



HAL
open science

Adaptive grid refinement for ship resistance computations

Jeroen Wackers, Ganbo Deng, Clémence Raymond, Emmanuel Guilmineau,
Alban Leroyer, Patrick Queutey, Michel Visonneau

► **To cite this version:**

Jeroen Wackers, Ganbo Deng, Clémence Raymond, Emmanuel Guilmineau, Alban Leroyer, et al..
Adaptive grid refinement for ship resistance computations. *Ocean Engineering*, 2022, 250, pp.110969.
10.1016/j.oceaneng.2022.110969 . hal-03609832

HAL Id: hal-03609832

<https://hal.science/hal-03609832v1>

Submitted on 15 Mar 2022

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.

Adaptive grid refinement for ship resistance computations

Jeroen Wackers*, Ganbo Deng, Clémence Raymond, Emmanuel Guilmineau,
Alban Leroyer, Patrick Queutey, Michel Visonneau

*LHEEA Lab, CNRS UMR 6598, Centrale Nantes
1 rue de la Noë, B.P. 92101, 44321 Nantes cedex 3, France*

Abstract

This paper studies the possibility of using adaptive grid refinement for routine, automated simulations of ship resistance in calm water. With the increase in maturity and reliability of mesh adaptation methods, the main remaining challenge is the creation of straightforward, universal user guidelines which allow these computations to be run correctly, without resorting to trial and error to set the parameters.

The paper uses the mesh adaptation in the flow solver ISIS-CFD. For this solver, a simulation protocol for resistance computation is proposed, which specifies for example the choice of the refinement criterion and the global mesh size. To investigate the reliability and generality of this protocol, it is fine-tuned on one test case and then applied unchanged to three different cases.

The tests show that the solutions have good behaviour and compare well with experiments. Furthermore, numerical uncertainty estimation works for these cases, which increases the trustworthiness of the solutions. Where this is tested, the mesh adaptation produces the same solutions as traditional meshing methods with reduced computational costs. As such, it is shown that mesh adaptation for resistance computations is possible today on a routine basis and that it is advantageous compared with other meshing techniques.

Keywords: Mesh adaptation, Ship resistance, Uncertainty estimation, Automated simulation

1. Introduction

For the continuing growth of computational fluid dynamics (CFD), mesh generation is considered as one of the main bottlenecks (Slotnick et al., 2014). Today, meshing requires more human intervention and more expert knowledge
5 from the users than any other aspect of simulation. Furthermore, these interventions often concern iterative corrections of the mesh, since meshes are created before the solution is known and may turn out to be ill suited for the flow. Thus, generating meshes for reliable simulations makes CFD both costly

*Corresponding author

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

and unpredictable in terms of user time, which is an obstacle for large-scale use
10 in industry (Michal, 2019).

Adaptive mesh refinement, which locally and automatically refines the mesh
during the simulation according to the requirements of the flow, is a natural
solution to this issue. After years of intensive development, mesh adaptation is
approaching maturity and its use for the simulation of complex flows is gaining
15 momentum. For example, the Unstructured Grid Adaptation Working Group
(Balan et al., 2020; Park et al., 2021) and the community around the open-
source (re)mesher MMG (e.g. Dapogny et al. (2014); Nardoni et al. (2019))
show the growing awareness of mesh adaptation for complex flow simulations in
academia and industry. For hydrodynamic simulations, mesh adaptation is used
20 with flow solvers such as OpenFOAM (Eskilsson and Bensow, 2012; Wang et al.,
2020), ReFresco (Windt and Klaij, 2012), StarCCM+ (Yilmaz et al., 2019) and
FINE/Marine (Hildebrandt and Reyer, 2015; Yvin and Muller, 2016; Wackers
et al., 2017).

The integration of adaptation in a routine meshing pipeline requires a para-
25 digm shift. Traditionally, adaptive refinement has been used to make simulations
locally more accurate or to reduce the costs for a given precision, i.e. to optimise
the use of computational resources. For large-scale routine meshing, resource
optimisation still matters, but the main purpose of the adaptation is to make
the computations easier to perform, more predictable and more robust, with less
30 need for user intervention. This implies first that a mesh refinement method
should work reliably under all circumstances, and second that it must operate as
a black box for the user (Wackers, 2021). Thus, precise guidelines are required to
automatically set the parameters for the mesh adaptation, such that it produces
the right meshes for a large range of cases. These guidelines tend to be based
35 on experience and establishing them is a long process (Wackers, 2019).

The purpose of this article is to show that for calm-water resistance of ships,
the required maturity of adaptive grid refinement to perform routine simulations
has been attained. The study uses the Navier-Stokes solver with integrated
mesh refinement ISIS-CFD (Queutey and Visonneau, 2007; Wackers et al., 2011)
40 which is developed by CNRS / Centrale Nantes and is included in FINE/Marine.
For this solver, a protocol of mesh adaptation guidelines is proposed based on our
earlier work. To study its generality, this protocol is fine-tuned on one particular
test case and then applied unchanged to three different cases of towed resistance
and self-propulsion. The accuracy of the adapted-grid solutions is evaluated
45 through uncertainty estimations based on mesh convergence, comparisons with
experiments, and studies of the difference with solutions on non-adapted meshes.
The work focuses on unappended displacement hulls modelled with standard
RANS turbulence models and wall-law boundary conditions.

The paper starts with a description of the flow solver ISIS-CFD and its mesh
50 refinement method (section 2). Section 3 presents the mesh adaptation protocol.
In section 4, the protocol is tested and adjusted for the KRISO Container Ship
(KCS) and the results are compared with non-adapted meshes. Section 5 then
applies the same protocol to the DTMB 5415 navy combatant, the KVLCC2
tanker, and the KCS in self-propulsion conditions using an actuator disk. In
55 section 6, we show how the simulations can be accelerated by computing the
flow on a series of coarse to fine meshes. To conclude, the range of applications
for the simulation protocol and the wider implications of the tests are discussed
in section 7.

2. ISIS-CFD and its adaptive mesh refinement

60 The flow solver ISIS-CFD is developed by CNRS / Centrale Nantes and distributed by Cadence – NUMECA as part of the FINE/Marine flow simulation suite. This section presents the flow solver and its integrated mesh adaptation method.

2.1. The flow solver

65 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 unstructured discretisation is face-based: fluxes are computed face by face, without a-priori assumptions about the cell topologies, which means
70 that cells with any number of arbitrarily shaped faces are accepted. 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 Rhie and Chow SIMPLE method (Rhie and Chow, 1983). Free-surface flow is simulated with a mixture flow approach (Wackers et al., 2011): the water surface is captured with a conservation equation for
75 the volume fraction of water, discretised with specific compressive discretisation schemes in order to keep the interface as sharp as possible. Both steady and unsteady free-surface flows are solved with a time integration technique.

The flow solver features classical Reynolds-averaged Navier-Stokes turbulence models and more sophisticated closures, such as the anisotropic EASM
80 model (Deng and Visonneau, 1999; Duvigneau and Visonneau, 2003). These are available with wall-resolved (no-slip) and law-of-the-wall boundary conditions. Several hybrid RANS/LES models are also available. 6 DOF resolved body motion (Leroyer and Visonneau, 2005) is combined with grid deformation
85 to simulate ship movement. Sliding interfaces and overset meshes allow a part of the mesh to move within the rest, these are used among others to simulate propellers and rudders. Ship propellers can also be modelled as an actuator disk or through the coupling with a panel code for propeller flow.

The solver is mostly used with unstructured hexahedral grids generated by
90 the Hexpress grid generator which is also part of FINE/Marine. These meshes combine structured-like regions and body-fitted viscous layer grids near the walls (see figure 2). The grids consist purely of hexahedral cells, with mesh size variations obtained by placing larger cells next to two or four smaller neighbour
95 cells. Due to its face-based nature, ISIS-CFD treats these grids just the same as any other type of mesh.

2.2. Adaptive refinement

The mesh adaptation method in ISIS-CFD (Wackers et al., 2012, 2014) performs anisotropic refinement of unstructured hexahedral meshes by dividing the
100 cells of an initial coarse mesh locally into finer cells, in one or more directions. This division is repeated until the desired cell sizes are obtained. Earlier refinement can be undone to accompany unsteady flows or changes in the solution as the computation converges. The refinement procedure is fully parallelised using the MPI protocol and includes an automatic adaptive load balancing.

