LAVA WMLES Results for the 5th High-Lift Prediction Workshop

Emre Sozer*, Michael F. Barad†, Jeffrey A. Housman‡, Francois Cadieux§, Victor C.B. Sousa¶, Ercan Dumlupinar‖, Kumar Saurabh**, Kiran Ravikumar†† and Jared Duensing‡‡
NASA Ames Research Center, Moffett Field, CA, 94035

Wall-Modeled Large Eddy Simulation (WMLES) studies for the 5th AIAA High-Lift Prediction Workshop (HLPW5), conducted using the Launch Ascent and Vehicle Aerodynamics (LAVA) unstructured solver, are presented. The numerical methods, modeling approaches, and meshing strategies are described. A subset of the results for workshop cases 2 and 3 is discussed, with comparisons to available experimental data and insights gained through grid sensitivity studies. Particular attention was focused on a WMLES modeling improvement in order to account for the leading edge regions, where the flow is typically transitional. A sensor-based approach for dynamically switching between laminar and turbulent wall stresses is explored. Computational resource requirements and performance across CPU and GPU systems are presented, highlighting the potential efficiency gains of GPU-based simulations.

I. Introduction

Scale-resolving simulation methods have become instrumental in accurately predicting $C_{L,max}$ for aircraft high-lift configurations [1–4]. Motivated by the CFD Vision 2030 Report [5], NASA has placed considerable emphasis on maturing high-fidelity Computational Fluid Dynamics (CFD) methods, with a particular focus on improving the accuracy, efficiency, and robustness of scale-resolving methods like Wall-Modeled Large-Eddy Simulations (WMLES) [6]. Consequently, routine modeling of the fluid dynamics that govern complex aerospace configurations has become a reality in recent years, positioning CFD capabilities to meet the overarching goal of enabling aircraft Certification and Qualification by Analysis (CQbA) [7]. WMLES can yield highly accurate aerodynamic predictions for flow regimes where other traditional CFD methods such as Reynolds-Averaged Navier-Stokes (RANS) tend to fall short, as evidenced through the vast spread of predictions computed for past AIAA High Lift Prediction Workshops [8–11].

The Launch, Ascent, and Vehicle Aerodynamics (LAVA) Framework [12] has been a regular contributor to the High Lift Prediction Workshop series. In support of the 3rd High-Lift Prediction Workshop, RANS-based CFD capabilities within the LAVA Curvilinear and LAVA Unstructured solvers were used to primarily assess turbulence model sensitivities for the high-lift configurations in question [13]. The considerable spread in predictions near aircraft $C_{L,max}$ and under-prediction of lift at near-stall conditions highlighted the extent to which RANS struggles to predict flow physics dominated by separation. The 4th High-Lift Prediction Workshop served as an opportunity to assess modeling sensitivities using a more comprehensive set of approaches. This involved thorough evaluations of high-lift configurations using RANS- [14], HRLES- [15], and WMLES-based [16] approaches, with rigorous sensitivity studies conducted to identify best practices for each of these three methods. Core findings regarding each method's performance for high-lift applications are summarized in [2]. The potential of stress-based WMLES approaches to accurately predict integrated aerodynamic coefficients near $C_{L,max}$ while side-stepping the issues of RANS-based methods is particularly attractive.

However, many challenges related to applying WMLES to high-lift aircraft configurations still persisted. Foremost among them was generating an appropriate computational mesh for this complex vehicle configuration. Another challenge was how to treat thin, attached and potentially laminar boundary layers near the leading edges of various

*Computational Aerosciences Branch, AIAA Member, emre.sozer@nasa.gov
†Computational Aerosciences Branch, Senior AIAA Member, michael.f.barad@nasa.gov
‡Computational Aerosciences Branch, Senior AIAA Member, jeffrey.a.housman@nasa.gov
§Computational Aerosciences Branch, AIAA Member, francois.cadieux@nasa.gov
¶Science and Technology Corporation, AIAA Member, victor.decarvalhobrittosousa@nasa.gov
‖Computational Aerosciences Branch, Senior AIAA Member, ercan.dumlupinar@nasa.gov
**Science and Technology Corporation, AIAA Member, kumar.saurabh@nasa.gov
††Science and Technology Corporation, kiran.ravikumar@nasa.gov
‡‡Computational Aerosciences Branch, AIAA Member, jared.c.duensing@nasa.gov

airplane components (slat, wing, flaps, etc). Key improvements were therefore critical for accurate and efficient simulations of the High Lift Common Research Model (CRM-HL) variants in the 5th High Lift Prediction Workshop (HLPW-5). The LAVA team addressed the first challenge by developing a Voronoi unstructured meshing methodology and a scalable mesh generator tool [17]. We then implemented a sensor-based approach to detect and appropriately treat attached laminar flow regions to address the second challenge. Other improvements such as the development of a low-dissipation, robust and efficient numerical scheme and porting the software to GPU were key enablers.

The HLPW-5 encompassed three test cases, which served as platforms to advance the state of the art toward the objectives of the workshop series [18]. While all three test cases and their required sub-cases were studied by our group during the workshop, we focus here on Cases 2.4 and 3. Both cases are slight variations of the CRM-HL model with leading and trailing edge high-lift devices deployed. Case 2.4 is supported by experimental data collected in the F1 low-speed wind tunnel [19] at a moderate Reynolds number of 5.9×10^6. Case 3 focuses on Reynolds number variations as high as 30×10^6.

The geometry definition for Case 2 is based on the as-designed ONERA 1/19.5-scale CRM-HL model tested in the F1 low-speed wind tunnel, while the definition for Case 3 represents the NASA 5.2% CRM-HL model tested in the National Transonic Facility (NTF) [20] at NASA Langley Research Center. The experimental data for the latter test has not yet been made available at the time of this writing. Some subtle geometric differences are observable across the two configurations, but key components are largely similar.

Early in the workshop, it became clear that the leading edge regions at high angles of attack, particularly for clean wing (without slat) configurations, were where code-to-code comparisons and numerical sensitivity tests tended to deviate the most. This highlighted one of the shortcomings of the WMLES method, especially in lower Reynolds number conditions. The typical WMLES approach imposes a wall stress calculated via a wall function with an underlying assumption of a fully turbulent boundary layer. For the high-lift problem, this assumption may break down near leading edges where the flow may be laminar or transitional.

In these regions where there is little to no resolved turbulent content, the application of fully turbulent wall stress may be excessive and may lead to premature loss of momentum and induce flow separation. Following these observations, our group adopted a sensor-based approach to detect resolved turbulent content at the wall and switch between turbulent wall stress (via wall model) or a no-slip boundary condition. This approach significantly improved our $C_{L,max}$ predictions (compared to experimental data for Case 2) as well as grid convergence characteristics. The method is based on work by Larsson et al. [21] with a few of our changes as described in detail in Section II.E. This simple approach works quite well with little added computational cost. The leading edge transitional behavior varies significantly across CFD codes since the onset of turbulent content development is dependent on the grid, numerical scheme, and the wall treatment. Sensitivity of this wall turbulence sensor to the unstructured mesh resolution and the sensor threshold value is explored in Sections IV.A.1 and IV.A.2.

In Section II, the numerical scheme used in the LAVA unstructured WMLES solver is presented in detail. Our Voronoi meshing strategy for the CRM-HL geometry, along with the resulting mesh statistics, is outlined in Section III. Section IV provides detailed discussions on a subset of results that were submitted to the workshop, presents various focused sensitivity studies, and provides comparisons to the experiments where available. Section IV.C provides an overview of the computational cost of the simulations on both CPU and GPU-based systems. Finally, Section V summarizes the core conclusions from this work and provides a description of the future research directions we intend to follow.

II. Numerical Scheme

The Navier-Stokes equations are solved using a 2nd order finite volume formulation which can be discretely expressed per control volume (computational cell) as:

$$J\frac{\Delta Q}{\Delta t} = \frac{1}{V}\sum_{f=1}^{N_f}\left(-F_{\mathrm{C}} + F_{\mathrm{D}}\right)A_f \tag{1}$$

where $Q = (p, u, v, w, T)$ is the primitive variable vector stored at cell centers, $W = (\rho, \rho u, \rho v, \rho w, \rho E)$ is the conservative variable vector, $J = \partial W/\partial Q$ is the conservative to primitive transformation Jacobian, V is the volume of the cell, f represents a cell face, and N_f is the total number of faces in a cell. Additionally, the convective and diffusive flux across f are described as,

$$F_{\mathrm{C}} = \begin{bmatrix} \rho(\mathbf{u}\cdot\mathbf{n}) \\ \rho\mathbf{u}(\mathbf{u}\cdot\mathbf{n}) + p\mathbf{n} \\ (\rho E + p)(\mathbf{u}\cdot\mathbf{n}) \end{bmatrix} \quad \text{and} \quad F_{\mathrm{D}} = \begin{bmatrix} 0 \\ \tau\cdot\mathbf{n} \\ (\tau\cdot\mathbf{u} - \mathbf{q})\cdot\mathbf{n} \end{bmatrix} \tag{2}$$

respectively. Here, $\mathbf{n}$ is the outward unit normal vector, $\mathbf{u}$ is the velocity vector, τ is the viscous stress tensor and $\mathbf{q}$ is the heat flux vector. In the discrete formulation currently used, the convective (F_{C}) and diffusive (F_{D}) fluxes across the surface f are, in general, a function of the primitive variable vector Q and its gradient ∇Q evaluated at centroid locations to the left ($Q_{\mathrm{L,c}}, \nabla Q_{\mathrm{L,c}}$) and to the right ($Q_{\mathrm{R,c}}, \nabla Q_{\mathrm{R,c}}$) of face f. The locations and the underlying notation are exemplified in Figure 1.

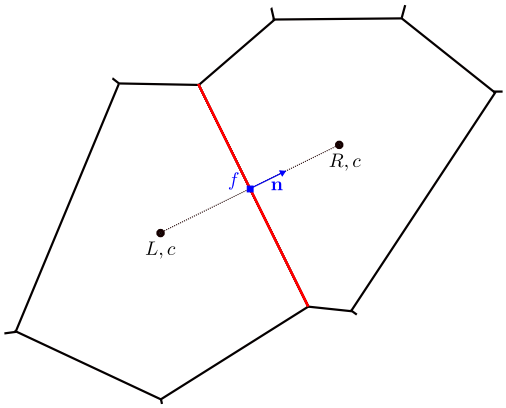

Fig. 1 Cell interface stencil notation.

A. Convective Flux

The convective flux scheme depends on the evaluation of two fluxes: a centered non-dissipative flux $F_{\mathrm{C}}^{\mathrm{C}}$ and an upwind-biased dissipative flux $F_{\mathrm{C}}^{\mathrm{U}}$. For the centered flux $F_{\mathrm{C}}^{\mathrm{C}}$, we implement an asymptotically entropy conserving and kinetic energy preserving (AEC KEP) flux based on the harmonic mean [22] using left and right cell center values $Q_{L,c}$ and $Q_{R,c}$ directly. The upwind flux leverages the MUSCL [23] reconstructed primitive variable vectors from either side of the face as:

$$Q_L = Q_{L,c} + \phi_L \nabla Q_{L,c}\cdot\vec{r}_L \quad \text{or} \quad Q_R = Q_{R,c} + \phi_R \nabla Q_{R,c}\cdot\vec{r}_R \tag{3}$$

where

$$\vec{r}_L = \vec{x}_f - \vec{x}_{L,c} \quad \text{and} \quad \vec{r}_R = \vec{x}_f - \vec{x}_{R,c} \tag{4}$$

$\nabla Q_{L,c}$ and $\nabla Q_{R,c}$ are, respectively, gradients at the left and right cell centroids calculated via the Green-Gauss method as described in Sozer et al. [24] and ϕ_L and ϕ_R are slope limiters evaluated for left and right reconstructions, with $\phi = 1$ representing no limiting and $\phi = 0$ representing first order evaluation. Limiting is typically not needed for subsonic flows and only becomes necessary in regions with flow discontinuities or with exceptionally problematic cell quality affecting the gradient evaluation. When a limiter is needed, a minmod [25] type function is used. After the MUSCL reconstruction, $F_{\mathrm{C}}^{\mathrm{U}}(Q_L, Q_R)$ is computed using the AUSMPW+ [26] flux function.

