Combined refinement criteria for anisotropic grid refinement in free-surface flow simulation

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

► To cite this version:

Jeroen Wackers, Ganbo Deng, Emmanuel Guilmineau, Alban Leroyer, Patrick Queutey, et al.. Combined refinement criteria for anisotropic grid refinement in free-surface flow simulation. *Computers and Fluids*, 2014, 92, pp.209-222. 10.1016/j.compfluid.2013.12.019 . hal-01155455

HAL Id: hal-01155455

<https://hal.science/hal-01155455>

Submitted on 22 Feb 2021

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

Combined refinement criteria for anisotropic grid refinement in free-surface flow simulation

Jeroen Wackers*, Ganbo Deng, Emmanuel Guilmineau,
Alban Leroyer, Patrick Queutey and Michel Visonneau
Laboratoire de recherche en Hydrodynamique, Energétique et Environnement Atmosphérique,
CNRS UMR 6598, Ecole Centrale de Nantes,
1 rue de la Noë, B.P. 92101, 44321 Nantes Cedex 03, France

Abstract

Anisotropic grid refinement is performed for the simulation of water flow with free-surface waves. For these flows, the refinement criterion must provide refinement at the water surface, to accurately resolve the conservation law which indicates the surface position, and below the surface to resolve the water flow. A combined criterion is presented, based on the free-surface position and on the Hessian of the pressure. Different forms of this criterion are presented, based on least-squares or Gaussian computation of the Hessian, in order to overcome irregularities in the computed pressure. The weighting of the two criteria for their combination is discussed; this weighting can be chosen independently of the Reynolds and Froude number of the flow. It is shown that the criterion creates suitable grids for two- and three-dimensional free-surface flows when starting from uniformly coarse original grids.

Keywords : Combined criteria, Hessian criteria, anisotropic refinement, free-surface waves.

1 Introduction

The simulation of two-fluid flows with a free surface between the fluids is inherently a multiphysics problem, since the motion of the free surface interacts with the turbulent viscous flows on both sides of the interface. In Navier-Stokes models for such flows, an equation for the water surface position has to be added to the standard flow equations [26], like a convection equation for either the volume fraction of water (volume-of-fluid methods) or for a smooth distance function to the surface (level set methods).

Automatic grid refinement is the technique of creating optimal meshes for a particular flow, by locally dividing the cells of a coarse original grid into smaller cells, so the grid resolution is increased there where greater precision is needed. Grid refinement methods are widespread, see for example [2, 11, 14]. Free-surface water flows have many features which are local in nature, so their precision can be increased with adaptive grid refinement. First, refinement around the surface strongly increases the resolution of the volume fraction equation [9, 8, 27, 28], so the modeling of the free surface is improved. But also other aspects of these flows, such as wakes and trailing vortices, are resolved with greater precision when grid refinement is used [10, 24, 29]. And finally, complex physical phenomena such as cavitation appear in general very locally and can be computed efficiently on adaptively refined grids.

The unstructured Reynolds-averaged Navier-Stokes solver ISIS-CFD which we develop contains an automatic grid refinement method [24, 27]. This flow solver, distributed by NUMECA Int. as part of the FINE/Marine computing suite, is aimed at the simulation of realistic flow problems in all branches of marine hydrodynamics. The grid refinement method is therefore developed to be general and flexible,

*Corresponding author: jeroen.wackers@ec-nantes.fr, +33 2 40 37 16 81

featuring anisotropic refinement on unstructured hexahedral grids, derefinement of previous refinements to enable unsteady flow computation, and full parallelisation including integrated dynamic load balancing. The anisotropic refinement is based on metric tensors.

In our earlier work on grid refinement for free-surface flows, the multiphysics character of the flows was not explicitly taken into account for the grid refinement. Instead, the original grid was chosen sufficiently fine to get a reasonable resolution of the flow, then automatic grid refinement was used to improve the accuracy of one particular flow feature. Thus, gravity waves at the water surface were computed with refinement based on the discontinuity in the volume fraction [28] and wake flows with refinement based on the pressure [24, 25].

The goal of the present paper is to obtain adapted fine grids for free-surface flows, starting from a coarse original grid without any initial refinement. To reach this goal, the focus is on the refinement criterion. In our mesh refinement technique, the criterion is a real field variable computed from the solution, which indicates everywhere in the flow domain the ideal sizes of the cells. For metric-based anisotropic refinement, the criteria are 3×3 symmetric tensors in each cell that indicate the local desired cell sizes independently in all directions. Given such a criterion, our existing mesh refinement module [27] automatically creates the best corresponding adapted mesh. Thus, the problem of producing refined free-surface meshes comes down to finding the right refinement criterion for free-surface flow.

The discussion in this paper follows the development of this refinement criterion, from the analysis of mesh requirements to practical details of the criterion computation. Good meshes for free-surface flow are studied first. We analyse the discretisation of the flow equations and the free surface model, as well as the physics of flows around bodies with free-surface effects and gravity waves, in order to find out what properties of the mesh are important for the accurate computation of such flows. The conclusions from this study are translated into requirements for the refinement criterion: it is shown that the multiphysics nature of two-fluid flows demands refinement criteria that are a combination of different physical sensors. We therefore discuss how multiple refinement criteria can be combined into one, how the different criteria should be weighted to achieve good accuracy in all equations, and which features are relevant as refinement criteria for hydrodynamic flows. The selected criterion combines directional refinement at the free surface with a pressure Hessian criterion that is modified to overcome the gravity-induced discontinuity in the pressure gradient at the water surface. Simulations of travelling wave trains and ship-generated wave patterns are presented to test these developments and to assess practical guidelines for their use.

The paper is organised as follows. Section 2 introduces the flow solver and the meshes used, concentrating on those aspects that are important for grid refinement. Section 3 gives an overview of the anisotropic mesh refinement method. Then, section 4 discusses the necessity of combined criteria for flows with a free water surface. Section 5 shows the construction of a criterion that combines directional refinement at the free water surface with a pressure Hessian criterion. Finally, section 6 analyses the weights that should be chosen for each criterion when two criteria are combined. Three test cases in section 7 indicate that the criterion generates effective meshes for a wide range of free-surface and cavitating flows. With these meshes, similar accuracies as on user-generated fine meshes are obtained with significantly less computational costs.

2 Surface capturing and finite-volume discretisation

The study of combined refinement criteria in this paper is conducted in the context of a finite-volume discretisation for the the Unsteady Reynolds-averaged Navier Stokes (URANS) equations on unstructured grids. This section describes the governing flow equations, the discretisation, and the type of meshes used, concentrating on those aspects that are most important for grid refinement and the construction of refinement criteria. Full details of the discretisation can be found in [19, 26].

2.1 Governing equations

The ISIS-CFD flow solver which we develop resolves the incompressible URANS equations in a mixture-fluid formulation to model water–air two-fluid flow. Here, the entire fluid is modelled as a numerical mixture of the pure fluids on the two sides of the interface. The system uses the conservation laws for momentum, total mass, and mass of each fluid. When the densities of the individual fluids are constant, the latter two reduce to $\nabla \cdot \mathbf{U} = 0$ and to a volume-of-fluid equation. In integral form, the equations are written as follows:

$$\frac{\partial}{\partial t} \int_V \rho \mathbf{U} dV + \int_S \rho \mathbf{U} (\mathbf{U} \cdot \mathbf{n}) dS = \int_S (\boldsymbol{\tau} - pI) \mathbf{n} dS + \int_V \rho \mathbf{g} dV, \quad (1)$$

$$\int_S \mathbf{U} \cdot \mathbf{n} dS = 0, \quad (2)$$

$$\frac{\partial}{\partial t} \int_V \alpha_i dV + \int_S \alpha_i \mathbf{U} \cdot \mathbf{n} dS = 0, \quad (3)$$

where V is a fixed volume, bounded by the closed surface S with a unit normal vector $\mathbf{n}$ directed outward. $\mathbf{U}$ and p represent, respectively, the velocity and pressure fields. $\boldsymbol{\tau}$ and $\mathbf{g}$ are the viscous stress tensor and the gravity vector, whereas I is a unit diagonal matrix. α_i is the volume fraction for fluid i and is used to distinguish the presence ($\alpha_i = 1$) or the absence ($\alpha_i = 0$) of fluid i . In the case of turbulent flows, additional transport equations are added to the system.

The effective flow physical properties (viscosity μ and density ρ) are obtained from the physical properties of the constituent fluids (μ_i and ρ_i) with the following constitutive relations. The last identity follows from the definition of the volume fraction.

$$\rho = \sum_i \alpha_i \rho_i, \quad \mu = \sum_i \alpha_i \mu_i, \quad 1 = \sum_i \alpha_i. \quad (4)$$

Thus, for two fluids, equation (3) only has to be solved for fluid 1; for this case, we denote the volume fraction of fluid 1 by α .

In this framework, free-surface water flows are modelled by specifying a discontinuous inflow condition for α ($\alpha = 1$ below the surface and $\alpha = 0$ above it). As equation (3) is a pure convection equation, the resulting solution for α in the whole domain is a discontinuity representing the free surface.