The refinement is called repeatedly during the simulation, to adapt the mesh
105 to the current flow. For physically steady flows, the flow and the mesh converge

together: if the grid is well adapted to the flow and the solution itself has converged, then calls to the refinement procedure no longer modify the mesh.

Metric-based refinement. Anisotropic grid refinement is based on metric tensors as refinement criteria (George et al., 1991). This approach is well known for the creation and adaptation of tetrahedral meshes (see for example Alauzet and Frazza (2021); Castro-Díaz et al. (1997); Loseille et al. (2010)), but it also works well for hexahedral mesh refinement. For metric-based adaptation, the refinement criterion is a continuous field which indicates the ideal size of the cells everywhere in the domain. Adaptive grid refinement is performed to get the actual cell sizes as close to these ideal sizes as possible. Thus, the refined mesh can be considered as a ‘discretisation’ of the criterion.

In ISIS-CFD, the refinement of the cells is decided as follows. First, the refinement criterion \mathcal{C} , which is a 3×3 symmetric tensor field (evaluated in the cell centres), is computed in some way from the flow solution. In each hexahedral cell i , the cell sizes $\mathbf{d}_{i,j}$ ($j = 1, 2, 3$) are defined as the vectors between the opposing face centres in the three cell directions j . Given a constant threshold parameter T_r , the goal of the grid refinement is to approach:

$$\|\mathcal{C}_i \mathbf{d}_{i,j}\| = T_r \quad \forall i, j. \quad (1)$$

This is obtained by refining a cell in the direction j when $\|\mathcal{C}_i \mathbf{d}_{i,j}\|$ exceeds T_r , while a previously refined group of cells is derefined in the direction j if $\|\mathcal{C}_i \mathbf{d}_{i,j}\|$ is lower than T_r/d for all cells in the group. The constant d is chosen slightly larger than 2, to prevent cells being alternately derefined and re-refined. 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 \mathcal{C} are based principally on the Hessian matrix of second spatial derivatives of the solution. Hessian matrices are linked to linear interpolation errors (Loseille et al., 2010), so they provide an indicator for the truncation error of a second-order finite-volume discretisation. The Hessian can be based on different variables. Here, to adapt the mesh to pressure-based flows but also to boundary layers, wakes, and shear layers, the criterion is based on the Hessians of both the pressure and the velocity, weighted in the way in which they appear in the flux (Wackers et al., 2017). A common weight ρV is assigned to all the velocity Hessians, where $V = \sqrt{u^2 + v^2 + w^2}$:

$$\mathcal{C}_H = \left(\max(\|\mathcal{H}(p)\|, \rho V \|\mathcal{H}(u)\|, \rho V \|\mathcal{H}(v)\|, \rho V \|\mathcal{H}(w)\|) \right)^\alpha. \quad (2)$$

\mathcal{H} is the Hessian operator; the absolute value of a matrix $\|\cdot\|$ corresponds to a matrix having the same eigenvectors as the original one and the absolute values of its eigenvalues. In the same way, the power α of a matrix is obtained by taking its eigenvalues to the power α while keeping the eigenvectors. We take $\alpha = \frac{1}{2}$. The maximum of two tensors is computed using the approximate procedure defined by (Wackers et al., 2012).

A second criterion creates the refined mesh at the free surface: this criterion is non-zero whenever the water volume fraction α (which indicates the free surface position) is neither 0 nor 1 (Wackers et al., 2012). Directional refinement normal to the surface is obtained from normal vectors $\mathbf{v} = \nabla \alpha_A / |\nabla \alpha_A|$, where

α_A is a smeared volume fraction field created with Laplacian smoothing. The criterion then becomes:

$$\mathcal{C}_S = \begin{cases} \mathbf{v} \otimes \mathbf{v} & \text{if } 0.1 \leq \alpha_A \leq 0.9, \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (3)$$

Here α_A is used again, to create a buffer layer of a few cells around the surface position for safety.

The two tensor criteria are combined into one (Wackers et al., 2014) by taking a weighted maximum of the tensors. Since the free-surface criterion always has a unit eigenvalue, a weighting constant c is applied only to the Hessian criterion:

$$\mathcal{C}_C = \max(\mathcal{C}_S, c\mathcal{C}_H). \quad (4)$$

Figure 1 shows a mesh generated with the combined refinement criterion.

Grid quality measures. Post-treatment of the refinement criterion is used to improve the regularity of the adaptively refined meshes. The first step is a protection of the viscous layer grid, which consists of several layers of wall-aligned cells (figure 2). For these layers, the refinement parallel to the wall is made the same throughout a column of cells from the wall to the outer layer: if one of the cells in a column needs to be refined, all the cells are refined. Thus, the column / layer structure of the boundary layer grid is preserved.

A second measure is to impose a minimum cell size: cells smaller than this size are no longer refined. This option prevents spurious refinement in case of locally large errors in the computation of the refinement criterion, which may appear for example in the high aspect-ratio cells of the near-wall boundary layer grid. Also, it prevents infinite refinement around flow singularities.

Finally, adaptive refinement can be forbidden outside of a 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.

Projection. Since the pressure on a wall depends on its curvature, the pressure reacts to any irregularity in the surface. To prevent spurious pressure peaks, a refined mesh must therefore be projected onto the true geometry (represented by a triangulation of the CAD file). Since existing nodes are already in the correct position, this projection is only applied to newly created nodes and the associated mesh deformation is handled in a local way. Ellipsoids are defined around each new boundary node, dimensioned such that they do not contain any other boundary node. The required displacement of the boundary node is then applied to all the nodes in the ellipsoid, in a weighted manner so that the displacement on the exterior of the ellipsoid goes to zero. With this approach, the displacement of all boundary nodes is handled independently.

3. Mesh adaptation protocol for ship resistance

The mesh refinement method in the previous section contains many free parameters. For mesh adaptation in routine simulation, the choice of these parameters should be mostly automated. This requires default values based on

165 guidelines which must be straightforward and correct every time, without a need
for trial and error. Establishing these is complicated, because the guidelines have
to be valid over a range of cases, so they must take into account a large number
of input variables: ship length, velocity, hull shape, etc.

170 Our approach for defining guidelines (Wackers, 2019, 2021) is to use physical
arguments, notably dimensional analysis, to reduce the number of input
variables that have to be considered and then to use experience to find sensi-
ble guidelines for the remaining parameters. This approach is applied here to
define standardised ISIS-CFD simulation guidelines for calm-water resistance
evaluation of displacement hulls. The guidelines are based on our earlier work
and have been fine-tuned using a test case, which is presented separately in the
following section 4 for clarity.




Figure 1: An example of an adapted mesh following the current protocol. Top to bottom: free surface, details of the Y -symmetry plane, X -cut at the stern. The case is the DTMB 5415 at $Fr = 0.41$ of section 5.1, with $T_{r,H} = 0.07$.

175 *Refinement criterion.* The first guideline concerns the choice of the refinement
criterion. To create meshes for predicting the forces on a hull, adjoint-based




Figure 2: Initial mesh corresponding to the bottom image of figure 1.

goal-oriented refinement criteria are the ideal choice (Alauzet and Frazza, 2021), since these create meshes which are directly optimised for the computation of the requested forces. However, for free-surface flows, evaluating an adjoint solution is difficult (Kühl et al., 2021). In the absence of an adjoint solver, an indicator for the local truncation errors (like a Hessian), which aims at an approximate equidistribution of the errors over the domain, is a suitable alternative. The question is, on which variables this criterion should be based.

To decrease the need for user intervention in the meshing process, the refinement criterion should react to all the important flow features, so that the meshes required to capture these features can be created entirely using adaptive refinement. Among others, this means that the fine mesh for the free surface and the ship wave pattern must be inserted with adaptation. The combination of the free-surface criterion and the Hessian of the pressure, which is a suitable indicator of the orbital flow fields in the waves, ensures this (Wackers et al., 2014). However, this is not enough: even if the forces on the hull are dominated by the pressure, the accurate resolution of the boundary layer and wake is required to correctly predict these forces (Wackers et al., 2017). Thus, refinement based on the velocity Hessians, which detect the wake, is also needed. The combined free-surface and flux-component Hessian criterion of section 2.2 is therefore a suitable choice.

