Can adaptive grid refinement produce grid-independent solutions for incompressible flows?

Jeroen Wackers, Ganbo Deng, Emmanuel Guilmineau, Alban Leroyer, Patrick Queutey, Michel Visonneau, Alexandro Palmieri, Alfredo Liverani

Can adaptive grid refinement produce grid-independent solutions for incompressible flows?

Jeroen Wackers[a], Ganbo Deng[a], Emmanuel Guilmineau[a], Alban Leroyer[a], Patrick Queutey[a], Michel Visonneau[a], Alexandro Palmieri[b], Alfredo Liverani[b]

[a] *LHEEA Lab, Ecole Centrale de Nantes, CNRS-UMR 6598, 44321 Nantes Cedex 3, France*
[b] *CIRI-MAM Tecnopolo della Nautica / UNIBO, Via Sant'Alberto 163, Ravenna, Italy*

Abstract

This paper studies if adaptive grid refinement combined with finite-volume simulation of the incompressible RANS equations can be used to obtain grid-independent solutions of realistic flow problems. It is shown that grid adaptation based on metric tensors can generate series of meshes for grid convergence studies in a straightforward way. For a two-dimensional airfoil and the flow around a tanker ship, the grid convergence of the observed forces is sufficiently smooth for numerical uncertainty estimation. Grid refinement captures the details of the local flow in the wake, which is shown to be grid converged on reasonably-sized meshes. Thus, grid convergence studies using automatic refinement are suitable for high-Reynolds incompressible flows.

Keywords: grid adaptation, grid convergence, uncertainty estimation, hydrodynamic flows

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

1. Introduction

One of the trends in computational fluid dynamics today is the high-fidelity simulation of more and more complex flows. For example, simulations for realistic geometries such as ships with their propellers and appendages have become possible. Physical flow features such as flow separation, vortex shedding and vortex breakup are simulated in unprecedented detail. Finally, multiphysics computations like fluid-structure interaction or the modelling of cavitation become common.

The results of such simulations depend on the physical models being used, such as the turbulence model in the Reynolds-averaged Navier-Stokes (RANS) equations. Often, such models are applied in situations which are far more complex than the ones for which they were developed and which may be outside their range of validity. Research of physical modelling, specifically for today's realistic simulations, is therefore of prime importance.

To accurately assess the precision of a physical model, we need to know a numerical solution in which the numerical errors are small with respect to the modelling errors: a solution that is close to grid convergence. In simple cases, it is possible for an experienced user to generate meshes which provide sufficiently small numerical errors. However, for more and more complex flows, the grid resolution needed to resolve the flow phenomena, as well as the precise position of these features, is uncertain. Therefore, it is impossible to know beforehand what mesh size is needed where, in order to obtain grid convergence.

The key to the study of physical flow modelling in complex cases may be the adaptive refinement of the grid. Mesh adaptation is the technique

of modifying locally the mesh density during a computation, in order to optimise the mesh for the flow being simulated. Such techniques are well established, see for example [1]. For the modelling of complex flows, adaptation techniques can detect where relevant flow features occur, so fine cells can be placed only around these features. This allows the creation of locally very fine meshes if these are needed to get grid convergence for flow details, while the total number of cells is kept reasonable.

In this paper, we see if today's grid refinement techniques can assure that the solution of an incompressible flow problem is grid-independent. This means not only obtaining a solution, but also proving that its numerical errors are small enough. For the paper, steady flows are considered and iterative convergence errors are minimised by setting strict tolerances for convergence. Thus, we concentrate on errors due to the spatial discretisation, i.e. interpolation errors and approximate integration.

In the absence of single-grid error estimation procedures like [2], estimating spatial discretisation errors requires simulations on several coarse to fine grids. Therefore, we wish to know if grid refinement studies can be performed with adaptive refinement. Specifically, we investigate the following: (a) Can grid adaptation create geometrically similar grids, where the local ratio of cell sizes between two grids is constant and the cell shapes and orientations are the same? (b) Are the results suitable for use with established uncertainty estimation methods? (c) Can the global and local flow be computed so accurately that the numerical errors become much smaller than the physical modelling errors?

Simulations and grid refinement are performed with the ISIS-CFD unstructured finite-volume incompressible RANS solver developed by the authors from ECN-CNRS. This code is used for the realistic simulation of industrial-type flows, it is available as the flow solver of the FINE^{TM}/Marine computing suite. Its adaptation method is based on the refinement of unstructured hexahedral grids by subdivision of the cells [3, 4]. Anisotropic refinement, where the cells are adapted to the flow both in size and in aspect ratio, is handled with the metric tensor approach introduced by George and Borouchaki [5], modified for the refinement of hexahedrals.

The tests in this paper are therefore based on the type of computations usually performed with ISIS-CFD: incompressible RANS flows, unstructured hexahedral grids, refinement by subdivision, and refinement criteria based on second spatial derivatives. However, the proposed method of uncertainty evaluation is not limited to this type of simulations. To complete the tests, we evaluate if the conclusions reached are also valid for other grid adaptation methods.

The flow solver is presented briefly in section 2, the refinement method and its use of metric tensors as refinement criteria are discussed in section 3. Section 4 investigates how such a refinement method can be used to produce geometrically similar series of meshes for convergence studies. A two-dimensional airfoil is simulated in section 5, to test whether these series can produce asymptotic convergence of the computed forces and useful estimations of the numerical uncertainty. Furthermore, we search the grid density needed to obtain grid independence for the solution in the wake. Finally, the flow around the KVLCC2 tanker is computed to see how close one

can get to grid convergence for a realistic ship flow test case (section 6). The conclusion (section 7) discusses the generality of the results and evaluates the limitations and perspectives of the proposed grid convergence method.

2. The ISIS-CFD flow solver

ISIS-CFD is an incompressible unsteady multifluid Navier-Stokes solver [6, 7, 8]. The solver is based on the finite volume method to build the spatial discretisation of the transport equations. Pressure-velocity coupling is obtained through a Rhie & Chow SIMPLE-type method [9]. The discretisation is face-based, so cells with an arbitrary number of arbitrarily-shaped faces are accepted; most computations are performed on unstructured hexahedral meshes. The code is fully parallel using the MPI (Message Passing Interface) protocol.

Turbulence is principally modelled with the Reynolds-averaged Navier-Stokes (RANS) equations and advanced turbulence closures, such as the anisotropic EASM model [6]. Furthermore, several Detached-Eddy Simulation (DES) models are available. If a free surface is present, it is captured with a mixture-model approach [7]. Finally, techniques such as mesh deformation allow the 6 DOF resolution of body motion, rotation of bodies is achieved through sliding interfaces, and coupling with other fluid or structure solvers is possible. For brevity, these options are not further described here.

3. Grid refinement method

An adaptive grid refinement technique is included in the solver ISIS-CFD [3, 4]. The method performs isotropic and anisotropic refinement of unstruc-

tured hexahedral meshes. Adapted meshes are created by dividing the cells of a coarse original grid into finer cells; this division can be repeated several times until the desired cell sizes are obtained. Anisotropy is introduced by splitting cells in one direction only. The grid refinement is performed in parallel and includes an automatic dynamic load balancing in order to redistribute the refined grid over the processors when some partitions have been refined more than the others. The mesh is adapted regularly during the computation in order to follow the convergence of the flow; earlier refinements can be undone to account for the changes in the flow as it converges. Recent developments of the technique include the treatment of this derefinement in an anisotropic way [10] and the flux-component Hessian criterion described below. The remainder of this section gives an overview of the refinement criteria which are used, followed by a description of two measures which ensure the quality of the refined grid.

3.1. Metric-based refinement

To obtain anisotropic grid refinement, we use metric tensors as refinement criteria. This technique was introduced for the generation of anisotropic tetrahedral cells [5], it has later been used successfully for the adaptive refinement of such meshes [11, 12]. The technique also provides a practical and flexible framework for the refinement of hexahedral meshes.

In our procedure, the refinement of the cells is decided as follows. First, the 3×3 criterion tensors $\mathcal{C}_i$ in each cell i are computed (in some way) from the flow solution. In a hexahedral cell, let the cell sizes $\mathbf{d}_{i,j}$ $(j = 1, 2, 3)$ be the vectors between the opposing face centres in the three cell directions. The goal of the grid refinement is then to create a grid which is uniform

under the transformation $\mathcal{C}$, which implies that:

$$\|\mathcal{C}_i \mathbf{d}_{i,j}\| = T_r \qquad \forall i, j, \tag{1}$$

where T_r is a constant. In the refinement procedure, this is obtained in the following way. Each time the procedure is called, the criterion $\mathcal{C}_i$ and the cell sizes $\mathbf{d}_{i,j}$ are computed on the current grid. A cell i is refined in the direction j when $\|\mathcal{C}_i \mathbf{d}_{i,j}\|$ exceeds the constant T_r, while a previously refined group of cells can be 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. Equation(1) implies that 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 $\mathcal{C}$.

3.2. Refinement criteria

The refinement criteria are based on the Hessian matrix of second spatial derivatives. Hessian matrices can be interpreted as rough error indicators, since they are linked to interpolation errors for linear interpolation [11]. Thus, they provide some measure of the truncation error for a second-order finite-volume discretisation.