2.2 Face-based discretisation

The flow equations of the previous subsection are discretised in a finite-volume framework, using pressure-velocity coupling obtained through implicit time integration with a Rhie & Chow SIMPLE-type method [20]. The discretisation is face-based. While all unknown state variables are cell-centered, the systems of equations used in the implicit time stepping procedure are constructed face by face. Fluxes are computed in a loop over the faces and the contribution of each face is then added to the two cells next to the face. This technique poses no specific requirements on the topology of the cells. Therefore, the grids can be completely unstructured, cells with an arbitrary number of arbitrarily-shaped faces are accepted.

2.2.1 Central discretisations

The core of the discretisation are the reconstructions of the state variables and their derivatives from the cell centres to the faces. For the diffusive fluxes and the pressure equation, these are basically central approximations of the normal derivatives. In case of misalignment, i.e. when the face normal vector is not aligned with the line between the neighbouring cell centres, extra correction terms are added using the cell-centered gradients computed with a Gauss method. These discretisations will be reused in the construction of refinement criteria (section 5.2).

As an example, the reconstruction of the pressure and its normal derivative is shown here for the case of constant density, i.e. below the surface. More details on the reconstruction, including the weighting

Figure 1: Cell face notations.

of the reconstruction with the density when it is non-constant, are given in [19, 26]. The pressure reconstruction used in the momentum equations is (see figure 1 for notations):

$$p_f = \frac{h^+}{h} p_L + \frac{h^-}{h} p_R + \left(\frac{h^- \mathbf{E}^+ - h^+ \mathbf{E}^-}{h} \right) \cdot \left(\frac{h^+}{h} \vec{\nabla} p_L + \frac{h^-}{h} \vec{\nabla} p_R \right). \quad (5)$$

Geometrical vectors $\mathbf{E}^\pm$ are introduced so that the framed term contribution goes to zero when the grid becomes orthogonal ($\mathbf{L}\mathbf{f}\cdot\mathbf{n} = \mathbf{f}\mathbf{R}\cdot\mathbf{n} = 0$):

$$\begin{aligned} \mathbf{E}^- &\triangleq (\mathbf{L}\mathbf{f}\cdot\mathbf{n}) \mathbf{n} - \mathbf{L}\mathbf{f}, \\ \mathbf{E}^+ &\triangleq (\mathbf{f}\mathbf{R}\cdot\mathbf{n}) \mathbf{n} - \mathbf{f}\mathbf{R}. \end{aligned} \quad (6)$$

Distances used are the projected distances to the face $h^\pm$ and the projected distance h between the L and R cell centers:

$$h^- = \mathbf{L}\mathbf{f}\cdot\mathbf{n}, \quad h^+ = \mathbf{f}\mathbf{R}\cdot\mathbf{n}, \quad h = h^- + h^+ = \mathbf{L}\mathbf{R}\cdot\mathbf{n}. \quad (7)$$

The reconstruction of the pressure normal derivative for the pressure equation is:

$$\left(\vec{\nabla} p \cdot \mathbf{n} \right)_f = \frac{p_R - p_L}{h} + \left(\frac{\vec{\nabla} p_L \cdot \mathbf{E}^- + \vec{\nabla} p_R \cdot \mathbf{E}^+}{h} \right). \quad (8)$$

Here again, the framed term contribution goes to zero when the grid becomes orthogonal.

2.2.2 Convective fluxes

The convective fluxes are computed using limited schemes in Leonard's Normalised Variable Diagram (NVD) [12]. These schemes use nonlinear interpolation between three points for the reconstruction at a cell face, depending on the flow direction, see figure 2. The points C and D are the cell centres of the two neighbour cells, in the upwind and downwind direction. The point U , for structured grids, is the cell centre upwind of point C ; for unstructured grids there usually is no suitable cell at this point, so the value in U is extrapolated from C using the gradient computed in C . Many different schemes exist, the standard scheme in ISIS-CFD is the AVLSMART scheme [18].

For the volume fraction equation, the solutions are known a-priori to be numerical approximations to a discontinuity at the water surface, with constant values elsewhere. For these special solutions, normal accuracy considerations do not hold. Therefore, downwind-biased NVD schemes are used which add artificial antidiffusion (negative diffusion) to the equation while preserving stability, see for example [23]. Antidiffusion continuously compresses the interface in α so it remains as sharp as possible. As

Figure 2: Points for NVD face reconstruction. The points C and D are the neighbouring cells of the face f , the point U is the mirror image of D with respect to C .

the discontinuous α does not have a well-defined gradient, the state U in our BRICS scheme [26] is not extrapolated from C but interpolated from the cell centre nearest to the point U and its neighbours.

Misalignment corrections, for the case where the face centre does not lie on the midpoint of the line CD (the point of interpolation) are not used. For the volume fraction equation, these are difficult to envisage because no useful gradient information exists. However, misalignments have a strong negative effect on the accuracy of the solution for α , so if possible the mesh should be made such that they are prevented as much as possible near the free surface.

2.3 Meshes

For this study, as usual for ISIS-CFD, unstructured hexahedral meshes generated with the HEXPRESS grid generator from NUMECA Int. are used (see figure 3 for an example). In these meshes, variations in cell size are handled by having small cells laying next to larger cells. This situation is called hanging nodes in the literature and solvers often need specific discretisations to handle these topologies. In ISIS-CFD, due to the face-based algorithm, these cells are treated in exactly the same way as all the others: the larger cells are simply seen as cells with more than 6 faces. Thus, no specific hanging-node treatment is included.

Figure 3: Cut through an unstructured hexahedral mesh.

Unstructured hexahedral grids are ideal for automatic grid refinement. Isotropic or anisotropic grid refinement can be applied to any of the hexahedral cells, the result will still be an unstructured hexahedral mesh. Therefore, locally refined meshes can be used directly in a flow solver that supports unstructured

hexahedral meshes, without requiring changes to the flow solver.

Due to the small cell – large cell transitions, strong cell misalignments exist in parts of the grid. When the grid is refined, these situations persist: no matter the size of the grid, there will always be cells that have neighbours twice smaller than themselves. Thus, misalignment problems may limit the accuracy of the solutions, certainly if the percentage of cells which have misaligned faces remains constant when the grid is refined. The only way to ensure good grid convergence is to make sure that the number of same-sized cells between the cell size transitions becomes larger and larger as the grid is refined, so the percentage of cells that have misalignments is reduced. For automatic mesh refinement, it is therefore essential to ensure grids with smoothly varying cell sizes.

3 Grid refinement procedure

The grid refinement procedure developed for ISIS-CFD [24, 27] is integrated completely in the flow solver. The method is entirely parallelised, including automatic redistribution of the grid over the processors. During a flow computation, the refinement procedure is called repeatedly. In such a call, first a refinement criterion is calculated, which is a real field variable based on the flow field that indicates where cells should be refined. Then, in a separate step of the procedure the grid is refined based on this criterion. These steps are kept separated so the criterion can be changed easily without modifying the rest of the refinement method. For steady flow, the refinement procedure eventually converges: once the grid is correctly refined according to the criterion, further calls to the procedure no longer cause any changes.

3.1 Anisotropic refinement

Grid refinement for hexahedral cells can be either isotropic, where a cell is always refined in all its directions at once, or anisotropic where division in only one or two directions is possible as well. For realistic applications, anisotropic refinement is essential. Isotropic grid refinement is very costly in three dimensions, since every refinement of a cell means a division in eight. Thus, creating very fine cells to accurately resolve a local flow phenomenon becomes far too expensive. However, by applying anisotropic refinement for flow features that require a fine grid in only one direction (notably, the water surface), the total number of cells required can be greatly reduced, or much finer flow details can be resolved.

A second reason for directional refinement is, that our refinement is based on unstructured hexahedral original grids as shown in figure 3. In these grids, cells of completely different aspect ratios lie side by side. Therefore, when refining, we need to control the size of the fine cells in all their directions independently, otherwise refined grids may have smoothly varying sizes in one direction, but repeated changes from fine to coarse and back to fine in another [27]. Isotropic refinement is not enough to prevent this, so directional refinement is the mandatory choice.

3.2 Tensor refinement criteria

For directional refinement, a way is needed to specify different cell sizes in different directions. The use of metric tensors as refinement criteria is such a way. This technique was first developed for the generation and refinement of unstructured tetrahedral meshes [1, 7, 14]. It is also an extremely useful and flexible framework for the refinement of unstructured hexahedral meshes.

In the metric context, the refinement criterion is a smoothly varying tensor field whose values at every point in the flow domain indicate what the ideal size for a cell in that position would be. As such, it can be thought of as the continuous equivalent of a mesh [1]. This ideal mesh depends on the flow field. There exists an ‘exact’ criterion which is computed from the exact solution; as the grid is refined the actual computed criterion converges to this exact criterion (this is different from the classical error indicating criteria where the criterion is halved when a cell is refined.) Adaptive grid refinement is performed to get the actual cell sizes in the refined grid as close to these ideal sizes as possible, so the refined mesh can be considered as a ‘discretisation’ of the criterion.