The upwind flux, by itself, is known to be too dissipative to be used in LES simulations [27]. The AEC KEP flux eliminates numerical dissipation in regions of regular stencils and drastically reduces it otherwise compared to the upwind flux. However, this desirable characteristic might not always be numerically stable in regions of significant compressibility effects and/or regions of questionable grid quality. To increase stability, the two fluxes previously described, $F_{\mathrm{C}}^{\mathrm{C}}$ and $F_{\mathrm{C}}^{\mathrm{U}}$ are blended with the aim to achieve low dissipation where turbulence is active and to selectively add appropriate dissipation where compressibility effects are prominent. The final convective flux is expressed as:

$$F_{\mathrm{C}} = (1-\zeta)F_{\mathrm{C}}^{\mathrm{C}} + \zeta F_{\mathrm{C}}^{\mathrm{U}} \tag{5}$$

where ζ is a Ducros [28] type upwind sensor that is computed locally at each face and is used to introduce numerical dissipation in dilatation-dominant regions of the domain. The sensor is formulated as:

$$\zeta = \frac{(\nabla \cdot \mathbf{V})^2}{(\nabla \cdot \mathbf{V})^2 + \alpha \left|\nabla \times \mathbf{V}\right|^2 + \epsilon} \tag{6}$$

Where ϵ is a small number to avoid division by zero. α is a factor that localizes the sensor further, i.e., the sensor becomes less sensitive to dilatation (and hence the compressibility effects) and reduces the upwind part of the convective flux. In this work, we used $\alpha = 100$.

B. Diffusive Flux

In order to evaluate the diffusive fluxes, primitive variable gradients at face centers are needed. The naive approach to get the face center gradients would be to simply take the average of left and right cell center gradients as:

$$\nabla \overline{Q}_f = \frac{1}{2}\left(\nabla Q_{L,c} + \nabla Q_{R,c}\right) \tag{7}$$

However, this approach is not stable and leads to the well-known even-odd decoupling issue, which can be remedied via various approaches as summarized by Jalali et al. [29]. Here, we correct the face-averaged gradient in the face normal direction as:

$$\begin{aligned} \nabla \overline{Q}'_f = \nabla \overline{Q}_f \\ - \left(\nabla \overline{Q}_f \cdot \hat{n}\right) \hat{n} \\ + \left(\frac{Q_{R,c} - Q_{L,c}}{d_\perp}\right) \hat{n} \end{aligned} \tag{8}$$

where $d_\perp$ is the distance between the left and right centroids in the face normal direction,

$$d_\perp = (\vec{x}_{R,c} - \vec{x}_{L,c}) \cdot \hat{n} \tag{9}$$

This treatment essentially replaces the face normal component of the face-averaged gradient $\nabla \overline{Q}_f$ with a directional derivative calculated using the two neighboring cells. The corrected face gradient $\nabla \overline{Q}'_f$ is then used to evaluate the stress tensor at the face and calculate the diffusive flux.

C. Time Integration

Time integration is carried out using an explicit, 3 stage, 3$^{\text{rd}}$ order strong stability preserving Runge-Kutta scheme [30]. For a Cartesian aligned grid with hexahedral cells, the CFL number can be defined as [31]:

$$\text{CFL} = \Delta t \left(\frac{|u| + a}{h_x} + \frac{|v| + a}{h_y} + \frac{|w| + a}{h_z}\right) \tag{10}$$

where u, v,w are the velocity components and h_x, h_y, h_z are the edge lengths in the 3 directions, while a is the local speed of sound and Δt is the time step size.

Instead of 3 distinct directions, for a general polyhedral cell we define CFL analogously for each cell face as:

$$\begin{aligned} \text{CFL}_L &= 3\Delta t \frac{\left|\vec{V}_L \cdot \hat{r}_L\right| + a_L}{h_L} \\ \text{CFL}_R &= 3\Delta t \frac{\left|\vec{V}_R \cdot \hat{r}_R\right| + a_R}{h_R} \\ \text{CFL}_f &= \max\left(\text{CFL}_L, \text{CFL}_R\right) \end{aligned} \tag{11}$$

where $\vec{V}$ is the velocity vector, $\vec{r}$ is the vector from cell center to face center as defined in Eq. 4, and $\hat{r}$ is $\vec{r}$ normalized by its magnitude. We define the unstructured cell length scale as $h = 2\|\vec{r}\|$, i.e. twice the magnitude of the vector from cell

center to face. The global $\text{CFL}_{\max}$ is then taken as the maximum of all face CFL_f values. The viscous contribution is omitted in the computation of the CFL number. The global $\text{CFL}_{\max}$ value is calculated at each time step and the step size Δt is adjusted to match the user-specified CFL limit.

D. Wall Model

Under-resolved boundary layers are modeled using an explicit, equilibrium wall function based on Musker [32]. However, our implementation utilizes optimizations that approximate the log functions for numerical efficiency while closely following Musker [32]. Wall function inputs are sampled at the first cell centroid off the wall. The wall stress computed by the wall model is directly imposed as the viscous flux at wall faces.

E. Wall Turbulence Sensor

WMLES for high-lift wing configurations has been demonstrated to be a promising approach throughout the high-lift prediction workshop series. However, one of the significant observed shortcomings of the approach is in the prediction of transitory flow regions such as wing and high-lift device leading edges, especially in lower Reynolds Number regimes.

In order to address this issue, we have adopted a wall turbulence sensor-based switch following Larsson et al. [21] with some key changes to the time-filtering and sensor threshold. It should be noted that this is not a true transition model but rather a crude, and importantly, a low computational cost model. The model utilizes a sensor measuring resolved turbulent kinetic energy at the wall and switches between the application of the fully turbulent wall-model shear stress and no-slip boundary conditions based on a sensor threshold. The premise of the model is that when there is little to no resolved turbulence, the application of the fully turbulent wall stress is excessive and can lead to premature loss of momentum and flow separation. In the context of this work, this leads to reduced $C_{L,max}$ as demonstrated in Section IV.A.2.

The sensor is defined as:

$$s_w(t) = \frac{\langle \rho_w \rangle(t) k(t)}{\langle \tau_w \rangle(t)} \tag{12}$$

where $k(t) = \langle u_k u_k \rangle - \langle u_k \rangle^2$ is the time-filtered turbulent kinetic energy, $\langle \rho_w \rangle(t)$ is the time-filtered density at the wall, and $\langle \tau_w \rangle(t)$ is the time-filtered wall shear stress as computed by the turbulent wall function.

Time-filtering of a quantity f is performed via an exponential moving averaging procedure which can easily be integrated by incrementing during each time step of the solver:

$$\langle f \rangle(t) = \alpha f(t) + (1 - \alpha) \langle f \rangle(t) \tag{13}$$

Where Δt is the time step size, $\alpha = \Delta t / T$ and T is the time scale. In Larsson et al. [21], the time scale is taken as $T(t) = 1/\sqrt{S_{ij} S_{ij}}$ where S_{ij} is the strain-rate tensor. In this work, however, we adopt a different time scale in order to avoid the computational cost of computing (or storing) the strain-rate tensor. Instead, we adopt a time scale as:

$$T(t) = \frac{h_w}{\kappa \sqrt{\langle \tau_w \rangle(t) / \langle \rho_w \rangle(t)} + \epsilon} \tag{14}$$

where $\kappa = 0.41$, h_w is the wall function sampling distance, and ϵ is a small number to avoid division by zero near stagnation points. This formulation essentially eliminates the computational cost of time scale evaluation by utilizing the shear velocity that is already computed via the wall function. In our experiments, we found this scale to provide a reasonable time averaging window.

Where sensor is greater than a set threshold $s_w(t) > s_{lim}$, the wall function evaluated shear stress is applied at the wall, whereas for $s_w(t) \leq s_{lim}$, a "no-slip" estimation of the wall stress is utilized assuming a linear velocity profile from the first off-wall cell center. While the threshold s_{lim} value suggested by Larsson et al. [21] is 0.25, we have found this value to be dependent on the numerical scheme of the particular solver that incorporates this model. For LAVA unstructured, $s_{lim} = 0.05$ was found appropriate based on experimentation.

F. Subgrid-Scale Turbulence Closure

Subgrid scale turbulence is modeled via the Vreman model [33]. LAVA unstructured uses a constant coefficient of 0.08 [34] which is derived based on numerical experiments using canonical test cases such as the homogeneous isotropic turbulence problem.

G. Numerical Boundary Layer Tripping

Penalty source-based numerical boundary layer tripping is applied at several locations chosen to be consistent with the experimental setup. The trip locations are shown for Case 2.4 on Figure 2. For Case 3, the tail is not present, and the small trip strips over the wing were not applied in our simulations.

The trip lines are used to mark cells for the application of source terms. Along the line, groups of cells are picked based on a regular on/off pattern that approximately activates for two consecutive cells and de-activates for a three-cell gap. The trip mechanism relies on momentum and energy penalty source terms as:

$$S_{u,i} = -\frac{\rho u_i}{\tau \Delta t} \tag{15}$$

$$S_E = -\frac{\rho u_i u_i}{\tau \Delta t} \tag{16}$$

where τ is a time relaxation factor. $\tau = 1$ would correspond to a sink strong enough to locally bring the flow to a halt. In this work we utilize $\tau = 40$, which is a mild application but has been found to effectively trigger the transition to turbulence where the grid resolution supports resolved content.

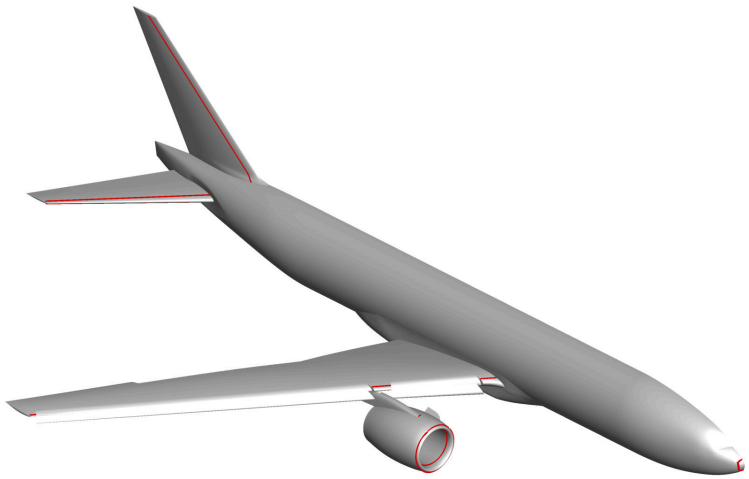

Fig. 2 Numerical trip locations for Case 2.4

III. Computational Grid

Voronoi unstructured meshes were created using the LAVA Voronoi tool [17]. Surface grid sizing was specified as shown in Figures 3 and 4 for Cases 2.4 and 3, respectively. The sizing regions are shown relative to the finest level, which spanned the slat, nacelle lip, and chine for Case 2.4.

In simulations of Case 2.4, relatively higher grid sensitivity was observed on the wing outboard upper surface. In addition, the flow topology over the flap at lower angles of attack was predicted differently between various workshop submissions. Taking these factors into consideration, we revised our surface sizing distribution for Case 3 by extending the finest region to cover the flap upper surfaces and the outboard section of the wing upper surface. In addition, fuselage was further refined to allow development of finer scale resolved turbulence content.

In the rest of this work, we refer to various grid levels by the finest surface spacing in millimeters for the flight-scale model, e.g. L4 level grid corresponding to the finest region having a 4 mm cell size near-wall. The LAVA Voronoi tool allows for the generation of consistently refined or coarsened grids by modifying this single reference sizing. The "L#" notation and "#mm" notation are used interchangeably throughout this work. The Voronoi grids used in this work are made up of isotropic cells only.

For volume meshing, a Body-Centered Cubic (BCC) initial seeding pattern was used, resulting in the majority of the grid having a truncated octahedron cell shape. No volume sizing fields were utilized. The volume grid is set to gradually coarsen with increasing distance from the surface. The coarsening is achieved via distinct telescoping regions, each of

which has a thickness corresponding to 8-cells in height, and each is consecutively coarser than the previous region by 2x. The finest level mesh only covers 4-cells in height. The behavior is demonstrated in Figure 5. Two layers of cells off the walls are aligned to the body (see Figure 5, bottom-right) and five layers of cells in the mesh size transition regions are marked for ten Lloyd smoothing iterations. Please see Sousa et al. [17] for further details of the Voronoi mesh generating techniques.

Resulting grid sizes are summarized in Table 1. LAVA Voronoi grids used for Case 2 are made available for download on the HLPW-5 webpage [35].