Initially, we based the criterion on the Hessian of the pressure [4], using the rationale that the pressure is important for the computation of forces on bodies and also a good indicator of waves and vortices. However, the pressure criterion is unable to track wakes. To adapt the mesh to pressure-based flows but also to boundary layers, wakes, and shear layers, a new criterion is introduced based on the Hessians of both the pressure and the

velocity. To give equal importance to the different Hessian matrices, they are weighted in the way in which they appear in the flux.

The Hessian criterion based on the pressure p (PH criterion) is computed as:

$$\mathbf{C}_{PH} = \left(\|\mathcal{H}(p)\|\right)^{\alpha}, \tag{2}$$

where $\mathcal{H}$ is the Hessian operator and 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. For second-order reconstructions, the interpolation errors are proportional to the Hessian, times the cell size squared. Thus, to get an approximate equidistribution of the interpolation errors, we choose $\alpha = \frac{1}{2}$ [3].

The new flux-component Hessian (FCH) criterion is computed from Hessians of the pressure and velocity components. To derive a general refinement criterion, we want to get as close as possible to an indication of the truncation error. For finite-volume discretisations, this error comes mainly from interpolation errors in the fluxes, so it contains all the different state variables (pressure, velocity, turbulent viscosity). However, from a computational point of view it is desirable to have a simple refinement criterion. Therefore, as a first step, only the convective and pressure part of the flux is used and the turbulence modelling is ignored. Furthermore, we do not take into account all the different products of velocity terms which appear in the fluxes but assign a common weight ρV to all the velocity Hessians, where ρ

is the density and $V = \sqrt{u^2 + v^2 + w^2}$. Thus, the criterion is chosen as:

$$\mathbf{C}_{FCH} = \Big(\max\big(\, \|\mathcal{H}(p)\|,\, \rho V\|\mathcal{H}(u)\|,\, \rho V\|\mathcal{H}(v)\|,\, \rho V\|\mathcal{H}(w)\|\,\big)\Big)^{\alpha}. \quad (3)$$

The maximum of two tensors is computed using the approximative procedure defined by [3]. Further testing has to determine if this criterion is adequate, or if diffusive and turbulence terms must also be added.

3.3. Boundary layer protection and minimum cell size

Two measures are available to improve the regularity of the adaptively refined meshes. The first (figure 1) is a protection of the boundary layer grid, which on the original grid consists of several layers of wall-aligned cells. For these layers, the refinement method contains an option to forbid any refinement in the wall normal direction (which would locally increase the number of layers). Furthermore, in all cases the refinement parallel to the wall is made the same in each 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.

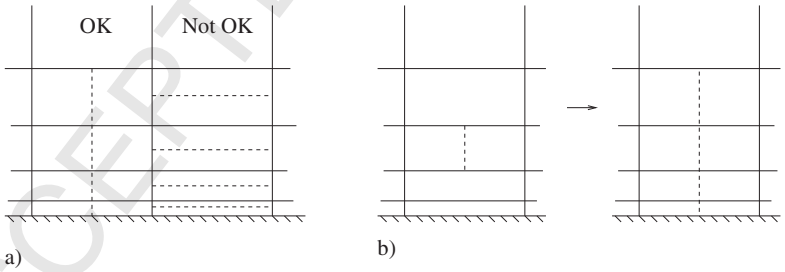

Figure 1: Boundary layer protection: preventing normal refinement (a) and copying parallel refinement in a column (b).

A second measure is to impose a minimum cell size: cells smaller than this size are no longer refined. This option prevents spurious refinement if locally large errors appear 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.

4. Using grid refinement for convergence studies

Metric-based grid adaptation methods like the one described in section 3 provide a simple and elegant way to create series of meshes for grid convergence studies. This section describes such grids and the way they can be created adaptively, plus it highlights certain theoretical limitations of the method proposed. The discussion is completed in section 5, which tests the effectiveness of the approach for obtaining grid-converged solutions in practice.

4.1. Geometrically similar grids

To allow extrapolation towards the grid-independent solution, grid convergence studies require a series of computations on different meshes for which the local truncation errors have the same spatial distribution and vary proportionally to a single global mesh size parameter (see for example [13, 14]). Since truncation errors in general depend on the shapes and orientations of the cells as well as on their size, this puts two requirements on the meshes:

1. The ratio of the cell sizes between two meshes must be constant throughout the mesh,

2. In any given position, the meshes need to have cells of similar shape and orientation.

For a user, creating such meshes is much easier with structured than with unstructured grid generators. In the former, the cell orientation is given by the shape of the grid block(s), the cell size can be varied freely by changing the number of cells in each direction. For the latter, the cell orientations are often not explicitly imposed. Also, it may require great care to obtain fine cells in exactly the same zones for all meshes, since the mesh generation algorithm does not ensure this naturally. In practice, this limits the use of refinement studies with unstructured grids.

4.2. Series of meshes with grid adaptation

Metric-based anisotropic grid refinement methods such as the one outlined in section 3 can create series of geometrically similar unstructured meshes naturally. Thanks to the use of metric tensors, the threshold T_r globally specifies the fineness of the grid: if the refinement criterion remains constant as the mesh is refined, the mesh size everywhere is proportional to T_r (equation 1). This is the case for any refinement criterion, local or integral-based, which is computed from the flow without explicitly taking into account the mesh. The Hessian criteria are an example. Furthermore, if the refined grids are created by cell division, the refined cells conserve the shape and orientation of the cells in the original grid, so two meshes refined from the same original grid have similar cell shapes. Thus, the two requirements of section 4.1 are satisfied. This means that series of geometrically similar meshes can be created by starting from the same original grid and simply varying T_r.

A limitation of this approach is that the refinement procedure cuts cells in half, so the cell sizes in the original grid can only be reduced by a power of two. Thus, if T_r is reduced by anything other than a power of two between subsequent meshes, the cell sizes are divided by the nearest power of two and the zones with cells of the same size become larger or smaller (see figure 3 for an example). While such meshes still resemble each other closely, they are not formally geometrically equivalent.

Furthermore, the mesh size is not proportional to T_r in those regions where the original grid is not refined at all. To obtain grids that are similar in these zones, one could make a different original grid for each threshold. However, if such grids are unstructured, they are rarely exactly similar close to the walls, where the grid is deformed to follow the walls. Thus, to improve the similarity in the unrefined parts of the mesh which are by definition the least interesting for the flow, one would reduce the similarity in the most critical parts of the mesh, the near-wall regions. Therefore, we base our convergence studies on the same original mesh for all grids.

4.3. Effects of protective measures

Protective measures also break the proportionality of the mesh size to T_r. Suppressing boundary layer refinement in the wall-normal direction results in the same cell sizes normal to the wall for all meshes. This is often the desired behaviour. For example in refinement studies with law-of-the-wall boundary conditions it is common practice to keep the boundary grids the same between meshes, since the gradient in the first layer of cells on the boundary (and therefore the solution) depends on the thickness of this layer. However, one must be aware of introducing a component to the numerical

error which does not disappear on grid refinement.

Imposing a minimum cell size has the same effect. Once the finest parts of the mesh have reached this size, they are no longer refined when T_r is reduced, so the truncation errors remain constant in these parts. A solution for this problem is to vary the minimum cell size on the different meshes, proportional to T_r. However, the minimum cell size only exists to counteract occasional large errors in the refinement criterion, it should have the least possible effect on the meshes. Therefore, we keep the minimum mesh size constant and we choose it as small as possible.

5. Two-dimensional case: Nakayama B airfoil

The objective of this section is to test the proposed method for grid-independence studies on a two-dimensional test case, for which very fine meshes can be produced if needed. Three different series of meshes with low-Reynolds and wall-law boundary conditions are created (section 5.1). The adaptively refined meshes are analysed in section 5.2, the convergence of the global forces and the related uncertainty estimation is studied in section 5.3. Section 5.4 considers the grid-independence of two local flow features, the pressure on the foil and the wake. The final section 5.5 gives estimations of the equivalent three-dimensional grid sizes for the grids obtained here.

5.1. Test case and computations

The test case is the Nakayama B airfoil [15]. This case is of particular interest since detailed measurements of the velocity and the turbulence intensity have been performed both in the boundary layers and in the wake, up to 2 chord lengths behind the trailing edge. The airfoil is a supercritical

profile with chord $c = 61$ cm, placed at $\alpha = 4^o$ in the centre of a 137 cm high test section. The flow is incompressible and $Re = 1.2 \cdot 10^6$. Simulations are performed with the $k - \omega$ SST turbulence model and with the AVLSMART scheme [16] for all convective terms.

Three series of computations have been performed. The first two series use a low-Reynolds (no-slip) boundary condition on the airfoil and serve to compare the pressure (LR PH) and flux-component (LR FCH) refinement criterion. The original mesh for these computations has 5755 cells and $y^+ \approx 1$ on the walls. A minimum cell size is imposed but this limit is kept low ($0.0002c$). The last series uses the FCH criterion with a wall-law boundary condition, for which the original grid has 2623 cells and $y^+ \approx 30$. For this series (WL FCH), the minimum cell size is $0.001c$, the size of the first layer of boundary cells. As seen in section 4.3, these cells cannot be refined.