Figure 4: Tensor refinement criterion. Cell Ω and unit circle (reference) in the physical space, deformed cell $\tilde{\Omega}$ and deformed circle after application of the transformation C , refinement decisions to create a uniform grid in the deformed space, and the resulting anisotropically refined grid.

The refinement criterion in each cell is a 3×3 symmetric positive definite matrix $\mathcal{C}$, which is interpreted as a geometric transformation of the cell in the physical space to a deformed space (figure 4). The refinement of the cells is decided as follows. Let the criterion tensors $\mathcal{C}$ in each cell be known (their computation from the flow solution is described in section 5). In each hexahedral cell, the cell size vectors $\mathbf{d}_j$ ($j = 1, 2, 3$), which are the vectors between the opposing face centres in the three cell directions, are determined. Next, the modified sizes are computed as:

$$\tilde{\mathbf{d}}_j = \mathcal{C} \mathbf{d}_j. \quad (9)$$

Finally, a cell is refined in the direction j when the modified size exceeds a given, constant threshold value T_r :

$$\|\tilde{\mathbf{d}}_j\| \geq T_r. \quad (10)$$

Thus, the objective of the refinement is to create a uniform grid in the deformed space. The tensors $\mathcal{C}$ are direct specifications of the desired cell sizes: in a converged refined grid, the cell sizes are inversely proportional to the magnitude of the $\mathcal{C}$. The threshold T_r functions as a global specification of the fineness of the grid; sensible choices for T_r in different situations are discussed in section 7 (see also sections 5.3 and 6.1).

3.3 Grid misalignment

The convection equation for the volume fraction is sensitive to the quality of the grid, as noted in section 2.2.2. Specifically, faces parallel to the water surface that have a large cell on one side and small cells on the other (see figure 5a) cause a distortion and diffusion of the water surface, since there is a strong misalignment between the face normal and the line connecting the centres of its two neighbour cells [28].

This problem is reduced by enforcing the refinement of extra cells to keep a good grid quality. For this, a quality criterion is applied (figure 5b): if refining the cell at one side of a face would cause too great an angle between the face normal and the line cell centre - face centre, the cell on the other side of the face is refined as well (figure 5c). The enforcement of this criterion improves the regularity of the grids at the water surface. However, on its own it is not enough.

Figure 5: Misalignment distorts the water surface (a), the curves are volume fraction isolines. Quality criterion (b): large angles between face normals and lines cell–face are forbidden. Added refinement (c).

4 The need for combined refinement criteria

Gravity waves in free-surface flows are usually generated by the flow passing around a foreign body, either a stationary object or a floating body such as a ship. The direct cause of the wave generation are the pressure and velocity disturbances created by the presence of this body. These disturbances are not only generated at the free surface, but also well below it; even a fully submerged object may create waves (see section 7.1 for an example). Then, once the waves are created, they propagate through a cyclic exchange of potential (gravity) and kinetic energy. Water particles in a travelling wave field describe an orbital motion; the velocities associated with this motion cause the propagation of the wave energy. Thus, to correctly resolve the generation and the propagation of a travelling surface wave, a good resolution must be obtained for the pressure and velocity fields below the surface, as given by equations (1) and (2).

Accurate computation of the volume fraction equation (3) is of prime importance as well. As the water surface is physically a discontinuity, the interface region for the volume fraction α must be as sharp as possible. This is the reason for the use of antidiffusive schemes for the volume fraction equation (section 2.2), which can be shown to give less numerical damping of waves than diffusive schemes [19]. It follows that a coarse grid at the free surface will also damp out waves. In our experience, the grid at the surface needs to be about twice as fine as the grid used in the vicinity of the surface, in order to resolve equation (3) correctly [28].

Thus, to create suitable grids, a grid refinement criterion for free-surface wave simulation must be based both on the pressure and velocity field and on the volume fraction. For these two, different indicators must be used. The reason for this is, that α is discontinuous at the free surface and constant everywhere else, while the pressure and the velocity are smooth in the whole flow field except at the surface. Therefore, spatial derivative based error indicators can be used to identify the regions of importance for the flow field below the surface, but not to find the water surface itself, since any derivative-based indicator applied to α would go to infinity when the grid is refined.

Also, grid misalignment must be avoided in the free-surface region, as it leads to large errors in the volume fraction (section 3.3). Therefore, the grid specified by the criterion must be as uniform as possible near the surface. Numerical evaluations of the derivatives of α can never be smooth, so they are not suitable for guaranteeing a uniform grid. Instead, a criterion can be based on the value of α itself.

Therefore, a suitable refinement criterion for water flow with waves is an error indicator for the flow field and a simpler criterion for α , combined into one.

5 Free surface – pressure Hessian criterion

For the simulation of flow with waves, we propose a criterion based on the Hessian matrix of second derivatives of the pressure, combined with a criterion that refines in the normal direction of the surface

for those cells where α is neither 0 nor 1. This section describes both of these criteria individually, then shows how they can be combined into one.

5.1 Free-surface criterion

To resolve accurately the solution of equation (3) which is a discontinuity for α that is convected with the flow, it is sufficient to refine the grid at and around the free surface, in the direction normal to the surface. When the surface is locally aligned with the cell directions, anisotropic refinement can be used to keep the total number of cells as low as possible. As noted in the previous sections, it is important that cells locally have the same size, to prevent misalignments near the surface.

The free-surface criterion $\mathcal{C}_S$ is based on α in the cells, it is non-zero when α is neither 0 nor 1. The normal direction to the surface is computed from a field α_s which corresponds to α , smoothed out by averaging over a cell and its neighbours a given number of times. The gradient of this field gives the normal directions. The criterion is then derived from vectors $\mathbf{v}_\alpha$ in each cell which are unit vectors in this normal direction for those cells where the smoothed α_s field is non-zero:

$$\mathbf{v}_\alpha = \begin{cases} \frac{\nabla \alpha_s}{\|\nabla \alpha_s\|} & \text{if } 0.1 \leq \alpha_s \leq 0.9, \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (11)$$

Using the smoothed field guarantees that the normals are well defined and also that the mesh is refined in a certain zone around the surface to create a margin of safety.

In tensor form, the free-surface criterion is computed as matrices having only one non-zero eigenvalue, associated with the direction of the vector $\mathbf{v}_\alpha$. The tensors $\mathcal{C}_S$ are computed as follows (with $\otimes$ representing the tensor product):

$$\mathcal{C}_S = \mathbf{v}_\alpha \otimes \mathbf{v}_\alpha. \quad (12)$$

In the directions normal to the vector $\mathbf{v}_\alpha$, the eigenvalues are zero. This implies a modified cell size of zero (equation (9)). As a consequence, the grid is not refined in these directions, the original cell sizes are kept. Since the $\mathbf{v}_\alpha$ are unit vectors near the surface, the only non-zero eigenvalues of $\mathcal{C}_S$ are equal to 1. Thus, from equation (10) it follows that the threshold value T_r directly indicates the desired cell size at the surface. We also see that the specified cell size normal to the surface is *exactly* the same in all surface cells, as required. The free-surface criterion has been used on its own, with good results, in our earlier work [28, 24, 27].

5.2 Computing the pressure Hessian

Hessian-based criteria are often used to control anisotropic grid refinement [1, 14, 15], since various demonstrations show that the Hessian matrix can be considered as an indicator of the local truncation error. Here this criterion is based on the pressure, because of our refinement strategy in boundary layers [27]. We consider that the number of layers in the boundary layer grid should be the same everywhere, to ensure the best grid quality. And since the approximate thickness of the boundary layer is known, these grid layers can be inserted on the original grid. Therefore, it is unnecessary to employ a criterion which has very high values in the boundary layer region. The pressure varies little over the thickness of a boundary layer, so its second derivatives there are limited.

To compute the Hessian matrix of a numerical solution, second-derivative operators must be discretised. A particular complication for this discretisation is that our meshes always contain places where the grid size changes abruptly, as small cells lie next to twice larger cells (see section 2.3). These places do not disappear when the mesh is refined, on the contrary their number increases significantly when automatic refinement is used. However, many discretisations of the second spatial derivatives depend on the mesh becoming smoother and smoother as it is refined, in order to obtain second-order accuracy. In particular, the well-known computation of the Hessian by using the Gauss theorem for finding the gradients of the quantity, then applying the Gauss theorem again to the gradients in order to compute the second derivatives, has an accuracy of order zero in places where the grid size changes abruptly. A

suitable technique for computing the Hessian, on the other hand, must be insensitive to these cell size changes. This section describes two possible ways of computing the Hessian matrix.