Figure 1 shows an adapted mesh of about 1M cells, created with the current protocol. Combined refinement at and below the surface can be seen in waves, while the mesh around the hull is refined at the bow and in the boundary layer. A comparison with the original grid in figure 2 shows that most of the cells have been inserted through adaptive refinement.

Threshold. The second question concerns the appropriate mesh density, specified here by the threshold T_r of equation (1). To manage computational resources, many authors adjust such mesh density parameters to obtain a specified number of cells. This adjustment can be automated (Pons and Ersoy, 2019) to remove the need for trial and error. Another approach, suitable for metric criteria, is to normalise the metric using its integral over the whole flow domain, in such a way that the mesh density parameter becomes a direct specification of the target number of cells (Loseille et al., 2010; Alauzet and Frazza, 2021).

To capture free-surface waves, we prefer a different approach. ISIS-CFD guidelines for non-adapted meshes specify target cell sizes at the free surface in terms of the ship length L (typically $L/1000$ for the vertical cell size). Thus, for

adaptive free-surface refinement, it is natural to link the mesh density parameter to the cell size at the surface rather than to the total number of cells (accepting that more deformed water surfaces require more cells to be captured correctly). Therefore, the free-surface refinement criterion is defined with unit vectors in equation (3), so the threshold T_r equals the desired cell size normal to the free surface.

To keep the criteria compatible, a similar behaviour is sought for the Hessian criterion of equation (2): the mesh density on the hull should be proportional to the threshold value, independent of the case parameters (velocity, density, etc.). To obtain this, the criterion field must have the same order of magnitude for all parameter values, which is achieved in equation (4) by nondimensionalising the Hessian criterion (Wackers et al., 2014; Wackers, 2019) with a weighting factor:

$$c = \bar{c} \left(\frac{L^2}{\rho U^2} \right)^a, \quad (5)$$

where L and U are the reference length and velocity, ρ is the water density, and the new weight \bar{c} is non-dimensional. The criterion $c\mathcal{C}_H$ is then independent of the reference length and velocity, it depends only on the Reynolds and Froude numbers and the hull geometry.

For selecting \bar{c} , it is convenient to treat the combined criterion as if the free-surface and Hessian criteria are used separately, with their own thresholds T_{rS} and T_{rH} . This is achieved by choosing:

$$T_r = T_{rS}, \quad \bar{c} = \frac{T_{rS}}{T_{rH}}. \quad (6)$$

Guidelines for these thresholds are based on experience. Since T_{rS} directly specifies the cell size at the free surface, this parameter is chosen equal to $L/1000$ which is the existing ISIS-CFD recommendation. For T_{rH} , an analysis in (Wackers, 2019) of several simulations including ships, wing profiles, and three-dimensional hydrofoils, finds a suitable range of $[0.4L, 0.04L]$. Serani et al. (2019) present computational settings which are valid over a family of ship hull geometries, in a context of automatic shape optimisation for displacement hulls. They suggest $T_{rH} \in [0.2L, 0.05L]$. A study included in (Demeester et al., 2021) applies the same values to a submarine. Here, based on the tests in section 4, we select $T_{rH} \in [0.2L, 0.025L]$.

Minimum cell size. Besides removing spurious refinement due to errors in the refinement criterion, the minimum cell size reduces the mesh refinement around the finest flow details. A large minimum cell size works like a cutoff filter: while the adaptive refinement still captures the main flow features, very fine refinement to resolve small details is not applied. Thus, the minimum cell size can be used to impose the level of resolution for the flow details that is required for a specific application.

For resistance computations, section 4.3 shows that details below the resolution of the free surface have a negligible influence on the computed forces and that not capturing these details reduces the total number of cells and the cost of the computations. Therefore, a large minimum cell size equal to $L/1000$ is applied.

245 *Limiting boxes.* To prevent outflow perturbations and to stabilise the computation, it is useful to damp out any waves before the flow reaches the outflow boundary. Coarse meshes near the outflow are commonly used for this. With adaptive refinement, these meshes can be created by imposing limiting boxes which forbid the refinement in x - and y -direction near the outflow boundary.
250 Refinement in z -direction is allowed throughout the domain, so the free surface remains sharply captured. The effect of a limiting box is visible to the left in the top view of figure 1.

While the minimum requirement for good damping is to forbid the refinement in the last L of the domain, the refinement can be suppressed much closer to
255 the stern of the ship, which reduces the number of cells. Serani et al. (2019) show that a limiting box at $0.25L$ behind the stern does not significantly change the computed forces, while a limit closer to the stern does. Since the tests in section 4.3 confirm this, a limit of $0.3L$ behind the stern is proposed here.

260 Additional limiting boxes are imposed at $0.3L$ in front of the bow and at $0.7L$ to the side of the ship. These do not alter the flow field but increase the stability of the computation by removing spurious waves radiated from the ship hull¹.

Initial mesh. The initial mesh, from which the computation is started, is created with Hexpress. Since all fine cells to capture the flow will be created with
265 adaptive refinement, the initial mesh only needs to resolve the geometry (figure 2), something that Hexpress can usually do automatically. In the rest of the domain, the mesh is left coarse to give as much manoeuvring room as possible to the adaptive refinement.

270 For metric-based refinement, the refinement criterion specifies the cell sizes in the adapted mesh throughout the domain. Thus, the initial mesh should not have a large influence on the final result of the computation. With division-based refinement however, there are two exceptions to this principle:

- The refined cell sizes cannot be chosen arbitrarily since division can only reduce the initial cell sizes by a power of two.
- 275 • The initial mesh imposes the orientation of the cells, because the refinement cannot turn them.

The first exception is relevant for the free-surface mesh. To obtain refined cells of $L/1024$ at the surface, all initial cells around the surface should have sizes of $L/2^j$, for integer $j \leq 10$. The second exception implies that the initial body-aligned viscous layer grid should be thick enough to fully contain the physical
280 boundary layer, so that the same will be true for the adapted meshes. In Hexpress, coarser meshes lead to thicker viscous layers; however, too coarse meshes may be unable to capture the hull geometry. As shown below, a good compromise is obtained with cubic cells of size $L/128$ around the hull.

¹These limitations were not used for all cases, which was deemed acceptable since no influence on the results is expected.

285 4. KCS fine-tuning

The test case in this section was used to fine-tune the simulation protocol in section 3. Results with the finished protocol are presented here. After an introduction of the case (section 4.1), the accuracy of the solution is evaluated in section 4.2 through a mesh convergence study and comparison with experiments. In section 4.3, the validity of the protocol is shown through systematic parameter variations. Finally, in section 4.4 the results are compared with classical meshing approaches to assess the efficiency of the mesh adaptation.

4.1. Test case and numerical setup

The test case is the KRISO Container Ship (KCS) in towed condition with free trim and sinkage, for which experimental results are available from KRISO (Van et al., 1998; Kim et al., 2001), see also (Hino, 2005). The test is in model scale ($L = 7.2786m$) at $U = 2.196m/s$, $\mu = 0.0012706kg/(ms)$, $\rho = 999.5kg/m^3$, $g = 9.81m/s$, which gives $Re = 1.257 \cdot 10^7$ and $Fr = 0.260$. The model weight is 1650.8kg with the centre of gravity at $x = 3.532m$ in front of the aft perpendicular. Half the ship is simulated and the domain runs from $1.5L$ in front of the ship to $3.5L$ behind the stern, sideways to $2L$ and vertically from $-1.75L$ to $1.25L$. The boundary conditions are imposed velocity and volume fraction on the side planes, imposed pressure on the top and bottom, a mirror condition on the centreplane, and wall law on the hull. Turbulence is modelled with $k - \omega$ SST.

Convective fluxes use the limited AVLSMART scheme (Pržulj and Basara, 2001), with central differences for diffusion. The time step is $0.0176s$ and 4000 time steps are run for convergence. The initial mesh has 128 cells/ L on the hull and $y^+ = 50$ for the first boundary layer cells; it has 79.9k cells in total. For the adaptive refinement, $T_{rS} = 1.3 \cdot L/1000$, where the safety factor of 1.3 prevents refinement below $L/1000$. A series of grids is produced by reducing the Hessian threshold T_{rH} with steps of a factor $\sqrt{2}$ from $0.2L$ to $0.025L$.