All series consist of ten meshes. For each mesh, the refinement threshold T_r is divided by $\sqrt{2}$ with respect to the previous mesh (thus, the cell sizes are halved every two meshes). The T_r go from 0.5000 to 0.0221.

5.2. Refined meshes

This section studies the series of adapted meshes to see if they have the characteristics described in sections 3 and 4. The difference between the three series is illustrated by figure 2. For the pressure criterion (figure 2a), the refinement is concentrated around the leading edge where the variations in the pressure are the largest. The same refinement is created by the flux-component criterion (figure 2b), since this criterion includes the pressure Hessian. However, the boundary layer and the wake are also refined, which is not the case for the pressure criterion. Contrary to the accepted practice

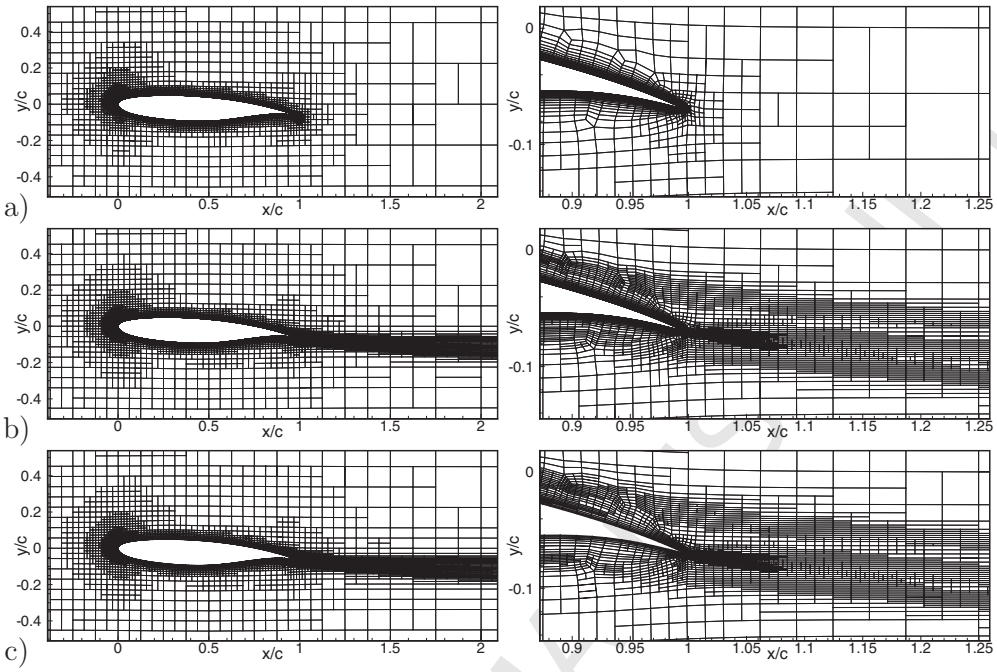

Figure 2: Refined meshes for the Nakayama airfoil with $T_r = 0.125$. Low-Reynolds conditions with the PH criterion (a) and FCH criterion (b), wall law with FCH (c). The right figures show a close-up around the trailing edge.

for boundary layer grids, fine cells are concentrated near the outside of the boundary layer. Finally, the wall-law FCH grids (figure 2c) are generated with the same criterion, the difference comes from the larger minimum cell size and the coarser near-wall grid in the original mesh. The figures show that even in the refined meshes, the difference with the LR FCH meshes remains limited to the near-wall region.

Are the meshes in one series geometrically similar? Figure 3 shows a detail of three successive meshes in the LR PH series. The threshold for the third mesh is twice smaller than for the first one and these two meshes indeed

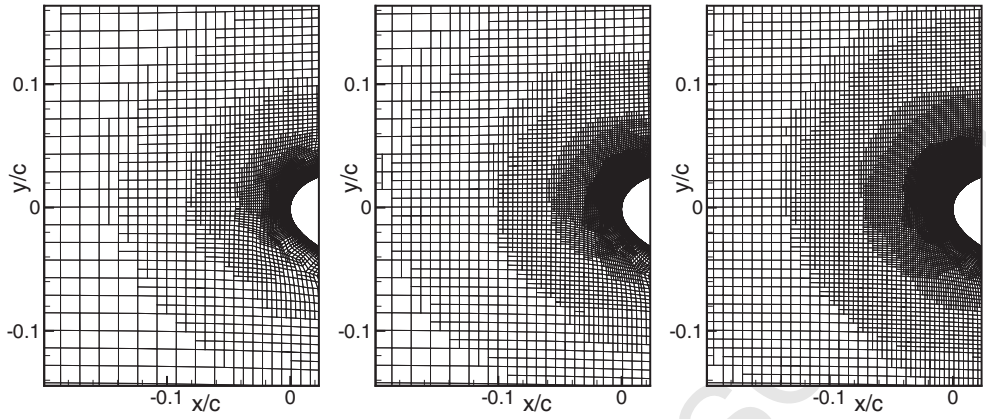

Figure 3: A detail of the LR PH meshes, for progressively finer thresholds. Left to right: $T_r = 0.0625$, $T_r = 0.0442$, and $T_r = 0.0313$.

have refinement in the same places (see section 4.2). For example, regions of cells with the same size start at $x/c = -0.14$ and $x/c = -0.07$ in both grids. The anisotropic refinement is also the same, as shown by the similar regions of vertically oriented cells above and below the regions of fine cells. Thus, the first and third grids are geometrically similar.

For the middle grid, the threshold has a ratio of $\sqrt{2}$ with respect to the other ones. Thus, the regions of fine cells are not the same, they start at $x/c = -0.1$ and $x/c = -0.05$ for example. Furthermore, the cells have the same size in the second mesh that they have in either the first or in the third mesh, so when going from a 'coarser' to a 'finer' mesh, in many places the cell sizes do not change! Locally, the second mesh is therefore not geometrically similar to the first and third ones. However, globally the three meshes resemble each other. The cell orientations are the same, the shape of the fine-cell regions is the same even down to the distribution of anisotropic refinement. The three meshes still form a series.

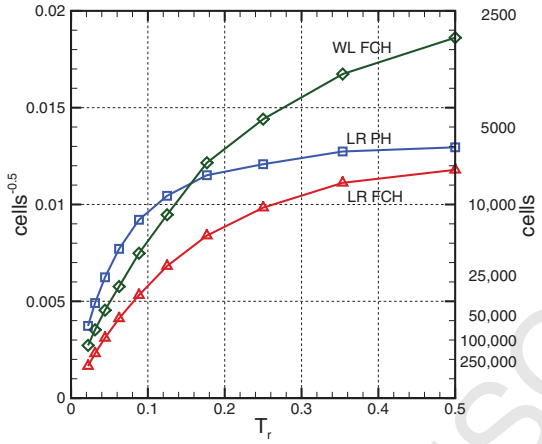

Figure 4: Number of cells in the three series of grids. The vertical axis shows the number of cells to the power -0.5, the corresponding number of cells is indicated to the right.

Finally, with the cell sizes proportional to the threshold (equation (1)), the number of cells should scale with T_r^{-2}. Figure 4 shows that for the Nakayama case, this is only asymptotically true as T_r goes to zero. On the coarser grids, in most of the domain the original grid is not refined. In these regions the number of cells does not change between grids; geometrical similarity is therefore not satisfied. However, even the coarsest grid is refined in those regions which are the most important for the simulation of the flow. Thus, in these regions the grids are geometrically similar. If this is enough for convergence studies, will be seen below.

5.3. Grid convergence of forces

For all computations, the convergence of the force coefficients $C_d = F_x/(\frac{1}{2}\rho V^2 cb)$ and $C_l = F_y/(\frac{1}{2}\rho V^2 cb)$ is studied, where the forces are computed on a wing segment of unit span $b = 1$. Uncertainty estimations are

performed following Eça and Hoekstra [14]. A power law:

$$C(T_r) = C_0 + \alpha(T_r)^p, \tag{4}$$

with an unknown order p is least-squares fitted through the force coefficients obtained on the five grids considered. This fit is used to compute the estimated error $\epsilon = \alpha(T_{r1})^p$ which is the difference between the value on the finest of these grids and the extrapolation to $T_r = 0$. Furthermore, the correlation error σ is computed as the standard deviation of the difference between the actual values on all grids and the power-law values obtained for the T_r of these grids. Then the uncertainty on the finest grid is estimated as $1.25\epsilon + \sigma$. The safety factor 1.25 due to Roache [17] is meant to ensure that the uncertainty interval contains the actual error with a 95% probability. When no satisfactory fit is found, new fits are computed with fixed-power laws and the factor of safety is increased. For details, the reader is referred to the original paper [14].

Other uncertainty estimation techniques such as those by Stern et al. [18] or Celik et al. [19] would probably also work with our data. However, the advantage of the Eça and Hoekstra approach for the current work is, that the standard deviation σ provides a measure of the quality of the fit. Thus, σ can be used to see how well our series of meshes produce asymptotic convergence.

