Nationaal Lucht- en Ruimtevaartlaboratorium

National Aerospace Laboratory NLR

NLR-TP-2015-015

A high-order finite-volume method with block-structured local grid refinement

J.C. Kok and H. van der Ven

Nationaal Lucht- en Ruimtevaartlaboratorium
National Aerospace Laboratory NLR

Anthony Fokkerweg 2
P.O. Box 90502
1006 BM Amsterdam
The Netherlands
Telephone +31 (0)88 511 31 13
Fax +31 (0)88 511 32 10
www.nlr.nl

Nationaal Lucht- en Ruimtevaartlaboratorium
National Aerospace Laboratory NLR

Executive summary

A high-order finite-volume method with block-structured local grid refinement

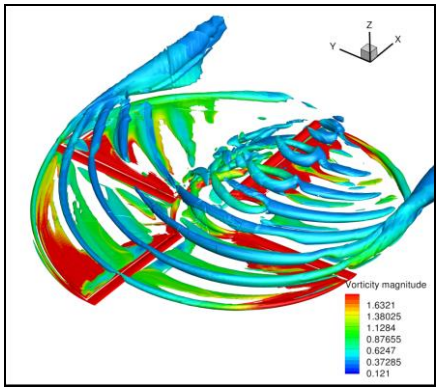

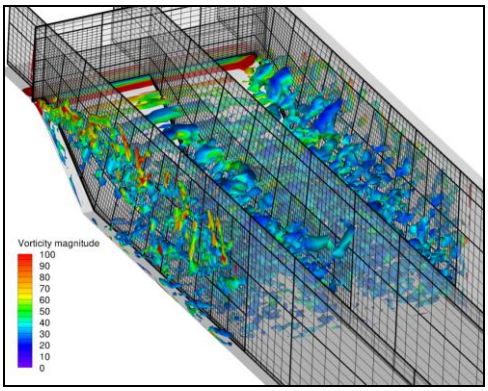

Problem area
Massively separated and vortex dominated flows play an important role in topics such as the design of silent landing gear, the study of stability and control properties of fighter aircraft in relation to vortex breakdown, the study of aerodynamic loads on structural aircraft components due to buffeting, and the study of helicopter vibrations induced by vortex-blade interaction.

Capturing these kind of flows is still a major challenge for Computational Fluid Dynamics. It requires advanced modelling approaches such as hybrid RANS-LES to capture the massively separated, turbulent flows as well as high-order numerical methods to accurately capture vortices. Additionally, local grid refinement may be employed, aiming to place the grid resolution where it is needed to capture the physics, while keeping the computations affordable.

Description of work
The high-order method employed in NLR's main CFD method ENSOLV has been extended with the capability to allow block-wise local grid refinement while maintaining the high-order accuracy. Furthermore, a method has been developed to automatically generate block-wise refined grids, placing uniform grid resolution where needed. First, the accuracy of the method has been verified for an elementary test case, consisting of the convection of an isentropic vortex. Next, the new capability has been applied to the time-accurate computation of helicopter vortex-blade interaction using the RANS equations. Finally, it has also been

Report no.
NLR-TP-2015-015

Author(s)
J.C. Kok
H. van der Ven

Report classification
UNCLASSIFIED

Date
January 2016

Knowledge area(s)
Computational Physics and Theoretical Aerodynamics

Descriptor(s)
CFD
high-order method
finite-volume method
local grid refinement

The contents of this report have been initially prepared for publication as an article in 'IDIHOM: Industrialization of High-Order Methods — A Top-Down Approach', Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 128, 2015 by Springer.

applied to the computation of the strongly separated, turbulent flows over a bump in a square duct and over a delta wing at high angle of attack using the hybrid RANS-LES method.

Results and conclusions
The performed computations clearly demonstrate that the newly developed grid refinement capability allows generating grids with the grid resolution where needed. For the helicopter case, this allowed improved vortex capturing without significantly increasing the computational costs. For the hybrid RANS-LES computations of strongly separated flows, the computational costs could be reduced without sacrificing accuracy.

Applicability
The new capability to perform accurate computations on locally refined grids, as well as to automatically generate such grids, can be applied to advanced CFD computations of strongly separated and vortex-dominated flows.

Nationaal Lucht- en Ruimtevaartlaboratorium, National Aerospace Laboratory NLR

Anthony Fokkerweg 2, 1059 CM Amsterdam,
P.O. Box 90502, 1006 BM Amsterdam, The Netherlands
Telephone +31 88 511 31 13, Fax +31 88 511 32 10, www.nlr.nl

Nationaal Lucht- en Ruimtevaartlaboratorium
National Aerospace Laboratory NLR

NLR-TP-2015-015

A high-order finite-volume method with block-structured local grid refinement

J.C. Kok and H. van der Ven

The contents of this report have been initially prepared for publication as an article in 'IDIHOM: Industrialization of High-Order Methods — A Top-Down Approach', Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 128, 2015 by Springer.

This publication has been refereed by the Advisory Committee AEROSPACE VEHICLES.

Customer	European Commission
Contract number	ACP0-GA-2010-265780
Owner	NLR
Division NLR	Aerospace Vehicles
Distribution	Unlimited
Classification of title	Unclassified
	Januari 2016

Approved by:

Author	Reviewer	Managing department
J.C. Kok	H. van der Ven	C. E. J. Wijnberg
Date: 24/11/2015	Date: 24/11/2015	Date: 30/11/15

Contents

This page is intentionally left blank.

A high-order finite-volume method with block-structured local grid refinement

Johan C. Kok and Harmen van der Ven

National Aerospace Laboratory NLR,
Flight Physics and Loads Department,
Amsterdam, the Netherlands,
johan.kok@nlr.nl, harmen.van.der.ven@nlr.nl

Abstract. Time-dependent vortex-dominated flows are computed accurately with a high-order finite-volume method on structured grids. In order to attain the required grid resolution in the vortex region, block-wise local grid refinement is employed. A new topology-based block-refinement algorithm allows the efficient generation of such block-wise refined meshes. The high-order finite-volume method is extended with high-order interpolation to deal with the partially continuous grids at block interfaces that result from the refinement. Results are presented for three applications: one time-accurate RANS computation of a helicopter flow case and two hybrid RANS–LES computations of strongly separated flows.

Keywords: CFD, high-order method, finite-volume method, local grid refinement

1 Introduction

Boundary-conforming block-structured finite volume solvers have several advantages. The block-structured mesh allows for accurate and efficient discretization of the flow equations. In particular, high-order methods can be devised by enlarging the stencil with only a moderate increase of computational effort. Using a single boundary-conforming mesh obviates the need for three-dimensional interpolation algorithms which are for instance necessary for Chimera methods. These advantages come at a price, though, as on the one hand the grid generation process is not fully automated and on the other hand generating an efficient grid in terms of number of grid cells is hampered by geometrical and topological restrictions, especially for complex geometries.

For time-dependent vortex-dominated flows, for example, the mesh should be as uniform as possible in the vortex regions. Within the multi-block structured framework such grids are difficult or even impossible to generate, unless some freedom in the grid structure is introduced. Here, this freedom consists of block-wise refinement of the mesh, where the grid in a block may be refined in one or more direction(s) by a factor 2^n. This results in partially continuous grids at block interfaces, where the ratio of mesh widths may be 2^n:1. Block-wise refined

meshes have the potential of attaining uniform grid spacing in the vortex regions. To generate these block-wise refined meshes efficiently, a topology-based block-refinement algorithm has been developed which refines the topology in such a way that the resulting block-refined mesh is as uniform as possible in, and only in, the vortex region.

The multi-block flow solver ENSOLV employs a low-dispersion symmetry-preserving fourth-order finite-volume method [7, 10]. This allows the accurate capturing of vortices, for example in helicopter rotor computations and in LES or hybrid RANS–LES of turbulent flow. This high-order method must be extended to deal with the partially continuous grids at block interfaces, resulting from the block-wise grid refinement. A new algorithm has been developed that can be both second-order and fourth-order accurate, and that can even be applied to completely discontinuous grids. Here, this algorithm is described as applicable to partially continuous grids. Furthermore, the algorithm is verified for a basic test case, consisting of the convection of an isentropic vortex in a uniform flow.