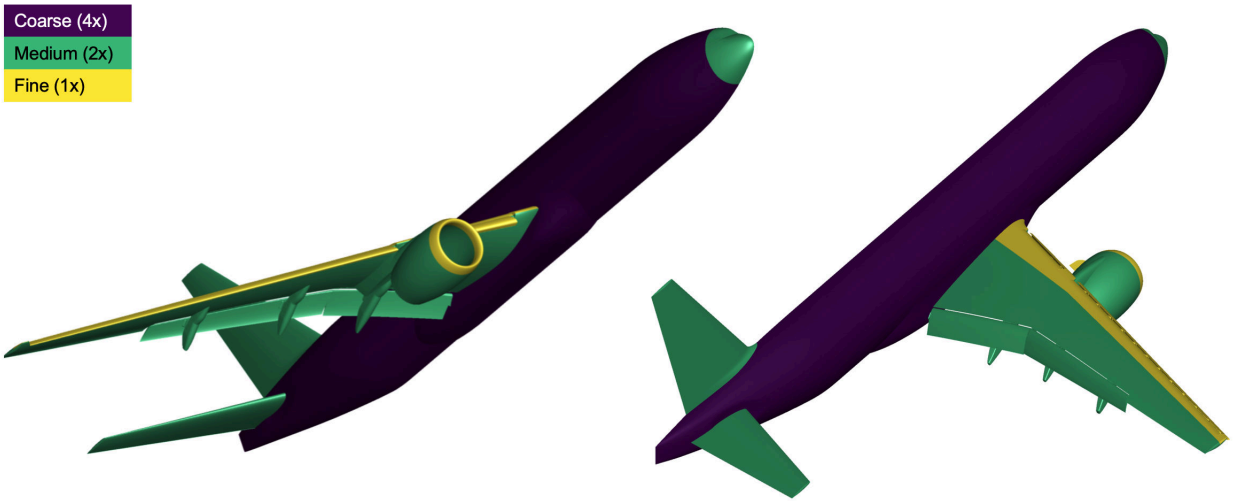

Fig. 3 Relative surface grid sizing for Case 2.4

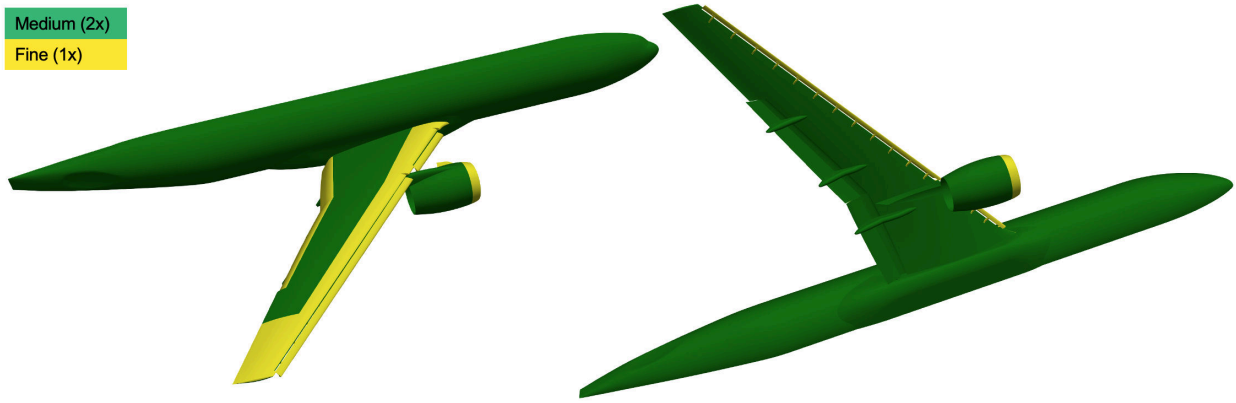

Fig. 4 Relative surface grid sizing for Case 3

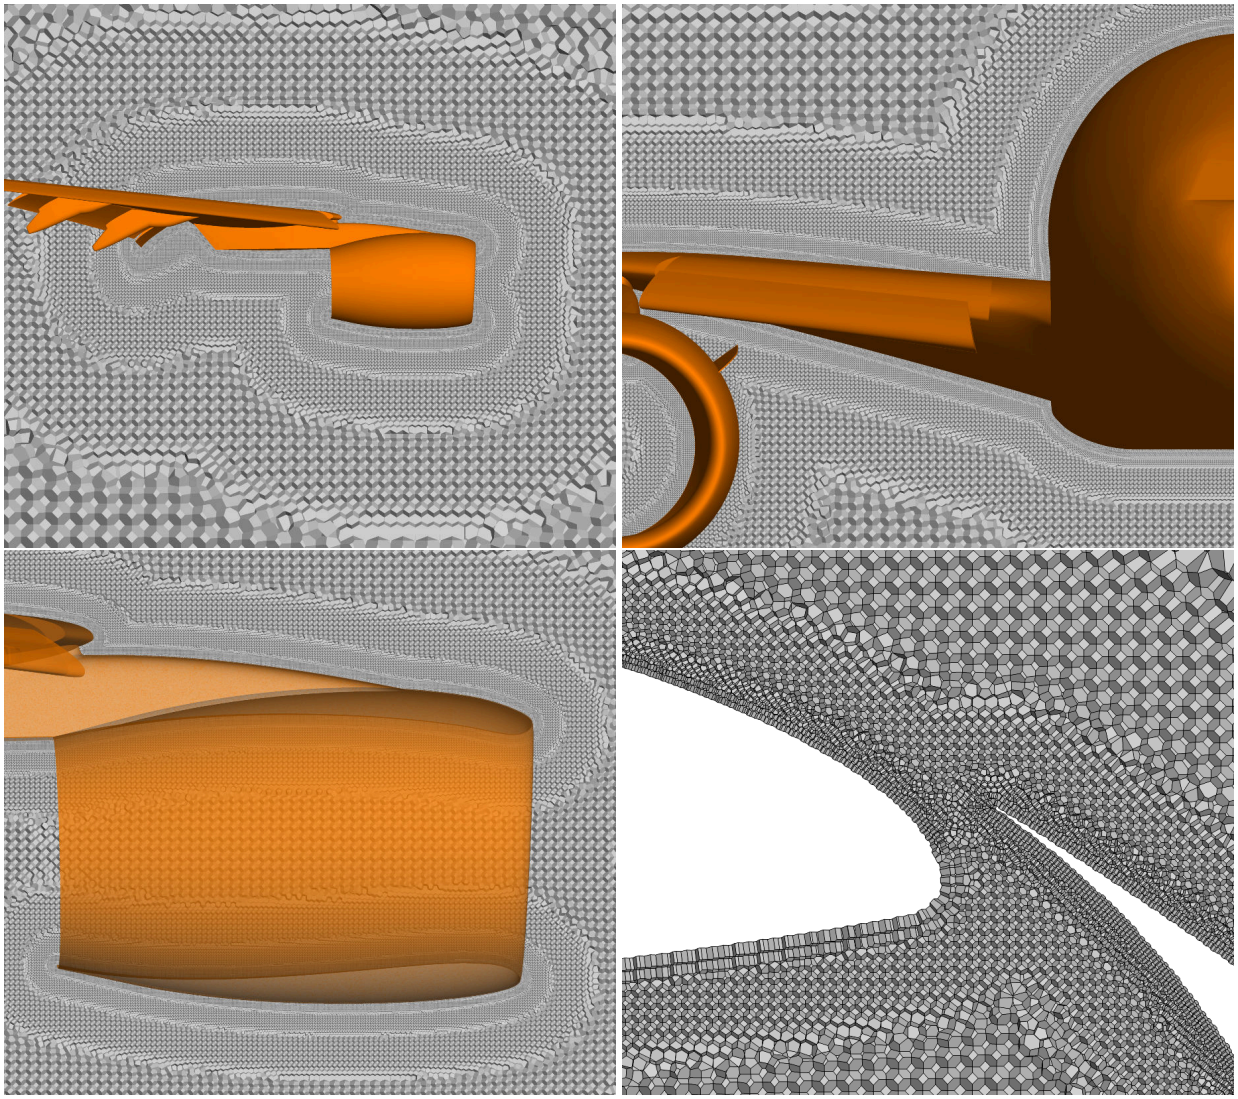

Fig. 5 Several views of the unstructured Voronoi volume grid for Case 3 at the coarse (L8) level.

Table 1 Voronoi grid statistics for each test case.

Case	L (mm)	# Cells	# Faces
2.4	8	84M	588M
2.4	6	147M	1,026M
2.4	4	325M	2,2068M
3	8	110M	769M
3	6	194M	1,351M
3	4	431M	3,004M

IV. Results

A subset of LAVA WMLES simulations for the 5th High-Lift Prediction Workshop will be discussed in this paper. In the workshop, all test cases utilized variants of the high-lift common research model, CRM-HL. Case 1 focused on a clean wing-body configuration with no high-lift devices deployed. Cases 2.1 through 2.4 represented a component build-up study with progressive additions of the slat, flaps, and the nacelle. Case 3 focused on several different Reynolds numbers as high as Re = 30×10^6. While we have participated in each of the workshop cases, this work focuses only on Cases 2.4 and 3 (see Figure 6), both of which highlight our most important conclusions from the workshop. Case 3 doesn't include horizontal and vertical tail surfaces. Otherwise, geometric differences between the cases are minimal. Further details of the test cases, as well as the submission data of the workshop participants and the available experimental data, can be found on the HLPW-5 webpage [35].

Our submissions for cases up to 2.4 were computed using the unstructured scheme described in our earlier work [34]. Later test cases were computed using the improved, lower dissipation numerical scheme as described in Section II, including the wall turbulence sensor-based boundary treatment.

This work expands on our workshop submissions by including additional simulations that explore the effect of the wall turbulence sensor (see Section II.E) with variations of grid resolution, Reynolds number, and the wall turbulence sensor threshold parameter s_{lim}. The test cases presented in this work and the corresponding flow conditions are listed in Table 2.

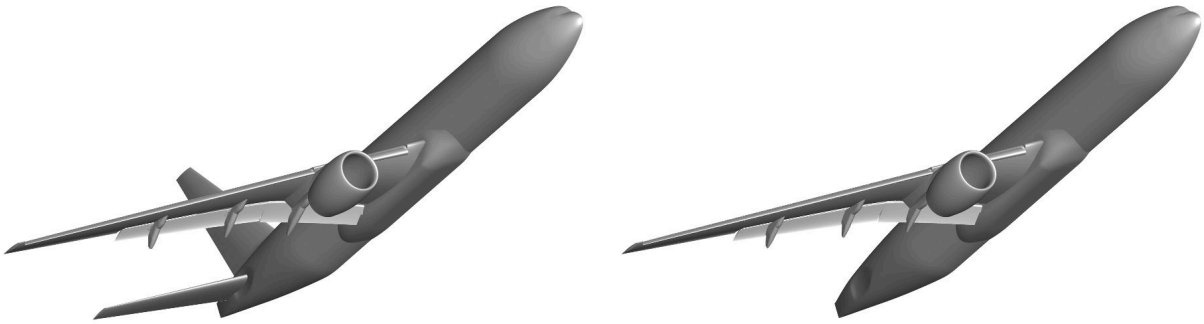

Fig. 6 Geometry for test Cases 2.4 (left) and 3 (right).

Table 2 Test case conditions.

Case	Mach	Re	Angles of Attack
2.4	0.2	5.9×10^6	7.6°, 10°, 14°, 16°, 17.7°, 19.7°, 23.6°
3.2	0.2	5.49×10^6	6°, 10°, 14°, 16°, 18°, 19°, 20°, 22°
3.3	0.2	16×10^6	6°, 10°, 14°, 16°, 18°, 19°, 20°, 22°
3.4	0.2	30×10^6	6°, 10°, 14°, 16°, 18°, 19°, 20°, 22°

All simulations have been run in free-air with a half-span model. The simulations have been impulsively started and carried out via explicit time integration until aerodynamic loads show meaningful statistical convergence. The integration time extent is shown in this work in terms of Convective Time Units (CTU), which is a non-dimensional time defined as:

$$\text{CTU} = \frac{t}{c_{\text{MAC}}/U_\infty} \tag{17}$$

where t is the time, c_{MAC} is the mean aerodynamic chord of the wing, and U_∞ is the free-stream velocity magnitude.

Example force and moment history for unstructured simulations of Case 2.4 at $\alpha = 7.6°$, $\alpha = 19.7°$ (near $C_{L,\max}$) and $\alpha = 23.6°$ (post-stall) are shown in Figures 7, 8 and 9, respectively. The figures show the raw data, along with the mean and standard deviation with a backward rolling averaging window of 30 CTUs. With increasing angles of attack, the aerodynamic loads exhibit an increased oscillation amplitude and a lower frequency content, indicating coherent

larger-scale unsteady flow structures. As a result, higher angle of attack simulations require a longer time integration window. In particular, the post-stall case at $\alpha = 23.6°$ has been integrated up to 80 CTUs, but could benefit from even further time integration.