5.3.1. Estimated uncertainties

The dependence of C_d and C_l on T_r is shown in figure 5; starting from the third grid, estimated uncertainties are included. To resemble the usual type of convergence study where a large number of grids is not always available, the uncertainties on each grid are computed from not more than five grids:

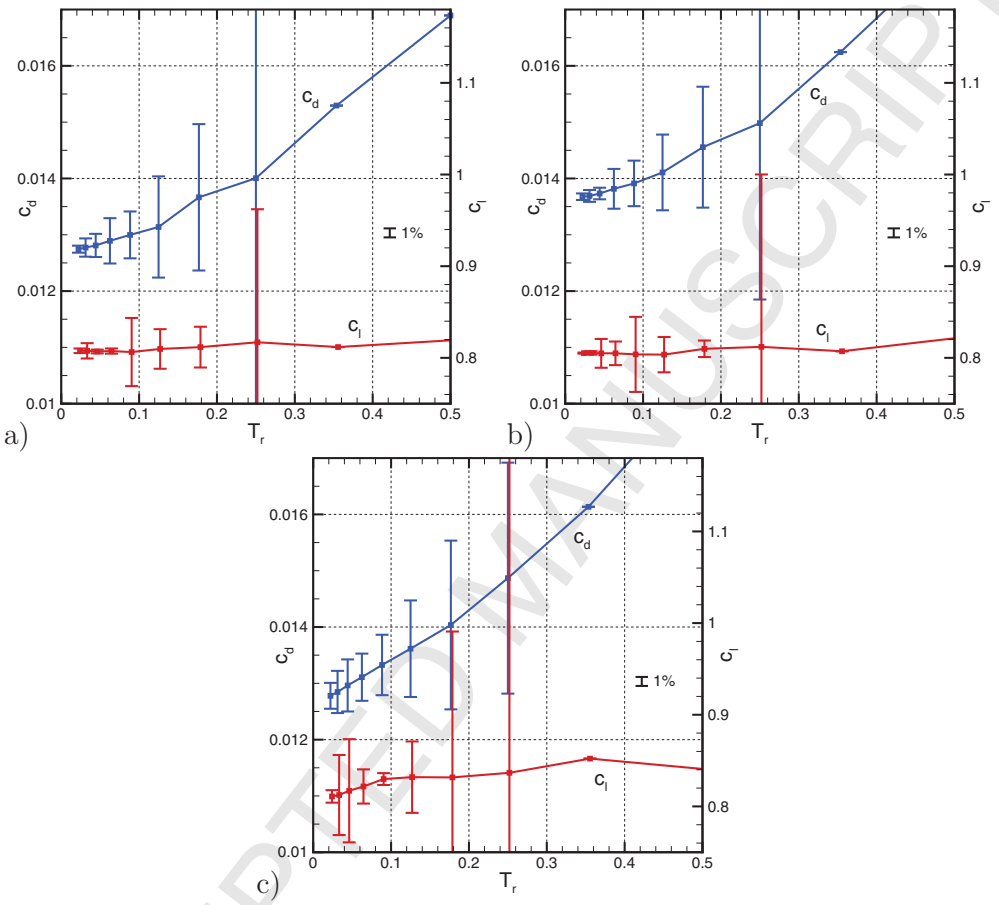

Figure 5: Convergence of drag coefficient C_d and lift coefficient C_l for the Nakayama airfoil, with respect to the threshold T_r. FCH with low-Reynolds grid (a) and wall-law grid (b), PH with low-Reynolds grid (c). The two vertical axes are scaled such that error bars of the same length for C_d and for C_l indicate the same relative uncertainty.

the grid itself and the previous four, if these exist.

The low-Reynolds grids with the FCH criterion provide the most complete simulation of the flow, since the wake and the boundary layer are fully resolved. Figure 5a shows a convergence for C_d that appears asymptotic with some superimposed high-frequency noise. The uncertainty estimation works well here: all the uncertainty intervals contain the actual converged C_d and those computed with five or four grids ($T_r \leq 0.1768$) are not excessively large. For C_l, even on the coarsest grids, the computed values are close to the solution on the finest grid. Unfortunately, this means that the convergence is dominated by noise so the estimated uncertainties are high; the estimation cannot show how good the coarse-grid solutions really are. Realistic error bars are obtained for $T_r \leq 0.0625$. Similar results are obtained for the WL FCH series (figure 5b), so the use of the wall law does not perturb the grid convergence.

However, the PH series shows a different convergence behaviour (figure 5c). First, C_l varies much more over the grids. And second, the shape of the convergence curve for both forces does not resemble a power law, which means that the uncertainty intervals on coarser grids do not contain the true converged values. So even though the lift is predominantly a pressure force, the pressure-based refinement criterion is not the optimal choice for computing it. On the contrary, a fine resolution of the boundary layers and the near wake is required.

Finally, the convergence behaviour on coarse grids is not much worse than on fine grids, even though these coarse grids are mostly equal to the original grid (see section 5.2): the convergence for C_d is smooth, while the values for C_l

rapidly approach the converged one. Thus, geometrically similar refinement of the most crucial regions is enough to get acceptable convergence.

5.3.2. Asymptotic convergence and noise

To see if the series of adaptively refined meshes produce asymptotically convergent solutions, we perform different power-law fits for the LR FCH series. For the least-squares method, the bias towards the finest grids suggested by Eça and Hoekstra is used. The coefficients p and C_0 for each fit, as well as the extrapolated error ϵ and correlation error σ (see equation (4) and its discussion) are gathered in table 1.

Table 1: Fits of the power law (4) to the forces on the Nakayama wing, using the results on different sets of grids. A symbol – indicates that no fit could be found. Percentages are expressed w.r.t. $D = 1.2714 \cdot 10^{-2}$ (C_d) and $D = 0.8070$ (C_l).

Finest T_r	grds.	$C_d.10^2$	p	$C_0.10^2$	$\epsilon\%D$	$\sigma\%D$	C_l	p	C_0	$\epsilon\%D$	$\sigma\%D$
All grids											
0.0884	6	1.3326	1.52	1.2659	2.38%	0.92%	0.8063	–			
0.0442	8	1.2963	1.56	1.2714	0.76%	0.57%	0.8069	0.68	0.8036	0.36%	0.23%
0.0221	10	1.2778	1.56	1.2714	0.26%	0.37%	0.8077	1.22	0.8070	0.03%	0.18%
5 grids											
0.0221	5	”	1.36	1.2698	0.36%	0.04%	”	1.00	0.8082	0.05%	0.05%
Every 2^{nd}											
0.0221	5	”	1.49	1.2697	0.33%	0.19%	”	1.48	0.8072	0.01%	0.19%

To eliminate noise in the convergence curves, Eça and Hoekstra recommend to use computations on as many grids as possible. Therefore, for the

first trial fits, the data on a grid and all the preceding ones is used. The resulting fit for C_d converges rapidly and is fully established with $T_r = 0.0442$ as the finest grid; adding the finest two grids does not change the fit, even though these points receive the largest weight in the fit. Furthermore a fit through the finest 5 grids, while not identical to the first fits, is good as shown by its low value of σ. This indicates that C_d actually behaves as a power law, perturbed only by small-scale noise.

This is different for C_l: the power-law fit has no solution for $T_r = 0.0884$, the coefficients vary for each added point, and C_{l0} is always lower than C_l, even though C_l on the finest grids increases with further refinement. The fit on the finest five grids looks more reasonable, however. It is plausible that the data for C_l contains a higher-order error term as well as high-frequency noise, so that fine grids are needed to reach the asymptotic range.

The noise in the data may be created because subsequent grids are not geometrically similar (section 5.2); indeed, figure 5 shows that most of the perturbations are positive for one point, negative for the next, etc. To eliminate this effect, a last fit was made using only every second grid, so each grid is twice finer than the previous one. For C_d, this fit is better than the one using all meshes since its σ is lower, but the improvement is small. For both forces, the fits are comparable. Thus, in practice it is possible to use the intermediate, not strictly similar meshes for data fitting.

5.3.3. Accuracy

What grid size is needed to accurately evaluate modelling errors? Figure 5 shows that a typical modelling error, switching from low-Reynolds to wall-law boundary conditions, creates a change in C_d of about 7%. For a

reliable comparison of the LR FCH and WL FCH computations, the numerical uncertainties for the two computations should be small compared to this modelling error. For example, an uncertainty of the order of 1% gives a 1.4% combined uncertainty for their difference; five times less than the modelling error. The results in figure 5 show that 1% estimated uncertainty for both C_l and C_d is reached around $T_r = 0.0313$, i.e. on very fine meshes.

However, 1% uncertainty is a very strict requirement which may not be necessary in practice for all computations. Comparing the computed values on each grid with the extrapolated value for the finest grid, we see that 1% *error* for C_d is reached around $T_r = 0.0442$ and for C_l around $T_r = 0.25$! And for C_d, using Richardson extrapolation, 1% error in C_{d0} is obtained from $T_r = 0.125$. Thus, modelling errors can be evaluated with results on coarser grids than the uncertainty estimates indicate, if one accepts these results without formal proof of their accuracy.