5.2.1 Definition of the Hessian criterion

The pressure Hessian matrix is:

$$\mathcal{H}(p) = \begin{bmatrix} (p)_{xx} & (p)_{xy} & (p)_{xz} \\ (p)_{xy} & (p)_{yy} & (p)_{yz} \\ (p)_{xz} & (p)_{yz} & (p)_{zz} \end{bmatrix}. \quad (13)$$

This matrix can be used directly as a tensor refinement criterion. However, if we assume heuristically that an indication of the local truncation error is given by $\mathcal{H}$ times the cell sizes to a power b (where b depends on the numerical method), equidistribution of this error indicator leads to a refinement criterion where the Hessian matrix is modified with a power law:

$$\mathcal{C}_H = (\mathcal{H}(p)^a), \quad (14)$$

where $\mathcal{H}^a$ has the same eigenvectors as $\mathcal{H}$ and eigenvalues that are those of $\mathcal{H}$ (in absolute value) to the power $a = \frac{1}{b}$. In general, we use $a = \frac{1}{2}$ which is appropriate for a second-order accurate discretisation.

5.2.2 Third-order least-squares (LS3) approximation

A first solution for the computation of the Hessian is to use a least-squares approximation [27]. In each cell, we construct a least-squares fit of a third-order polynomial to the solution in the cell, its neighbour cells and its neighbours' neighbours. The approximated Hessian is then computed from the second derivatives of this polynomial. Let $P_j(\mathbf{x})$, $j = 1 \dots 20$ be the set of basic three-dimensional polynomial functions in $\mathbf{x}$ of up to third order (i.e. $P_1(\mathbf{x}) = 1$, $P_2(\mathbf{x}) = x$, $P_3(\mathbf{x}) = y$, $\dots$, $P_5(\mathbf{x}) = x^2$, $\dots$, $P_{11}(\mathbf{x}) = x^3$, $\dots$, $P_{20}(\mathbf{x}) = xyz$). Furthermore, $\mathbf{I}$ is the vector of cell indices of a cell i , its neighbours and its neighbours' neighbours. Then we shall search coefficients β such that the polynomial:

$$\phi(\mathbf{x}) = \sum_{j=1}^{20} \beta_j P_j(\mathbf{x} - \mathbf{x}_i), \quad (15)$$

is the closest fit to the values of p in the cell centres of $\mathbf{I}$, within the space defined by the set $P_j(\mathbf{x} - \mathbf{x}_i)$. Defining the matrix A and vector $\mathbf{b}$ as:

$$A_{jk} = P_j(\mathbf{x}_{\mathbf{I}_k} - \mathbf{x}_i), \quad b_k = p_{\mathbf{I}_k}, \quad j = 1 \dots 20, k = 1 \dots \text{size}(\mathbf{I}), \quad (16)$$

the coefficients β are found as:

$$\beta = (A^T A)^{-1} A^T \mathbf{b}. \quad (17)$$

According to the definition of the least-squares procedure, there is no better third-order polynomial fit to the points $\mathbf{I}$, so the error in the fit is at least fourth-order. Therefore, if p is a sufficiently smooth function, the approximated Hessian $\mathcal{H}_{LS3}(p)$ is a second-order accurate approximation to $\mathcal{H}(p)$ (two orders are lost by the double differentiation). Tests with manufactured solutions in [27] confirm this on our refined grids.

5.2.3 Smoothed Gauss (SG) method

Unfortunately, the numerically evaluated pressure p_h is *not* a smooth function. Our SIMPLE-based pressure equation contains a Laplace-type operator in finite-volume form, for which the fluxes over the faces are based on the normal derivatives of the pressure computed with equation (8). On arbitrary meshes, these are first-order accurate. Therefore, the truncation error of the Laplace equation which contains the derivatives of the fluxes is formally of order zero. The solutions for the pressure still have at least first-order accuracy ($p_h = p + \mathcal{O}(h)$ where h is a measure of the grid size), because these local truncation errors depend on the relative sizes of a cell and its neighbours so they have opposite signs in

small and large cells lying next to each other, which means that they mostly cancel globally. However, the second derivatives of the pressure appear directly in the pressure equation so they have the same order of accuracy as the truncation error, i.e. $\mathcal{H}(p_h) = \mathcal{H}(p) + \mathcal{O}(1)$. It has been numerically confirmed for an 1D case that the LS3 Hessian gives errors of order zero where small cells lie next to larger cells.

The consequence for grid refinement is, that refining cells creates large errors in the Hessian on the boundaries between finer and coarser cells. Thus, the grid is not only refined where the solution dictates it, but also in places where it has already been refined. This spurious refinement leads to irregular meshes.

As the error in the Hessian is related to small-scale irregularities in the pressure field, it can be reduced by smoothing. Therefore, we define a smoothed Gauss (SG) Hessian. Let the Gauss approximation to the gradient of a field q be given as:

$$\vec{\nabla}_G(q) = \frac{1}{V} \sum_f q_f S_f \mathbf{n}_f, \quad (18)$$

where the face values q_f are computed with the expression (5) and V is the volume of the cell, S_f are the areas of the faces. With the same face reconstructions, a Laplacian smoothing $\mathcal{L}$ is defined as:

$$\mathcal{L}(q) = \frac{\sum_f q_f S_f}{\sum_f S_f}. \quad (19)$$

Then the SG Hessian is computed as follows:

1. Compute the gradient of p using $\vec{\nabla}_G$,
2. Smooth each component of the gradient by applying N times the smoothing $\mathcal{L}$, where $N = 4$ is sufficient in most cases,
3. Compute the gradients of the smoothed gradient components using $\vec{\nabla}_G$,
4. Symmetrize the resulting Hessian matrix by setting $(\mathcal{H})_{ij} = \frac{1}{2}((\mathcal{H})_{ij} + (\mathcal{H})_{ji})$,
5. Smooth the Hessian by applying N times $\mathcal{L}$ to each component.

Since the error in the pressure p_h has an oscillatory component of $\mathcal{O}(h^2)$, differentiating this solution creates an oscillatory error of $\mathcal{O}(h)$ in the first and $\mathcal{O}(1)$ in the second derivatives. The smoother $\mathcal{L}$ uses the same type of interpolation to the faces as the Laplace operator in the pressure equation, so it produces $\mathcal{O}(h^2)$ oscillations itself. Therefore, $\mathcal{L}$ cannot increase the smoothness of p_h , which is the reason why the pressure is not smoothed. On the other hand, the smoothing of the gradients (2) is essential. The $\mathcal{O}(h)$ wiggles in $\vec{\nabla} p_h$ are small compared to the gradient of p so smoothing is very effective for removing these wiggles. The new $\mathcal{O}(h^2)$ oscillations introduced by $\mathcal{L}$ are inoffensive since the gradients are only differentiated once more. In the Hessian, the remaining oscillations of the original solution are of the same order as the solution itself so smoothing cannot improve the accuracy. The step (5) is only applied to create better mesh quality through a smoother criterion.

The resulting Hessian is not second-order accurate but its smoothness makes it interesting as a refinement criterion, since smooth criteria provide good mesh quality. However, smoothing decreases the spurious oscillations in the refinement criterion but also reduces the intensity of physical small-scale features. This limitation of the criterion is the reason that all smoothing should be kept to a minimum.

5.2.4 Hessian at the free surface

The Hessian criterion cannot be directly evaluated at the free surface. The pressure derivatives are proportional to the density ρ according to equation (1), so the exact pressure gradient has a discontinuity in the normal direction at the free surface and the second derivative is a Dirac δ function. For numerical solutions, the second derivative in the zone of varying α has a peak which grows as the grid becomes finer. Numerical differentiation produces large errors in this case.

As a result, no correct values can be computed for the pressure Hessian around the surface so an approximative procedure is needed. For LS3 (section 5.2.2) it is considered that, while the Hessian has a peak at the surface, this peak is associated with an eigenvalue normal to the surface. The pressure gradient parallel to the surface is approximately zero, so the second derivative parallel to the surface is close to zero as well. Therefore, we compute the Hessian at the surface from the unmodified pressure, then we limit all eigenvalues of the matrix $\mathcal{C}_H$ (equation (14)) at the surface to a constant value. The LS3 Hessian criterion then behaves like the free-surface criterion around the surface.

For the smoothed Gauss evaluation (section 5.2.3), this procedure is impossible since the smoothing spreads out the discontinuity in the first and second derivatives, thus creating high values of the Hessian in a large zone around the surface. Therefore, the gradient smoothing (step (2) of the algorithm in section 5.2.3) is not performed in those cells where $0.0001 \leq \alpha \leq 0.9999$. The Hessian is smoothed in three steps:

- 5a. Smooth the Hessian in all cells except those where $0.0001 \leq \alpha \leq 0.9999$ plus two layers of cells around those, to take into account that the perturbed pressure gradient in the cells at the surface influences the Hessian in their neighbours as well.
- 5b. Copy the computed values of the Hessian from outside the zone in (5a) across the zone, following the vertical direction (this removes the peak at the surface). The criterion is copied in the upward direction, so the Hessian values computed in the water are used across the surface region.
- 5c. Smooth the Hessian only in the zone of (5a).

The idea of this approach is, that the SG Hessian at the surface must not be used. Therefore, sensible values must be copied from elsewhere. While both this procedure and the one for the LS3 criterion are heuristic, they work well in practice as will be shown in section 7.