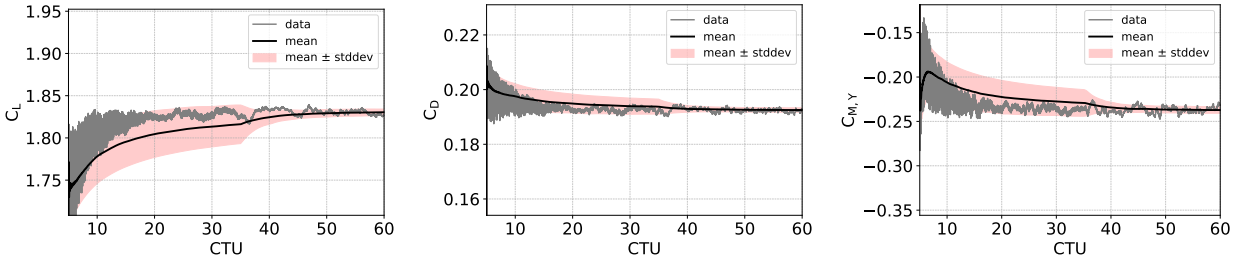

Fig. 7 Force and moment convergence history for Case 2.4 at $\alpha = 7.6°$. The mean and standard deviation shown is computed using a backward rolling window of 30 CTUs.

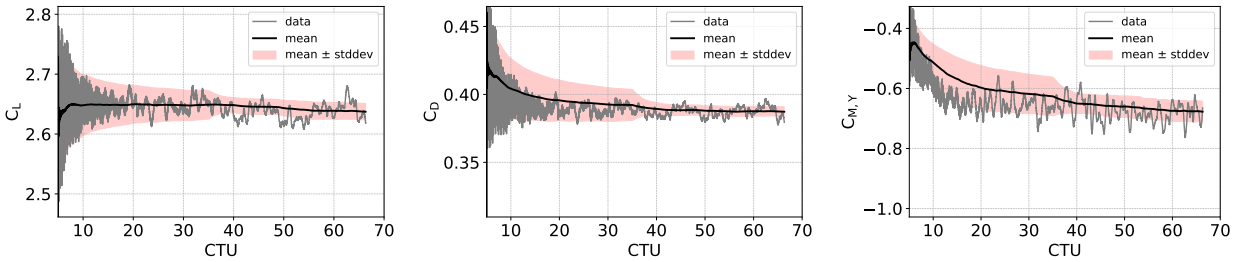

Fig. 8 Force and moment convergence history for Case 2.4 at $\alpha = 19.7°$. The mean and standard deviation shown is computed using a backward rolling window of 30 CTUs.

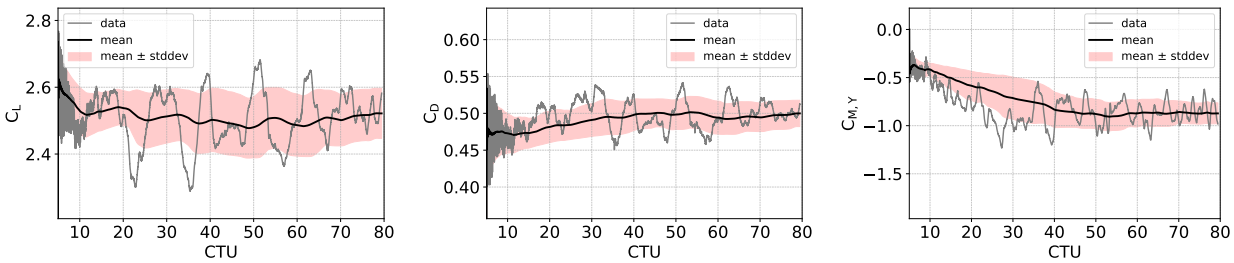

Fig. 9 Force and moment convergence history for Case 2.4 at $\alpha = 23.6°$ (post-stall). The mean and standard deviation shown is computed using a backward rolling window of 30 CTUs.

A. Test Case 2.4

1. Grid Sensitivity

Unstructured simulations of Case 2.4 at the angles of attack listed in Table 2 have been conducted using three different grid resolution levels corresponding to 8 mm, 6 mm and 4 mm finest grid spacings on the surface, denoted as L8, L6, L4 grids, respectively. Time-averaged force and moment results are shown in Figure 10. Note that a wall turbulence sensor threshold of $s_{lim} = 0.05$ is used for all simulations unless stated otherwise.

The finest (4 mm) grid results show excellent agreement with the experimental data, with the exception of a slight under-prediction of $C_{L,max}$ and an under-prediction of $C_{M,Y}$ at $\alpha = 19.7°$. The medium level (6 mm) grid produces nearly identical results except for the lift coefficient at the post-stall angle and the pitching moment coefficient at the

highest two angles of attack. The coarsest level (8 mm) is generally under-predicting lift and over-predicting the pitching moment across the angle of attack range.

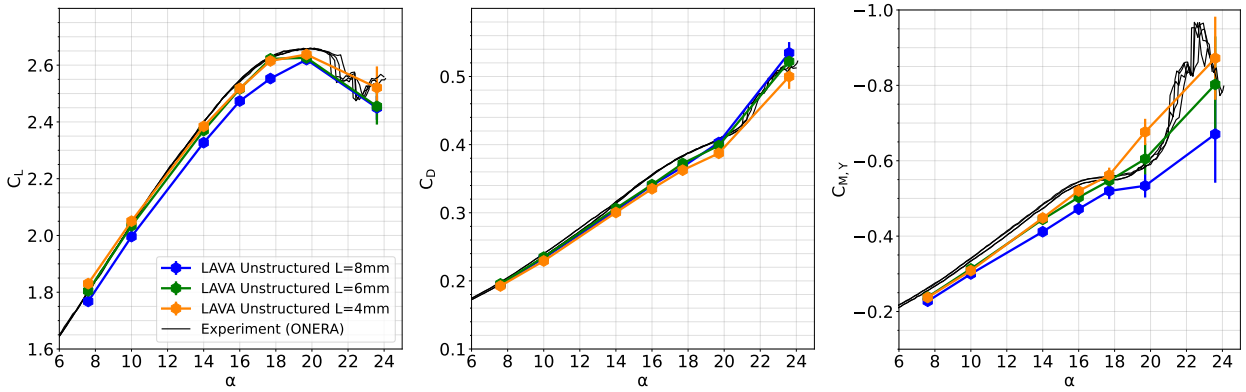

Fig. 10 Time-averaged aerodynamic loads at different grid resolution levels.

Figure 11 shows the unstructured surface solution over the wing for all three grid resolution levels at $\alpha = 19.7°$ near $C_{L,max}$. In each row, going left to right corresponds to coarse (L8), medium (L6) and fine (L4) grid solutions, respectively. It is useful to first discuss the wall turbulence sensor field (see Section II.E) as shown in the bottom row. The values below the threshold of $s_{lim} = 0.05$ have been blacked out to indicate where the sensor detects little resolved turbulence content and hence dictates the switch to a no-slip boundary condition (which results in much lower skin friction levels, as can be observed in the top two rows). A notable and encouraging observation is that the no-slip (blacked-out) regions are highly consistent across different grid resolutions with a few exceptions: the inboard slat surface and the outboard side of the nacelle are developing more resolved turbulence with refinement. The outboard slat seems to exhibit an opposite trend where it is tending towards laminarization with grid refinement. We can also observe that the sensor detects resolved turbulence immediately downstream of numerical tripping locations (see Figure 2) such as the nacelle lip. While the threshold location is the only influence of the sensor on the simulation, the rest of the sensor field is informative to observe regions of high turbulent kinetic energy which largely correspond to separated flow regions.

Inspecting the instantaneous and time-averaged surface skin friction distributions in the top two rows, we can make several immediate observations: As expected with LES methods, finer scale eddy structures are resolved with grid refinement. At the coarsest level, the separation at the wing outboard is prominent, but with refinement, separation regions gradually narrow into distinct, triangular shapes (often referred to as the "pizza slices"). In contrast, the mid-board region seems insensitive to grid refinement in terms of flow topology. Wing-fuselage junction also exhibits noticeable grid sensitivity between L8 and L6 levels where the latter displays a relative weakening of the boundary layer (and consequently a lower skin friction), although a separation is not observed in the time-averaged surface skin friction lines. Between L6 and L4 levels, this trend continues but is more subdued. At the L4 level, a small corner separation appears at the wing-body juncture trailing edge.

Figure 12 offers a similar view but at the lowest angle of attack of $\alpha = 7.6°$. At this angle, the solution is much less grid sensitive compared to the higher angles. Apart from the development of finer scale structures with refinement, as expected, the grid sensitivity regarding the flow topology seems to be limited to the outboard region. At this angle of attack, various WMLES results at the HLPW-5 were divided with respect to the occurrence of flap separation. In our simulations, the flap separation was consistently predicted at all grid levels.

Comparisons of simulation results (using fine L4 grid) with experimental oil flow images [19, 35] are shown in Figure 13 for $\alpha = 7.6°$ (top row) and $\alpha = 19.7°$ (bottom row). Please note that the color scheme for the skin friction contours in this figure has been inverted (with respect to the skin friction contour plots in the rest of the paper) in order to be able to distinguish skin friction lines in separated regions more clearly. At the low angle of attack, outboard flap separation is remarkably well reproduced in the simulations. Over the inboard flap, agreement holds well near the wing junction edge as well as over the pylon location and the junction of the two flaps. Mid-board of the inner flap shows attached flow in the simulation, whereas it seems separated in the oil flow image. However, the existence of near-separation conditions is evidenced by the turning of the surface skin friction lines. Over the main wing element, the simulated flow topology is consistent with the oil flow. The only significant difference at this angle is observed at the

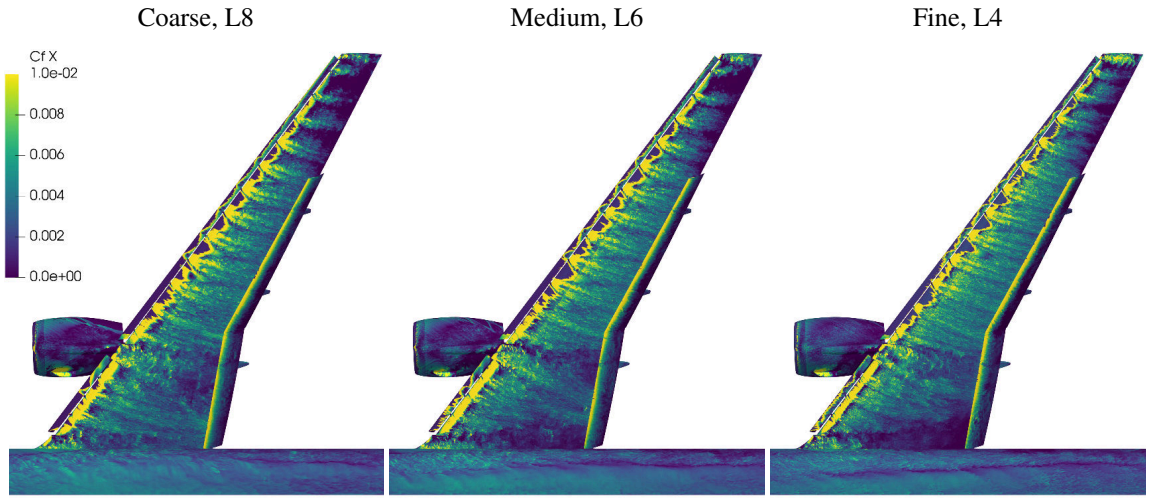

(a) Instantaneous axial skin friction coefficient.

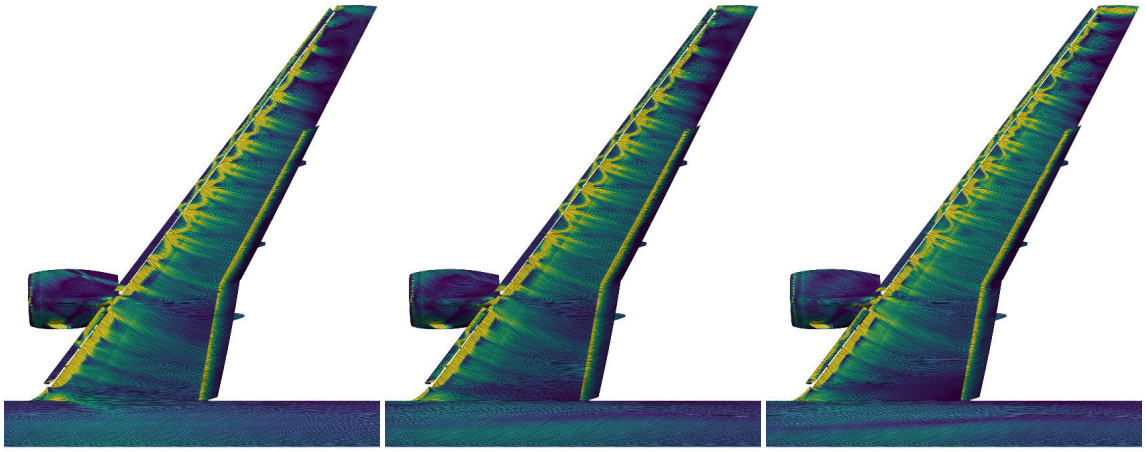

(b) Time-averaged skin friction coefficient magnitude with surface skin friction lines.