This paper presents results for the application of local grid refinement to both helicopter flow simulations and X-LES computations. The helicopter flow concerns the baseline case of the HART-II experiment. The X-LES method [8, 10] is a DES-type [12] hybrid RANS–LES method. Two cases are considered: the flow over a bump in a square duct and the flow over a delta wing at high angle of attack. For both cases, due to the simple geometry, a baseline grid is available that is practically uniform in the separated flow region where an LES is performed. Therefore, local grid refinement is not needed to obtain the desired grid resolution in the region of interest. Instead, the grid outside the regions of interest is coarsened to investigate whether a reduction of computational costs can be obtained in this way, without loss of accuracy.

2 Basic method description

The flow solver ENSOLV has the following capabilities which are relevant to the subject of this paper.

A fourth-order, symmetry-preserving, low-dispersion finite-volume scheme [7] is used to discretize convection. A central (instead of upwind) discretization is used, so that the method contains no numerical dissipation. The finite-volume method maintains its properties (fourth-order accurate, low numerical dispersion, no numerical dissipation) on non-uniform, curvilinear grids. On uniform, Cartesian grids, it is equivalent to the DRP scheme of Tam and Webb [13]. An important property of the symmetry-preserving finite-volume method is that it ensures that kinetic energy is conserved by convection (see e.g. Verstappen and Veldman [16]). This improves the stability properties of the scheme significantly, requiring only a low level of artificial diffusion, if any at all. For compressible flow, sixth-order artificial diffusion is added to the equations, maintaining the fourth-order accuracy [10].

For aeroelastic simulations, the flow solver is coupled in a strong way (exchange at each time step) to a modal representation of the structural model. At

each time step, the full 3D grid must be deformed in order to remain boundary conforming with the aeroelastically deformed geometry. This is done using a combination of a 3D volume spline (block vertices and edges) and transfinite interpolation (block faces and interior). The fourth-order finite-volume method has been extended to remain fourth-order accurate as well as fully conservative on such deforming grids.

The X-LES method [8] is a DES-type hybrid RANS–LES method based on the k–ω turbulence model. A single set of turbulence-model equations is used to model both the Reynolds stresses in RANS zones and the subgrid stresses in LES zones. The X-LES method in particular is based on the TNT k–ω model. The method switches to LES when the RANS length scale ($l = \sqrt{k}/\omega$) exceeds the LES length scale ($C_1\Delta$, with Δ the filter width and $C_1 = 0.05$). The RANS length scale is then replaced by the LES length scale in the expression for the eddy viscosity as well as in the expression for the dissipation of turbulent kinetic energy. The filter width Δ is defined at each grid point as the maximum of the mesh size in all directions.

To improve the capturing of free shear layers, two modifications have been added to the X-LES method: a stochastic subgrid-scale (SGS) [9] model and a high-pass filtered SGS model [10]. For the present computations, only the high-pass filtered SGS model has been employed. In this SGS model, the SGS stresses are computed from the velocity fluctuations instead of the instantaneous velocity, where the velocity fluctuations are obtained by applying a temporal high-pass filter to the velocity field.

To prevent the X-LES method from inadvertently switching to LES inside attached boundary layers (resulting in so-called shear-stress depletion), the shielding function of DDES [14] is included.

3 Block-structured local grid refinement

3.1 Algorithm

The new topology-based local grid refinement algorithm is described briefly below. The interested reader may find more details in Van der Ven et al. [15]. The algorithm consists of the following steps:

1. uniform refinement of the block topology, splitting all blocks into subblocks of a fixed block size (measured as the number of cells within a block); typically the target block size is 8^3 or 16^3;
2. further topology-refinement of those blocks which are targeted for grid refinement; the subblock topology is such that each grid-refined subblock is as close as possible to the fixed block size used in step 1;
3. based on the user-defined maximum refinement ratio (2:1, 4:1, etc.) blocks bounding the refinement region are targeted for refinement, and step 2 is repeated for those blocks;
4. generate the block grids within the refined topology.

In step 2 a refinement sensor is used which indicates how the grid within a block should be refined. The sensor consists of two parts: 1) the target mesh width; 2) the region where this mesh width should be attained. The target mesh width should be attained for all cells within a block. The refinement region can be specified in different ways: distance to a geometric object; specification of specific blocks; or a region described by a simple geometric object such as a sphere, cylinder, or cube. As the aim of the topology refinement is a uniform mesh, the grid refinement within a block is allowed to be anisotropic.

The algorithm is demonstrated for the NACA0012 airfoil, using a refinement region defined by the distance to a line segment. This example serves as an illustration of the algorithm only; in the next section the algorithm will be applied to the isolated HART II rotor.

The original mesh and topology is shown in Figure 1a. With a specified block size of 8^2, the refined topology of step 1 of the algorithm is shown in Figure 1b. The line segment used to define the refinement region is shown in Figure 1c. A region at a distance of 10% chord to this segment is defined as the refinement region. Within this region a mesh width of 0.001 chord should be attained. Figure 1d shows the refined topology which is the result of step 2 of the algorithm. Note that at this stage the grid has not been refined yet and that some of the blocks in the refined topology consist of a single cell of the original mesh. Subsequently, the user-defined refinement ratio is applied. Figure 1e shows the refined topology based on a maximum ratio of 2:1. In the last step of the algorithm the refined grid is generated, which is shown in Figure 1f.

It is worthy to note that the final topology consists of 982 blocks. Clearly, it is unfeasible to expect from a user to generate such topologies by hand.

3.2 HART II rotor

For rotor flows the convection of the tip vortices in the wake of the blades is important to capture the blade-vortex interaction. As the vortices move through the wake, the easiest way of obtaining a mesh with sufficient resolution in the vortex regions is to uniformly refine the mesh in a cylinder around the rotor blades. So the definition of the refinement region in this case is a cylinder centred at the rotor hub with radius equal to the rotor radius and sufficient height to contain the vortices.

The target mesh width is $0.01R$, where R is the rotor radius. This corresponds to 16% chord, which is a rather coarse resolution, but nevertheless leads to a refined mesh of 26 million cells, whereas the original mesh had 13 million cells. Figure 2a shows a detail of the original mesh in a horizontal plane through the hub. The same detail of the refined mesh is shown in Figure 2c. The uniform resolution in the wake is clearly visible. The imposed refinement ratio of 2:1 is also visible in the resolution outside the rotor disk. Figures 2b and 2d show the refinement in a vertical plane through the hub bisecting the rotor disk between two blades.

Figures 2e and 2f demonstrate that the target mesh width is actually obtained in the specified refinement region. The figure shows the ratio of the maximum