5.3 The combined criterion

The final criterion is a combination of the two criteria above. Considering the problems of the Hessian criterion at the surface, it is tempting to select the free-surface criterion there and the Hessian criterion everywhere else. However, the free-surface criterion specifies no refinement in the direction parallel to the surface, while this refinement may be needed if only to ensure that the grid at the surface is not less refined than just below it. Therefore, the Hessian criterion must be considered everywhere (this explains why the approximative procedures detailed above are necessary). On the other hand, while the LS3 Hessian has a behaviour similar to a free-surface criterion, the real free-surface criterion is used as well because it guarantees that the grid at the surface is absolutely regular and that a safety zone of refined cells is generated around the surface. The criterion in each cell is thus computed from both criteria.

The tensor criteria are combined into one by taking a weighted maximum of the two tensors. We want the threshold T_r to indicate directly the desired cell size at the surface (as for the free-surface criterion), so a weighting factor c is applied only to the Hessian criterion:

$$\mathcal{C}_C = \max(\mathcal{C}_S, c\mathcal{C}_H). \quad (20)$$

The approximate maximum of the two tensors is computed using the procedure described in [27], which is an improvement of the one in [7]. First, the eigenvalues and eigenvectors of the two tensors are computed. Then new eigenvalues for each tensor are set as the maximum of the original eigenvalue and the length of the corresponding eigenvector when it is multiplied by the *other* tensor. This gives two approximations to the maximum tensor; the final tensor in each cell is the element-by-element average of these two.

6 Proportionality and scaling rules

Inherent in the choice of a combined refinement criterion is the need to define the proportionality between the criteria being combined. In our criterion, the proportionality is given by the factor c in equation (20). No universally valid expression for c exists, since the choice depends on the problem and on the output

quantities of interest. However, it is possible to give guidelines for c purely in terms of the non-dimensional characteristics of the flow, since the scaling of the flow does not influence the optimal mesh. This section first derives a scale-independent form for the weighting, then discusses which non-dimensional parameters should be taken into account when constructing guidelines for c . Some practical advice for choosing c will be discussed during the analysis of the test cases in section 7.

6.1 Scale-independent form of c

Let a flow problem be characterised by a geometry-based reference length L , velocity V_∞ , and gravity g , as well as a density ρ_∞ and a kinematic viscosity μ_∞ . In a liquid-gas two-fluid flow the latter two will generally be the properties of the liquid, since this is the dominant fluid in the interaction. With these, we define non-dimensional coordinates and pressure:

$$\bar{\mathbf{x}} = \frac{\mathbf{x}}{L}, \quad c_p = \frac{p - p_\infty}{\frac{1}{2}\rho_\infty V_\infty^2}, \quad (21)$$

as well as the Reynolds and Froude number:

$$Re = \frac{\rho_\infty V_\infty L}{\mu_\infty}, \quad Fr = \frac{V_\infty}{\sqrt{gL}}. \quad (22)$$

Now consider two flows that are identical up to a scaling factor, i.e. they have the same non-dimensional pressure distribution $c_p(\bar{\mathbf{x}})$. It is logical to require that these flows get the same refined mesh in $\bar{\mathbf{x}}$ -coordinates; one mesh should be a scaled version of the other. This also guarantees that identical meshes are produced when a problem is computed in dimensional or in non-dimensional form. For two such scaled flows, the free-surface part of the refinement criterion has identical values, since its eigenvalues are equal to 1 at the surface (see section 5.1). It follows that the pressure-Hessian part of the criterion should also be the same for the two flows, which is obtained when the product $c\mathcal{C}_H$ is non-dimensional. The Hessian matrix $\mathcal{H}$ has the dimension of the second spatial derivative of the pressure and scales therefore with $\frac{1}{2}\rho_\infty V_\infty^2/L^2$. Taking into account the power a from equation (14), a non-dimensional criterion is obtained by setting c as:

$$c = \bar{c} \left(\frac{L^2}{\frac{1}{2}\rho_\infty V_\infty^2} \right)^a. \quad (23)$$

The new proportionality factor $\bar{c}$ is non-dimensional. For constant $\bar{c}$ the entire refinement criterion is identical for the two flows so similar meshes can be obtained, as in the case for the pure free-surface criterion, by setting the threshold T_r proportional to L (see equation (10)). The parameter $\bar{c}$ will be used to formulate guidelines for the proportionality which do not depend on the scaling of the problem.

6.2 Dependence on Fr

The question remains, how $\bar{c}$ should depend on the non-dimensional parameters of the problem. For typical water flows, the influence of Re is weak, since the boundary layers are thin and have little effect on the pressure. Thus, the criterion with scaling (23) is already more or less independent of Re , so there is no reason to modify $\bar{c}$ with this parameter. However, the influence of Fr is significant. To analyse the effect of the Froude number on the pressure Hessian for waves around objects in steady motion, we will study a simple example: 2D linear waves.

6.2.1 Hessian criterion in linear waves

Consider a travelling wave of moderate amplitude in deep water. Supposing irrotational flow gives the following velocity potential [22]:

$$\Phi = A \exp(kz) \cos(k(x - vt)), \quad (24)$$

for the amplitude A , where the wave number $k = \frac{2\pi}{\lambda}$. The wave length λ and the celerity v are linked through the dispersion relation $v^2 = \frac{g\lambda}{2\pi}$ or, in non-dimensional form:

$$\bar{v} = \frac{\bar{\lambda}}{2\pi Fr^2}, \quad (25)$$

where the non-dimensional quantities are $\bar{v} = \frac{v}{V_\infty}$ and $\bar{\lambda} = \frac{\lambda}{L}$. Furthermore, let $\bar{k} = kL$, $\bar{A} = \frac{A}{LV_\infty}$ and $\bar{t} = \frac{tV_\infty}{L}$. With the unsteady form of Bernoulli's equation we find the pressure:

$$c_p = -\frac{2\bar{z}}{Fr^2} + \frac{2\bar{v}\bar{A}}{Fr^2} \exp\left(\frac{\bar{z}}{Fr^2}\right) \sin\left(\frac{\bar{x} - \bar{v}\bar{t}}{Fr^2}\right), \quad (26)$$

The (dimensional) 2D Hessian becomes:

$$H = \frac{\frac{1}{2}\rho_\infty V_\infty^2}{L^2} \frac{2\bar{v}\bar{A}}{Fr^6} \exp\left(\frac{\bar{z}}{Fr^2}\right) \begin{bmatrix} -\sin\left(\frac{\bar{x} - \bar{v}\bar{t}}{Fr^2}\right) & \cos\left(\frac{\bar{x} - \bar{v}\bar{t}}{Fr^2}\right) \\ \cos\left(\frac{\bar{x} - \bar{v}\bar{t}}{Fr^2}\right) & \sin\left(\frac{\bar{x} - \bar{v}\bar{t}}{Fr^2}\right) \end{bmatrix}. \quad (27)$$

Finally, we find for the Hessian contribution to the combined criterion (cf. equations (14), (20) and (23)):

$$c\mathcal{C}_H = \bar{c} \left(\frac{2\bar{v}\bar{A}}{Fr^6}\right)^a \exp\left(\frac{a\bar{z}}{Fr^2}\right) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (28)$$

This result indicates that the refined cell sizes in waves will increase exponentially with the depth and that the cells will be square, even when using directional refinement.

In a steady flow, the waves created by an immersed object have $v = V_\infty$ or $\bar{v} = 1$. For these waves, the pressure at the surface ($z = 0$ in linear theory) becomes:

$$c_p = \frac{2\bar{A}}{Fr^2} \sin\left(\frac{\bar{x} - \bar{t}}{Fr^2}\right). \quad (29)$$

For the criterion at the surface, we get:

$$c\mathcal{C}_H = \bar{c} \left(\frac{2\bar{A}}{Fr^6}\right)^a \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (30)$$

This expression can be used to study the grid refinement obtained near the surface due to the contribution of the Hessian criterion. Two extreme cases exist. For long waves at high Fr , i.e. when $\lambda \approx L$, the extreme values of c_p in the wave will be proportional to the extremes on the body, which are more or less constant with Fr (c_p on a streamlined body varies from 1 to about -3). Thus, from (29) it follows that $\frac{2\bar{A}}{Fr^2}$ is independent of Fr , so (30) shows that $c\mathcal{C}_H \sim \bar{c}Fr^{-4a}$ for long waves. In the typical case where $a = \frac{1}{2}$ we get $c\mathcal{C}_H \sim \bar{c}Fr^{-2}$. Given the dispersion relation (25) which reduces to $\bar{\lambda} \sim Fr^2$ for $\bar{v} = 1$, this implies that the number of cells per wavelength is constant for varying Fr when $\bar{c}$ is constant.