(c) Wall turbulence sensor s_w.

Fig. 11 Effect of grid resolution on the wing surface solution for Case 2.4 at α = 19.7°. L8 grid (left), L6 grid (middle), L4 grid (right).

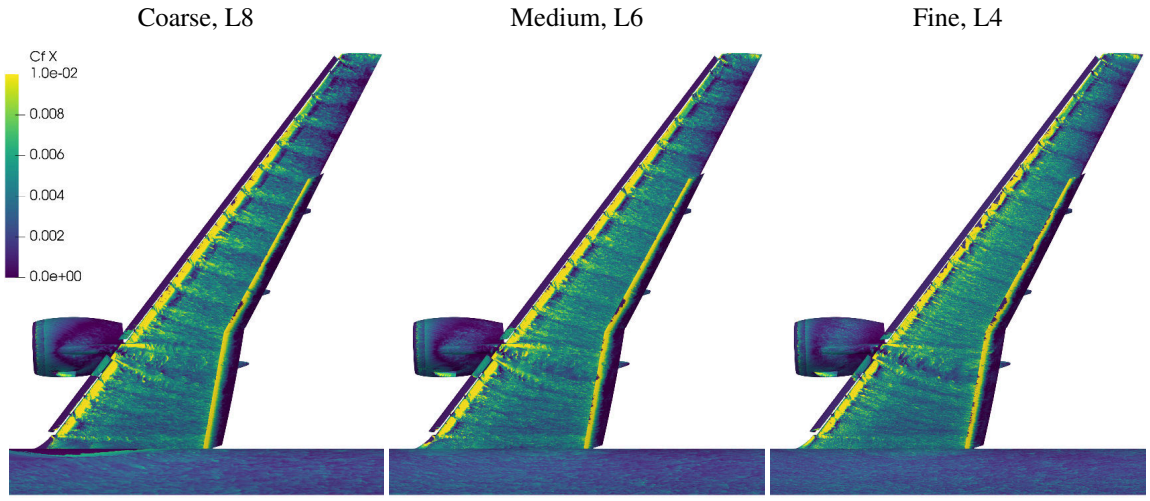

(a) Instantaneous axial skin friction coefficient.

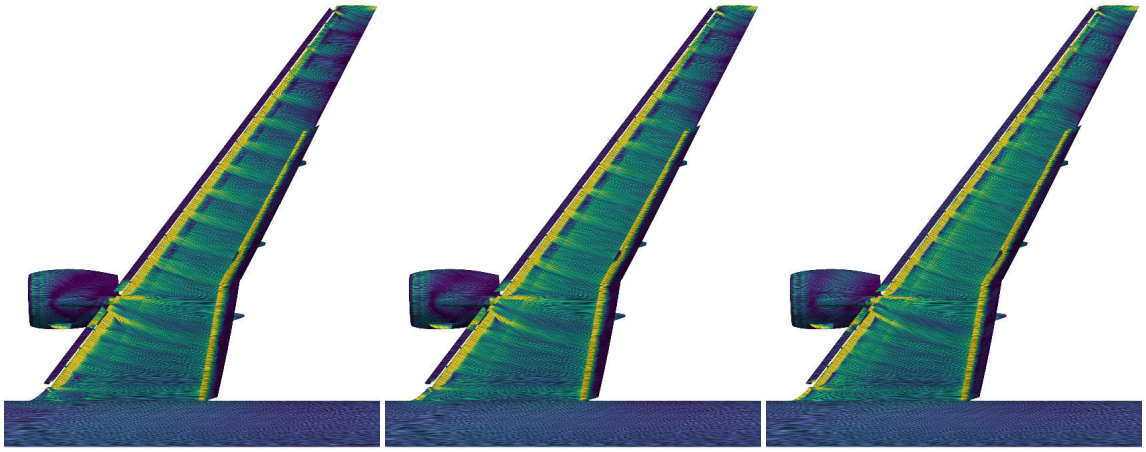

(b) Time-averaged skin friction coefficient magnitude with surface skin friction lines.

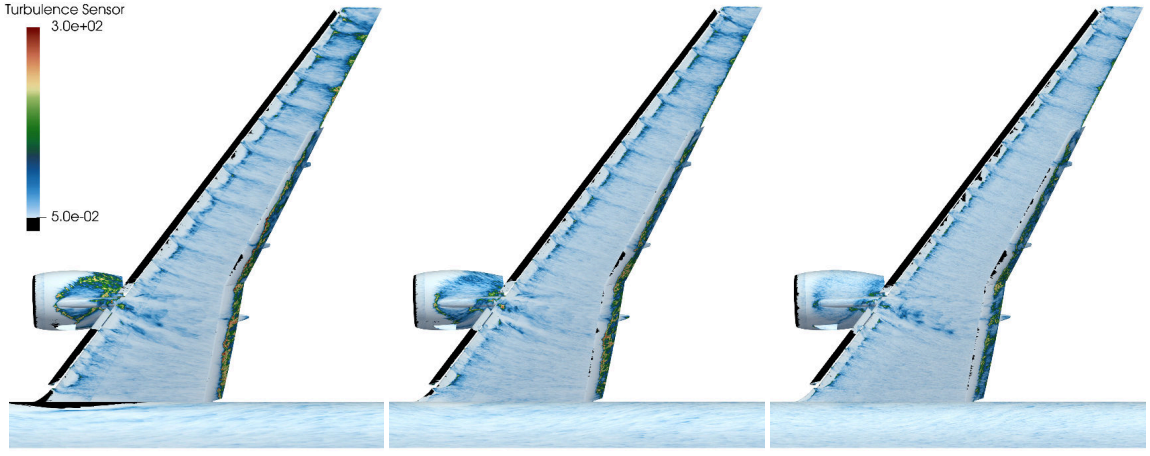

(c) Wall turbulence sensor s_w.

Fig. 12 Effect of grid resolution on the wing surface solution for Case 2.4 at $\alpha = 7.6°$. L8 grid (left), L6 grid (middle), L4 grid (right).

outboard side of the nacelle pylon, where the recirculating flow region is narrower in the simulation. Near the $C_{L,max}$ angle of $\alpha = 19.7°$, there is excellent agreement in flow topology over the flaps. Over the main wing, the outboard incipient separation wedges and the wing tip region are also very consistent with the oil flow image. The spread of the slat attachment bracket wakes throughout the span follows the experiment closely. Focusing on the wing root, the direction of the flow is well reproduced in the simulation. However, the small corner flow separation that is seen in the simulation is not present in the oil flow image. Note that at this angle, the model is close to the stall onset and slight variations in the angle of attack could affect the corner flow topology significantly. Over the outboard side of the nacelle pylon, the oil flow image seems to have degraded results, but the flow direction seems consistent with the simulation. Lastly, the entirety of the slat upper surface seems to have attached flow, which is consistent with the simulation thanks to the wall turbulence sensor-based boundary treatment.

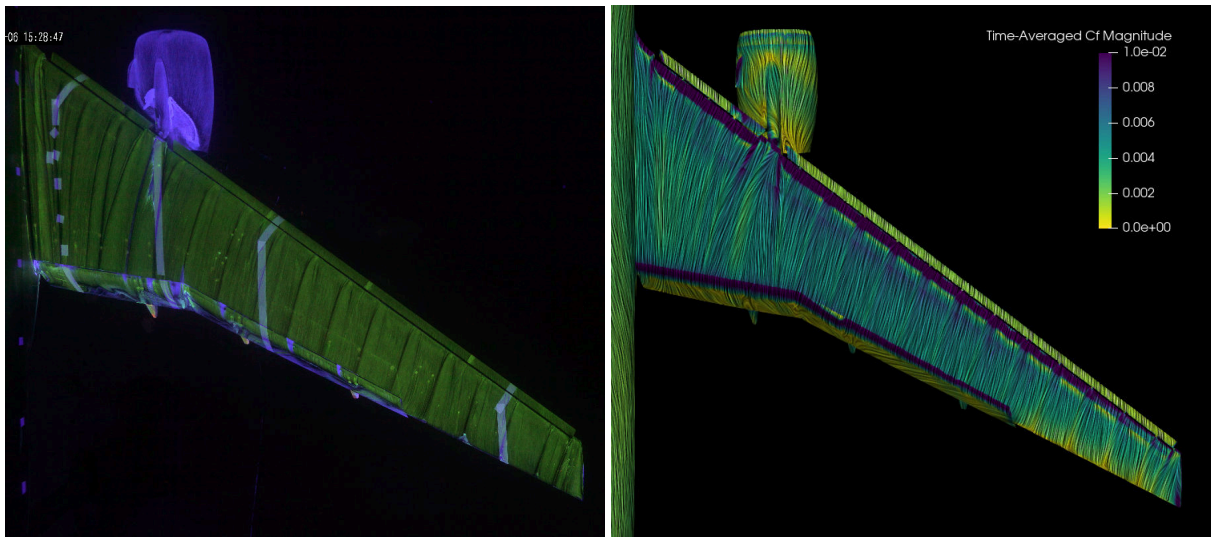

(a) $\alpha = 7.6°$

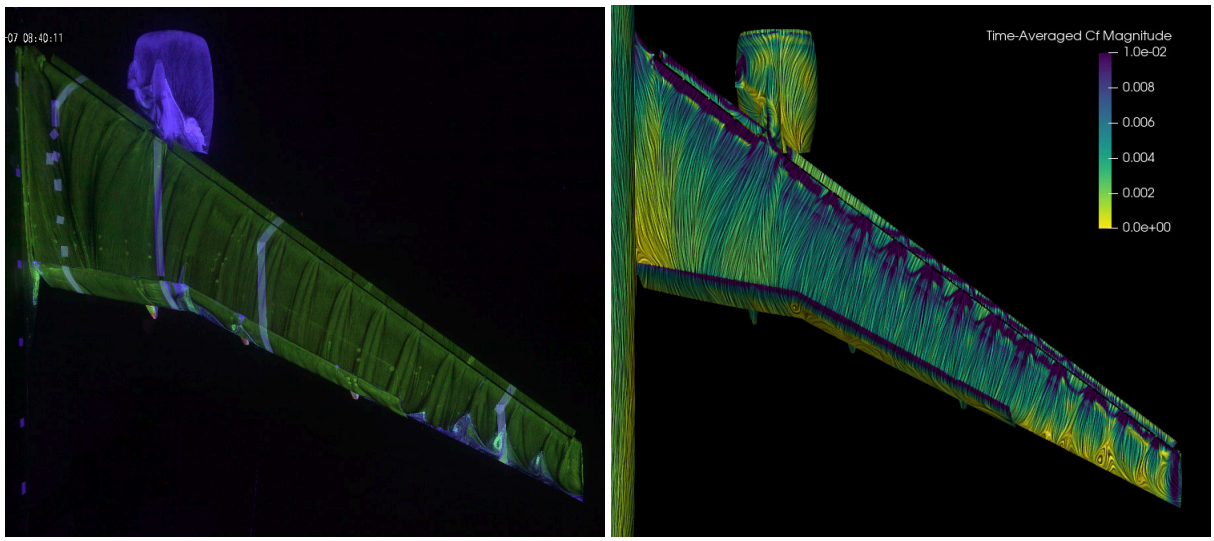

(b) $\alpha = 19.7°$

Fig. 13 Comparison of experimental (ONERA [19, 35]) oil flow images (left) to fine grid (L4) simulation results for Case 2.4 (right).

Based on the grid sensitivity observations discussed in this section, the fine (L4) Voronoi grid was considered to be sufficiently grid converged and adopted as the best practice. However, on account of the relatively higher grid sensitivity observed at the wing outboard and the wing-fuselage juncture, along with the workshop's code-to-code disagreements

on flap separation at the lowest angle, we made further targeted surface refinements for Case 3 as described in Section III.