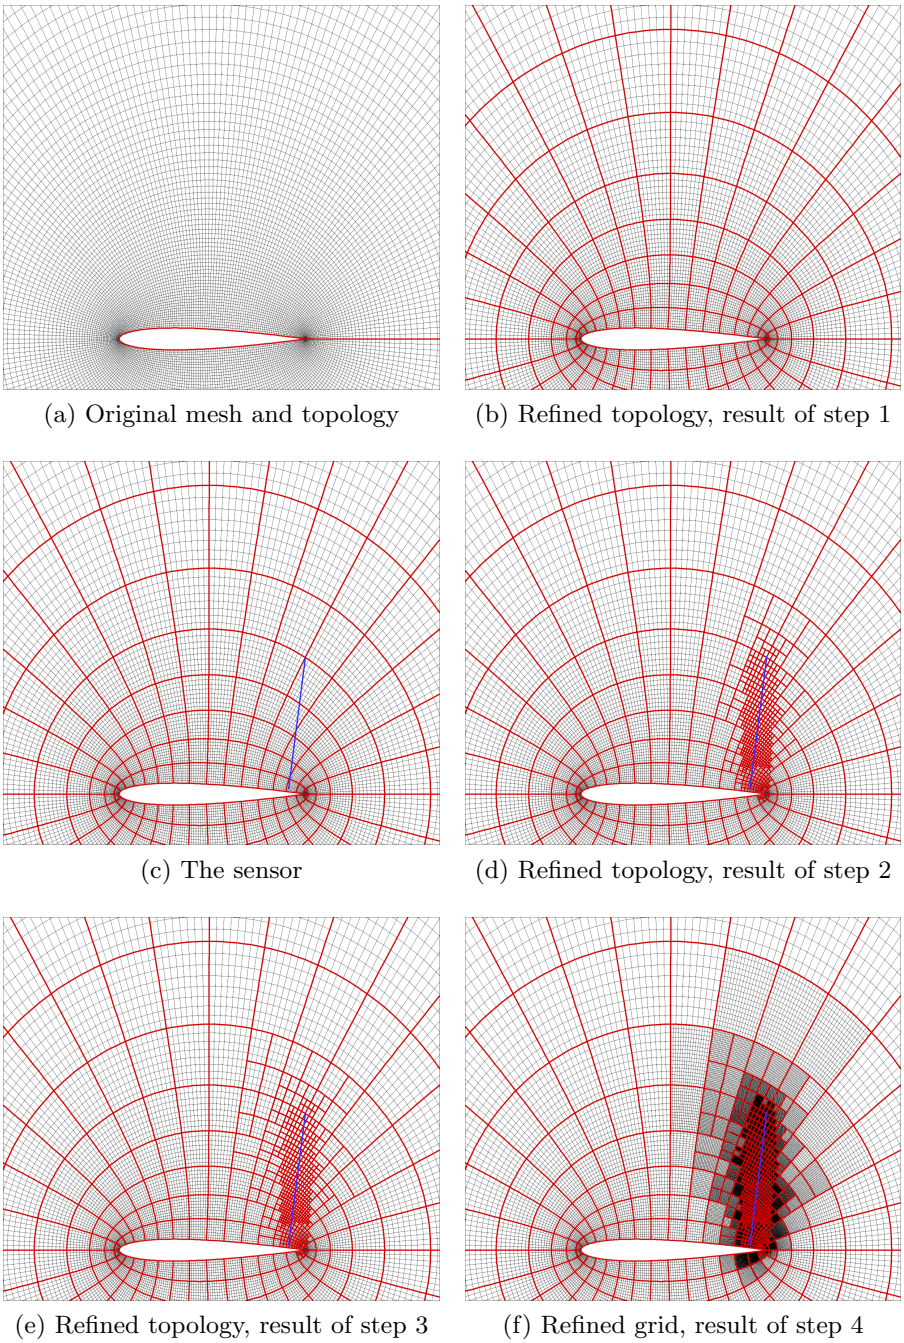

(a) Original mesh and topology

(b) Refined topology, result of step 1

(c) The sensor

(d) Refined topology, result of step 2

(e) Refined topology, result of step 3

(f) Refined grid, result of step 4

Fig. 1: Illustration of the grid generation algorithm. Block boundaries are shown in red; grid lines in black; the line segment used in the sensor is shown in blue.

mesh width in a cell over the target mesh width. As can be seen the ratio is less than one in the refinement region, demonstrating the correct functioning of the algorithm. The figure also shows that the algorithm is relatively efficient: the region where the target mesh width is attained is not that much bigger than the refinement region.

4 High-order method at partially continuous grids

4.1 Algorithm

ENSOLV employs a cell-centred finite-volume method. For the second-order scheme, a five-point stencil is used in each computational direction. For the fourth-order scheme, a seven-point stencil is used. To couple the flow computations in different blocks of the multi-block grid, dummy cells are used. Considering the size of the stencil, the second-order scheme requires two rows of dummy cells along each interface, and the fourth-order scheme requires three rows. The flow solution at the cell-centres of the dummy cells are determined by interpolating the conservative flow variables of the adjacent block. The interpolation algorithm is described for the case of a 2:1 ratio in mesh size, both along and normal to the interface.

A generic interpolation algorithm has been developed, applicable both to partially continuous and full discontinuous grids. Main requirement has been that the algorithm is robust and accurate for time-accurate computations, maintaining the properties of the high-order finite-volume method as much as possible. This was obtained by initially dropping the requirement that the scheme is fully conservative at the interface. Future extension towards a fully conservative scheme, however, is possible.

First, the second-order variant is described. In this case, the two rows of dummy-cell values are obtained by linear interpolation in computational space, which is second-order accurate. The interpolation is performed in each computational direction separately. A first option is to apply this interpolation only along the interface, as shown in figure 3, ignoring the jump in mesh size in the direction normal to the interface. This has as advantage that the symmetry of the discretization normal to the interface is maintained, requiring two rows of source cells on each side of the interface. Furthermore, this scheme is generic enough to be applied to a fully discontinuous grid (with appropriately defined interpolation coefficients). A second option is to interpolate the solution in the normal direction as well (after first interpolating along the interface), as shown in figure 4. Note that this scheme is now asymmetric, requiring four instead of two rows of source cells from the fine grid.

This generic algorithm can be readily extended to fourth-order accuracy by increasing the order of the interpolation from second to fourth. This entails a larger interpolation stencil. Furthermore, three rows of dummy cells are needed for the fourth-order finite-volume method. Again, the first option consists of interpolating along the interface only, as shown in figure 5, maintaining the

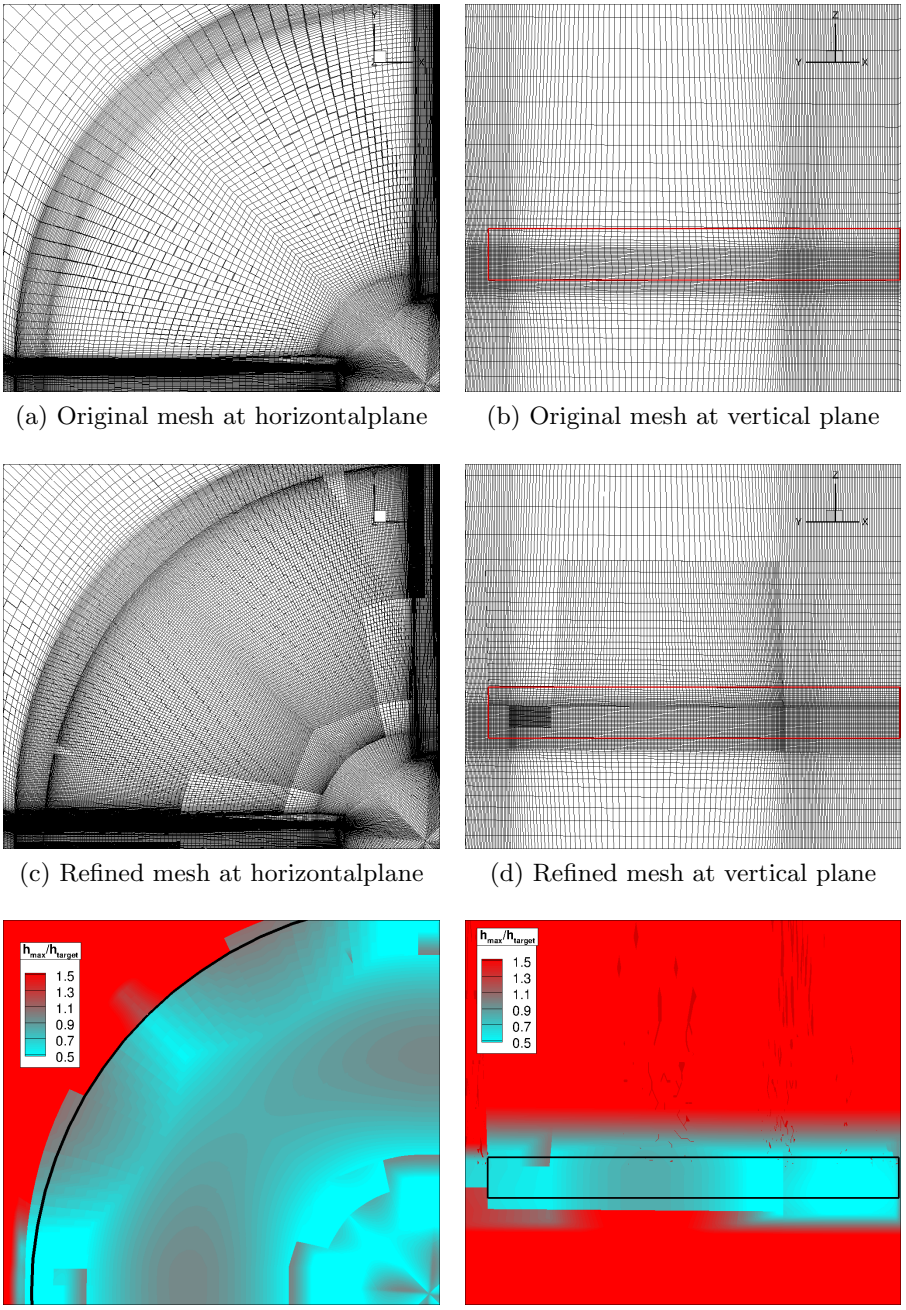

(a) Original mesh at horizontalplane (b) Original mesh at vertical plane

(c) Refined mesh at horizontalplane (d) Refined mesh at vertical plane

(e) Mest width ratio at horizontal plane (f) Mest width ratio at vertical plane

Fig. 2: Original and refined mesh for the HART II rotor at a horizontal and a vertical plane through the hub; the border of the refinement region is shown in red. Ratio of the maximum mesh width in a cell over the target mesh width in the same two planes; the border of the refinement region is shown in black.