The other extreme situation is the creation of very short waves ($\lambda \ll L$) at low Fr . In this case the pressure field of the body, which creates the wave, can be supposed to vary linearly over the length of the wave. As a result, the extreme values of c_p in the wave will be proportional to the wave length $\bar{\lambda}$. Therefore, we get $\frac{2\bar{A}}{Fr^2} \sim \bar{\lambda} \sim Fr^2$ (from (25)) so $c\mathcal{C}_H \sim \bar{c}Fr^{-2a}$. For $a = \frac{1}{2}$ this results in $c\mathcal{C}_H \sim \bar{c}Fr^{-1}$ so the cell sizes scale with Fr when $\bar{c}$ is constant; as the wave length itself scales with Fr^2 the number of cells per wavelength diminishes when Fr goes to zero.

6.2.2 Variation of $\bar{c}$ with Fr

Given the result of section 6.2.1 it makes sense to choose $\bar{c}$, for a given type of geometry, independent of the Froude number. First of all, this means that flow phenomena near the body far below the surface

(such as trailing vortices) are resolved with the same mesh density for all Fr , which is good because these phenomena do not depend much on the Froude number, especially for low Fr . Also, the accuracy of a computed wave is related to the number of cells per wave length, so waves at various Froude numbers are resolved with the same accuracy if $\bar{c}$ is independent of Fr . Only the waves at very small Fr are resolved less accurately, but even this is natural since the influence of these waves on the flow is small.

However, it may also be sensible to increase $\bar{c}$ slightly with Fr . A consequence of choosing $\bar{c}$ constant is, that low- Fr flows require more cells than high- Fr flows. Also, *better* accuracy may be required for waves at high Fr since these waves have a stronger influence on the flow. For longer waves, the number of cells per wave length is proportional to $\bar{c}$ (see above), so setting $\bar{c} \sim Fr^{1/2}$ for example results in a number of cells per wave length that varies with $Fr^{1/2}$. The test case in section 7.2 specifically addresses the choice of $\bar{c}$ depending on Fr .

7 Test cases

This section presents three test cases which analyse the meshes created by the combined refinement criterion, the choice of the weighting factor c , and the different types of flow that can be computed with the criterion. The cases are the two-dimensional immersed wing of Duncan (section 7.1), the Series 60 ship (section 7.2), and the cavitating INSEAN E779A propeller (section 7.3).

7.1 Immersed NACA0012 wing

The first test case is meant to evaluate the behaviour of the combined refinement criterion for a two-dimensional wave field. The case is the geometry studied by Duncan [5], a wave train generated by an immersed NACA0012 profile of chord 0.203 m at 5 degrees angle of attack, with its centre point at 0.236 m below the surface. The inflow velocity is 0.8 m/s and $Re = 1.42 \cdot 10^5$.

Figure 6: Refined mesh around the immersed profile for the Duncan case with $c = 0.002$. The Hessian is evaluated with SG (a) and LS3 (b).

The figure 6 shows the refined meshes created with the LS3 and the SG Hessian criterion of section 5.2, starting from an original mesh that has some refinement around the wing but none at all around the free surface (figure 7). The figures show that the refinement criterion detects the free surface, the pressure peaks at the leading and trailing edges, and the zone where the pressure disturbance is transferred from the wing to the surface. Little refinement is created in the air zone, since the pressure is nearly constant there; for this case, a fine grid in the air is in fact not really necessary. The SG criterion gives by far the better mesh: it is smoother around the wing and below the surface and the spurious refinement in the wake is absent. The refined wake for the LS3 criterion comes from small oscillating numerical errors

Figure 7: Original mesh for the Duncan case.

in the pressure, as discussed in section 5.2.3; it is irregular and does not increase the precision of the solution. Also, the SG mesh around the surface has a continuous horizontal cell size, while the LS3 mesh is too fine at the surface. Thus, the explicit removal of the peak (step 5b in section 5.2.4) is a more robust procedure than the limiting of the LS3 criterion.

Figures 8a–c show the influence of varying T_r and c for the SG Hessian. Figure 8a gives the first wave crest for a target cell size $T_r = 0.0025m$ and $c = 0.002$. The grid here is four times finer at the surface than just below it. When T_r is increased, this ratio does not change but the entire grid becomes coarser (figure 8b). Reducing c produces a coarser grid only below the surface (figure 8c); the grid at the surface does not change. All these figures show directional refinement at the surface, the grids are coarser in horizontal direction when the surface is horizontal. The pressure Hessian criterion creates only square cells below the waves, as is predicted in section 6.2.1.

Table 1: Number of cells in the refined meshes, Duncan test case with $T_r = 0.0025$.

c	0	0.001	0.002	0.004
$\bar{c}$	0	0.088	0.176	0.352
cells	8151	14392	25848	54890

In figure 8d, the position of the free surface is given for $T_r = 0.0025m$ and four different values of c . Except for the case with pure free-surface refinement ($c = 0$), the agreement with experiments is as good as can be expected in this case (RANS computations tend to underpredict the wave height for Duncan’s case). Moreover, the results on the two finest grids produced with the combined criterion are very similar. Thus, it is not necessary to refine the grid below the surface to twice the size of the grid at the surface, four times coarser cells as obtained with $c = 0.002$ (corresponding to $\bar{c} = 0.176$) are acceptable. As the total number of cells increases strongly with the parameter c (table 1), this information can be used to keep the total number of refined cells low.

7.2 Series 60 wave pattern

The goal of the second test case is to study the dependence of the refined mesh and the flow solution on the weighting factor $\bar{c}$ for a typical ship flow case at different Froude numbers. Thus, we can obtain practical guidelines for the choice of $\bar{c}$ depending on Fr . The test case is the Series 60 hull in straight-ahead motion and calm water. While this hull form is outdated, it has the advantage of producing very

Figure 8: Refined meshes at the first wave crest for the Duncan case with $T_r = 0.0025$, $c = 0.002$ (a), with $T_r = 0.005$, $c = 0.002$ (b), and with $T_r = 0.0025$, $c = 0.001$ (c). Figure (d) compares the wave elevation for $T_r = 0.0025$ and four different c with Duncan's experiment [5].

clear and distinct wave patterns that are ideal for evaluating refined meshes. Detailed experiments for this case are available from IIHR [13] at $Fr = 0.316$ and $Re = 5.3 \cdot 10^6$. Apart from this Froude number, we also compute the flow at $Fr = 0.16$ and $Fr = 0.41$ with $Re = 2.7 \cdot 10^6$ and $Re = 6.9 \cdot 10^6$ respectively.

Like the previous case, the computations are started from a mesh that has no initial refinement at all around the free surface, to show that a sensible refined mesh for free-surface ship flows can be obtained entirely with automatic grid refinement. This original grid has 242k cells. Based on the Duncan test case, the SG Hessian is chosen for the refinement criterion. The threshold for all computations (equivalent to the desired cell size at the surface) is kept constant, $T_r = 0.001L$ which is the usual cell size at the surface for ISIS-CFD computations without grid refinement. The weighting factor is varied: for each Froude number, the flow is computed with $\bar{c} = 0.016$, 0.024 and 0.032 which gives the grid sizes in table 2. Reference results come from computations without grid refinement on a fine mesh of 3.45M cells. The turbulence is modelled with the Menter $k - \omega$ SST model.

Table 2: Number of cells in the refined meshes for the Series 60 test cases.

$\bar{c}$	0.016	0.024	0.032
$Fr = 0.16$	698k	1169k	1793k
$Fr = 0.316$	597k	892k	1360k
$Fr = 0.41$	757k	973k	1348k

To analyse the flow, the wave pattern at the three Froude numbers is shown in figure 9. While the

Figure 9: Wave patterns for the Series 60 at $Fr = 0.16$ (a), $Fr = 0.316$ (b), and $Fr = 0.41$ (c). The isoline distance is the same for all figures and corresponds to $L/1000$. The $Fr = 0.316$ result is compared with experiments from IIHR [13].

Figure 10: Nondimensional stern wave length, bow wave height, and horizontal cell size at the stern wave in the automatically refined meshes ($c = 0.032$) for the Series 60 test case at different Fr .

Figure 11: Series 60, cross-sections of the grid at $Y/L = 0.118$ for $Fr = 0.16$ (a), $Fr = 0.316$ (b), and $Fr = 0.41$ (c).

wave strength varies strongly with Fr , the main diverging waves in all cases are created at the bow and the stern; for the two highest Fr , the bow wave breaks. Transverse waves are visible behind the bow wave for $Fr = 0.16$ and behind the stern for the other Fr . To verify the approximations of wave length and pressure variation used to study the pressure Hessian in waves (section 6.2.1), the nondimensional length of the transverse stern wave $\bar{\lambda} = \lambda/L$ and the height of the bow wave $\bar{h} = h/L$ (first trough to second crest) measured from these plots are given in figure 10. As expected, the wave length varies linearly with Fr^2 , confirming the dispersion relation (25). The difference in c_p over the wave height (Δc_p) is computed by noting that $\Delta c_p \sim g\bar{h} \sim \bar{h}/Fr^2$, so this quantity is plotted. Δc_p was assumed in section 6.2.1 to be constant for high Fr and proportional to $\bar{\lambda} \sim Fr^2$ for lower Fr . The graph shows a behaviour which is neither constant nor linear in Fr^2 , but increases faster for lower than for higher Fr . Thus, the two approximations are reasonable as limit values for very small or large Fr .