2. Wall Turbulence Sensor Sensitivity

In the previous section, we established that the wall turbulence sensor behavior, particularly where it detects the switch to a no-slip boundary condition, is largely insensitive to the grid resolution. In this section, we probe further by testing the sensitivity of the solution to the sensor threshold value s_{lim}. Figure 14 shows the integrated aerodynamic loads at the fine (L4) grid resolution level for sensor threshold values of 0 (sensor off, wall function stress is applied everywhere), 0.025, 0.05 (current best practice) and 0.075. The sensor threshold variation sweep was only evaluated at $\alpha = 7.6°$ and $\alpha = 19.7°$. For the former, little difference is observed in the loads. At $\alpha = 19.7°$, however, the results are highly sensitive to the sensor threshold value. When the sensor is turned off with $s_{lim} = 0$, $C_{L,max}$ is under-predicted by a significant margin. The lift value progressively increases with increasing sensor threshold values but seems to converge at $s_{lim} = 0.075$, showing excellent agreement with the experimental value.

Considering that further grid refinement could slightly raise the predicted $C_{L,max}$ values, the sensor threshold value of $s_{lim} = 0.05$ was chosen to be used as the best practice. The effect of the sensor threshold on drag is negligible. The pitching moment, on the other hand, is sensitive at $\alpha = 19.7°$ where, unlike lift, it exhibits a non-monotonic behavior, i.e. the pitching moment seems to oscillate with the $s_{lim} = 0$ producing the closest result to the experiment, $s_{lim} = 0.05$ being the furthest away and the other s_{lim} values producing identical values in between. This phenomenon, and our deviation from the experimental pitching moment near $C_{L,max}$, are still under investigation, while we suspect an under-resolved interaction between the wing wake/juncture vortex and the horizontal tail as a possible explanation.

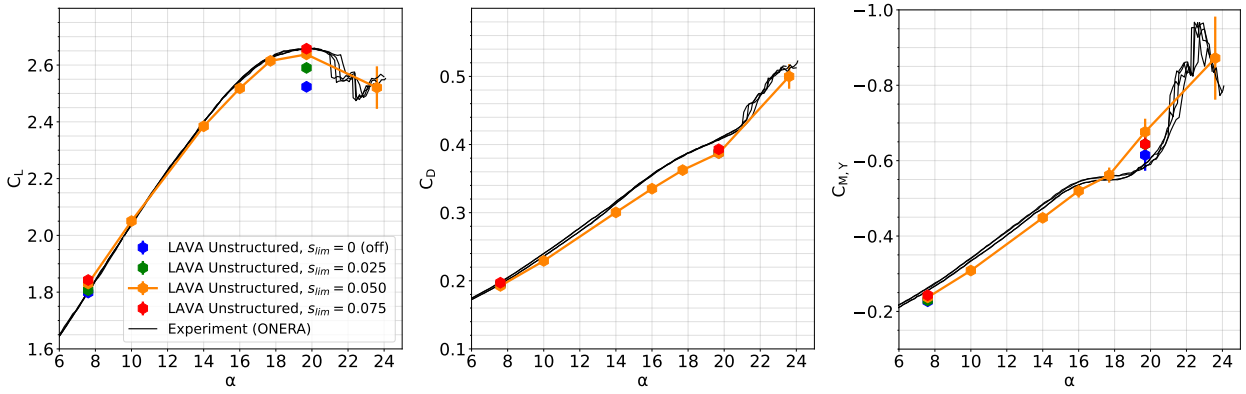

Fig. 14 Effect of wall turbulence sensor threshold, s_{lim}, on aerodynamic loads.

Figure 15 offers another view of the sensor sensitivity by showing the surface solution for different threshold values. As expected, the effect of the sensor is focused on the leading edge areas, with increasing threshold values leading to larger pockets of no-slip boundary condition application (indicated as the blacked-out regions). A major impact of the sensor is seen over the slat which exhibits flow separation when the sensor is turned off ($s_{lim} = 0$). This also leads to premature momentum loss and a larger separation region at the wing outboard. Both of these contribute to the lift under-prediction as seen in Figure 14. The activation of the sensor with any of the non-zero threshold values immediately leads to an attached flow over the slat and improves the lift prediction. Note that the experimental oil flow images show no evidence of slat separation (see Figure 13. With the threshold values of 0.05 and 0.075, the sensor begins to affect the main element leading edge where it detects no resolved turbulent content between the slat attachment brackets and hence leads to pockets of no-slip zones, further increasing the lift.

Our use of the wall turbulence sensor is a recent development with promising improvements for predictions of this case. We continue to evaluate this model for various other cases with different geometries and flow conditions to deduce how universally applicable it is. Our preliminary experience indicates that the sensor threshold value, s_{lim}, is dependent on the numerical scheme of the solver, but it doesn't need to be adjusted for grid resolution, and the value calibrated for this case may be directly applicable to other problems.

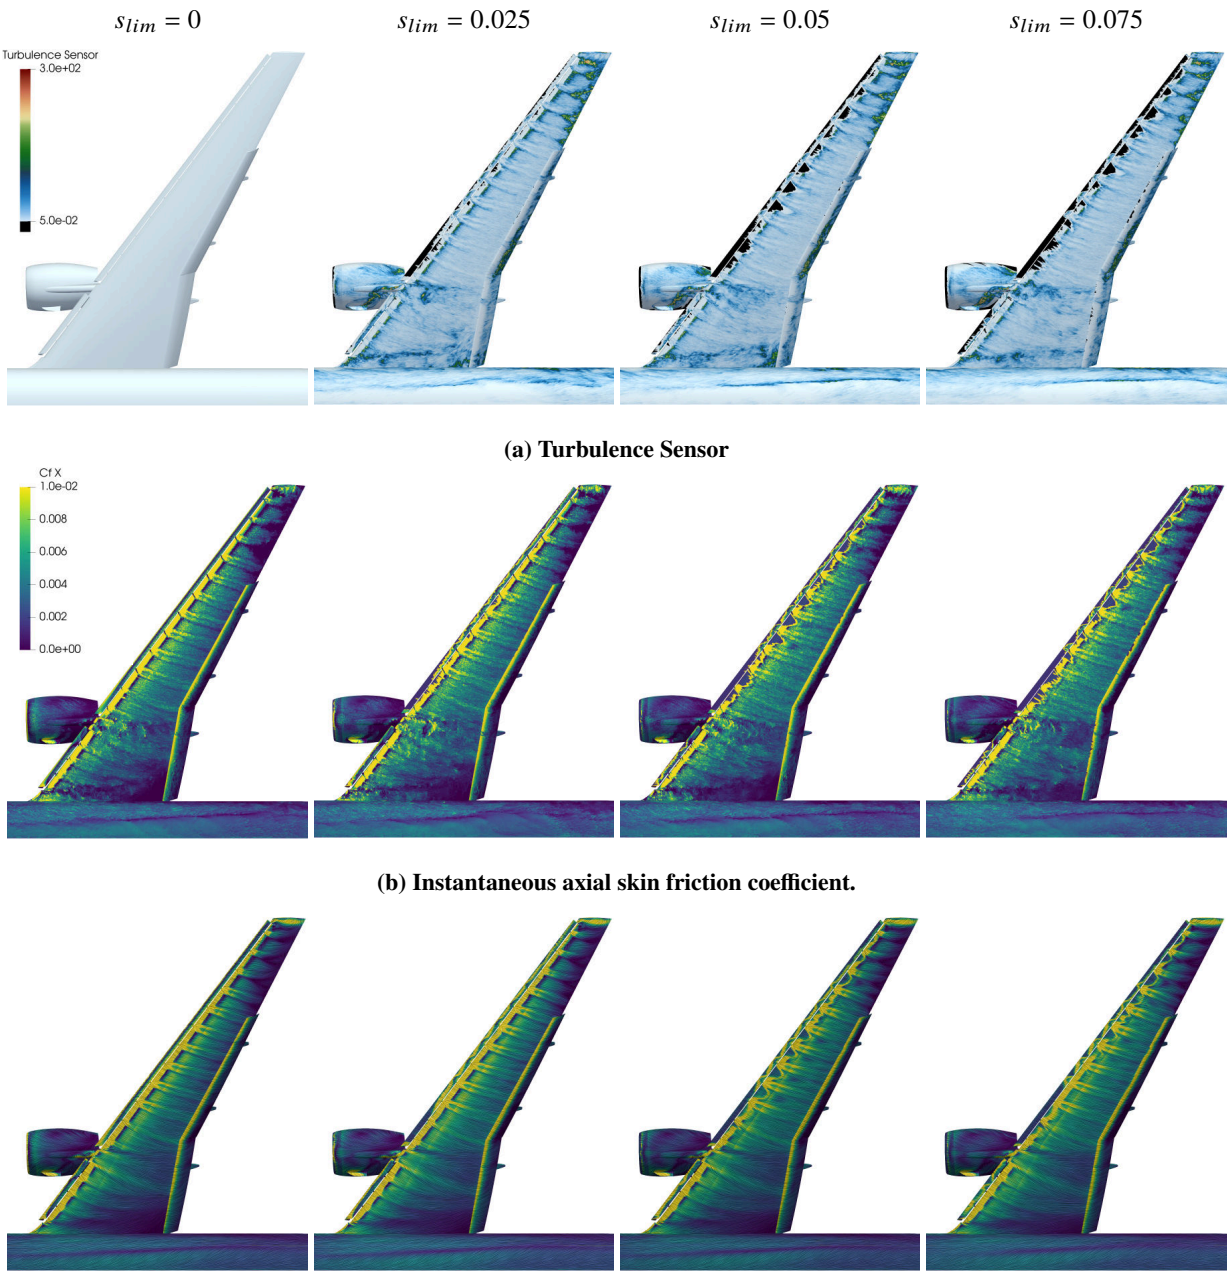

(a) Turbulence Sensor

(b) Instantaneous axial skin friction coefficient.

(c) Time-averaged skin friction coefficient magnitude with surface skin friction lines.

Fig. 15 Effect of wall turbulence sensor threshold, s_{lim}, for Case 2.4 at $\alpha = 19.7°$ using the fine (L4) grid.

B. Test Case 3

Case 3 utilized the CRM-HL model at 5.2% scale without horizontal or vertical tails, as seen in Figure 6. It focused on a range of Reynolds numbers, starting with the optional Case 3.1 with the lowest Re = 1.05×10^6. Our group only participated in Cases 3.2, 3.3, and 3.4 with Re = 5.49×10^6, 16×10^6, and 30×10^6 respectively, as listed in Table 2. Unfortunately, the experimental data have not yet been made available for Case 3.

Same family of grids as described in Section III were used for each of the Reynolds numbers simulated. We might expect a finer grid resolution requirement for the higher Reynolds number cases, particularly within the boundary layer. However, at the time of the workshop, our Voronoi meshing tool [17] implementation was serial and had limitations on the scale of mesh we could create. We have since overcome that limitation with distributed memory parallelism, but we have not yet revisited this test case with finer meshes. That said, the following grid sensitivity discussions reveal satisfactory grid convergence for all Reynolds numbers tested.

Figure 16 shows the sensitivity of time-averaged C_L to grid resolution (left) and Reynolds number (right). All of these sensitivity cases were computed at a single angle of attack of $\alpha = 18°$. On the left plot, the decreasing L number corresponds to finer grids. At Re = 5.49×10^6, lift seems to increase nearly linearly with grid refinement, albeit at a shallow slope. Looking back at the Case 2.4 grid sensitivity (see Figure 10), we might be seeing a similar behavior where the coarsest solution (L8) exhibits higher lift near $C_{L,max}$ angle than the rest of the α sweep trend would suggest. For the higher Reynolds numbers, a more satisfactory grid convergence is observed with a significantly shallower slope between L6 and L4 levels compared to L8 and L6 levels. For all Reynolds numbers, the change of lift between the two finest grid levels is less than 25 counts and hence the fine (L4) grid level is deemed satisfactory. The plot on the right shows the variation of lift with changing Reynolds numbers. All three grid levels exhibit a consistent trend of increasing lift with increasing Reynolds number, with greater sensitivity between Re = 5.49×10^6 and Re = 16×10^6.

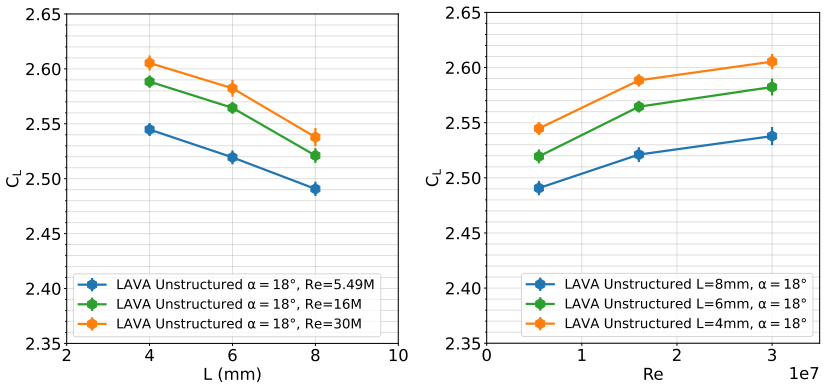

Fig. 16 Time-averaged aerodynamic loads for Case 3, showing the grid sensitivity (left) and Reynolds number sensitivity (right) effects at $\alpha = 18°$.