4.2. Grid convergence and accuracy

To determine the accuracy of the numerical solution, the grid convergence study for the Hessian threshold is analysed. Since the mesh size everywhere is proportional to T_r (equation (1)), varying this threshold produces a series of similar coarse to fine meshes (Wackers et al., 2017). The evolution of the resistance with the mesh size can be used to evaluate its numerical uncertainty through standard procedures. We use the estimation of Eça and Hoekstra (2014). Starting from the third grid, the uncertainty on each grid is evaluated using the result on the grid and all the coarser ones; T_{rH} is used as a convenient measure of the grid size.

Figure 3 shows the convergence of the resistance (averaged over the last 25% time steps). Visually, the force converges smoothly, which is confirmed by the uncertainty estimation. Variable power-law behaviour (the best possible result for Eça and Hoekstra's procedure) was obtained, with similar results for all meshes. Also, most uncertainty intervals overlap, which shows that the uncertainty estimation is effective even on the coarser grids. The estimated uncertainty is low (0.4% on the finest grid). Plots of the free surface (figure 4) show that the flow field is visually converged around $T_{rH} = 0.05L$. Finally, a comparison with the experimental results from Van et al. (1998); Kim et al.




Figure 3: KCS resistance, grid convergence with T_{rH} and estimated numerical uncertainty. The blue lines show the uncertainty interval of the experimental results (Van et al., 1998; Kim et al., 2001).

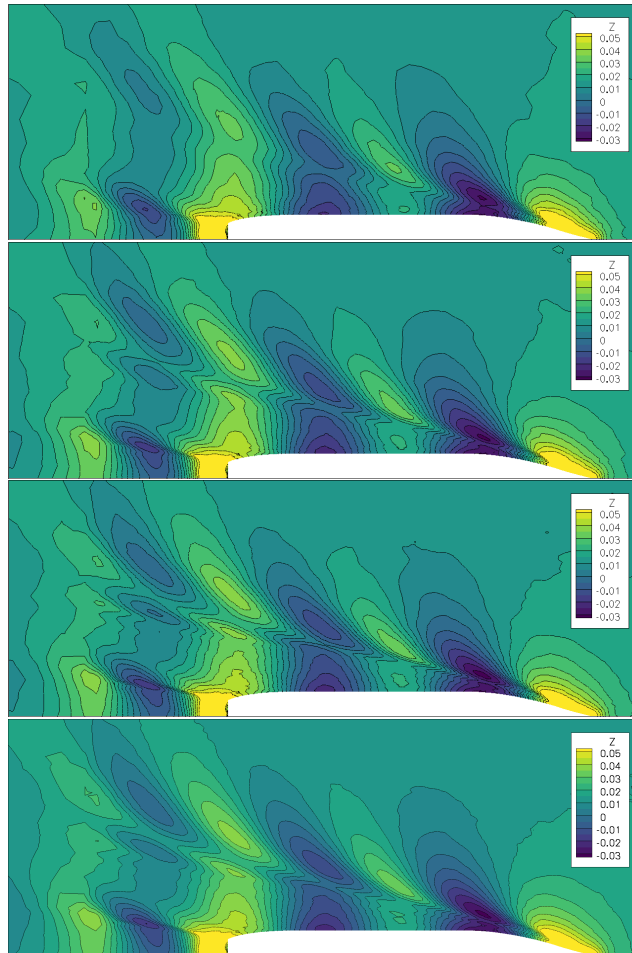


Figure 4: KCS, free-surface elevation. Top to bottom: $T_{rH} = 0.2L$, $0.1L$, $0.05L$, and $0.025L$.

(2001) in figure 3 shows that the converged numerical resistance lies within the experimental uncertainty interval.

4.3. Parameter study

335 The settings of the simulation protocol are checked here through systematic parameter variation. The first test concerns the effect of the minimum cell size. This parameter (chosen relatively small in earlier tests, around $10^{-4}L$) has a much more pronounced influence than we expected. Figure 5 shows the grid convergence for three different minimum cell sizes. For the smaller values of T_{rH} , a small minimum cell size will lead to the resolution of many flow details, which may not be essential for capturing the drag. Thus, the figure shows that for a larger minimum cell size, the number of cells is strongly reduced, for only a modest loss of precision. Also, the solutions for the different minimum cell sizes converge to the same value when T_{rH} is reduced. This means that for resistance simulations, the choice of the largest minimum cell size, i.e. $1.3 \cdot L/1000$ like the free-surface threshold, is preferable.

340

345




Figure 5: KCS, influence of the minimum cell size on the convergence of the resistance.

The limiting box proposed at $0.3L$ behind the stern is restrictive: it removes most of the trailing wave field behind the ship and also the far-field wake. Therefore, a test is performed with a much longer limiting box which ends at $2L$ behind the stern. Figure 6 compares the free-surface elevations for the two tests: in the interior of the smaller limiting box, the solutions are visually identical. Furthermore (figure 7), apart from some small perturbations, the forces are similar for the same T_{rH} so there is no systematic advantage to using the larger limiting box. On the other hand, the larger box significantly increases the number of cells (up to 75% for $Tr=0.05L$). This again confirms that the smaller box is acceptable. However, Serani et al. (2019) tested an even shorter limiting box ($0.05L$) and found significant perturbations. This leads us to adopt the box at $0.3L$ as the best choice.

350

355

To see the influence of the original mesh, the standard mesh (where the cells have a size of $L/128$ on the hull) is compared with two other Hexpress meshes that have $L/96$ and $L/192$ on the hull. Figure 8 shows the boundary layer on the aft hull. The figure shows that the different initial meshes lead to varying

360




Figure 6: KCS, Free-surface elevation with limiting box $2L$ behind the stern (top) and $0.3L$ behind the stern (bottom). Hessian threshold $T_{rH} = 0.05L$.

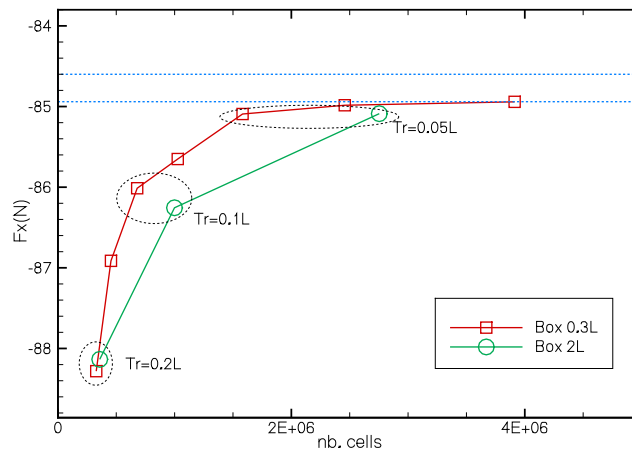


Figure 7: KCS, convergence as a function of T_{rH} for two different limiting box sizes.

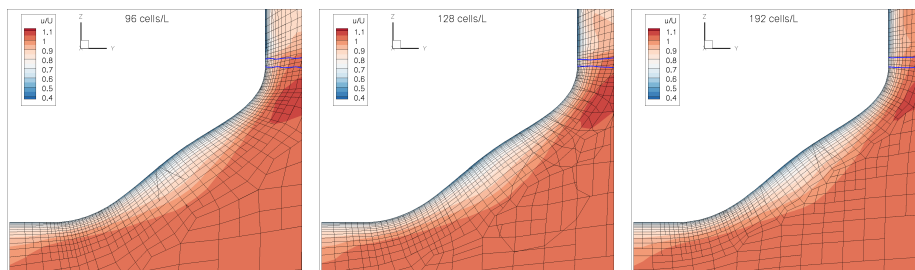


Figure 8: KCS, axial velocity, free surface, and mesh at $x = 0.86L$. Left to right: 96 cells/ L initial grid, standard grid, and 192 cells/ L grid. $T_{rH} = 0.1L$.

