time make the sock-mesh approach interesting in the scope of industrial use.

7.4. FSI: flexible NACRA foil

800 Finally, the fluid-structure interaction capability and the quasi-static modal procedure of section 5.1 are tested on a Z-shaped foil of the Olympic sailing catamaran Nacra 17. This foil was extensively studied in Marimon Giovannetti et al. (2022) and Knudsen et al. (2024), including experimental testing of the foil in a cavitation tunnel. The case has no free surface; the foil is attached to
805 one of the tunnel walls, which are represented in the simulation.

The foil is fixed at 3° pitch, 0° roll and velocity $V = 7\text{m/s}$, which gives a foil deformation $0.890c$ (with maximum chord $c = 0.2\text{m}$) in the experiments. A sock mesh (figure 8) is used to accompany this fairly large deformation. Turbulence is modelled using $k-\omega$ SST with a wall-modelled boundary condition and y^+ close
810 to 50. For all quasi-static and RBF interpolation parameters, the default values of sections 2.2, 4.3, and 5.1 are chosen. The foil travels 1 chord length in 10 time steps and 2500 steps are run in total. The forces are non-dimensionalised with $\frac{1}{2}\rho V^2 A$, where A is half the wetted surface.




Figure 23: Eigenmodes $\phi_1 - \phi_5$ (left to right) of the Nacra 17 foil. The colours indicate the magnitude of the deformation; the undeformed shape is shown in grey. The eigenfrequencies $f_1 - f_5$ are 12.98, 31.23, 59.95, 141.11, and 217.95Hz.

815 The Abaqus model of Marimon Giovannetti et al. (2022) has been used here to conduct the modal analysis (section 4.1). Since the exact structure and mechanical properties of the NACRA foil are kept secret by the manufacturer, the model assumes a shell with a single spar, made of isotropic material with density $\rho = 3000\text{kg/m}^3$, Young's modulus² $E = 65\text{GPa}$ and Poisson coefficient

²The model we received has a higher Young's modulus than the article ($E = 30\text{--}55\text{GPa}$).

820 $\sigma = 0.35$. It is meshed using shell elements of typical size $d_{elmts,FEA} = 0.004m$.
 The internal structure is regarded as highly uncertain and later, Knudsen
 et al. (2024) changed to a variable wall-thickness model with $\rho = 1600\text{kg/m}^3$,
 $E = 61\text{GPa}$ and $\sigma = 0.33$.

825 *FSI convergence.* Figure 23 shows the shapes of the eigenmodes which form the
 basis of the reduced-order structural model. To reduce the computational time,
 only the first five modes obtained in this FEM simulation were used, since in our
 experience, for typical sailboat foils, eigenmodes with frequencies above 200Hz
 contribute little to the global deformation. The mode ψ_1 is the main bending
 mode, ψ_2 is a twisting of the upper foil which rotates the lower part fore-and-aft,
 ψ_3 and ψ_4 are higher-order bending modes, and ψ_5 is a pitch-changing twist
 830 of the lower foil.




Figure 24: Time-convergence of the Nacra 17 simulations for $T_{rH} = c/8$. Top: modal amplitudes $q_1 - q_5$ for the flexible foil (centre: zoomed), bottom: resistance coefficient C_{Fx} .

835 As can be seen in figure 24, the first mode is dominant, with a small contribu-
 tion of the second bending mode q_3 . The effect of the other modes is negligible,
 which validates the choice to include not more than 5 modes in the model. The
 convergence of the modal amplitudes is smooth, thanks to the under-relaxation
 of 0.5. The resistance however has strong oscillations, like for the quasi-static
 positioning (figure 11). The smooth convergence of the rigid-body resistance
 confirms that these peaks are due to the foil deformation.

840 However, the quasi-static procedure keeps the modal FSI stable and the
 convergence to a steady state takes the same time as for the rigid case (about
 4s). Since the time steps are the same for the two cases, this shows that FSI
 can be taken into account for a low additional computational cost, thanks to
 the quasi-static modal approach.