Figure 11 shows cross-sections of the mesh for the three Froude numbers at the largest weighting factor $\bar{c} = 0.032$. Around the position of the free surface, which clearly shows the difference in wave amplitude between the cases, the meshes have directional refinement. The cell size below the surface decreases gradually from the bottom up, the finest cells are concentrated in the bow wave (left) and the stern wave (right). As for the Duncan test case, the refined cells below the waves are predominantly

square, although some cells near the surface are smaller in the vertical direction. The size of the cells in the original grid can be seen in the upper right corner of figure 11b, so the entire fine grid is effectively created by automatic refinement.

To study the dependence of the mesh size on Fr , the horizontal cell size in the stern wave just below the surface for the meshes of figure 11 is also given in figure 10. As the refined cells are created by division of the original cells, there is only a limited number of possible cell sizes, so the sizes for the $Fr = 0.316$ and $Fr = 0.41$ case are the same. Therefore, it is impossible to conclude whether the number of cells per wave length stays constant for high Fr . When Fr is reduced, the cell size decreases slower than the wave length so the number of cells per wave decreases. Still, for low Fr the cells are smaller than for high Fr . On one hand this is logical, since figure 9 reveals that the waves at $Fr = 0.16$ are extremely fine so a good grid resolution is needed to resolve them. On the other hand, the total grid sizes are larger for small Fr than for large Fr (see table 2), which may seem counter-intuitive since the waves at low Fr have little influence on the solution.

Figure 12: Series 60, comparisons of the water surface in three X -cuts at $Fr = 0.16$ (a), $Fr = 0.316$ (b), and $Fr = 0.41$ (c). Comparison of the velocities on the line $X/L = 1$, $Z/L = -0.02$ (d) for $Fr = 0.316$. Experiments from IIHR [13].

Finally, the solutions for all cases are compared in figure 12, which shows the evolution of the free surface in three X -cuts when $\bar{c}$ is varied, compared with the non-adapted fine grid. For the highest Froude number (figure 12c), even the solution with $\bar{c} = 0.016$ is in close agreement with the fine-grid solution, only some discrepancy can be noted for the $\bar{c} = 0.024$ solution behind the ship. This is no longer the case at $Fr = 0.316$ (figure 12b), where notable differences exist for $\bar{c} = 0.016$. However, $\bar{c} = 0.024$ gives sufficient accuracy. Finally, at $Fr = 0.16$ (figure 12a) the computation for $\bar{c} = 0.016$ is the closest to the fine-grid solution. Thus, this solution is probably wrong! The waves in this case have a height of less than one grid cell, a situation for which wave heights may be overpredicted on too coarse grids [3]. The solutions for $\bar{c} = 0.024$ and $\bar{c} = 0.032$ are reasonably close, so again $\bar{c} = 0.024$ is satisfactory. The figure 12d shows velocity profiles on a horizontal line below the water at the stern. Only $Fr = 0.316$ is shown here, since the results for all Froude numbers are similar. All $\bar{c}$ give results that are close to the fine-grid solution, the discrepancy with the experimental results may be due to the isotropic $k - \omega$ SST turbulence model which is not always well adapted to the simulation of ship wake flows [4, 6].

In conclusion, the analysis of the meshes for the Series 60 test case confirms the findings of section 6.2 on the behaviour of the pressure Hessian in waves. To accurately model the wave pattern, higher values for $\bar{c}$ are needed at low Fr than at high Fr . However, considering the weak influence of the waves on the rest of the flow field at low Fr , a constant choice for $\bar{c}$ is justified, which gives the added advantage that the grid on the hull below the surface is refined in the same way for all Fr . For slender hulls like the Series 60, setting $T_r = 0.001L$ to obtain the standard ISIS-CFD cell size at the surface with a value of $\bar{c}$ around 0.024 gives sufficient accuracy. Higher values of $\bar{c}$ should be used with care, since they increase significantly the total number of cells (table 2).

7.3 Cavitating INSEAN E779A propeller

The two-fluid mixture model framework can also be used for the modelling of cavitation. For these computations, water and its vapour are considered as two separate fluids so evaporation and condensation are transformations from one of the fluids into the other. These transformations are produced by a source term in equation (3). Different approximative cavitation models exist in the literature, we use the model of Merkle [16] which expresses the production or destruction of vapour as a function of the difference between the pressure and the vapour pressure.

To demonstrate that combined grid refinement may also be useful for this particular kind of flow, an initial test case is presented: the INSEAN E779A propeller is simulated in open-water conditions, for a regime with steady sheet cavitation on the propeller blades and cavitating cores in the trailing vortices. The propeller with diameter $D = 0.227\text{m}$ runs at $n = 36\text{rps}$ with an advance coefficient $J = \frac{V_\infty}{nD} = 0.71$ and a cavitation number $\sigma = \frac{p_\infty - p_{\text{vap}}}{1/2\rho_\infty V_\infty^2} = 3.5$. Experiments for this case have been performed by CNR-INSEAN [17, 21]. The flow requires a very fine resolution of the vortices, in order to capture the low pressure peaks in the cores. The advantage of the combined refinement criterion here is that the Hessian part detects the entire structure of the vortices, while the free-surface criterion ensures that the cavitating core has a regular, fine grid.

The computation is started on a coarse grid that has 770k cells, with some refinement around the blades but very little in the wake zone. The free-surface refinement for the combined criterion is based on the volume fraction of the vapour. A single computation is performed with a threshold $T_r = 0.0012$ and a scaling factor $c = 0.000012$. The refined mesh has 1.23M cells. As a reference, the computation is also performed on the coarse grid without adaptive refinement. The LS3 Hessian is used, rather than the smoother and more robust SG Hessian, since it gives a sharper definition of the vortex core. In fact, the pressure gradient in the core varies rapidly in space, so any smoothing of the gradient affects the Hessian, lessening the amount of grid refinement in the core.

The computed cavitation zones are given in figure 13. Compared with the cavitation observed experimentally (figure 14), the outline of the vapour bubble computed on the refined grid comes close. The pocket on the blade has the right shape, both the tip and propeller hub vortices are clearly seen and the characteristic rollup of the tip vortices is reproduced correctly in the simulation. This is different for the original-grid solution. While the sheet cavitation on the blades is reasonable and similar to the different

Figure 13: Computed surface of the cavitating pockets (isovalue $\alpha = 0.5$) for the INSEAN E779A propeller, on the original grid (a) and the adaptively refined grid with $T_r = 0.0012$ (b). (The small spots on the surface are plotting artefacts.)

numerical results presented by [21], the cavitation in the trailing vortices is not simulated at all.

The main difference between the adapted-grid simulation and the experiments is, that the cavitating tip vortices do not persist over a very long distance. We have also observed this effect for non-cavitating propellers, it is due to the damping out of the pressure peak in the vortex cores caused mainly by the RANS turbulence models, which overpredict the turbulent viscosity in the core. Numerical diffusion from the spatial and temporal discretisation also contributes. However, the vortices do not disappear when the cavitating cores collapse. A cross-section of the mesh (figure 15) shows grid refinement in the cores of the first and second vortices downstream of the blade tips, which means that a low pressure peak persists over this distance. The influence of these vortices on the flow near the propeller is therefore taken into account for the simulation of the cavitation near the tips. As these non-cavitating vortex cores cannot be detected by a free-surface refinement criterion alone, this clearly demonstrates the value of a combined criterion.

Figure 14: Experimental view of the cavitation on the E779A propeller, from CNR-INSEAN [17].

Figure 15: Refined grid in the $Y = 0$ plane for the INSEAN E779A propeller.

This cavitation study is preliminary, since the influence of different models, model settings, and indeed the interaction between the cavitation model and the grid refinement were not tested. Also, sensible guidelines for choosing c and the threshold T_r are not yet available. However, these first results indicate that continuing such research is worthwhile.

8 Conclusion

Grid refinement criteria are studied for automatic mesh adaptation in free-surface flow simulation. It is shown that these criteria must refine both around the surface, to resolve the convection equation for the volume fraction, and in the region below it in order to capture the orbital flow fields. While derivative-based criteria suitable for detecting the velocity and pressure fields can in principle locate the free surface, a dedicated robust free-surface capturing criterion is preferred in order to ensure a regular mesh at the surface and to prevent errors in the volume fraction coming from cell misalignment. Thus, the refinement criterion must be a combination of a free-surface criterion with a derivative-based sensor.