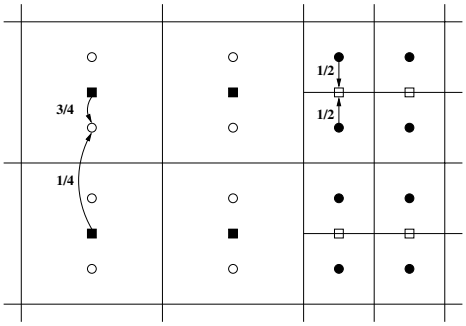

Fig. 3: Stencil to determine dummy-cell values using second-order interpolation along interface only (solid symbols: source points (cell centres); open symbols: target points (dummy cells))

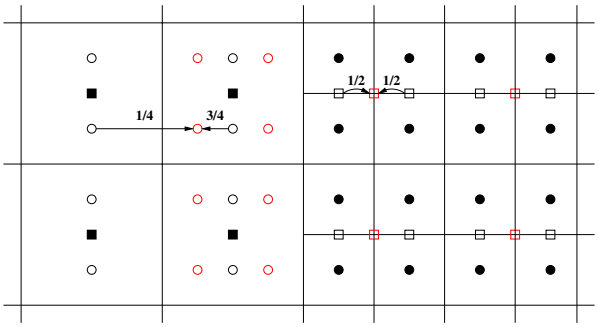

Fig. 4: Stencil to determine dummy-cell values using second-order interpolation both along and normal to interface (solid symbols: source points (cell centres); open black symbols: intermediate points after interpolation along interface; open red symbols: target points (dummy cells))

symmetry of the scheme normal to the interface. Now, three rows of source cells are required on each side of the interface. The second option, consisting of subsequently interpolating in normal direction as well, is shown in figure 6. This algorithm requires seven rows of source cells of the fine grid. Thus, the fine-grid block is required to have a minimum dimension of seven cells.

The second option is more accurate than the first, because it takes the jump in mesh size in normal direction into account. It is less robust, however, due to the loss of symmetry. This reduced robustness is felt the strongest when a multi-grid scheme is applied to solve the non-linear system of equations (per time step). Therefore, in the multi-grid scheme, the first option is applied on the coarse grid levels and the second option on the finest level only. This has as additional advantage that on the coarse levels, the blocks have a minimum grid size of only three cells instead of seven cells.

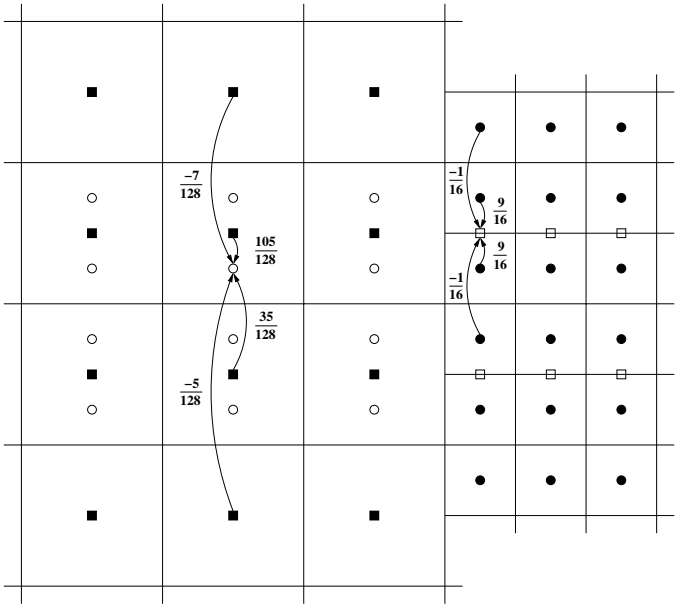

Fig. 5: Stencil to determine dummy-cell values using fourth-order interpolation along interface only (solid symbols: source points (cell centres); open symbols: target points (dummy cells))

4.2 Verification

In order to test the capability of the high-order finite-volume scheme to accurately capture vortices without significant dissipation or dispersion on locally refined grids, the convection of a 2D isentropic vortex in a uniform flow is considered. This test case has been used to verify to high accuracy of the low-dispersion symmetry-preserving fourth-order finite-volume method on curvilinear grids [7].

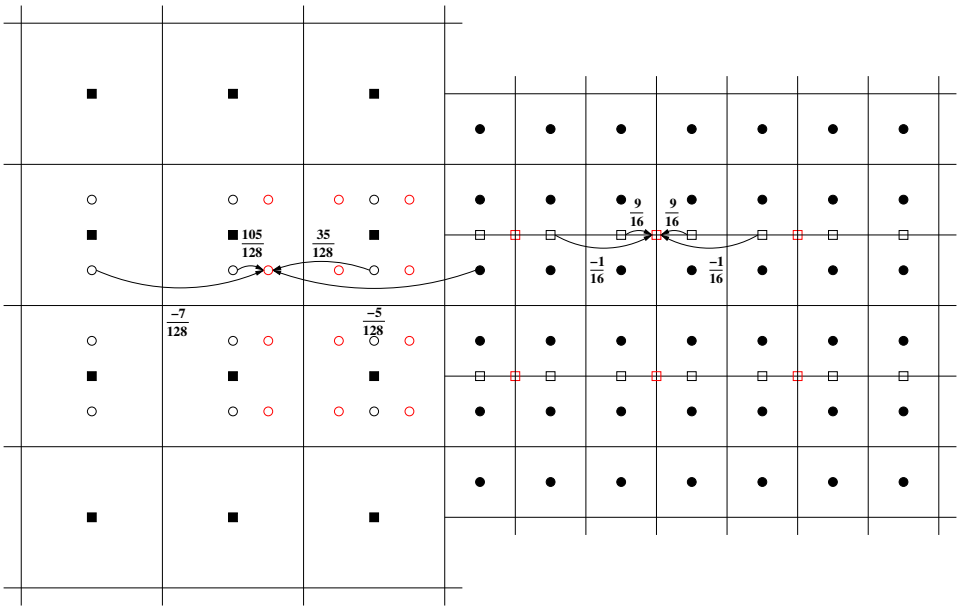

Fig. 6: Stencil to determine dummy-cell values using fourth-order interpolation both along and normal to interface (solid symbols: source points (cell centres); open black symbols: intermediate points after interpolation along interface; open red symbols: target points (dummy cells))

Computations are performed on a 10×10 multi-block grid with local grid refinement in a checker-board pattern, as illustrated in figure 7. This figure also shows the initial location of the vortex; the final location is just before the downstream boundary. The number of cells in the coarse blocks equals 20×20 and in the fine blocks 40×40. The equations are integrated in time by a low-storage 4-stage Runge–Kutta scheme.

The interpolation algorithm at partially continuous grid interfaces is verified for the fourth-order finite-volume method. Figure 8 shows the final temperature field of the vortex with three approaches: using second-order interpolation along the interface only, using fourth-order interpolation along the interface only, and using fourth-order interpolation both along and normal to the interface. First of all, for all results the vortex is located in the correct position and only small differences with the analytic solution can be seen. This was expected, as the fourth-order finite-volume method was particularly designed to have low dispersion and dissipation in order to accurately capture vortices (and waves). The differences between the first two approaches (with interpolation only along the interface) are only small, both showing some small disturbances compared to the analytic solution and a weak increase of the temperature in the vortex core. The highest accuracy is obtained when the fourth-order interpolation is applied both along and normal to the interface. In that case, the solution is practically identical to the analytic solution (the differences in the vortex core essentially due

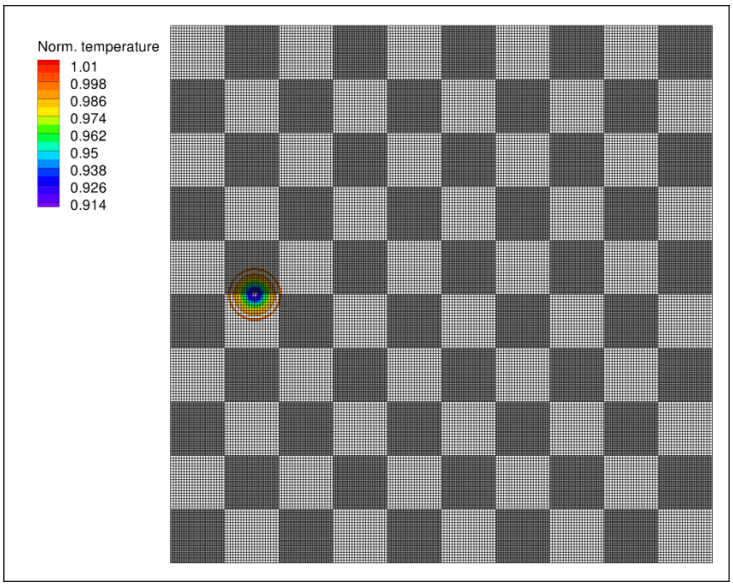

Fig. 7: Initial temperature field of strong isentropic vortex on locally refined grid

to the limited accuracy of the plotting program). Note that these computations are stable with a very low level of sixth-order artificial diffusion ($k^{(6)} = \frac{1}{8}$).