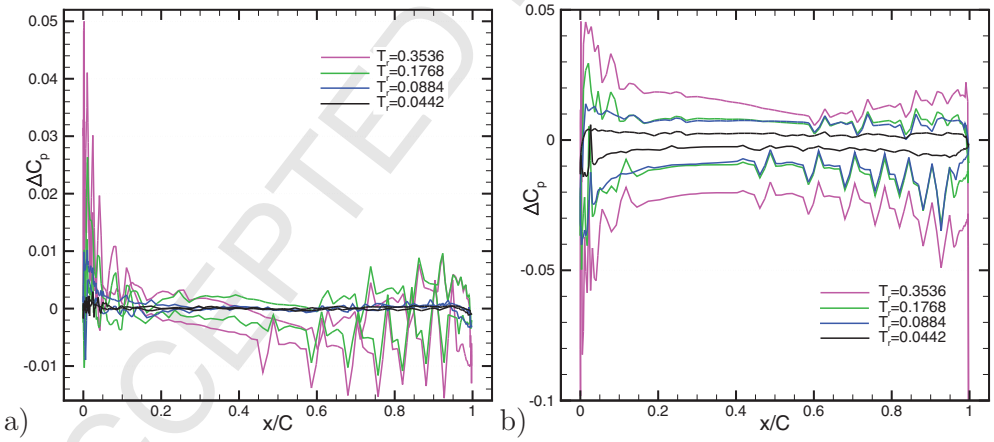

Figure 6: The difference in pressure coefficient C_p with respect to the finest grid $T_r = 0.0221$, for four coarser grids. LR FCH (a) and LR PH (b). Note that the vertical scale of the two figures is different.

5.4. Grid-independence of the local flow

Much of the interest of grid adaptation is in the precise computation of local flow details. This section studies the grid convergence of two such features: the pressure on the airfoil and the flow in the wake. For the distribution of $C_p = p/(\frac{1}{2}\rho V^2)$ on the profile, figure 6 shows the difference in the solution between four coarser grids and the finest one. For the LR FCH series (figure 6a), although the pressure signal is noisy, the distance between solutions decreases as the grids become finer. The main differences are found in the stagnation point and the suction peak on the leading edge, so the regions of fine cells created at the leading edge (see figures 2 and 3) are justified. From $T_r = 0.0884$ on, the pressure difference with the finest grid is less than 1% of the stagnation pressure everywhere.

The result for the LR PH series is different (figure 6b). The distance between solutions decreases for finer grids, but the difference with the finest grid is almost uniform over the entire profile and the differences between grids are larger. This confirms what was found in section 5.3.1: even for computing the pressure, the accurate resolution of the boundary layer flow is a necessity.

The second feature studied is the near and far wake. Figure 7 gives the velocity and one component of the turbulent fluctuation for the two FCH series (with the PH criterion, the wake is completely diffused). Judging visually from the distance between the curves, on the $T_r = 0.0442$ grid the solution is converged, the $T_r = 0.0884$ grid is close, and the $T_r = 0.1768$ grid already gives the right tendencies.

The results are not the same for low-Reynolds and wall-law boundary

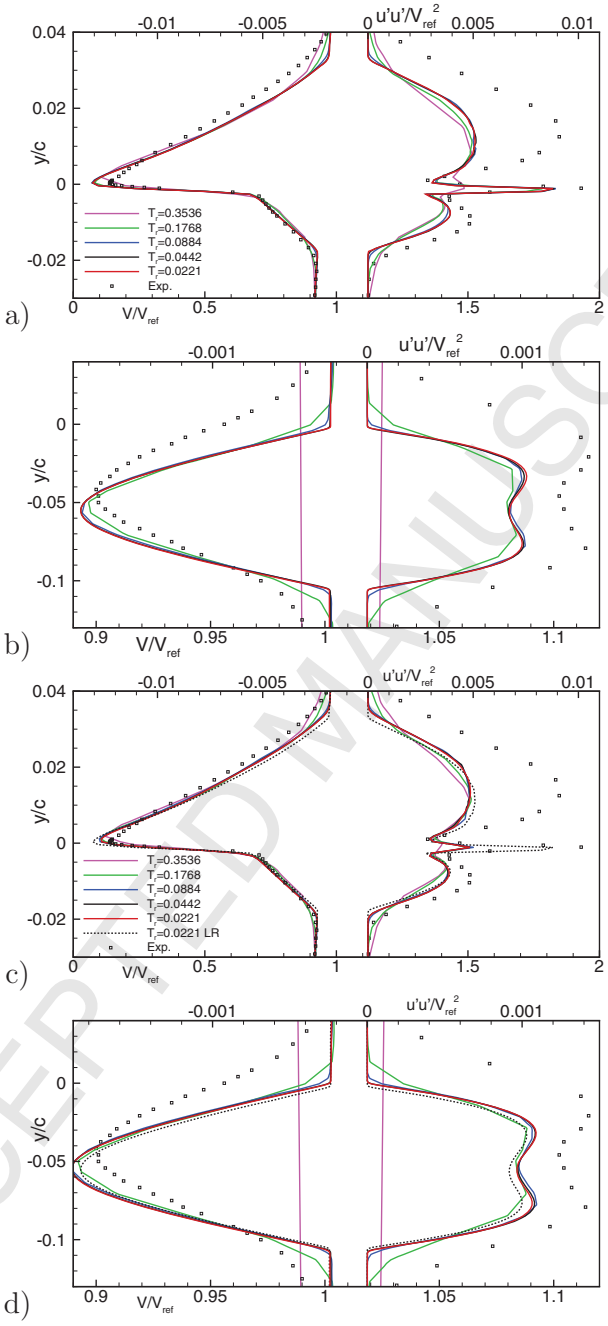

Figure 7: Velocity V (left lines) and fluctuation correlation $\overline{u'u'}$ (right lines) in the wake of the Nakayama airfoil, at $0.01c$ (a, c) and $2.0c$ (b, d) behind the trailing edge. LR FCH (a, b) and WL FCH (c, d).

conditions. The most notable effect of the wall law is the attenuation of the turbulence peak at the trailing edge in the near wake, which corresponds to the position where the grids differ most (figures 2b and c). However, even in the far wake a reduction in the velocity defect and the turbulence intensity can be seen due to the wall law.

Compared with the experiments [15], the shape of the turbulence profiles is well captured in both cases but the intensity is too low, so the wake and the boundary layer are too thin. Also, unlike the experiments, the edges of the wake are sharp in the numerical solutions. This is only observed because of the extra-fine grid around the edge of the wake (figure 2b); on coarser refined grids and therefore on standard boundary layer grids also, the numerical solution is closer to the experiments. This modelling error is therefore only visible thanks to the adaptive refinement.

Thus, from the $T_r = 0.0884$ grids on, the differences between the solutions in one series are small compared with the difference between low-Reynolds and wall-law solutions. Both these differences are small with respect to the distance between the simulations and the experiments. Thus, the $T_r = 0.0884$ grid and all finer grids indicate modelling errors in the wake reliably.

5.5. Equivalent three-dimensional grid sizes

The preceding sections have established that grid convergence for the Nakayama wing can be obtained using adaptive refinement. 1% proven uncertainty in the forces is obtained for $T_r = 0.0313$, while the far wake is converged for $T_r = 0.0884$. A reasonable approximation of both is already obtained when $T_r = 0.1768$. To generalise these conclusions, the thresholds should be translated to numbers of cells, or even better, to an order of

magnitude for the number of cells in an equivalent three-dimensional grid.

To estimate the number of cells in a 3D case which resembles the Nakayama airfoil, we suppose that the number of cells in the third direction (the spanwise direction for a wing) is of the same order as the number of cells in the flow direction. For want of something better, this number is approximated by the number of cells on the surface of the airfoil. Multiplying the number of surface cells by the total number of 2D cells therefore gives a rough indication of an equivalent 3D grid size. These numbers are given in table 2; they indicate that the grids for $T_r = 0.0884$ correspond to habitually used grid sizes, while those for $T_r = 0.0221$ are much finer than what is used in practice.

Table 2: Nakayama airfoil: estimated number of cells in three-dimensional grids.

T_r	0.3536	0.1768	0.0884	0.0442	0.0221
Low-Re FCH	2.1M	5.6M	26M	164M	983M
Low-Re PH	1.1M	1.6M	3.4M	12M	60M
WL FCH	0.7M	1.8M	8.1M	38M	188M

For low-Reynolds FCH computations, getting less than 1% uncertainty for the forces with our procedure requires more than 100M cells. However, these grids resolve the wake with very fine cells, which may not be necessary for obtaining only the forces; also, 1% uncertainty is often too strict (section 5.3.3). For reliable LR FCH predictions of the near and far wake, 20M to 50M cells are needed, so such simulations can be easily envisaged.

With the wall law and FCH, less than 1% uncertainty requires around 50M cells while the far wake can be captured accurately with about 10M

cells; both objectives are therefore possible. The LR PH computations are efficient, but to obtain 1% uncertainty in the forces they need even finer thresholds than those tested here. Furthermore, the wake is not resolved. Therefore, these computations may be of little interest despite their apparent low number of cells. Note that these are indications only, actual numbers of cells will depend on the real 3D geometry and the flow characteristics, as well as the quality of the original grid and the exact choice of the refinement strategy.