Flexible versus rigid simulations. The deformation of the NACRA foil and the
 pressure on the suction side are shown in figure 25. Since the foil has a pitched-

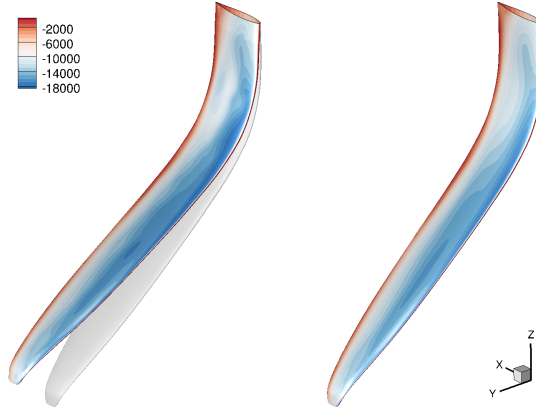


Figure 25: Comparison of foil shapes and suction-side pressure for the deformed (left) and rigid (right) Nacra 17 hydrofoil.

845 up attitude, bending increases the angle of attack, as shown by the stronger suction peak compared with the rigid case. This, and the bottom section which becomes more horizontal, increase the lift F_z (+23% vs. rigid) and slightly reduce the sideforce F_y (-5% vs. rigid). Due to the higher loading, the (induced) drag increases (+6% vs. rigid). See table 2 for the results.

850 The mesh convergence in figure 26 is oscillating for the flexible foil, while the rigid foil has a smooth convergence. This may be due to the noisier time signals for the flexible foil. However, the numerical uncertainties are similar for the two cases and the convergence is globally the same. The foil tip displacement u_{tip} converges exactly like F_z , which proves that the (linear) structure model is
855 solved accurately. Overall, the precision of the simulation results is not affected by the FSI. However, the significant effect of FSI on the forces is confirmed.

Table 2: Comparison of simulation ($T_{rH} = c/16$) and experimental results for the Nacra 17.

	C_{Fx}	C_{Fy}	C_{Fz}	u_{tip}/c
Experimental	0.0132	0.252	0.246	0.890
Simulation	0.0168	0.268	0.279	0.866
Simulation (rigid)	0.0156	0.284	0.227	–
Simulation (Knudsen et al., 2024)	0.0158	0.284	0.297	1.095

860 *Comparison with experiments.* For the flexible foil at $T_{rH} = c/16$, table 2 shows a deformation at the tip of the foil within 4% of the experimental result, but the difference is 27% for the drag and 13% for the lift. Since the numerical uncertainty for this mesh is 2.5% for the drag and around 1% for the side and lift forces and the deformation, this could indicate a significant modelling error.

865 However, the results are coherent. Our simulation overestimates the experimental forces and underestimates the deformation, which indicates that the structure model is too rigid, due to the choice of the material properties, the constant shell thickness over the span, the simplistic spar, or the boundary conditions representing the attachment in the tunnel. Based on different tests, Marimon Giovannetti et al. (2022) reach the same conclusion, even using




Figure 26: Convergence of the resistance for the Naca 17 hydrofoil, with uncertainty.

$E = 55\text{GPa}$. The reworked model of Knudsen et al. (2024) is probably more accurate, since their result overestimates both the forces and the deformation.

870 Furthermore, all simulations overestimate the lift, probably due to an alignment issue in the experiments, but the ratio 1.04 between the lift and side-force is close to the experiments (0.98) and almost the same as Knudsen et al. (2024)'s simulation (1.05). Compared to the rigid-foil ratio of 0.80, this again underscores the importance of FSI for accurate force predictions. Finally, no
875 transition model is used, while the actual flow is probably partially laminar. This, and the higher induced drag, explain the overestimation of C_{Fx} , which Knudsen et al. (2024) also observe.

The linearised deformation does not preserve the structure exactly and actually increases its surface area, but for the current case, this increase is only
880 1% so it is not the main cause of the differences. Thus, this test cannot be seen as a formal validation but, given the uncertainty in the structure model, the agreement with experiments is satisfactory.

8. Conclusion and discussion

This paper discusses numerical techniques for the efficient RANS compu-
885 tation of the flow around lifting hydrofoils. These require a fine resolution of the flow around the foil: above $10M$ cells may be needed to get high precision. Also, fluid-structure interaction has a significant influence on the forces. While this does not make hydrofoil simulation impossible, it shows that computational efficiency is crucial to make these simulations possible for daily use in a ship
890 design environment.