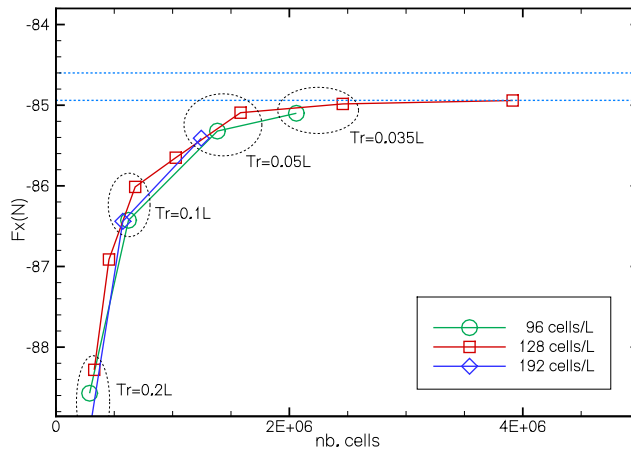


Figure 9: KCS, influence of the initial Hexpress mesh.

viscous layer mesh thicknesses. However, all the meshes are able to capture the full aftship boundary layer in the viscous layer mesh (although the $L/192$ mesh is at its limit). Furthermore, the free surface discontinuity is thinner in the standard mesh. The reason for this is that the adaptive refinement cuts cells in two, so it can only divide the initial cell sizes by a power of 2. Therefore, for the standard initial mesh, the grid at the free surface ends up at $\Delta z = L/1024$, while for the new grids it ends up at $\Delta z = L/768$. The resistance is shown in figure 9. The two new meshes give similar results, but there is a difference with the standard mesh, which gives faster and better convergence towards the experimental results. Thus, the coarser grid at the surface is enough to perturb the computed forces.

4.4. Comparison of meshing methods

To conclude the KCS test, the results with adaptive refinement are compared with two non-adaptive techniques for generating series of meshes. The first approach is to create a series of meshes with Hexpress using the C-Wizard automatic setup tool included in FINE/Marine. These meshes follow the standard FINE/Marine guidelines, with uniform cell sizes over the hull. As indicated in section 4.3, the viscous layers of such meshes become thinner as the mesh is refined, so the series of meshes is not strictly similar. The second series uses a combined Hexpress and systematic grid refinement (SGR) approach: the three coarsest grids are generated by Hexpress, using manually increased refinement at the bow and stern to capture the pressure fields. The three finest grids are created from these coarser grids by refining all cells once (which is possible with adaptive refinement). This approach is costly in terms of the number of cells, but it guarantees that the viscous layers remain thick enough as the mesh is refined.

Figure 10 compares the three results. The figures show that for very coarse meshes (300k cells) where the adaptive refinement is mainly used to capture the free surface, the C-Wizard series gives similar or slightly better results than the adaptive series. However, from 600k cells on, the adaptive series is clearly better than the other two, using on the order of 40% less cells than the C-Wizard and




Figure 10: KCS, convergence in terms of the number of cells for three meshing strategies and estimated uncertainties. The x -axis uses a log scale.

60% less than SGR for the same accuracy. The adaptive series also performs
 395 better for the uncertainty estimation: as noted in section 4.2, variable power-law fits are obtained with similar behaviour for all meshes. Also, the estimated uncertainty is low (0.4% on the finest grid). The SGR series is perturbed by the switch from pure Hexpress meshes to SGR, but it also produces power-law convergence for the two finest grids, with low uncertainty (0.7% on the finest
 400 grid). The convergence of the C-Wizard series is perturbed by the change in thickness of the viscous layer: because the convergence is not asymptotic, the uncertainty is high (13% – 18%). Finally, the figure confirms that the adaptive and SGR series converge to the same value for the resistance, which underlines the reliability of both series. The C-Wizard series may not converge to the same
 405 value however, because the meshes are not similar. Thus, for this case, Hessian-based adaptive refinement is the best meshing strategy, both for coarse-mesh and fine-mesh simulations.

5. Blind application tests

When considering the good results of section 4, it must be taken into account
 410 that the protocol was calibrated for the KCS. To test if the mesh refinement settings are generally valid, the protocol of section 3 is applied here, without any adjustments, to three cases which span the range of bare displacement hulls: a slender hull (the DTMB 5415 naval combatant), the high-block-coefficient KVLCC2, and the KCS in self-propulsion with an actuator disk.

5.1. DTMB 5415

For this slender ship, three speeds are simulated, from $Fr = 0.1$ where the waves are barely perceptible to $Fr = 0.41$ where heavy wave breaking occurs both at the bow and the stern. The test is based on the measurements at INSEAN by Olivieri et al. (2001), using the INSEAN model 2340 with a length
 420 of $L = 5.72m$. The density is $\rho = 1000kg/m^3$, the gravity $9.8058m/s^2$ and the dynamic viscosity is $0.0010080kg/(ms)$. The three velocities are $0.749m/s$

$(Fr = 0.10, Re = 4.25 \cdot 10^6)$, $2.097m/s$ ($Fr = 0.28, Re = 1.19 \cdot 10^7$), and $3.071m/s$ ($Fr = 0.41, Re = 1.74 \cdot 10^7$). Trim and sinkage are fixed to the experimentally obtained values. The boundary conditions and turbulence model
 425 are the same as for the KCS. The domain size is from $1.5L$ in front of the ship to $3L$ behind it, sideways to $y = 2L$ and vertically from $z = -1.5L$ to $z = L$.

The same initial mesh is used for all three velocities. It has 128 cells/ L on the hull and $y^+ = 50$ for the first boundary layer cells at $Fr = 0.28$, with 77.8k cells in total. The flux discretisation and the mesh adaptation settings are the same as for the KCS. The time steps are chosen such that the ship advances
 430 one reference length in 100 time steps; 2000 time steps are run for convergence.

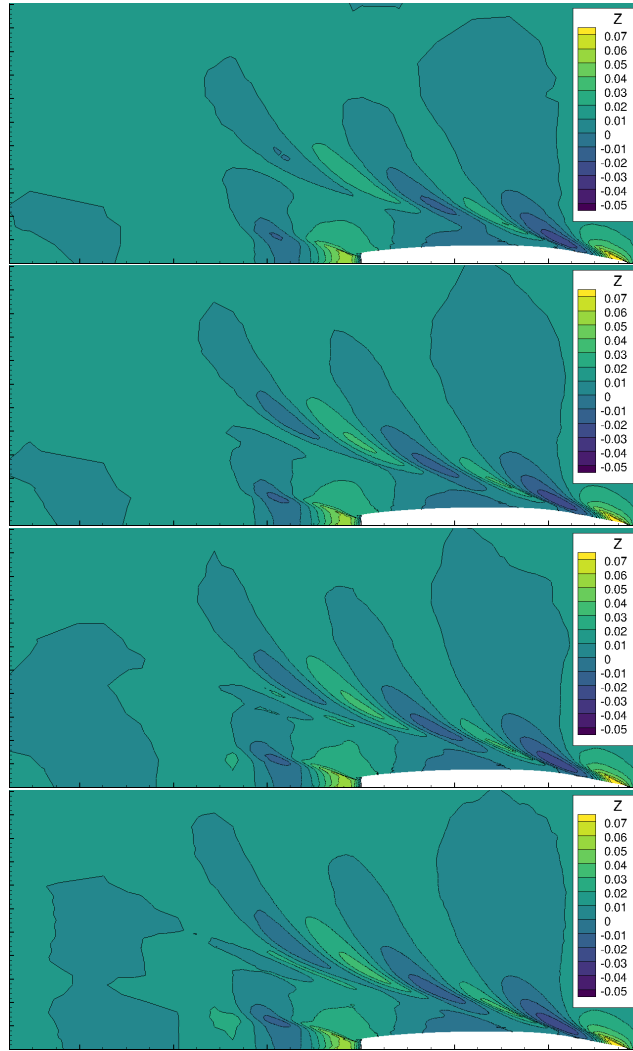


Figure 11: DTMB 5415, convergence of free-surface elevation at $Fr = 0.28$. Top to bottom: $T_{rH} = 0.2L, 0.1L, 0.05L$, and $0.025L$.