6. The step to 3D: KVLCC2 tanker

A representative three-dimensional case is the flow around the KVLCC2 tanker ship. We wish to see if this three-dimensional case confirms the findings for the Nakayama airfoil. The KVLCC2 hull has a complex aftbody flow with boundary layer separation and different vortical structures; figure 8 gives a global view of the flow field, with the thickening and separation of the boundary layer.

The flow is computed for a model of the ship with length $L = 5.571m$, which gives $Re = 4.6 \cdot 10^6$. Since the flow is symmetric, only half the hull is simulated. Free-surface effects are ignored: the water surface is treated as a symmetry plane, like the vertical centreplane. Windtunnel measurements are available from Postech [20], Larsson et al. [21] give a complete description of the test case.

The flow is simulated here with a wall-law approach and the anisotropic EASM turbulence model including rotation correction, since this model represents flows dominated by longitudinal vorticity better than the standard

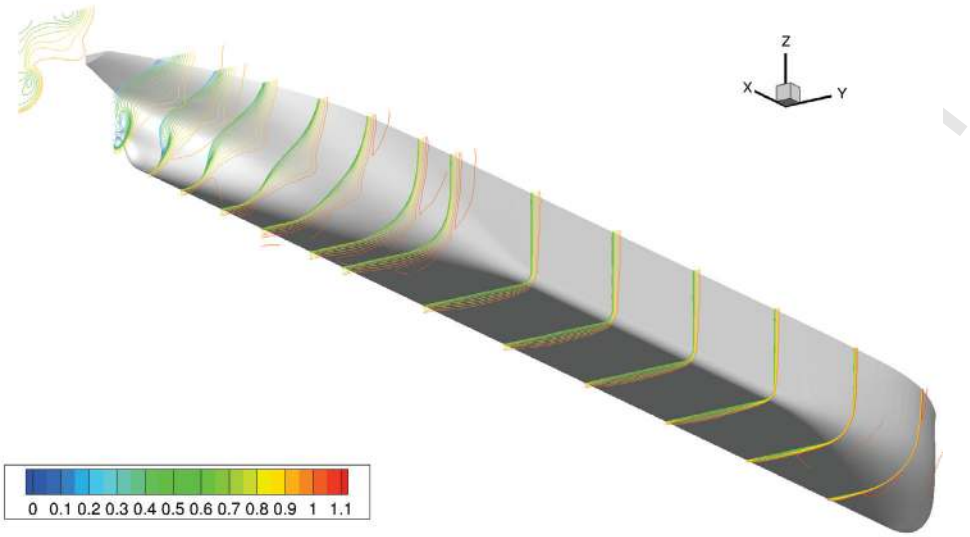

Figure 8: KVLCC2: the (half) hull and the axial velocity field.

$k-\omega$ SST model [6]. Convective terms are discretised with a blended central / upwind scheme. FCH is used as refinement criterion with the minimum cell size equivalent to the first-layer thickness in the boundary layer grid (see sections 4.3 and 5.1).

Table 3: KVLCC2: thresholds and mesh sizes.

T_r	3.0	2.0	1.5	1.0	0.75
nb. cells	247k	613k	1.34M	5.02M	10.18M

6.1. Refined meshes

For the convergence study five different thresholds are used (table 3), with a factor 4 between the cell sizes of the coarsest and finest mesh. This

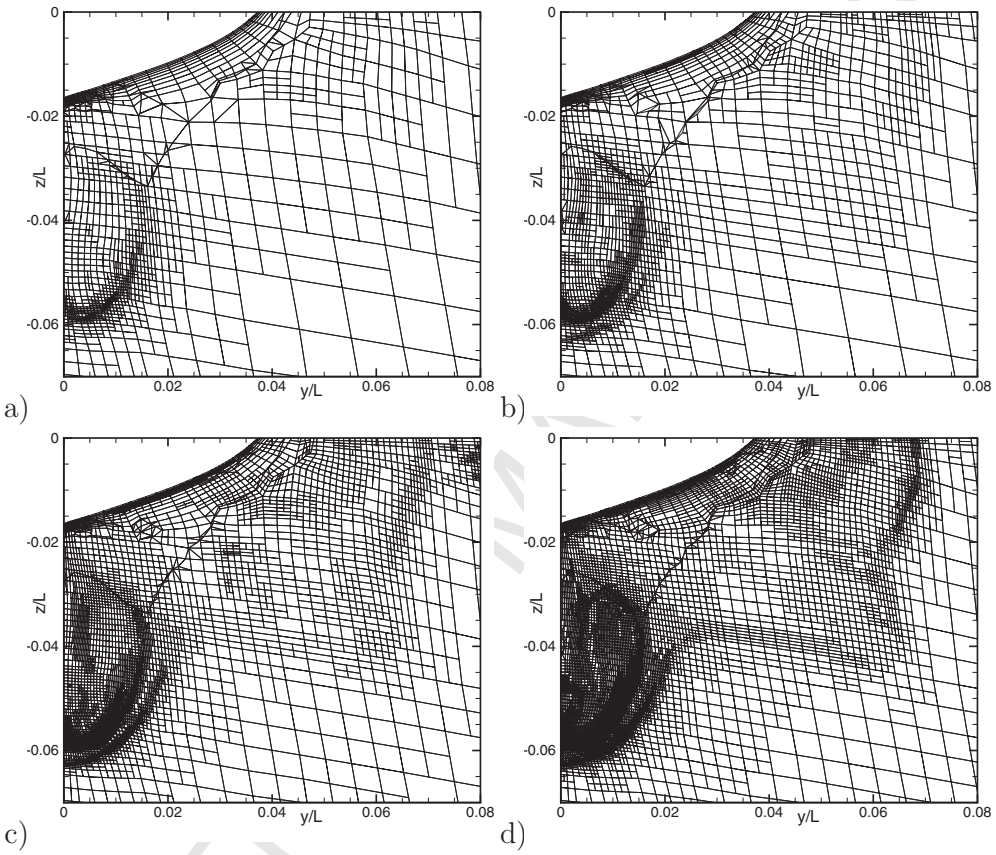

Figure 9: Meshes in the propeller plane $x/L = 0.9825$ of the KVLCC2, $T_r = 2.0$ (a), $T_r = 1.5$ (b), $T_r = 1.0$ (c), and $T_r = 0.75$ (d).

represents a more realistic study than the 10-grid Nakayama case, whose coarsest grids would be insufficient to capture the details of the more complex KVLCC2 flow, while the finest grids would be too expensive in three dimensions.

The refined meshes in the propeller plane (the aftmost cut plane on the hull in figure 8) are shown in figure 9 for the four finest thresholds. These meshes are refined around the principal features of the flow, such as the main aft-body vortex whose centre is near (0.01, -0.04). This vortex creates an open-type separation on the hull, with an associated ear-shaped shear layer that is captured with fine cells. A second shear layer starting at (0.0, -0.065) separates the boundary layer flow around the vortex from the flow outside. Finally, like for the 2D wing, the outer edge of the boundary layer attracts grid refinement everywhere.

The pairs of grids in figures 9a and c, as well as b and d, have thresholds T_r which differ by a factor two, so the pairs should be geometrically similar. This is mostly the case; the centre of the vortex with its vertically elongated cells is an example. However, there is a perturbing effect here that was not encountered before: as the mesh is refined, the flow field changes so the refinement criterion is modified, which leads to non-similar meshes. This effect is stronger than for the 2D case, because the KVLCC2 flow has more shear layers and discontinuities in the velocity gradients, which become sharper on fine meshes. Thus, especially on the outer edge of the boundary the grid size decreases more than linearly with T_r. This is not necessarily a problem, since the sharp resolution of these features improves the quality of the local flow on the fine meshes, while globally seen the four grids form a reasonable

series.

6.2. Convergence of the flow

The convergence of the resistance force $C_T = F_x/(\frac{1}{2}\rho V^2 S)$, with $S = 0.2682L^2$, is monotone (figure 10a). Numerical uncertainties were again estimated with the method of Eça & Hoekstra (section 5.3). The error bars computed with 3, 4 and 5 grids overlap, which is an indication that the procedure works well. On the finest grid, the uncertainty is 1.69% with an order $p = 1.7$. However, more noise is present than for the Nakayama case, which may be due to the irregularities in the mesh noted above. Due to the complexity of the case, it is also possible that a higher-order error is present.

The convergence of the pressure force coefficient C_{Tp} and the viscous force coefficient C_{Tv} are given in figures 10b and 10c respectively. While the viscous force is about four times larger than the pressure force, their variation across the grids and the computed uncertainties are of the same magnitude. Both figures show that the forces are not in the asymptotic range, but the estimated uncertainties are reasonable, especially for the viscous forces where the uncertainty intervals overlap. The pressure force has some more irregularity. The estimated uncertainties are 6.47% with order $p = 1.51$ for C_{Tp} and 1.11% with $p = 2.26$ for C_{Tv}.