5 Applications

5.1 HART II rotor

As a first application, simulations are performed for an isolated four-bladed rotor in one of the flow conditions of the HART-II experiment. The HART-II experiment is described in detail in Van der Wall et al.[17]. The chosen flow condition is a slow descent flight where multiple BVI events take place during a rotor revolution. The rotor radius of the windtunnel model is 2 meters. Rotational frequency is 109.12 rad/s. Forward speed of the rotor is 32.9 m/s. Effective shaft angle is 4.5 degrees (tilted backward).

The turbulent flow is modeled using the Reynolds-averaged Navier-Stokes equations with the TNT k-ω turbulence model [6]. All simulations have been run using the fourth-order accurate finite volume scheme. Simulations have been performed on the original mesh of 13 million elements, and the locally refined mesh of 26 million elements. The discretization algorithm on locally refined meshes is described above, with second order interpolation at the partially continuous grid interfaces. The time step corresponds to 0.5° azimuth for all simulations.

For the simulation on the original mesh, rigid blades are considered. For the simulation on the refined mesh, both rigid and flexible blades are considered. In all simulations the rotor is trimmed to the experimental time-averaged forces.

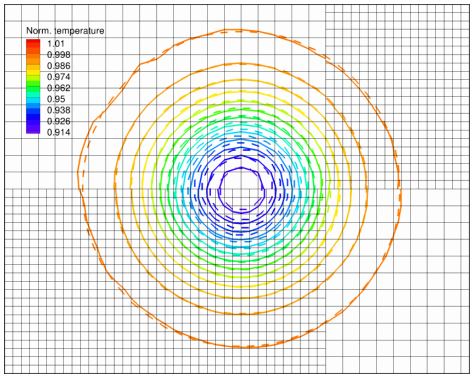

(a) Second-order interpolation along interface only

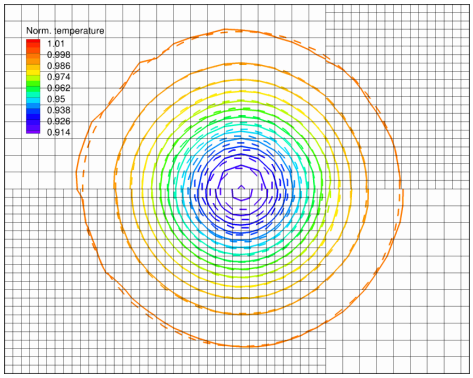

(b) Fourth-order interpolation along interface only

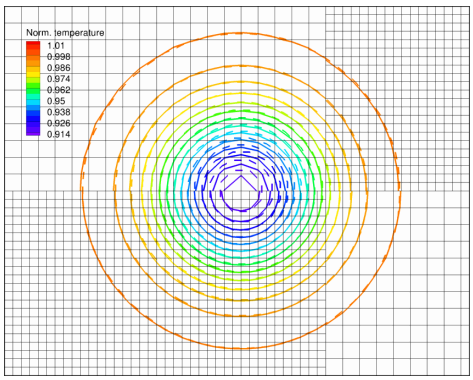

(c) Fourth-order interpolation along and normal to interface

Fig. 8: Convection of strong isentropic vortex on locally refined grid with fourth-order finite-volume method: final temperature field (dashed lines: analytic solution; solid lines: computed solution)

The aeroelastic simulation setup is described in detail in the chapter on aeroelastic test cases in this book [11], Section ??. In the current paper the vortex resolution of the two simulations is compared.

Figure 9a shows the instantaneous iso-contour of the Q-criterion for the original mesh. The Q-criterion distinguishes flow under shear stress from rotational flow, and is used here to focus on the tip vortices rather than on the boundary layer vorticity. The value of Q is equal to $0.72(U_\infty/R)^2$; the vorticity is scaled with $2U_\infty/R$. Figure 9b shows the instantaneous iso-contour of the Q-criterion for the refined mesh at the same values of Q and vorticity. The increase in vortex resolution is evident. Hence, the local grid refinement performs as expected.

Continuing with the simulation using flexible blades, the instantaneous iso-contour of the Q-criterion is shown on the refined mesh for the trimmed solution in Figure 10. Compared to the simulation on the original grid, the lift of the blades on the advancing side (azimuth angle between 90 and 180 degrees) is smaller. Hence the strength of the tip vortex shed in this quadrant is less. In order to visualise the vortex system the value of Q is reduced by a factor of two with respect to the previous figures.

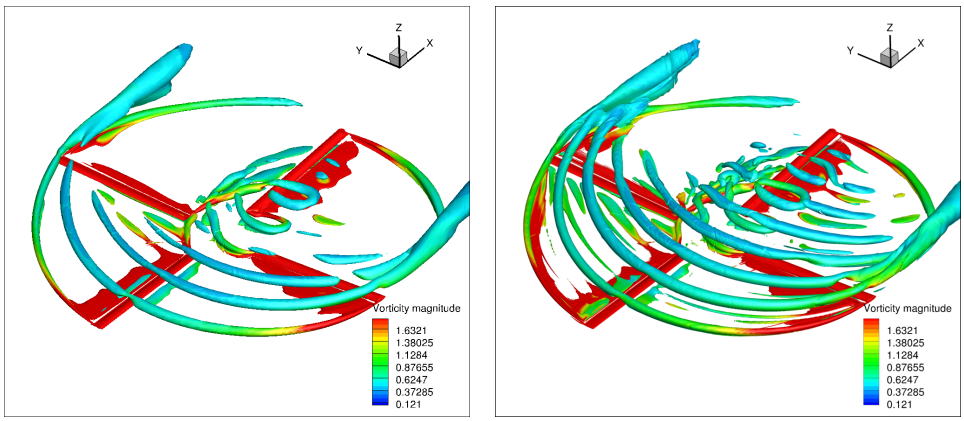

(a) Original mesh (b) Refined mesh

Fig. 9: Instantaneous iso-contour of the Q-criterion (blade at 0° azimuth) coloured with vorticity magnitude. Advancing side of the rotor is on the top-left side of the figures (positive y), retreating on the bottom-right side (negative y).

5.2 Results for bump in square duct

As a second application, the turbulent, separated flow over a rounded bump in a square duct is considered. Experiments have been performed within the EU-project DESider by ONERA [1] in a hydraulic channel. The duct has a height of 0.3 m, a width of 0.5 m, and a length of 2.367 m. The bump has a height of 0.138 m, starts at the inflow plane at $x = 0.367$ m, and ends at $x = 0$ m. At

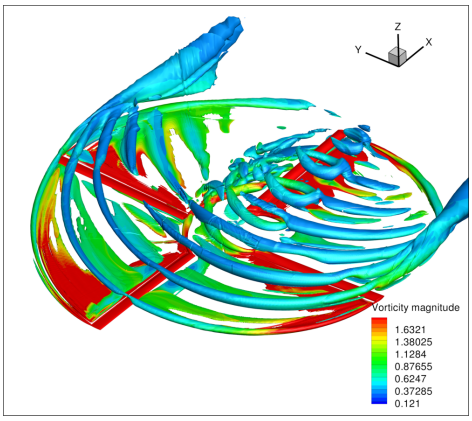

Fig. 10: Instantaneous iso-contour of the Q-criterion (blade at 0° azimuth) coloured with vorticity magnitude for the refined mesh and flexible blades. Advancing side of the rotor is on the top-left side of the figure (positive y), retreating on the bottom-right side (negative y).

the entrance, velocity profiles from the experiment are prescribed, which have a centre velocity of approximately 7 m/s. Furthermore, water with a density of 997 kg/m^3 and a dynamic viscosity of $0.89{\cdot}10^3$ Pa·s is considered. As a compressible flow solver is used for solving this incompressible flow, an inflow Mach number of $M = 0.1$ is chosen.