Results. Figure 11 shows the convergence of the free surface for $Fr = 0.28$. The level of convergence is similar to the KCS, even though the flow is different:

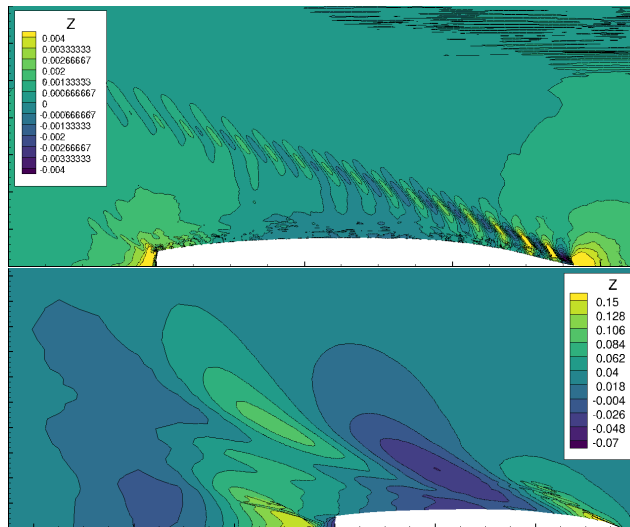


Figure 12: DTMB 5415, free-surface elevation on fine grids ($T_{rH} = 0.025L$). Top: $Fr = 0.1$, bottom: $Fr = 0.41$.

there is stronger breaking in the bow wave system, especially the first shoulder
 435 wave. The flow has converged around $T_{rH} = 0.05L$. The results on the finest
 meshes for $Fr = 0.1$ and $Fr = 0.41$ are given in figure 12. For brevity the
 grid convergence is not shown; the results at $Fr = 0.41$ are converged around
 $T_{rH} = 0.05L$ despite the much stronger breaking waves. However, at $Fr = 0.1$
 440 the bow wave is not converged even at $T_{rH} = 0.025L$, although the flow in zones
 close to the ship has converged. Still, the waves are so small (about 5 times
 smaller than the vertical cell size) that their influence is probably negligible.
 Note that the range of the images is different: for $Fr = 0.1$, the waves are 17
 times smaller than for $Fr = 0.28$ and 37 times smaller than for $Fr = 0.41$.

Figure 13 gives the drag force (averaged over the last 25% time steps) ac-
 445 companied by the numerical uncertainty, evaluated following Eça and Hoekstra
 (2014). The blue lines represent the experimental results. The vertical scale is
 chosen such that a 1% uncertainty is the same height in the 3 graphs. These fig-
 ures show that the convergence behaviour is similar for $Fr = 0.1$ and $Fr = 0.41$,
 with the resistance varying with T_{rH} to the power of about 0.6. This is normal
 450 behaviour for pressure-dominated flows. The finest-grid uncertainties are
 around 1.8%. For $Fr = 0.28$ however, the convergence is less regular but the
 solution is more accurate, with 0.6% uncertainty on the finest grid. The two fast
 cases agree most closely with the experiments. The overestimation for $Fr = 0.1$
 may be due to the boundary layer grid, which was created for $Fr = 0.28$; at the
 455 lowest velocity, y^+ is 20 which is close to the transition layer around $y^+ = 11$
 where the wall law is the least accurate. Still, the agreement with experiments
 is acceptable for all cases.

Figure 13 also shows the evolution of the number of cells with T_{rH} . The
 results are similar until $T_{rH} = 0.025L$, where the mesh for $Fr = 0.1$ is clearly
 460 finer than for the other two. This phenomenon is due to the smaller waves for
 the reduced speed, which require many fine cells to be resolved correctly. The
 number of cells agrees well with the KCS case (figure 3).




Figure 13: DTMB 5415 resistance, grid convergence with T_{rH} and estimated numerical uncertainty for (top to bottom) $Fr = 0.1$, $Fr = 0.28$, and $Fr = 0.41$. The blue lines represent the experimental results (Olivieri et al., 2001).

5.2. KVLCC2

To test a high block-coefficient hull, the KVLCC2 tanker is simulated at $Fr = 0.14232$ and $Re = 4.6 \cdot 10^6$, following the experiments given by (Van et al., 1998; Kim et al., 2001). The density is $\rho = 999.25 \text{ kg/m}^3$, the gravity 9.8058 m/s^2 and the dynamic viscosity is $0.0012547 \text{ kg/(ms)}$. The velocity is 1.047 m/s and the reference length $L = 5.5172 \text{ m}$. All settings (motion, turbulence model, etc.) are the same as for the KCS; the time step is 0.0263 s and 2000 time steps are run. The initial mesh has 58k cells.




Figure 14: KVLCC2 resistance, grid convergence with $T_{r,H}$ and estimated numerical uncertainty, compared with experimental results (Van et al., 1998; Kim et al., 2001).

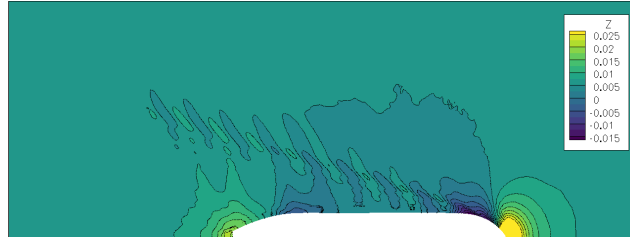


Figure 15: KVLCC2, free-surface elevation for $T_{r,H} = 0.025L$.

The wave pattern is similar to the DTMB 5415 at $Fr = 0.1$ (figure 15). Likewise, the finer grids have similar numbers of cells (figure 14). And while the convergence of the resistance has some slight oscillations, the global behaviour is smooth and the final uncertainty is low (0.61%). The converged difference with the experiments is about 1.5%.

5.3. KCS self-propulsion

The final test case is the KCS in self-propulsion, where the propeller is modelled by an actuator disk. Compared with the towed-resistance cases for which the mesh adaptation protocol was developed, this poses three additional challenges. First, the global flow may be different. Furthermore, it is not known if the mesh refinement reacts to a force field like the actuator disk. And finally,

a ship propeller is tiny with respect to the ship itself. Can something so small be captured by mesh adaptation that is configured for the flow around the ship?

The test case as described by (Hino, 2005) is the same as in section 4.1, except that the model has free sinkage only. Also, $\mu = 0.0011379 \text{ kg}/(\text{ms})$ and $\rho = 999.1 \text{ kg}/\text{m}^3$ which gives $Re = 1.40 \cdot 10^7$ and $Fr = 0.260$. The model weight is 1648.5 kg . The actuator disk has an outer and inner radius of 0.125 m and 0.025 m , with a thickness of 0.05 m . Ship-point self propulsion is obtained by applying a forward skin-friction correction force of 30.3 N to the hull; the actuator disk thrust is automatically updated to compensate for the remainder of the resistance. The corresponding torque is interpolated from a $K_T - K_Q$ curve obtained using 21 open water simulations of the propeller. To accommodate the flow rotation behind the disk, the entire ship is simulated. Otherwise, the initial Hexpress mesh is equal to the one from section 4; no particular refinement is applied around the actuator disk.




Figure 16: Actuator disk: open-water axial velocity in the y -symmetry plane for $T_{rH} = 0.1L$. The box shows the disk position.

As a first test, the actuator disk is simulated in an open water setting. A coarse initial mesh is made, with cell sizes equal to the largest cells that appear on the KCS hull. Furthermore, even though the actuator disk is simulated on its own, the adaptive refinement is configured using the reference length of the ship. The results are good (figure 16): even on a mesh which is coarse for the ship ($T_{rH} = 0.1L$), both the actuator disk and its wake are well represented by the refined mesh. This implies that no particular initial fine mesh in the actuator disk region is needed: self-propulsion can be computed using initial meshes made for towed resistance.

For the actual ship, this conclusion is confirmed: the meshes for two different thresholds (figure 17) capture the propeller flow well. Furthermore, the meshes correspond visually to what the thresholds imply for a ship: $T_{rH} = 0.1L$ is a rather coarse mesh, while $T_{rH} = 0.025L$ gives a fine mesh. Thus, refinement based on the ship length scale is appropriate for a detail like the actuator disk.