The grid convergence of the local flow is good (figure 11). Even though the overlapping isolines of the Nakayama case are not obtained, the velocity and turbulence isolines for the finest grids are very close throughout the aft-body flow and also in the near wake. Even the 0.9 axial velocity isoline which represents the outer boundary layer, often underresolved in KVLCC2 simulations, is close to convergence thanks to the adaptive refinement. The

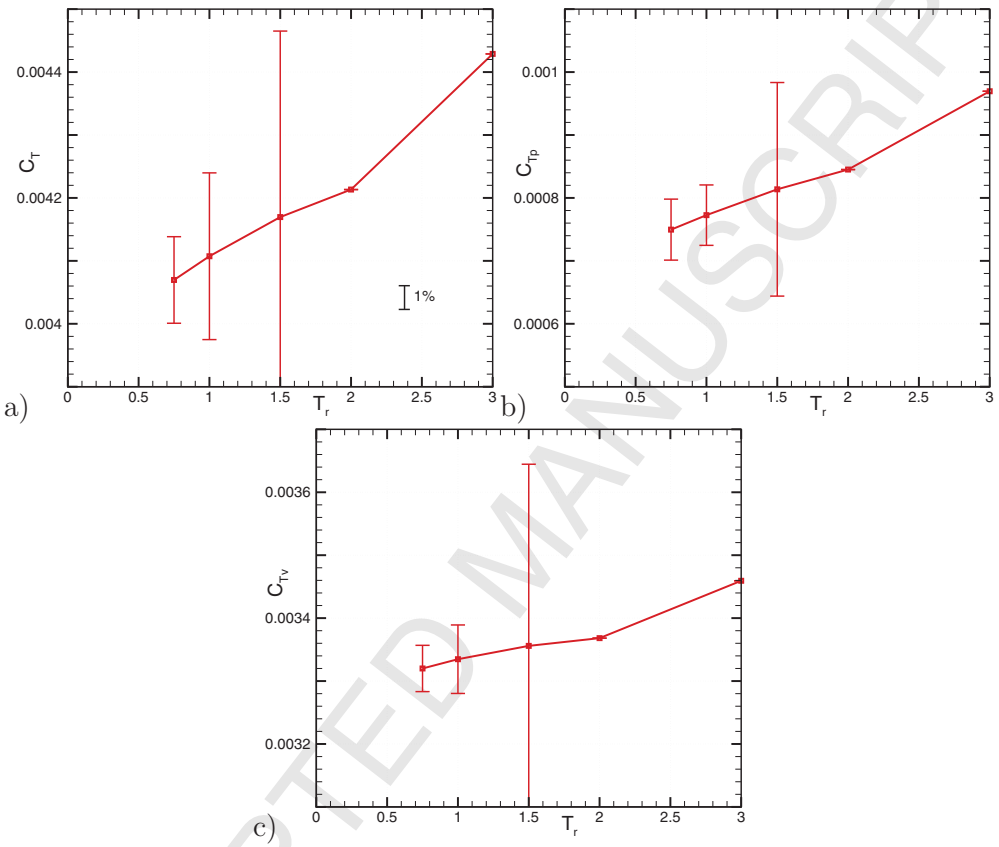

Figure 10: Grid convergence for the KVLCC2: resistance coefficient C_T (a), pressure and viscous resistance coefficients C_{Tp} (b) and C_{Tv} (c), with uncertainty computed as in section 5.3, using all the coarser grids available. The scale of the axes is the same for the three graphs.

turbulence kinetic energy in the shear layer is the same on the two finest grids. And finally, the differences between the two finest meshes are small compared to their difference with the wind-tunnel measurements from [20]. Thus, the solution on the finest mesh is sufficiently precise to assess modelling errors accurately.

These results agree with the findings of the Nakayama case (section 5.5). Convergence of the local flow in the wake is indeed obtained at 10M cells. Furthermore, we expect a 1% uncertainty in the forces around $T_r = 0.5$ judging from figure 10. Such a grid would have around 40M cells (eight times the number of cells in the $T_r = 1.0$ grid), which also corresponds to the Nakayama estimate. This confirms the soundness of the 2D-based orders of magnitude for the number of cells, although the KVLCC2 is not a wing so the nearly exact agreement is due to coincidence.

7. Conclusion

7.1. Evaluation of the results

To use adaptive grid refinement for computing grid-independent solutions, the introduction identified three requirements: grid adaptation should create geometrically similar grids, these have to work with established uncertainty estimation methods, and the resulting global and local computed quantities must indeed be accurate. Can these requirements be met?

In principle, the adaptive refinement of a coarse, unstructured hexahedral grid using a procedure based on metric tensors is ideal for grid convergence studies. It ensures that cells in all meshes have the same shapes and orientations, while the cell sizes in a mesh are proportional to the global refinement

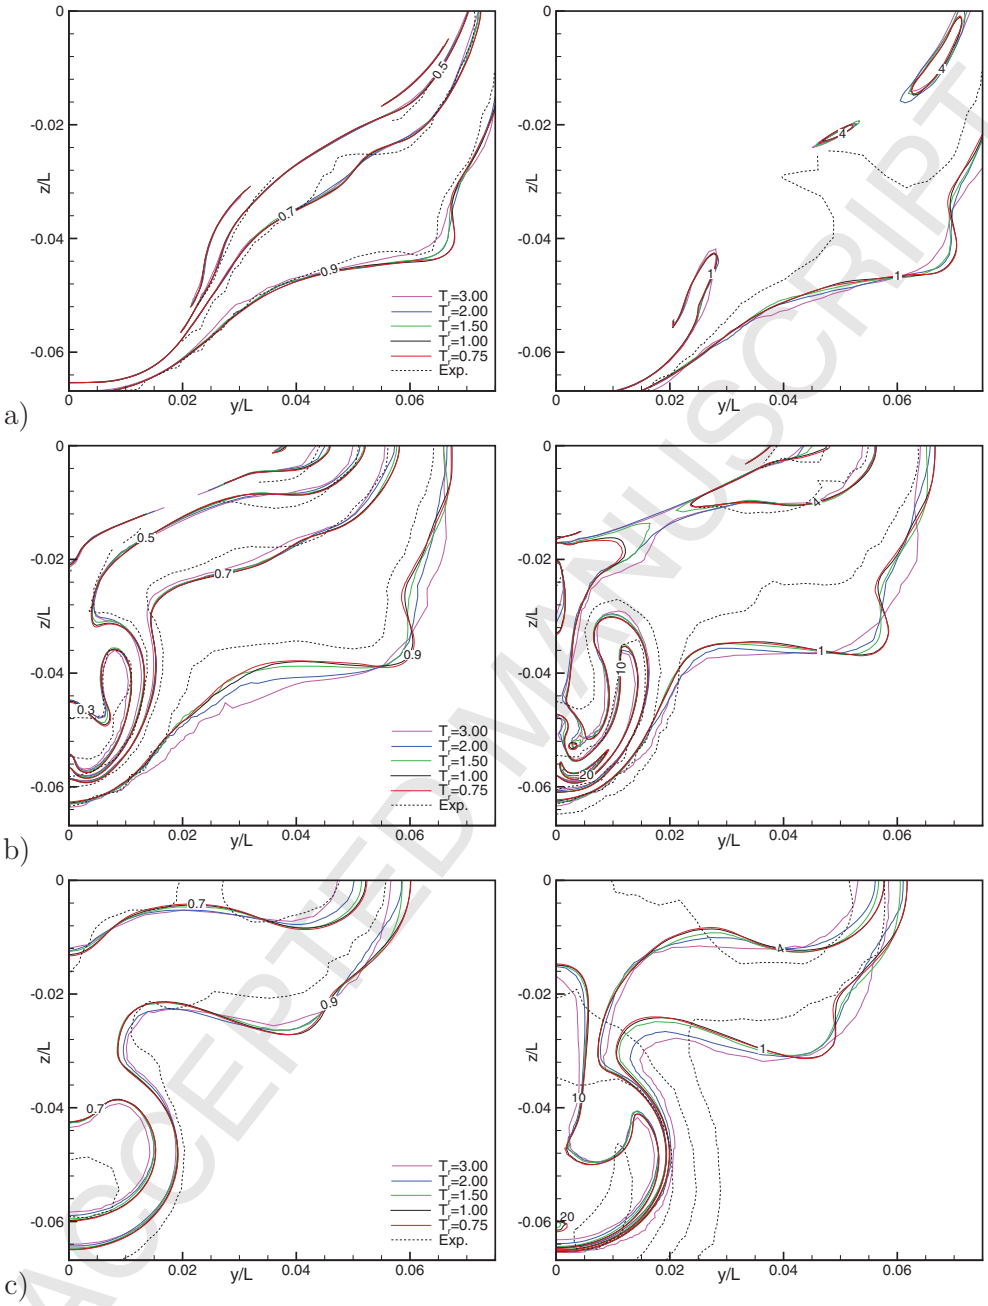

Figure 11: Grid convergence for the KVLCC2 in cut planes at $x/L = 0.9$ (a), in the propeller plane $x/L = 0.9825$ (b), and behind the ship at $x/L = 1.1$ (c): axial velocity (left) and turbulence kinetic energy (right).

threshold T_r. The refined meshes should therefore be similar.