Mesh adaptation. In our view, efficient hydrofoil simulation is impossible without adaptive mesh refinement. For free-surface capturing, due to the dynamic positioning, the foil deformation and the large waves that hydrofoils generate, the water surface position cannot be predicted in advance. Thus, without adap-
895 tation, large boxes of fine cells are required: in the L-foil test of section 7.1, the finest non-adapted mesh contains $13.9M$ cells (62% of the total) to capture the surface. For comparison, the $T_{rH} = 4c$ adapted mesh, which has full free-surface refinement, has only $700k$ cells. Furthermore, without adaptation, computations may be erroneous if the refinement zone is not placed correctly.
900 For these reasons, free-surface mesh adaptation is already standard practice in the industrial use of FINE/Marine.

Adaptation around the hydrofoil is equally important, and anisotropy is crucial: the ‘sock’ mesh, a flow- and geometry-aligned base mesh that enables anisotropic refinement, is required for good uncertainty at reasonable costs.
905 The IMOCA case in section 7.3 shows 4.8% estimated uncertainty for $T_{rH} = c/8$, which is three times better than the adapted non-sock mesh for the L-foil (section 7.1) at the same threshold. Hessian-based adaptation provides not only accuracy but also safety, as shown by the kitefoil of section 7.2 where the non-adapted meshes ‘converge’ to an incorrect solution.

910 The sock approach requires small overset domains with boundaries close to the foil surface, so accurate interpolation between the overset and background domain is needed. Mesh adaptation is the only way to ensure equal cell sizes in the interpolation regions of both domains without requiring a large box of fine cells in the background grid.

915 This work demonstrated a standardised mesh adaptation protocol for hydro-
foils: despite their different natures, all the test cases in section 7 were performed
with the same guidelines for the computational settings. This paves the way
to automated, less expensive, and less technical hydrofoil CFD. De facto, this
makes simulation accessible to a larger number of users.

920 *Modal FSI.* The modal approach for FSI combines fluid simulation with a linearised
eigenvalue decomposition of the structural finite-element model, created
off-line with an external structural solver. For hydrofoil simulation, this offers
several advantages. First, there is no direct coupling between the fluid and
structure solvers; any structure solver and any model can be used. Thus, there
925 is no limitation for the use of anisotropic material behaviour, which is essential
to model composite structures correctly. Furthermore, once the linearised
reduced-order model is constructed, it can be evaluated almost instantly which
means that the computational overhead for the model FSI is close to zero. The
main impact on the fluid simulation cost is actually the RBF deformation of
930 the mesh, which can be sped up by the greedy approach. In addition, when
using quasi-static deformation, the mesh is updated at most once per time step
(instead of each non-linear iteration), which drastically limits the overhead.

The main disadvantage of the modal approach is, that it is linear without any
coupling between the modes, and therefore limited to moderate deformations.
935 Modern foils can bend a lot (up to 20% of their span for an IMOCA foil), in
which case the linearised model may be inaccurate. For this reason, we study a
one-dimensional beam solver as an alternative for the modal approach. Still, the
modal FSI provides good results for small to medium deformations (the Nacra
17 case of section 7.4 has an error of only 1% in the foil surface area) and since
940 the computation is not much more expensive than for rigid bodies, the approach
can be recommended for industrial applications.

Quasi-static positioning and deformation. To efficiently find equilibrium solu-
tions for the equations of rigid and flexible body motion, we propose quasi-static
fixed-point algorithms, where in each step an approximate equilibrium position
945 is estimated from the current flow. The body is then gradually moved to this
new position. The central part of the algorithm is estimating the dependence of
the forces on the position, to deduce the equilibrium position or shape. A wide
range of options exist: here, the estimation for the rigid-body positioning is
based on the 2D analytical lift slope, integrated over the hydrofoil surface. For
950 the foil deformation, the dependence is extracted with finite differences from the
two previous solutions. These approximations do not have to be very accurate;
for example, the rigid-body estimation ignores the presence of the free surface.

The quasi-static procedure eliminates the time step constraints for the time-
accurate solution of body motions and the tests in section 7.1 and 7.4 show that
955 the simulations converge as fast as for rigid bodies, which means again that the
computational overhead for body motion is negligible.

Overset meshes. Two meshing approaches to handle hydrofoil motion are pre-
sented: mesh deformation and overset meshes. Deformation is successful for
smaller foil motions and it has the advantage that the mesh is one block, so no
interpolation is needed and the discretisation is fully conservative. However, the
960 mesh deformation limits the amplitudes and the speed with which the quasi-
static procedures can move the foil, since too fast deformations smear out the

water surface. Also, the single-block Hexpress meshes have diagonal ‘diamond’ cells wherever the foil is not aligned with the mesh directions. These cells reduce the accuracy of the solution.