Figure 17 offers a broader look into the Reynolds number sensitivity by examining both lift, drag, and pitching moment for an α sweep at the finest grid resolution (L4). The offsets in lift and pitching moment coefficients due to Reynolds number variations appear to be nearly uniform across the angle of attack range, with notable differences emerging at the post-stall angle. The drag coefficient demonstrates remarkable Reynolds number independence, with only subtle variations observed at post-stall conditions, where higher Reynolds numbers correspond to marginally reduced drag. Intriguingly, the pitching moment characteristics exhibit a distinct behavior shift in the post-stall regime, with lower Reynolds number cases displaying a more pronounced pitch break. Another notable observation is that the lowest Re case exhibits $C_{L,max}$ at a slightly lower angle of attack. For the higher Reynolds numbers, we could benefit from additional simulations between $\alpha = 20°$ and $\alpha = 22°$ to identify a more precise location of stall onset.

Figure 18 shows the instantaneous axial skin friction coefficient over the wing top surface for different grid resolutions at each Reynolds number. Both increasing grid resolution and increasing Re result in the appearance of finer-scale turbulent structures. Across all Re conditions, topological grid sensitivity effects are largely concentrated on the outboard section, showing narrowing "pizza slice" separations with grid refinement. This observation is consistent with Case 2.4 (see Figure 11) although due to additional surface refinement in this region for Case 3, the sensitivity is less pronounced. At Re = 5.49×10^6, the nacelle wake over the wing breaks down into finer structures with grid refinement, but this effect is not as obvious for Re = 30×10^6, potentially indicating a need for further targeted refinement in this

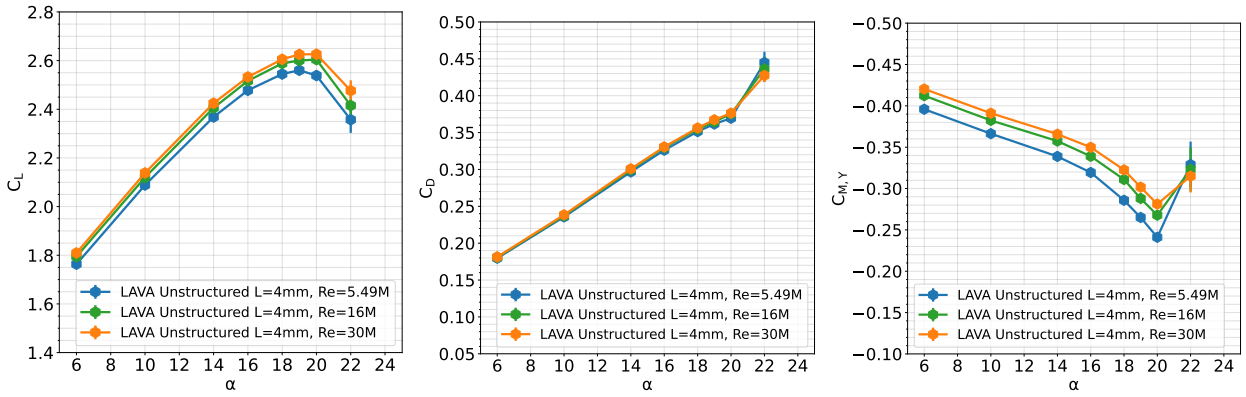

Fig. 17 Time-averaged aerodynamic loads for Case 3 at different Reynolds numbers.

area. We do not observe significant topological sensitivities with refinement over the wing-body junction, nacelle, flaps, or the mid-board section of the slat.

Figure 19 shows the effect of Reynolds number on instantaneous axial skin friction coefficient over the wing at three different angles of attack. All simulations shown in the figure are computed with the fine (L4) mesh. As observed in the previous figure, increasing Reynolds number results in finer scale turbulent structures across the angle of attack range. At $\alpha = 6°$, we do not observe any topological changes in the flow, but we notice a reduction of peak skin friction coefficient at the main element leading edge with increasing Re. At $\alpha = 19°$, we observe a slight decrease in the extent of the outboard incipient separation wedges with Re change from 5.49×10^6 to 16×10^6, along with a reduction in skin friction coefficient in the nacelle wake. Between Reynolds number of 16×10^6 and 30×10^6, no noticeable change in flow topology is observed via inspection of the skin friction coefficient. At $\alpha = 22°$, which corresponds to post-stall, the dominant sensitivity to Re is realized over the extent of the large inboard separation. At the lower Re, the inboard separation is extensive enough to interact and merge with the nacelle wake over the wing. The inboard separation region progressively narrows with increasing Re. At the highest Re, the nacelle wake and the inboard separations are more distinct.

Figure 20 shows the wall turbulence sensor field computed on the fine (L4) mesh at $\alpha = 19°$ at each of the Reynolds numbers. Focusing on the leading edge, the extent of the no-slip zones (blacked out regions, corresponding to regions of little to no resolved turbulence content) detected by the sensor is slightly shrinking with the increase of Re from 5.49×10^6 to 16×10^6. This indicates earlier development of turbulent content with increased Reynolds number as expected. Going from Re = 16×10^6 to 30×10^6, the sensor field seems much less sensitive. Besides the leading edge behavior, the sensor field indicates stronger resolved content in the nacelle wake with increasing Re.

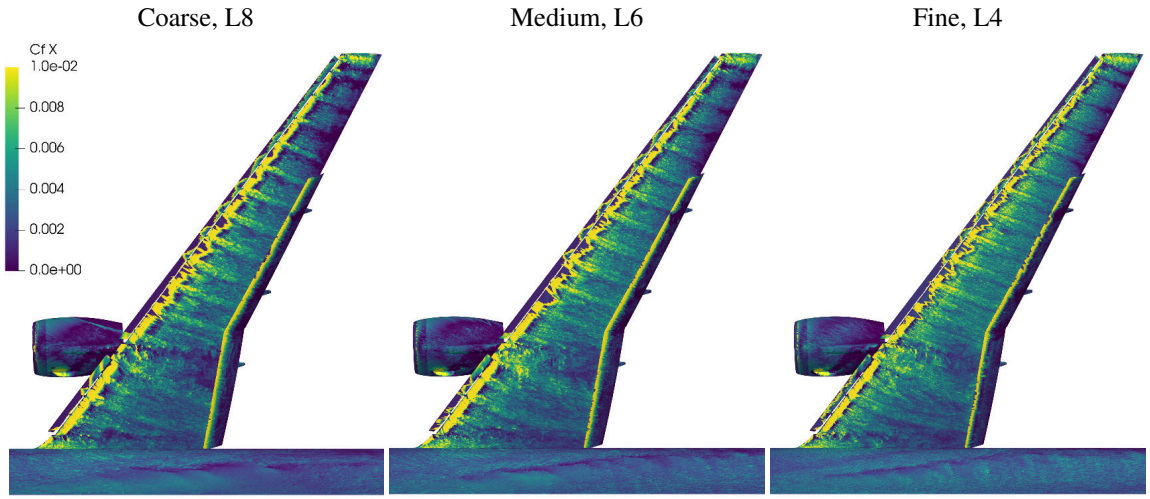

(a) Re = 5.49×10^6

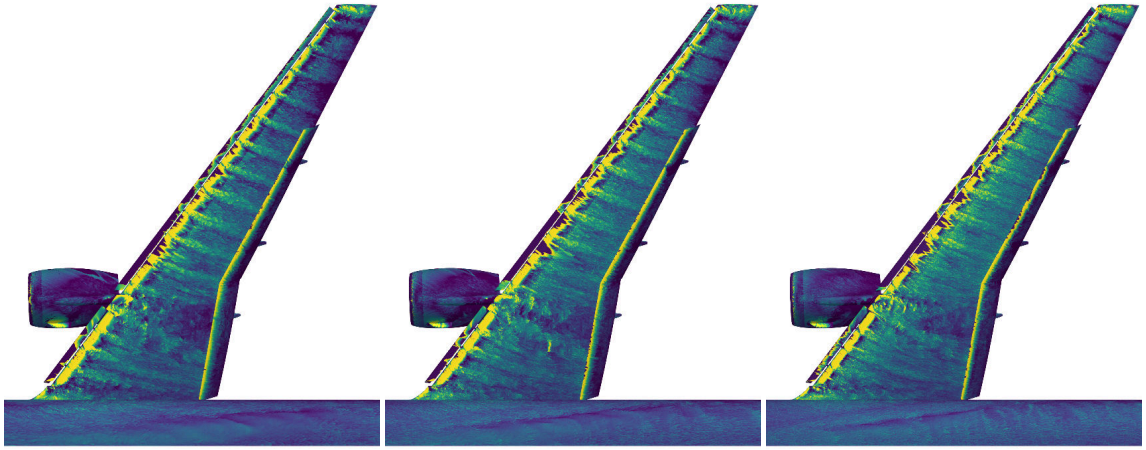

(b) Re = 16×10^6

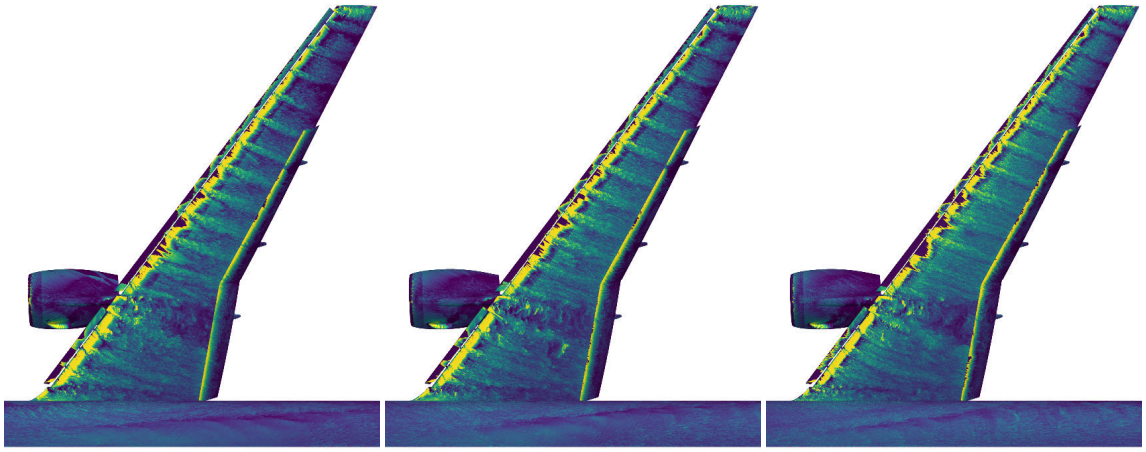

(c) Re = 30×10^6

Fig. 18 Effect of grid resolution on the solutions of Case 3 at $\alpha = 18°$ for all three Reynolds numbers, showing instantaneous axial skin friction coefficient for L8 grid (left), L6 grid (middle), and L4 grid (right).