However, the possible grid sizes are limited by the refinement method, since refined cells are created by dividing existing cells in half. Thus, if the threshold is reduced by a factor less than 2, the cell sizes are not reduced by this factor; instead, the same cell sizes are kept and the zones of fine cells are enlarged. Furthermore, geometric similarity is perturbed when the flow contains discontinuities and sharp features. These are better resolved on finer grids so the refinement criterion is increased locally, which leads to more refinement than expected on the finest grids. Finally, on coarse meshes the original grid is conserved over much of the domain, while on fine meshes, the refinement may be limited by an imposed minimum cell size.

Despite these limitations, series of grids are produced which are globally similar. The consistent orientations of the cells, as well as the equal position of the boundary layer mesh in each grid, are useful to ensure smooth convergence. Finally, for cases where the flow is well resolved and the threshold is varied by a factor 2, local geometric similarity has been confirmed by our tests.

A least-squares based uncertainty estimation technique was applied to the computed forces, with success. Oscillations in the results are observed when the thresholds between meshes do not vary by a factor 2. However, the least-squares approach removes these oscillations if enough meshes are used. Using only meshes which vary by a factor 2 improves the fits, but the improvement is small. Thus, series of adaptively refined meshes that are not locally similar can be used for uncertainty estimation.

The original grid which is not refined everywhere for coarse meshes, does not deteriorate the convergence. The mesh is refined there where this is crucial for the flow, which is enough to ensure grid convergence of the forces. The lift of an airfoil converges quickly with grid adaptation, because the leading-edge pressure peak and the boundary-layer displacement are well resolved on adapted meshes. Finally, the KVLCC2 case shows that the forces are indeed perturbed when the refinement criterion is resolved better on fine meshes. However, this probably means that the flow is not yet in the asymptotic range, so perturbations in the convergence can be expected even if the grids were perfectly similar.

The convergence of local flow features, studied by visual inspection of isoline plots, is good. For the two-dimensional airfoil, the computed wake on the finest grids is identical, even far behind the trailing edge. In the more complicated flow around the KVLCC2, the velocity and the turbulence are nearly identical on the finest grids; the differences are small with respect to the distance between computations and experiments.

7.2. Limitations and perspectives

This final section discusses the practical usefulness of grid convergence studies using adaptive refinement. First of all, how large is their field of application? The tests performed here use the ISIS-CFD flow solver and grid adaptation, so they are limited to high-Reynolds incompressible RANS solutions, hexahedral grids, and Hessian-based refinement criteria. However, the idea that series of grids with geometrically similar cell sizes can be obtained by varying a refinement threshold is valid for all metric-based adaptation techniques which use solution-based refinement criteria (i.e. crite-

ria with weak mesh dependence). This approach is not even limited to fluid dynamics, it could be used for the simulation of structures etc.

Obtaining identical cell orientations between grids requires adaptive grid refinement by subdivision of an original grid, which is typically associated with mesh adaptation for hexahedral cells. However, as shown above, exact local similarity may not be required for convergence studies. Therefore, the proposed technique could also work for tetrahedral mesh generation, especially for anisotropic meshes where the element orientation between meshes is similar, since it is guided by the anisotropy.

Thus, any combination of a flow solver, an adaptive mesher and a metric-based refinement criterion can potentially be used to perform convergence studies. The exact procedure needs to be determined for each method.

For incompressible flows, we have shown that the technique can produce converged local-flow solutions: the numerical accuracy of the computed wake flows is sufficient to assess modelling errors due to turbulence models, wall laws, etc., on meshes with acceptable numbers of cells. Therefore, the technique can be used for research in turbulence modelling of complex flows.

For the computation of forces, useful uncertainty estimations are produced. However, some perturbations are generated because the meshes are not perfectly similar. Low uncertainties for wall-law meshes can be obtained with reasonable numbers of cells, while the same is costly for wall-resolved boundary layers. This is because the grid refinement method resolves all the details of the wake, which may not be necessary if the only objective is to compute forces. Thus, convergence studies on structured, non-adapted

grids are to be preferred for force computations, if these grids can actually be generated. Otherwise, the proposed procedure is a valid option.

Using adaptive meshing for grid convergence studies has a final advantage: it is simple to perform. Once a first computation is set up, simulations on different meshes are obtained by changing only the refinement threshold. Furthermore, a grid convergence study is not expensive, since one is interested in the results on the finest grid. Therefore, with respect to a one-grid computation only coarser grids are added. In three dimensions, the computations on all coarser grids together need little more than half the time for the finest grid. Finally, unstructured grids can be used without added difficulty.

It is perhaps this aspect, as well as the ability to resolve the details of complex flows, which makes the idea of fundamental importance. Adaptive refinement can make grid convergence studies so straightforward that they become accessible for everyday CFD simulation.

Acknowledgements

This work was granted access to the HPC resources under the allocation 2015-2a1308 made by GENCI (Grand Equipement National de Calcul Intensif), which is gratefully acknowledged.

References

[1] J. P. Moitinho de Almeida, P. Díez, C. Tiago, N. Parés (Eds.), Adaptive Modelling and Simulation 2013, Lissabon, Portugal, 2013.

[2] A. Hay, M. Visonneau, Error estimation using the error transport equation for finite-volume methods and arbitrary meshes, Int J Comput Fluid Dyn 20 (7) (2006) 463–479.

[3] J. Wackers, A. Leroyer, G. B. Deng, P. Queutey, M. Visonneau, Adaptive grid refinement for hydrodynamic flows, Comput Fluids 55 (2012) 85–100.

[4] J. Wackers, G. B. Deng, E. Guilmineau, A. Leroyer, P. Queutey, M. Visonneau, Combined refinement criteria for anisotropic grid refinement in free-surface flow simulation, Comput Fluids 92 (2014) 209–222.

[5] P. L. George, H. Borouchaki, Delaunay Triangulation and Meshing - Application to Finite Elements, Hermes, 1998.

[6] R. Duvigneau, M. Visonneau, On the role played by turbulence closures in hull shape optimization at model and full scale, J Mar Sci Tech 8 (2003) 11–25.

[7] P. Queutey, M. Visonneau, An interface capturing method for free-surface hydrodynamic flows, Comput Fluids 36 (9) (2007) 1481–1510.

[8] J. Wackers, B. Koren, H. C. Raven, A. v. d. Ploeg, A. Starke, G. B. Deng, P. Queutey, M. Visonneau, T. Hino, K. Ohashi, Free-surface viscous flow solution methods for ship hydrodynamics, Arch Comput Meth Eng 18 (2011) 1–41.

[9] C. M. Rhie, W. L. Chow, A numerical study of the turbulent flow past an isolated airfoil with trailing edge separation, AIAA J 17 (1983) 1525–1532.

[10] J. Wackers, G. B. Deng, E. Guilmineau, A. Leroyer, P. Queutey, M. Visonneau, Creating free-surface flow grids with automatic grid refinement, in: S. Perotto, L. Formaggia (Eds.), New Challenges in Grid Generation and Adaptivity for Scientific Computing, Springer, 2015.

[11] A. Loseille, A. Dervieux, F. Alauzet, Fully anisotropic goal-oriented mesh adaptation for 3D steady Euler equations, J Comput Phys 229 (2010) 2866–2897.

[12] A. Belme, A. Dervieux, F. Alauzet, Time accurate anisotropic goal-oriented mesh adaptation for unsteady flows, J Comput Phys 231 (2012) 6323–6348.

[13] R. V. Wilson, F. Stern, H. W. Coleman, E. G. Paterson, Comprehensive approach to verification and validation of CFD simulations – Part 2: application for Rans simulation of a cargo/container ship, J Fluids Eng 123 (2001) 803–810.

[14] L. Eça, M. Hoekstra, A procedure for the estimation of the numerical uncertainty of CFD calculations based on grid refinement studies, J Comput Phys 262 (2014) 104–130.

[15] A. Nakayama, Characteristics of the flow around conventional and supercritical airfoils, J Fluid Mech 160 (1985) 155–179.

[16] V. Pržulj, B. Basara, Bounded convection schemes for unstructured grids, in: 15th AIAA Computational fluid dynamics conference, AIAA paper 2001-2593, Anaheim, CA, 2001.

[17] P. J. Roache, Fundamentals of Verification and Validation, Hermosa Publishers, 2009.

[18] F. Stern, R. V. Wilson, H. W. Coleman, E. G. Paterson, Comprehensive approach to verification and validation of CFD simulations – Part 1: methodology and procedures, J Fluids Eng 123 (2001) 793–802.

[19] I. B. Celik, U. Ghia, P. J. Roache, C. J. Freitas, H. Coleman, P. Raad, Procedure for estimation and reporting of uncertainty due to discretization in CFD applications, J Fluids Eng 130 (2008) 1–4.

[20] S. J. Lee, H. R. Kim, W. J. Kim, S. H. Van, Wind tunnel tests on flow characteristics of the KRISO 3,600 TEU containership and 300K VLCC double-deck ship models, J Ship Res 47 (1) (2003) 24–38.

[21] L. Larsson, F. Stern, M. Visonneau, Numerical Ship Hydrodynamics, An assessment of the Gothenburg 2010 Workshop, Springer, 2014.