Still, figure 17 also shows that the mesh size throughout the actuator disk region is equal to the minimum cell size (chosen the same as the size at the free surface), even for $T_{rH} = 0.1L$. It is likely that the mesh at $T_{rH} = 0.025L$ would become much finer locally if the minimum cell size were smaller. The mesh is fine enough for the actuator disk: it contains 40 cells over the disk diameter, which is twice more than the recommended minimum for ISIS-CFD. However, it is possible here that the minimum cell size prevents the convergence to a mesh-independent solution.

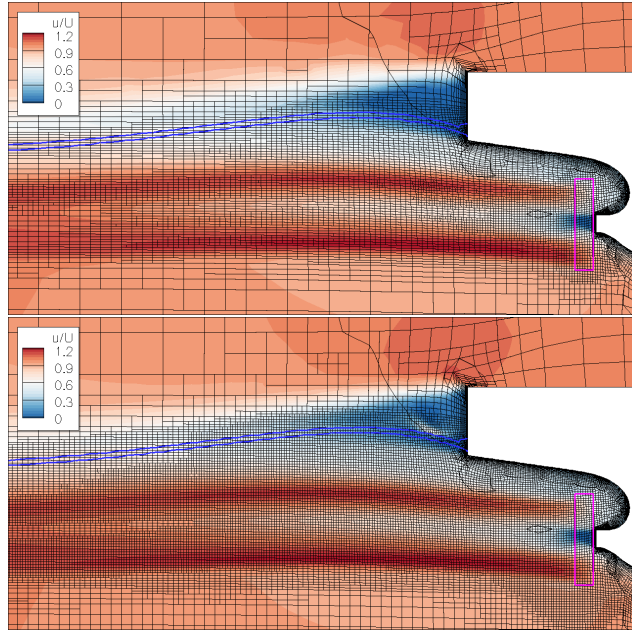


Figure 17: KCS self-propulsion: axial velocity in the y -symmetry plane. Top: $T_{rH} = 0.1L$, bottom: $0.025L$.




Figure 18: KCS self-propulsion: axial velocity in the propeller plane $x/L = 0.9911$. Left to right: $T_{rH} = 0.1L$, $0.025L$, and experiments.

To assess the convergence of the local flow, figure 18 shows the total wake field in the propeller plane. The flow is similar between the two solutions shown. This is partially due to the minimum cell size – the meshes are equal in the centre of the actuator disk. But even outside the actuator disk region, where the minimum cell size plays a smaller role, the agreement is good. The experimental wake is somewhat different, since it comes from an actual propeller, but the boundary layer near the hull is captured well. Thus, the simulation of the local flow is satisfactory.

The same holds for the resistance (figure 19). While some oscillations are present, the force behaves globally like the towed case (figure 3) and the estimated uncertainties are of the same order². The resistance result confirms that the mesh adaptation protocol can be used for self-propulsion.

²The convergence of other self-propulsion parameters such as thrust and torque has also been studied, but since these are proportional to the resistance (either directly or via the $K_T - K_Q$ lookup table), they behave the same and are therefore not shown here.




Figure 19: KCS self-propulsion resistance, grid convergence with T_{rH} , estimated numerical uncertainty, and experimental result from (Hino, 2005).

530 6. Convergence acceleration

Adaptively refined meshes have a significant final advantage (Loseille et al., 2007): it is trivial to initialise a fine-grid solution from the converged solution on a coarser grid, by simply restarting the simulation with a smaller T_r . Since coarse-grid simulations are cheap, this can accelerate the solution on the fine grid. Furthermore, all the coarse-grid results are available for free when the fine-grid simulation is finished.




Figure 20: DTMB 5415 at $Fr = 0.28$, grid convergence and uncertainty for three simulation approaches.

An example is shown here³ for the DTMB 5415 at $Fr = 0.28$ of section 5.1. Instead of running 2000 time steps, the solutions are stopped when no force

³This is an illustrative example only. We do not claim that the settings are universally correct, nor that they are optimal for this case.

Table 1: Computation time with 28 cores (in minutes) for three simulation approaches. The number of time steps is given between parentheses. For multilevel, the time is summed from the coarsest level to the current level, while the number of time steps is for the current level only.

$T_{r,H}$	Full simulation	Single level	Multilevel
0.2L	47.5 (2000)	20.4 (915)	20.4 (915)
0.1L	94.8 (2000)	35.1 (857)	36.1 (+356)
0.05L	296.4 (2000)	144.6 (1107)	90.4 (+358)
0.025L	767.6 (2000)	279.9 (910)	274.3 (+466)

540 value in the last 150 time steps differs by more than 0.25% from the average over the last 250 time steps. Computations are run for $T_{r,H} = 0.2L$, $0.1L$, $0.05L$, and $0.025L$. Two series are performed: a single-level series where each grid is started from zero, and a multi-level one where each simulation starts from the previous coarser solution.

545 Figure 20 shows the convergence of the solutions and the estimated uncertainties for each series, compared with the full 2000 time step simulations. The number of time steps and the computation time on a 28-core Intel Xeon node is reported in table 1. While differences can be observed, the behaviour is globally similar and the multilevel approach is the fastest for the two finest thresholds.

550 We have tested the same procedure on other cases and obtained an acceleration everywhere. This means that, by correctly using multilevel simulations, one can obtain a fine-grid solution plus an uncertainty estimation faster than just the solution on a single mesh.

7. Discussion and conclusion

555 The previous sections presented four test cases where combined free-surface and Hessian-based mesh adaptation is used to create the meshes for ship resistance evaluation. This section attempts to generalise the findings of these test cases by discussing two points. First, is routine automated use of adaptive grid refinement to perform resistance computations possible? And second, is it advantageous with respect to more traditional meshing methods?

560 *Is it possible?* As discussed in section 1, routine mesh refinement requires a protocol without trial and error to set precise values for all simulation parameters, that is general enough to apply to a wide range of ships and conditions. The protocol in section 3 is precise. But how general is it?

565 For bare hulls, the protocol is applied in this paper to displacement hulls which range from slender to full. The DTMB 5415 test shows that the procedure is effective over a range of velocities from slow steaming to military full speed. In the course of this test campaign, several other hulls such as the Duisburg Test Case (DTC) and the Virtue tanker have been tested, which are not shown here because they are similar to the cases presented. In all cases, the agreement with experiments is reasonable or better. Furthermore, the adapted-mesh solutions 570 converge well and where this was tested, they converge to the same results as non-adapted meshes. Thus, we consider that the protocol is applicable to all bare displacement hulls in fixed position, free in trim and sinkage, and in self-propulsion with a body force or an actuator disk model.

575 Can it also be used with appendices or actual propellers? The choice of
the criterion and the threshold T_{rH} is probably correct: the KCS actuator disk
test suggests that an appendix which generates enough force to move or turn
the ship, creates enough pressure variations to be captured well by adaptive
refinement tuned for the ship. And if ‘inactive’ appendices such as rudders
580 in a neutral position do not attract much refinement, their influence on the
ship’s forces is small so they may not require an accurate representation by the
grid. However, the minimum cell size for the ship may be too coarse to capture
the flow around the appendices. While the current settings are acceptable for
actuator disks, rudders and propellers may require locally reduced minimum cell
585 sizes or other modifications of the criterion. This is a subject for future study.

Further questions are the extension of combined Hessian and free-surface
refinement to other hull forms such as planing hulls, other approaches such as
wall-resolved turbulence or full-scale simulation, or unsteady simulations like
added resistance in waves. For all these applications, FINE/Marine users em-
590 ploy free-surface refinement (with the fine mesh around the hull created by
Hexpress) on a routine basis. Thus, the only open question is the choice of the
Hessian parameters. We have little experience with unsteady simulations, so
more testing is required for these applications. For all steady cases, it is likely
that the correct range for T_{rH} is more or less the same as the one given here,
595 since a similar range is even applicable to wings and hydrofoils (Wackers, 2019)).
And other steady cases are close enough to what was tested here, that even if
the protocol in section 3 has to be adapted for them, it is likely that a similar
protocol exists. This answers the main question: routine mesh adaptation for
resistance simulations is realistic today.

600 *Is it advantageous?* If Hessian-based adaptation is to be introduced widely
for resistance computations, it has to offer advantages with respect to existing
meshing approaches. A first advantage, efficiency, is discussed in section 4.4:
figure 10 shows that adapted meshes provide similar performance as a pure
Hexpress mesh for situations where the Hessian refinement plays almost no
605 role (300k cells). For all finer meshes, the adaptive approach reaches the same
accuracy with significantly less cells and thus faster computations.