By combining overset meshing with foil-aligned ‘sock’ initial meshes, only the small overset domain has to move, which means that larger and faster motions are possible. Furthermore, the body-aligned sock mesh makes efficient anisotropic refinement possible, as noted above. For these reasons, we recommend the sock approach as part of a standardised hydrofoil simulation method.

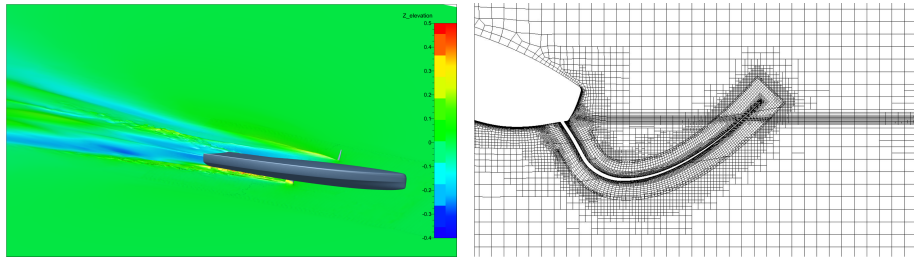


Figure 27: Simulation of a full yacht hull with hydrofoils.

The paper shows the strong synergy of these approaches. Thus, the efficiency of a hydrofoil simulation does not come from one technique, but from a careful coupling of methods that help each other to be efficient. Combined with a well-tested, reliable Navier-Stokes solver, these methods result in a simulation approach which can provide high-accuracy simulation of hydrofoils in an industrial design environment.

Upcoming challenges are the simulation of a full sailing yacht with foils and appendages (Figure 27) and the modelling of dynamic motion in waves. For the former, simulation accuracy must be preserved at the junction between the foil overset domain and the hull in the background domain. Furthermore, the mesh adaptation protocols for hulls and hydrofoils (intended for vastly different length scales) must be combined into one approach. For the latter, stable and efficient time-accurate simulation is crucial. The basis of rapid FSI through well-chosen efficient structure models, and efficiency and reliability provided by mesh adaptation, remains important for these advanced applications.

Acknowledgements

We thank the Institut Carnot MERS and the research directorate of Centrale Nantes for funding the ORUP project, in which the hydrofoil adaptation protocol was developed.

Furthermore, we gratefully acknowledge Dr. Patrick Ploé whose thesis forms the basis of section 5.2, Dr. Paolo Perali who provided the PUFFIn results in section 7.2, Mr. Antoine Koch who co-designed the IMOCA foils, and Dr. Laura Marimon Giovannetti and her colleagues who shared with us their Nacra 17 geometry and its structure model.