Computations have been performed on a baseline grid of $142 \times 60 \times 76 =$ 647,520 cells and on a grid with LGR (local grid refinement) with 355,560 cells, see figure 11. To be more precise, the latter grid actually contains local grid coarsening: the grid is coarsened by a factor 2 in all directions outside the flow separation region where the computation is in LES mode (i.e., outside the red surface in figure 11b), saving about 45% of grid points. A time step has been used equal to $\Delta t = 1.89{\cdot}10^{-4}$ s. A total of 50,000 time steps has been taken and flow statistics have been computed over the last 45,000 steps.

Figure 12 gives an instantaneous impression of the flow in terms of the Q-criterion, revealing the vortical structures in the separated flow region. For the grid with LGR, the figure shows that the separated flow is mostly contained within the fine-grid region, although occasionally a vortex may escape, in particular in the centre of the tunnel.

The mean pressure coefficient along the bottom wall is shown for three stations in figure 13. Compared to the experiment, the separation region is too small: the pressure recovers upstream of the experiment. This finding is consistent with the DES results from the DESider project [5]. More important here is the comparison between the computations on the baseline (BSL) grid and the grid with LGR. They are found to lie close together, with the LGR computation having a slightly smaller separation region.

The resolved turbulence level is compared on both grids in terms of the RMS value of the x-component of velocity (figure 15). The level of resolved turbulence

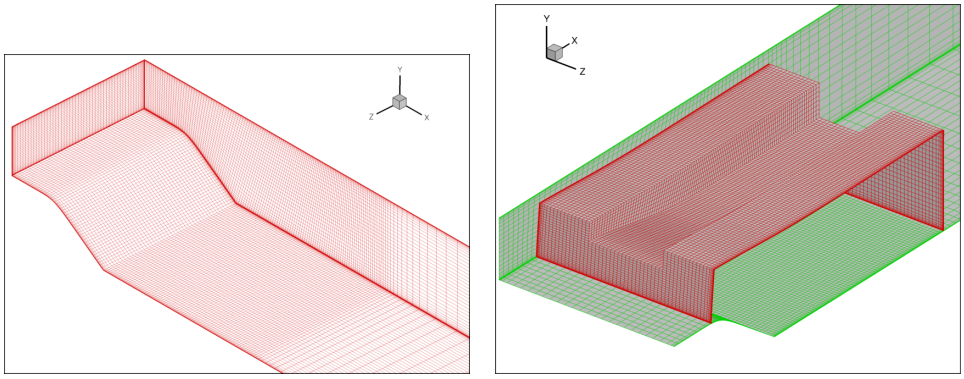

(a) Baseline grid ($142 \times 60 \times 76$ cells)

(b) Grid with LGR: coarsened outside red surface, i.e., outside LES region

Fig. 11: Baseline grid and grid with LGR for bump in square duct

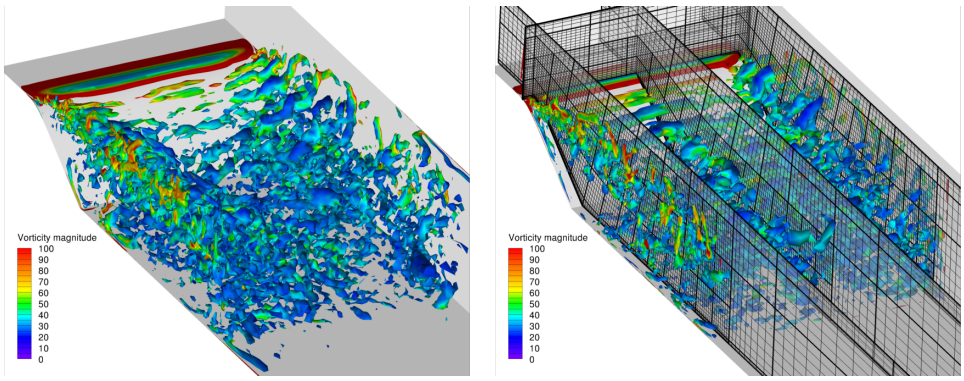

(a) Baseline grid

(b) Grid with LGR

Fig. 12: Instantaneous isocontours of Q-criterion, coloured with vorticity magnitude

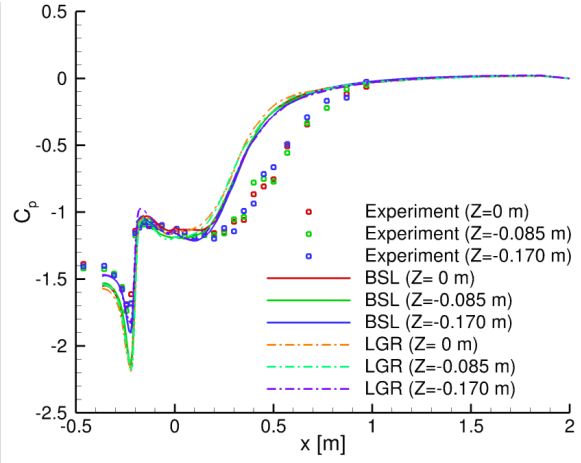

Fig. 13: Mean pressure coefficient along bottom wall

in the shear layer is slightly lower on the grid with LGR. This may be related to the slightly smaller separation region.

A possible explanation for the differences between the solutions on the baseline grid and the grid with LGR is that the refined grid region has been defined too narrow around the separated flow. Although for the time-averaged flow solution, the separation region is well within the refined region, this is not the case instantaneously. In particular, at the downstream boundary (large) vortical structure leave the refined region and are then suppressed on the coarser grid. Occasionally, this may also happen at the downstream part of the upper boundary of the refined region. This may reduce the level of resolved turbulence. Furthermore, at the upstream boundary, the incoming, attached boundary layer is also represented on a coarser grid and this may affect the initial development of the shear layer after separation.

5.3 Results for delta wing at high angle of attack

As a third application, the flow around a delta wing with a sharp leading edge at high angle of attack and high Reynolds number is considered. This flow is characterized by the main vortex developing above the wing. The vortex is formed as the shear layer emanating from the leading edge rolls up, starting immediately at the apex. At high Reynolds numbers, the shear layer rapidly becomes unstable and a turbulent vortex is formed. At a sufficiently high angle of attack, the vortex breaks down: the high axial velocity in the vortex core drops rapidly to a value close to zero.

The NASA delta wing geometry of Chu and Luckring [2] is considered, for which experiments that include measurements of velocity fluctuations have been performed by Furman and Breitsamter [3, 4]. Vortex breakdown occurs at the