We choose a criterion which combines directional refinement in the region where the volume fraction is between 0.1 and 0.9 with refinement based on the pressure Hessian. The difficulty in obtaining an accurate and smooth Hessian criterion lies in the numerical evaluation of second spatial derivatives on unstructured meshes. An evaluation using least-squares interpolation with third-order polynomials (LS3) is presented, which is insensitive to the mesh geometry. However, this formulation still produces irregular criteria since the pressure solution itself is non-smooth. This leads us to a second type of computation (SG) where Gauss derivation is used twice to compute first the gradient of the pressure and then its Hessian. Both the gradient and the Hessian are smoothed after their computation, to remove irregularities. Numerical tests show that SG gives much better mesh quality in waves than LS3, so it is the preferred choice for combined refinement. However, the smoothing may damp out small-scale structures such as vortex cores, so the LS3 Hessian is better adapted for these flows; improving the vortex-capturing performance of the SG Hessian is a subject for further study.

The numerical pressure Hessian has a peak at the free surface, due to the discontinuity in the pressure gradient. For LS3, this peak is conserved since it is associated with an eigenvector normal to the surface; the Hessian therefore acts like the free-surface criterion at the surface. The peak eigenvalue is simply limited to a given maximum. This is not possible for SG, since the smoothing diffuses the peak. Therefore, the free-surface region is not smoothed and the computed Hessian outside the surface zone is extrapolated

through this zone. This technique performs best, since it is shown to produce no unnecessary refinement parallel to the surface.

When combining two refinement criteria, the relative weight of each criterion must be chosen. Since the free-surface criterion is non-dimensional, it is useful to introduce a non-dimensional scale-independent form for the Hessian criterion as well, with a weight $\bar{c}$. This weight can be set independently of the Reynolds number for the flow, since the pressure for typical hydrodynamic flows varies little with Re . Analysing the dependence of the Hessian on the Froude number, we find that in gravity waves for a constant $\bar{c}$, the number of cells per wave length remains constant for high Fr and diminishes for low Fr . A numerical test on the Series 60 hull shows that, to get good accuracy for the free surface, a higher value for $\bar{c}$ is needed at low Fr than at high Fr . However, due to the limited influence of the free surface on the flow at low Fr , it is acceptable to choose $\bar{c}$ independently of Fr . Sensible values for slender ships are $\bar{c} = 0.02 - 0.025$ with a refinement threshold set around $T_r = L/1000$.

The ultimate goal of combined refinement criteria for free-surface flows is to generate fine meshes of good quality entirely with grid refinement, starting from uniformly coarse original meshes. Thus, a user does not need to take into account the free surface when generating this original mesh, which greatly simplifies the setup of a computation and even forms an important step towards the automatization of simulation. For three different free-surface test cases, including a case with cavitating flow, the current paper has demonstrated such mesh generation. Creating complete free-surface meshes with grid refinement therefore seems entirely possible.

Acknowledgements

The authors thank Master students A. Jebali, for his contribution to the validation of the combined criterion, and F. Lemoine, for his work on the application to cavitation. Computations were performed using HPC resources from GENCI (Grant2010-x2010021308), which is gratefully acknowledged.

References

- [1] F. Alauzet and A. Loseille. High-order sonic boom modeling based on adaptive methods. *J Comput Phys*, 229(3):561–593, 2010.
- [2] D. Aubry, P. Díez, B. Tie, and N. Parés, editors. *Adaptive Modelling and Simulation 2011*. Paris, France, 2011.
- [3] G. B. Deng, E. Guilmineau, A. Leroyer, P. Queutey, M. Visonneau, and J. Wackers. Verification and validation for unsteady computation. In L. Larsson, F. Stern, and M. Visonneau, editors, *A Workshop on Numerical Ship Hydrodynamics*, Gothenburg, Sweden, 2010. Chalmers University of Technology.
- [4] G. B. Deng and M. Visonneau. Comparison of explicit algebraic stress models and second-order turbulence closures for steady flows around ships. In *Proceedings of the 7th Int. Conf. on Numerical Ship Hydrodynamics*, Nantes, France, 1999.
- [5] J. H. Duncan. The breaking and non-breaking wave resistance of a two-dimensional hydrofoil. *J Fluid Mech*, 126:509–520, 1983.
- [6] R. Duvigneau and M. Visonneau. On the role played by turbulence closures in hull shape optimization at model and full scale. *J Mar Sci Tech*, 8:11–25, 2003.
- [7] P. L. George and H. Borouchaki. *Delaunay Triangulation and Meshing - Application to Finite Elements*. Hermes, 1998.
- [8] A. Hay, A. Leroyer, and M. Visonneau. H-adaptive Navier-Stokes simulations of free-surface flows around moving bodies. *J Mar Sci Tech*, 11:1–18, 2006.

- [9] A. Hay, P. Queutey, and M. Visonneau. Computation of three-dimensional free-surface flows with an automatic adaptive mesh refinement and coarsening strategy. In *Proceedings of the 25th Symposium on Naval Hydrodynamics*, St. John's, Newfoundland and Labrador, 2004.
- [10] J. Karlsson, A. Feymark, and C. Eskilsson. Implementation of anisotropic mesh refinement in OpenFOAM. In *Proceedings of the 15th Numerical Towing Tank Symposium (NuTTS 2009)*, Cortona, Italy, 2012.
- [11] N. Kroll. Achievements of the European project ADIGMA on adaptive high-order methods for aerospace applications. In *Proceedings of ECCOMAS CFD 2010*, Lisbon, Portugal, June 2010.
- [12] B. P. Leonard. The ULTIMATE conservative difference scheme applied to unsteady one-dimensional advection. *Comp Meth Appl Mech Engrg*, 88:17–74, 1991.
- [13] J. Longo and F. Stern. Effects of drift angle on model ship flow. *Exp Fluids*, 32:558–569, 2002.
- [14] A. Loseille, A. Dervieux, and F. Alauzet. Fully anisotropic goal-oriented mesh adaptation for 3D steady Euler equations. *J Comput Phys*, 229:2866–2897, 2010.
- [15] J. Majewski. Anisotropic adaptation for flow simulations in complex geometries. In *36th CFD/ADIGMA Course on hp-adaptive and hp-multigrid methods*, Lecture Series 2010-01. Von Karman Institute, 2009.
- [16] C. L. Merkle, J. Feng, and P. E. O. Buelow. Computational modeling of the dynamics of sheet cavitation. In *Proceedings of the 3rd International Symposium on Cavitation*, Grenoble, France, 1998.
- [17] F. Pereira, F. Salvatore, and F. Di Felice. Measurement and modelling of propeller cavitation in uniform inflow. *J Fluids Eng*, 126:671–679, 2004.
- [18] V. Pržulj and B. Basara. Bounded convection schemes for unstructured grids. In *15th AIAA Computational fluid dynamics conference*, AIAA paper 2001-2593, Anaheim, CA, 11-14 June 2001.
- [19] P. Queutey and M. Visonneau. An interface capturing method for free-surface hydrodynamic flows. *Comput Fluids*, 36(9):1481–1510, November 2007.
- [20] C. M. Rhie and W. L. Chow. A numerical study of the turbulent flow past an isolated airfoil with trailing edge separation. *AIAA J*, 17:1525–1532, 1983.
- [21] F. Salvatore, H. Streckwall, and T. van Terwisga. Propeller cavitation modelling by CFD - Results from the VIRTUE 2008 Rome workshop. In *Proceedings of the First Int. Symposium on Marine Propulsion, SMP 2009*, Trondheim, Norway, June 2009.
- [22] J. J. Stoker. *Water Waves – The Mathematical Theory with Applications*. Interscience, New York, USA, 1957.
- [23] O. Ubbink. *Numerical predictions of two fluid systems with sharp interfaces*. PhD thesis, Imperial College of Science, Technology & Medicine. University of London., January 1997.
- [24] J. Wackers, K. Ait Said, G. B. Deng, P. Queutey, M. Visonneau, and I. Mizine. Adaptive grid refinement applied to RANS ship flow computation. In *Proceedings of the 28th Symposium on Naval Hydrodynamics*, Pasadena, California, 2010.
- [25] J. Wackers, G. B. Deng, P. Queutey, M. Visonneau, E. Guilmineau, and A. Leroyer. Sliding grids and adaptive grid refinement applied to ship hydrodynamics. In *Proceedings of the 29th Symposium on Naval Hydrodynamics*, Gothenburg, Sweden, 2012.

- [26] J. Wackers, B. Koren, H. C. Raven, A. van der Ploeg, A.R. Starke, G. B. Deng, P. Queutey, M. Visonneau, T. Hino, and K. Ohashi. Free-surface viscous flow solution methods for ship hydrodynamics. *Arch Comput Meth Eng*, 18:1–41, 2011.
- [27] J. Wackers, A. Leroyer, G. B. Deng, P. Queutey, and M. Visonneau. Adaptive grid refinement for hydrodynamic flows. *Comput Fluids*, 55:85–100, 2012.
- [28] J. Wackers, P. Queutey, and M. Visonneau. Adaptive grid refinement for unsteady ship flow and ship motion simulation. In *12th Numerical Towing Tank Symposium (NuTTS 2009)*, Cortona, Italy, 2009.
- [29] J. Windt and C. Klaij. Adaptive mesh refinement in MARIN’s viscous flow solver ReFRESCO: implementation and application to steady flow. In *Proceedings of MARINE 2011*, Lisbon, Portugal, 2012.