995 **References**

- Alauzet, F., Frazza, L., 2021. Feature-based and goal-oriented anisotropic mesh adaptation for RANS applications in aeronautics and aerospace. *J Comput Phys* 439, 110340.
- Balan, A., Park, M., Anderson, K., Kamenetskiy, D., Krakos, J., Michal, T.,
1000 Alauzet, F., 2020. Verification of anisotropic mesh adaptation for turbulent simulations over ONERA M6 wing. *AIAA J* 58, 1–16.
- Banks, J., Marimon Giovannetti, L., Turnock, S.R., Boyd, S.W., 2014. Developing tools for assessing bend-twist coupled foils, in: 17th Numerical Towing Tank Symposium, Marstrand, Sweden.
- 1005 Bayona, V., 2019. An insight into RBF-FD approximations augmented with polynomials. *Comput Math Applic* 77, 2337–2353.
- Biancolini, M.E., Viola, I.M., Riotte, M., 2014. Sails trim optimisation using CFD and RBF mesh morphing. *Comput Fluids* 93, 46 – 60.
- de Boer, A., van der Schoot, M.S., Bijl, H., 2007. Mesh deformation based on
1010 radial basis function interpolation. *Comput Structures* 85, 784–795.
- Bos, F.M., van Oudheusden, B.W., Bijl, H., 2013. Radial basis function based mesh deformation applied to simulation of flow around flapping wings. *Comput Fluids* 79, 167–177.
- Brizzolara, S., 2019. CFD investigation of cavitation-ventilation coupling on
1015 surface-piercing supercavitating hydrofoils, in: MARINE 2019, Gothenburg, Sweden.
- Duvigneau, R., Visonneau, M., 2003. On the role played by turbulence closures in hull shape optimization at model and full scale. *J Mar Sci Tech* 8, 11–25.
- Eça, L., Hoekstra, M., 2014. A procedure for the estimation of the numerical
1020 uncertainty of CFD calculations based on grid refinement studies. *J Comput Phys* 262, 104–130.
- Eskilsson, C., Bensow, R.E., 2012. A mesh adaptive compressible Euler model for the simulation of cavitating flow, in: MARINE 2011, Lisbon, Portugal.
- Franke, R., 1982. Scattered data interpolation: tests of some methods. *Math
1025 Comput* 38, 181–200.
- George, P.L., Borouchaki, H., 1998. *Delaunay Triangulation and Meshing - Application to Finite Elements*. Hermes.
- George, P.L., Hecht, F., Vallet, M.G., 1991. Creation of internal points in Voronoi’s type method. Control adaptation. *Adv Eng Software* 13, 303–312.
- 1030 Groth, C., Cella, U., Costa, E., Biancolini, M.E., 2019. Fast high fidelity CFD/CSM fluid structure interaction using RBF mesh morphing and modal superposition method. *Aircraft Eng Aerospace Techn* 91, 893–904.
- Hardy, R.L., 1990. Theory and applications of the multiquadric-biharmonic method 20 years of discovery 1968–1988. *Comput Math Applic* 19, 163–208.

- 1035 Harwood, C.M., Young, Y.L., Ceccio, S.L., 2016. Ventilated cavities on a surface-piercing hydrofoil at moderate Froude numbers: cavity formation, elimination and stability. *J Fluid Mech* 800, 5–56.
- Knudsen, S., Marimon Giovannetti, L., Legarth, B., Walther, J., 2024. Dynamic fluid structure interaction of NACRA 17 foil. *J Mar Sci Eng* 12, 237.
- 1040 Korkmaz, K.B., Werner, S., Bensow, R., 2023. Investigations on experimental and computational trim optimisation methods. *Ocean Eng* 288, 116098.
- Leroyer, A., Barré, S., Kobus, J., Visonneau, M., 2008. Experimental and numerical investigations of the flow around an oar blade. *J Mar Sci Techn* 13, 1–15.
- 1045 Leroyer, A., Visonneau, M., 2005. Numerical methods for RANSE simulations of a self-propelled fish-like body. *J Fluid Struct* 20, 975–991.
- Loseille, A., Dervieux, A., Alauzet, F., 2010. Fully anisotropic goal-oriented mesh adaptation for 3D steady Euler equations. *J Comput Phys* 229, 2866–2897.
- 1050 Mallol, B., Guisnel, X., Leroyer, A., Wackers, J., 2019. Several leaps forward for CFD simulations of hydrofoils, in: *MARINE 2019*, Gothenburg, Sweden.
- Marimon Giovannetti, L., Farousi, A., Ebbesson, F., Thollot, A., Shiri, A., Eslamdoost, A., 2022. Fluid-structure interaction of a foiling craft. *J Mar Sci Eng* 10, 372.
- 1055 Matveev, K., Wheeler, M., Xing, T., 2019. Numerical simulation of air ventilation and its suppression on inclined surface-piercing hydrofoils. *Ocean Eng* 175, 251–261.
- Menter, F., 1994. Two equation eddy-viscosity turbulence modeling for engineering applications. *AIAA J* 32, 1598–1605.
- 1060 Mouton, L., Leroyer, A., Deng, G.B., Queutey, P., Soler, T., Ward, B., 2018. Towards unsteady approach for future flutter calculations. *J Sailing Tech* 7, 1–19.
- Park, M., Balan, A., Clerici, F., Alauzet, F., Loseille, A., Kamenetskiy, D., Krakos, J., Michal, T., Galbraith, M., 2021. Verification of viscous goal-based anisotropic mesh adaptation, in: *AIAA SciTech Forum*, Anaheim, CA.
- 1065 Perali, P., 2023. Etude numérique et expérimentale des écoulements autour d’un profil flexible, incluant le régime cavitant. Ph.D. thesis. Ecole Centrale de Nantes. France.
- Perali, P., Sacher, M., Leroux, J.B., Wackers, J., Augier, B., Hauville, F., Bot, P., 2024. Performance prediction of a hydrofoil near the free surface using low (BEM) and high (RANS) fidelity methods. Submitted , 1–29.
- 1070 Perić, M., 2022. Prediction of tip vortex flows on propeller blades using adaptive local grid refinement, in: *24th Numerical Towing Tank Symposium*, Zagreb, Croatia.