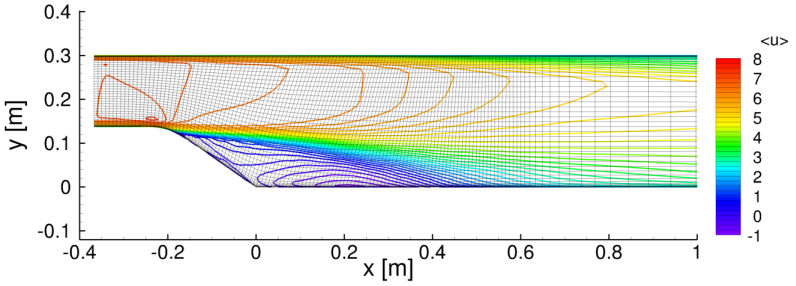

(a) Baseline grid

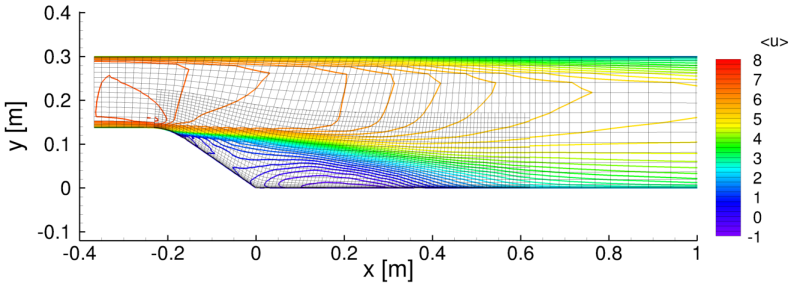

(b) Grid with LGR

Fig. 14: Mean x-component of velocity in plane $z = 0$

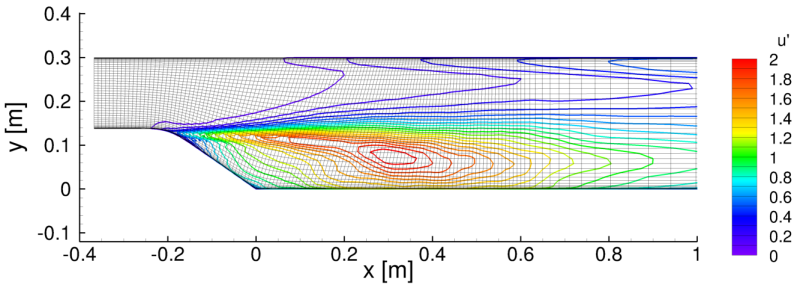

(a) Baseline grid

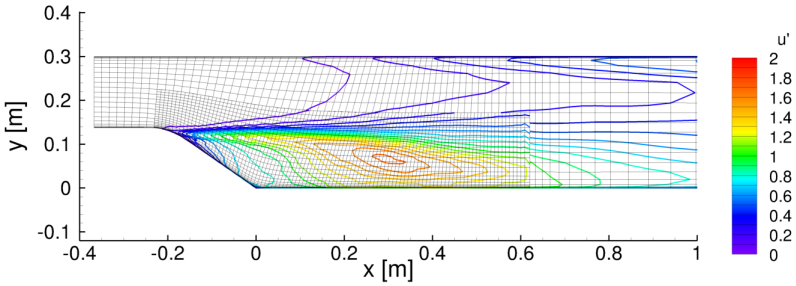

(b) Grid with LGR

Fig. 15: RMS of x-component of velocity in plane $z = 0$

flow conditions $M = 0.07$, $Re_{\text{mac}} = 1{\cdot}10^6$, and $\alpha = 23°$ (with the Reynolds number based on the mean aerodynamic chord c_{mac}).

A baseline multi-block structured grid has been generated (within the ATAAC project), consisting of 22 blocks and 6.3 million grid cells (figure 16a). The grid has a conical structure over a large part of the wing: the grid covering the main vortex is essentially isotropic at each chordwise station (outside the boundary layer) and the mesh width grows in all directions (including the streamwise direction) together with the main vortex, going from approximately $0.003c_{\text{mac}}$ to $0.011c_{\text{mac}}$. In other words, the grid resolution relative to the main vortex is kept constant. Only in a small region near the apex, the conical structure is not fully maintained, avoiding a grid singularity. The far-field boundary is located at three root chord lengths from the wing. To study grid sensitivity, also a finer grid with the mesh width reduced by a factor 2/3 in all directions (21.4 million grid cells) has been generated.

As for the previous test case, an LGR grid with local coarsening is created (figure 16b), containing 3.7 million grid cells (reduction of 41%). The refined grid region contains the complete wing, including the boundary layers, as well as the main vortices above the wing.

X-LES computations have been performed with a time step of $3.75{\cdot}10^{-4}$ CTU (convective time units: c_{mac}/u_∞) on the baseline and LGR grids and a time step of $2.5{\cdot}10^{-4}$ CTU on the fine grid. The flow statistics have been obtained by averaging over 21 CTU for the baseline grid, 18 CTU for the LGR grid, and 8 CTU for the fine grid (after transients of about 6 CTU).

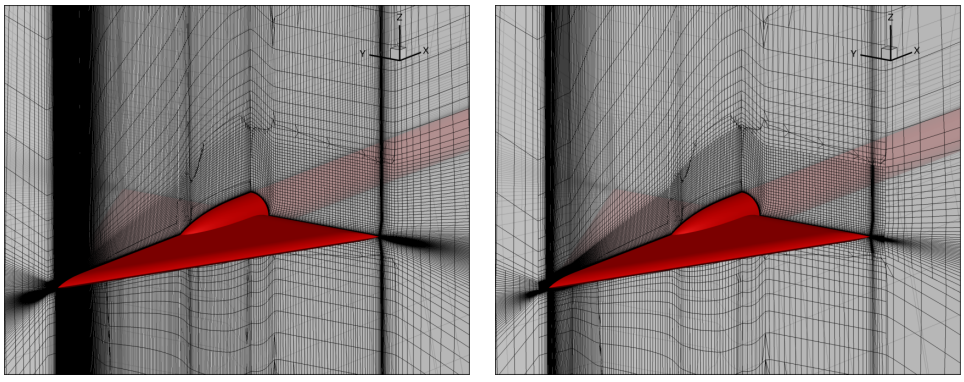

(a) Baseline grid (b) Grid with LGR

Fig. 16: Grids for VFE2 delta wing

An impression of the instantaneous results is given in figure 17 in terms of isocontours of the Q-criterion. Similar fine-scale turbulent structures can be observed on the baseline and LGR grids.

The mean and RMS values of the pressure coefficient are given at two representative chordwise station in figures 18 and 19. For the mean values, there are

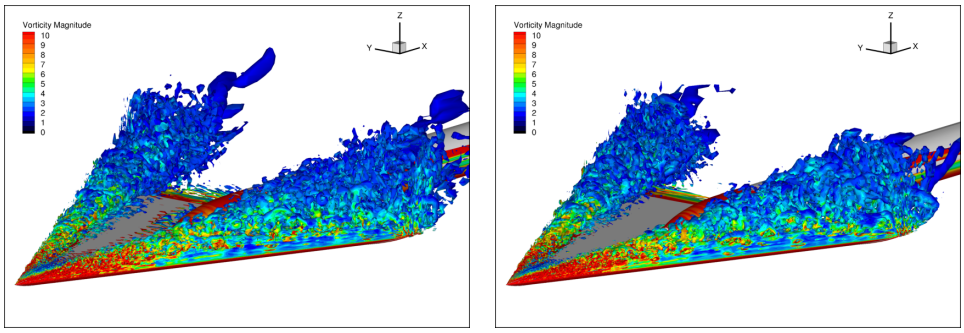

(a) Baseline grid (b) Grid with LGR

Fig. 17: Instantaneous isocontours of Q-criterion, coloured with vorticity magnitude

no differences visible between the baseline and LGR grids, while there are only small differences with the fine grid. Larger, but still relatively small, differences are seen for the RMS values.

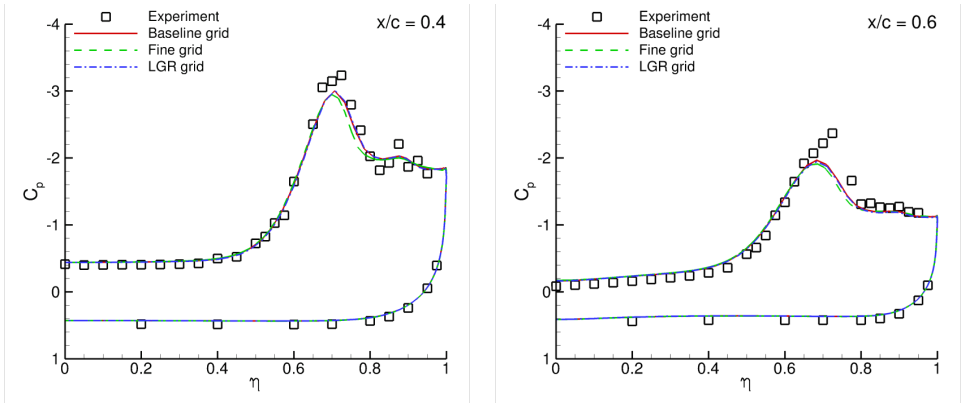

(a) $x/c = 0.40$ (b) $x/c = 0.60$

Fig. 18: Mean pressure coefficient at five chordwise stations

Finally, a comparison of the level of resolved turbulence is shown in figure 20 in terms of the resolved turbulent kinetic energy at a single station (40% root chord). Practically the same levels of resolved turbulence are obtained on the baseline and LGR grids, and also similar levels are seen on the fine grid and in the experiment.