Figure 21: Water surface (volume fraction α) for the DTMB 5415 at $Fr=0.28$ with standard non-adapted mesh (left) and adapted mesh (right).

Furthermore, mesh adaptation makes it easier to obtain reliable simulations, since the adaptation automatically ensures that the mesh is suitable for the flow. For example, figure 21 shows the DTMB 5415 at $Fr = 0.28$, which is a moderate

610 Froude number for a destroyer. However, the free-surface refinement in a state-
of-the-art non-adapted mesh fails to capture the bow wave, which may introduce
errors in the computed forces; this mesh would require manual modification and
a new simulation. The adapted mesh has the free-surface refinement directly in
the right position to capture the waves.

615 A third advantage is that adaptive-refinement based forces behave remark-
ably well for uncertainty estimation. Sharp, meaningful uncertainty estimations
increase the confidence in the results by providing a quantitative measure of their
reliability, which is considered more and more important (see for example Stern
et al. (2017)). And yet, section 4.4 shows that uncertainty estimation is not a
620 trivial matter for general unstructured grids. For adapted meshes however, the
grid convergence in all cases presented here is smooth, which leads to accurate
power-law fits to the data; this was also observed by (Wackers et al., 2017)
for mono-fluid flows. Furthermore, Alauzet and Frazza (2021) observed similar
good convergence for a completely different setting: simulations of wings and
625 aircraft using compressible flow and tetrahedral mesh adaptation. This means
that the good convergence behaviour is likely a property of the metric-based
refinement. The capacity of adaptive refinement to provide uncertainty esti-
mations for routine simulations on unstructured grids could be a major step
forward for practical CFD simulations.

630 A final point is that multilevel simulations like in section 6 are easy to per-
form with adaptive refinement. This technique can accelerate the computations
to a point where both the result and the uncertainty estimation are produced
for less than the cost of a single-level simulation.

A disadvantage of adaptive meshing for resistance simulations is that it is a
635 new technique, contrary to fixed meshing which has the benefit of years of user
experience in research and industry settings. However, for all simulations which
require reliable (nearly-)automatic meshing, such as automatic shape optimisa-
tion or large series of computations, adaptive meshing is one day going to be
a necessity. The current results show that the technique is mature enough to
640 start gaining experience with large-scale practical application.

Acknowledgements

This work was granted access to HPC resources of CINES under the allo-
cation A0092A01308 made by GENCI (Grand Equipement National de Calcul
Intensif), which is gratefully acknowledged.

645 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.,
650 Alauzet, F., 2020. Verification of anisotropic mesh adaptation for turbulent
simulations over ONERA M6 wing. *AIAA J* 58, 1–16.
- Castro-Díaz, M.J., Hecht, F., Mohammadi, B., Pironneau, O., 1997. Anisotropic
unstructured mesh adaptation for flow simulations. *Int J Num Meth Fluids*
25, 475–491.

- 655 Dapogny, C., Dobrzynski, C., Frey, P., 2014. Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems. *J Comput Phys* 262, 358–378.
- Demeester, T., van Brummelen, E.H., Degroote, J., 2021. An efficient quasi-Newton method for three-dimensional steady free surface flow. *Int J Num Meth Fluids* 93, 2581–2610.
660
- Deng, G.B., Visonneau, M., 1999. Comparison of explicit algebraic stress models and second-order turbulence closures for steady flows around ships, in: 7th Int. Conf. on Numerical Ship Hydrodynamics, Nantes, France.
- 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.
665
- Eça, L., Hoekstra, M., 2014. A procedure for the estimation of the numerical 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.
670
- 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.
- Hildebrandt, T., Reyer, M., 2015. Business and technical adaptivity in marine CFD simulations: Bridging the gap, in: COMPIT ’15, Ulrichshusen, Germany. pp. 394–405.
675
- Hino, T. (Ed.), 2005. CFD Workshop Tokyo 2005. National Maritime Research Institute, Tokyo, Japan.
- Kim, W.J., Van, D.H., Kim, D.H., 2001. Measurement of flows around modern commercial ship models. *Exp Fluids* 31, 567–578.
- 680 Kühn, N., Kröger, J., Siebenborn, M., Hinze, M., Rung, T., 2021. Adjoint complement to the volume-of-fluid method for immiscible flows. *J Comput Phys* 440, 110411.
- Leroyer, A., Visonneau, M., 2005. Numerical methods for RANSE simulations of a self-propelled fish-like body. *J Fluid Struct* 20, 975–991.
- 685 Loseille, A., Dervieux, A., Alauzet, F., 2010. Fully anisotropic goal-oriented mesh adaptation for 3D steady Euler equations. *J Comput Phys* 229, 2866–2897.
- Loseille, A., Dervieux, A., Frey, P., Alauzet, F., 2007. Achievement of global second-order mesh convergence for discontinuous flows with adapted unstructured meshes, in: 18th AIAA Conference, Miami, FL.
690
- Michal, T., 2019. Anisotropic solution adaptive meshing for production aerospace applications, in: Tetrahedron VI, Paris Saclay, France.
- Nardoni, C., Bordeu, F., Cortia, J., 2019. Body-fitted discretization, accurate interface tracking and geometrical constraints in topology optimization, in: Adaptive Modelling and Simulation 2019, El Campello, Spain.
695

- Olivieri, A., Pistani, F., Avanzini, A., Stern, F., Penna, R., 2001. Towing tank experiments of resistance, sinkage and trim, boundary layer, wake, and free surface flow around a naval combatant INSEAN 2340 model. Technical Report 421. IIHR. Iowa.
- 700 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.
- Pons, K., Ersoy, M., 2019. Adaptive mesh refinement method. Part 1: Automatic thresholding based on a distribution function. URL: <https://hal.archives-ouvertes.fr/hal-01330679>. unpublished report.
- 705 Pržulj, V., Basara, B., 2001. Bounded convection schemes for unstructured grids, in: 15th AIAA Computational Fluid Dynamics Conference, Anaheim, CA.
- Queutey, P., Visonneau, M., 2007. An interface capturing method for free-surface hydrodynamic flows. *Comput Fluids* 36, 1481–1510.
- 710 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.
- 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.
- 715 Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., Mavriplis, D., 2014. CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences. Technical Report NASA/CR–2014-218178. NASA.
- 720 Stern, F., Diez, M., Sadat-Hosseini, H., Yoon, H., Quadvlieg, F., 2017. Statistical approach for computational fluid dynamics state-of-the-art assessment: N-version verification and validation. *J Verification, Validation and Uncertainty Quantification* 2, 031004.
- Van, S.H., Kim, W.J., Yim, G.T., Kim, D.H., Lee, C.J., 1998. Experimental investigation of the flow characteristics around practical hull forms, in: 3rd Osaka Colloquium on Advanced CFD Applications to Ship Flow and Hull Form Design, Osaka, Japan.
- 725 Wackers, J., 2019. Adaptivity for complex flows. HDR thesis, Université de Nantes, France.
- 730 Wackers, J., 2021. Industry application of adaptive grid refinement – the case of FINE/Marine, in: Adaptive Modelling and Simulation 2021, online.
- Wackers, J., Deng, G.B., Guilmineau, E., Leroyer, A., Queutey, P., Visonneau, M., 2014. Combined refinement criteria for anisotropic grid refinement in free-surface flow simulation. *Comput Fluids* 92, 209–222.

- 735 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.
- 740 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.
- 745 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.
- 750 Windt, J., Klaij, C., 2012. Adaptive mesh refinement in MARIN’s viscous flow solver ReFRESCO: implementation and application to steady flow, in: *MARINE 2011*, Lisbon, Portugal.
- Yilmaz, N., Atlar, M., Khorasanchi, M., 2019. An improved mesh adaption and refinement approach to cavitation simulation (MARCS) of propellers. *Ocean Eng* 171, 139–150.
- 755 Yvin, C., Muller, P., 2016. Tip vortex cavitation inception without a cavitation model, in: *19th Numerical Towing Tank Symposium*, St Pierre d’Oléron, France.