- 1075 Ploé, P., 2018. Surrogate-based optimization of hydrofoil shapes using RANS simulations. Ph.D. thesis. Ecole Centrale de Nantes. France.
- Prasad, B., Hino, T., Suzuki, K., 2015. Numerical simulation of free surface flows around shallowly submerged hydrofoil by OpenFOAM. *Ocean Eng* 102, 87–94.
- 1080 Queutey, P., Visonneau, M., 2007. An interface capturing method for free-surface hydrodynamic flows. *Comput Fluids* 36, 1481–1510.
- Rendall, T.C.S., Allen, C.B., 2009. Efficient mesh motion using radial basis functions with data reduction algorithms. *J Comput Phys* 228, 6231–6249.
- Rhie, C.M., Chow, W.L., 1983. A numerical study of the turbulent flow past an isolated airfoil with trailing edge separation. *AIAA J* 17, 1525–1532.
- 1085 Richeux, J., 2022. Towards the automation of adaptive RANS simulations for hydrofoils. M.Sc. thesis. Ecole Centrale de Nantes. France.
- Sacher, M., Durand, M., Berrini, É., Hauville, F., Duvigneau, R., Le Maître, O., Astolfi, J.A., 2018. Flexible hydrofoil optimization for the 35th America’s Cup with constrained EGO method. *Ocean Eng* 157, 62–72.
- 1090 Schaback, R., 1995. Creating surfaces from scattered data using radial basis functions, in: Dæhlen, M., Lyche, T., Schumaker, L.L. (Eds.), *Mathematical Methods in Computer Aided Geometric Design III*, Vanderbilt University Press. pp. 477–496.
- 1095 Serani, A., Pellegrini, R., Diez, M., Wackers, J., Jeanson, C.E., Queutey, P., Visonneau, M., 2019. Adaptive multi-fidelity sampling for CFD-based optimization via radial basis function metamodels. *Int J Comp Fluid Dyn* 33, 237–255.
- Wackers, J., 2019. Adaptivity for complex flows. HDR thesis, Université de Nantes, France.
- 1100 Wackers, J., Deng, G., Raymond, C., Guilmineau, E., Leroyer, A., Queutey, P., Visonneau, M., 2022. Adaptive grid refinement for ship resistance computations. *Ocean Eng* 250, 110969.
- Wackers, J., Deng, G.B., Guilmineau, E., Leroyer, A., Queutey, P., Visonneau, M., Palmieri, A., Liverani, A., 2017. Can adaptive grid refinement produce grid-independent solutions for incompressible flows? *J Comput Phys* 344, 364–380.
- 1105 Wackers, J., Koren, B., Raven, H.C., van der Ploeg, A., Starke, A., Deng, G.B., Queutey, P., Visonneau, M., Hino, T., Ohashi, K., 2011. Free-surface viscous flow solution methods for ship hydrodynamics. *Arch Comput Meth Eng* 18, 1–41.
- Wackers, J., Leroyer, A., Deng, G.B., Queutey, P., Visonneau, M., 2012. Adaptive grid refinement for hydrodynamic flows. *Comput Fluids* 55, 85–100.

- 1115 Wang, Z., Li, L., Cheng, H., Ji, B., 2020. Numerical investigation of unsteady cloud cavitating flow around the Clark-Y hydrofoil with adaptive mesh refinement using OpenFOAM. *Ocean Eng* 206, 107349.
- Yvin, C., Leroyer, A., Visonneau, M., Queutey, P., 2018. Added mass evaluation with a finite-volume solver for applications in fluid–structure interaction problems solved with co-simulation. *J Fluid Struct* 81, 528–546.
- 1120 Zhi, Y., Zhan, J., Huang, R., Qiu, R., Wang, Y., 2022. Numerical investigations into the ventilation elimination mechanism of a surface-piercing hydrofoil. *Ocean Eng* 243, 110225.