Thus, the impact of coarsening the grid away from the region of interest is found to be small for this case; smaller than for the previous case. The main reason is that it is easier to fully encompass the region of interest (i.e., the

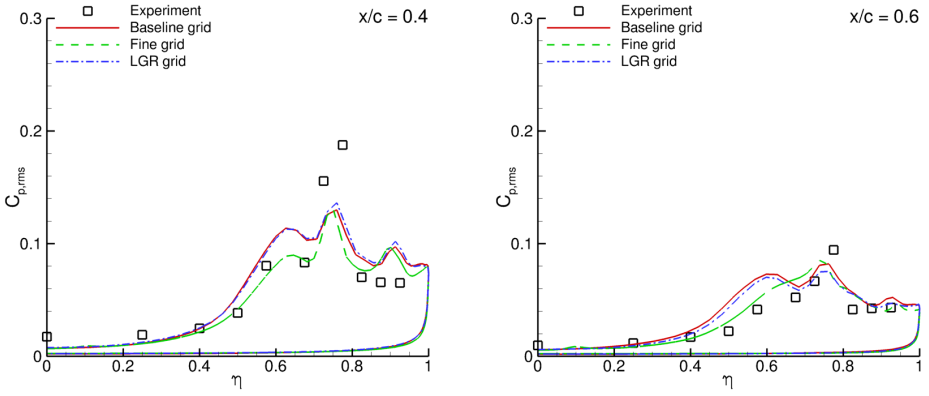

(a) $x/c = 0.40$ (b) $x/c = 0.60$

Fig. 19: RMS of pressure coefficient at five chordwise stations

separated flow captured with LES) in the refined grid region for an external flow than for the internal flow of the previous case.

6 Conclusions

An automatic topology-based grid refinement algorithm has been developed and demonstrated with the specific aim of obtaining block-structured meshes with uniform resolution in user-specified subdomains.

A generic algorithm has been developed to deal with partially continuous grids in a fourth-order finite-volume method. The stability and accuracy of the fourth-order finite-volume method have been verified for the convection of an isentropic vortex on a locally refined grid. The highest accuracy is obtained when the dummy-cell values are computed by fourth-order interpolation both along and normal to the partially continuous grid interfaces.

The simulation of an isolated helicopter rotor in forward flight clearly demonstrated improved vortex resolution on the locally refined mesh. This is essential for capturing blade-vortex interaction.

X-LES computations were performed for two tests cases on grids using local grid refinement (LGR). Compared to the baseline grids, the LGR grids were coarsened away from the separated flow regions that were captured with LES, saving 40 to 45% of the grid points. For the bump in a square duct, the results on the LGR grid showed a slightly smaller separation region than on the baseline grid. Possibly, the refined grid region was chosen too tightly around the region of interest. For the delta wing, no significant differences were found between the LGR and baseline grids. Thus, local grid refinement can be used to reduce the computational costs of X-LES computations, provided the grid is coarsened well outside the region of interest, which is easier to achieve for external flows than for internal flows.

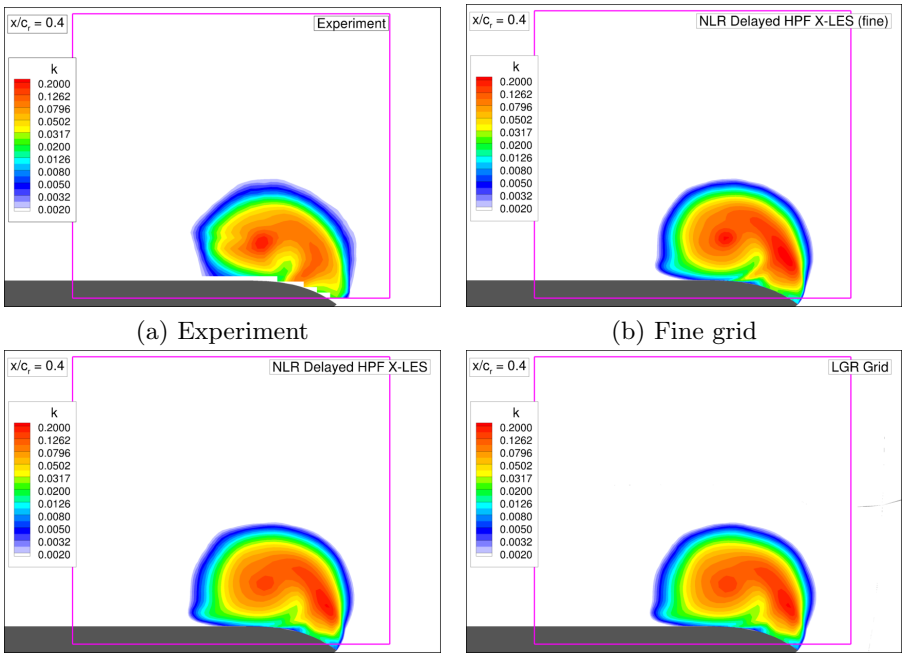

(a) Experiment (b) Fine grid

(c) Baseline grid (d) Grid with LGR

Fig. 20: Resolved turbulent kinetic energy at chordwise station $x/c = 0.4$

References

1. Aupoix, B.: The DESider bump experiment for hybrid modelling validation. In: EWA Workshop. Prague (5–6 June 2007)
2. Chu, J., Luckring, J.M.: Experimental surface pressure data obtained on 65° delta wing across Reynolds number and Mach number ranges. TM 4645, NASA (1996)
3. Furman, A., Breitsamter, C.: Turbulent and unsteady flow characteristics of delta wing vortex systems. In: 46$^{\text{th}}$ AIAA Aerospace Sciences Meeting and Exhibit. Reno (NV), USA (7–10 January 2008), AIAA 2008-381
4. Furman, A., Breitsamter, C.: Experimental investigations on the VFE-2 configuration at TU Munich. In: Understanding and Modeling Vortical Flows to Improve the Technology Readiness Level for Military Aircraft, chap. 21. NATO RTO (2009), RTO-TR-AVT-113
5. Haase, W., Braza, M., Revell, A. (eds.): DESider — A European Effort on Hybrid RANS–LES Modelling, Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 103. Springer (2007)
6. Kok, J.C.: Resolving the dependence on freestream values for the k–ω turbulence model. AIAA Journal 38(7), 1292–1294 (2000)
7. Kok, J.C.: A high-order low-dispersion symmetry-preserving finite-volume method for compressible flow on curvilinear grids. Journal of Computational Physics 228, 6811–6832 (2009), NLR-TP-2008-775
8. Kok, J.C., Dol, H.S., Oskam, B., van der Ven, H.: Extra-large eddy simulation of massively separated flows. In: 42$^{\text{nd}}$ AIAA Aerospace Sciences Meeting. Reno, NV (5–8 January 2004), AIAA paper 2004-264

9. Kok, J.C., van der Ven, H.: Destabilizing free shear layers in X-LES using a stochastic subgrid-scale model. In: Peng, S.H., Doerffer, P., Haase, W. (eds.) Progress in Hybrid RANS–LES Modelling, Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 111, pp. 179–189. Springer (2009), NLR-TP-2009-327
10. Kok, J.C., van der Ven, H.: Capturing free shear layers in hybrid RANS–LES simulations of separated flow. In: Third Symposium 'Simulation of Wing and Nacelle Stall'. Braunschweig, Germany (21–22 June 2012), http://hdl.handle.net/10921/914, NLR-TP-2012-333
11. Morzynski, M., Nowak, M., Kok, J.C., van der Ven, H.: Assessment of high-order methods – aeroelastic test cases. In: Kroll, N., Hirsch, C., Bassi, F., Johnston, C., Hillewaert, K. (eds.) IDIHOM – Industrialisation of High-Order Methods – A Top Down Approach, Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 1xx. Springer (2014)
12. Spalart, P.R.: Detached-eddy simulation. Annual Review of Fluid Mechanics 41, 181–202 (January 2009)
13. Tam, C.K.W., Webb, J.C.: Dispersion-relation-preserving finite difference schemes for computational acoustics. Journal of Computational Physics 107, 262–281 (1993)
14. Travin, A.K., Shur, M.L., Spalart, P.R., Strelets, M.K.: Improvement of delayed detached-eddy simulation for LES with wall modelling. In: Wesseling, P., Oñate, E., Périaux, J. (eds.) ECCOMAS CFD 2006. Egmond aan Zee, The Netherlands (5–8 September 2006)
15. van der Ven, H., Boelens, O.J., Kok, J.C., Prananta, B.B., van Rooij, M.P.C.: Simulation of a rotor in forward flight using topology-based refinement of multiblock structured meshes. In: 38th European Rotorcraft Forum. Amsterdam, the Netherlands (September 4–6 2012), NLR-TP-2012-485
16. Verstappen, R.W.C.P., Veldman, A.E.P.: Symmetry-preserving discretization of turbulent flow. Journal of Computational Physics 187(1), 343–368 (2003)
17. van der Wall, B.G.: 2nd HCC aeroacoustics rotor test (HART II) Part I: Test documentation. Tech. Rep. IB 111-2003/19, DLR (2003)