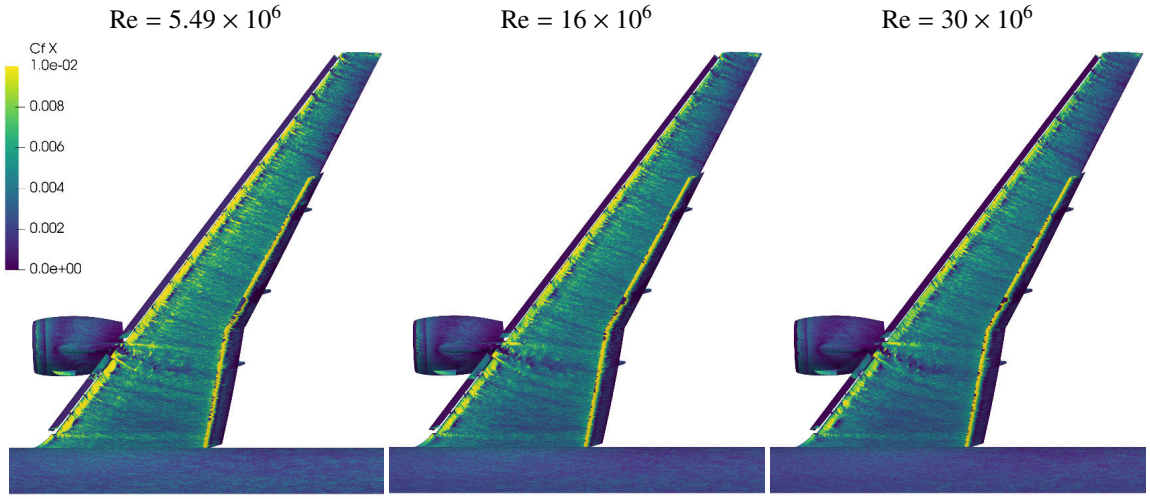

(a) $\alpha = 6°$

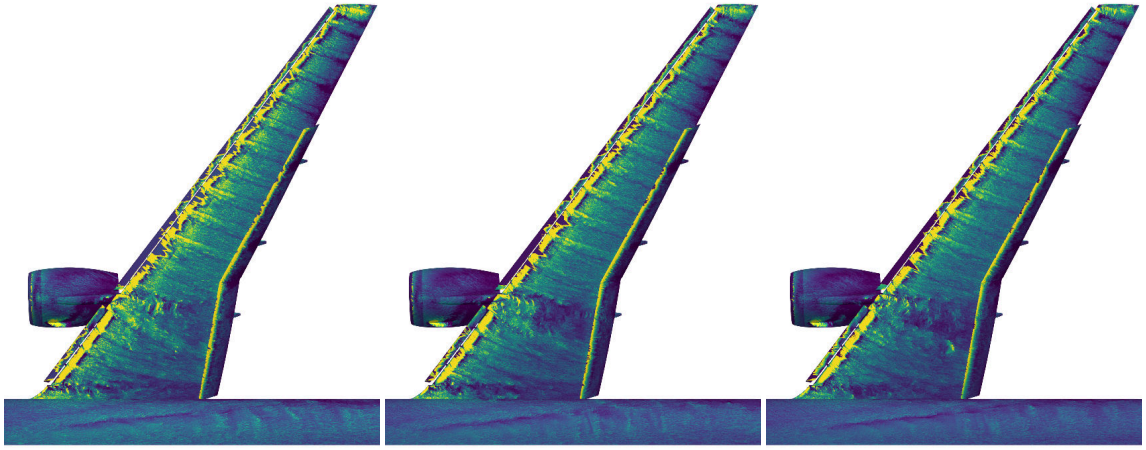

(b) $\alpha = 19°$

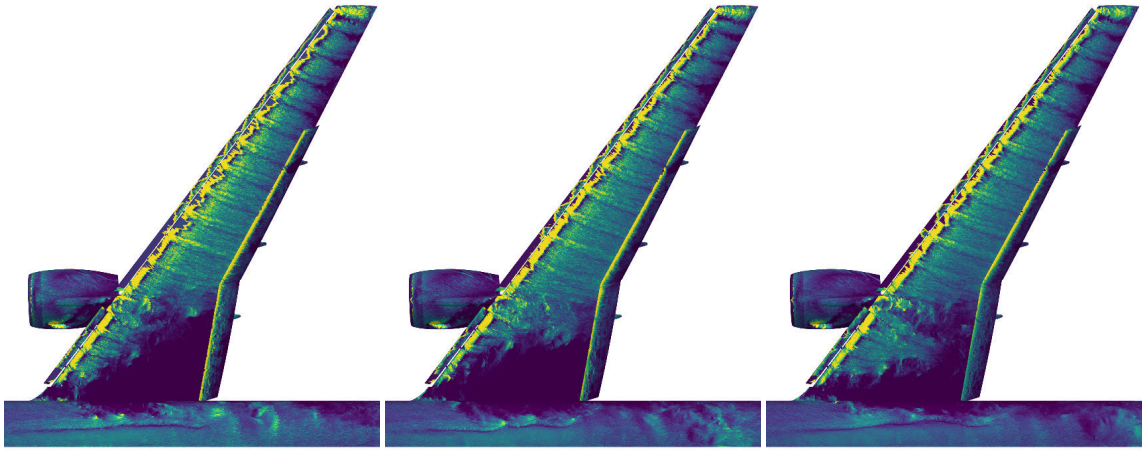

(c) $\alpha = 22°$

Fig. 19 Effect of Reynolds number on the solutions of Case 3 with the fine (L4) grid resolution level for three different angles of attack, showing instantaneous axial skin friction coefficient for $\mathrm{Re} = 5.49 \times 10^6$ **(left),** $\mathrm{Re} = 16 \times 10^6$ **(middle), and** $\mathrm{Re} = 30 \times 10^6$ **(right).**

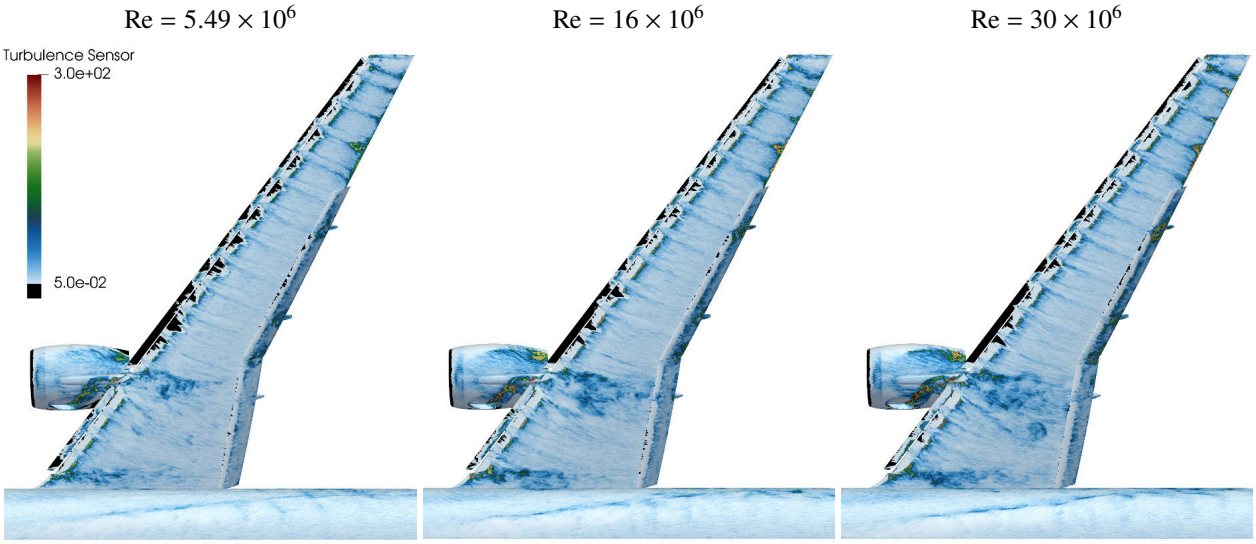

Fig. 20 Wall turbulence sensor field over the wing surface for all three Reynolds numbers at $\alpha = 19°$, using the fine (L4) grid.

C. Computational Cost

The advent of Graphics Processing Units (GPUs) as accelerators for high-performance computing has revolutionized computational science, enabling simulations of increasingly complex geometries with high resolutions. Heterogeneous computing, where GPUs are used as accelerators alongside CPUs, is a common feature in modern supercomputers. Therefore, it is critical that codes can efficiently exploit the capabilities of these diverse architectures.

In this section, we evaluate the performance of the LAVA unstructured WMLES solver on Cabeus and Aitken supercomputers, which are housed in NASA's Advanced Supercomputing Division in Moffett Field, California. The Cabeus cluster consists of 128 nodes, each with an AMD Milan CPU and 4 NVIDIA A100 GPUs. The Aitken cluster is a CPU-only machine with 2048 nodes, each with 2 AMD Rome CPUs.

Table 3 Computational cost of simulations for Case 2.4 using LAVA unstructured WMLES solver. Performance normalized by grid count is listed in terms of Millions of Updates Per Second (MUPS). Energy consumption for the simulations is provided in megajoules per CTU (MJ/CTU). All measures are reported for double-precision computations. Each Cabeus node listed here has 4 NVIDIA A100 GPU cards.

Metric	Units	CPU	GPU
Solve Cells		325M	325M
Resources		100 AMD Rome nodes	16 NVIDIA A100 nodes
Performance	MUPS / node	19.1	192
Wall Time	hours / CTU	1.62	1.01
Energy Consumption	MJ / CTU	467	46.4

Table 3 shows various performance metrics of the LAVA unstructured solver using the numerical scheme and the best practices for the test cases described in this work. In particular, Case 2.4 simulations using the fine (L4) Voronoi grid are employed as a benchmark for performance.

One measure of performance listed in the table is the Millions of Updates Per Second (MUPS), which quantifies the number of cells (in millions) that can be advanced a single time step in one second of wall time. The MUPS figure in the table is given per compute node (i.e. millions of cells updated per second per node). While MUPS is helpful in quantifying the performance, it does not serve well for comparisons across different solvers that utilize different

computational cell types and/or utilize different time integration schemes. The BCC seeded Voronoi grids employed in this study are primarily comprised of truncated octahedra with 14 faces per cell. In contrast to traditional structured and unstructured element types such as hexahedra (6 faces/cell) and tetrahedra (4 faces/cell), these Voronoi grids require 2-3 times more flux evaluations. These additional evaluations can potentially mitigate directional bias, and enhance computational accuracy, as demonstrated by Sozer et al. [34].

The energy consumption, measuring the energy spent for time integration of one CTU, serves as a more consistent metric to compare performance across solvers and compute platforms. The energy consumption is estimated based on peak power recordings of 450 W per node for CPU simulations and 800 W per node for GPU simulations listed in the table.

It would be useful at this point to briefly comment on the coding paradigms employed. The LAVA unstructured is written predominantly in C++, utilizing MPI for distributed memory parallelism, and OpenMP for shared memory parallelism on the CPUs. Careful consideration has gone into vectorizing and optimizing the CPU implementation. The GPU implementation with CUDA is relatively new and further optimizations are in progress.

A comparison between CPU and GPU energy consumptions reveals that the GPU implementation is approximately 10x more energy-efficient per CTU. This efficiency translates to achieving the same computational results while consuming only one-tenth the energy. Further significant performance improvements (as much as 2x increase in speed/reduction of energy consumption) may be realized by switching to a single-precision or a mixed-precision approach. However, our early experience with this approach revealed a significant sensitivity of the $C_{L,max}$ value with respect to the numerical precision for the cases in this work. Hence, double-precision storage and arithmetic remain our current best practice while a mixed-precision approach continues to be actively investigated.

V. Conclusions

This work summarized the results of our multi-year effort to develop an accurate and efficient unstructured WMLES capability for the prediction of high-lift aerodynamics. The work was motivated by NASA's CFD Vision 2030 and driven by our participation in AIAA's 5th CFD High Lift Prediction Workshop. We have described our numerical approach, gridding approach, and best practices in detail.

With Case 2.4, we demonstrated WMLES results in near-perfect agreement with experimental measurements. Integrated aerodynamic forces and moments through a wide range of angles of attack were accurately predicted by WMLES simulations. Of particular interest, capability to accurately predict both the value of $C_{L,max}$, and the angle of attack at which it occurs was demonstrated. During the workshop, a critical improvement to the wall modeling approach was implemented and tested in the LAVA Unstructured solver. This model, based on a wall turbulence sensor, detects the level of resolved turbulence content at the wall and appropriately switches between no-slip and wall function stresses based on the selected sensor threshold. This approach was shown to be effective in removing excessive momentum loss and hence preventing premature separation at the slats and other near-leading edge features. The approach significantly improved the accuracy of $C_{L,max}$ predictions. In this work, the sensitivity of this approach, particularly the choice of the sensor threshold, has been examined in detail with respect to grid resolution and Reynolds number variations. The sensor threshold location was found to be largely insensitive to grid resolution and to have a weak dependence on Reynolds number, as expected.

With Case 3, we focused on simulations of the CRM-HL at Reynolds numbers ranging from 5.49×10^6 up to 30×10^6 and demonstrated the applicability of the WMLES approach at flight scale conditions. While experimental data for this case have not yet been made available, the simulations across the Reynolds number variations demonstrated generally coherent and consistent trends that are in line with expected flow physics.

In addition to accuracy considerations, this work demonstrated the computational feasibility of the WMLES approach for complex geometries at realistic flight conditions. The computational performance as reported in this work has demonstrated that the routine use of WMLES is becoming more and more accessible. The progress in this regard has been greatly accelerated with the advent of GPU-based systems for which we demonstrated an order of magnitude increase in energy efficiency for the same computations, in comparison to modern CPU-based systems. This makes WMLES a powerful technology for aiding the design of higher performance, more efficient, and safer air vehicles.

Going forward, we plan to extend this work by investigating other relevant challenge problems and evaluating potential improvements to the subgrid-scale and wall models. Recognizing the time-consuming nature of iterating between tests of manually specified grid distributions, adaptive mesh refinement WMLES techniques will also be an important area of focus in our future research.

Acknowledgments

This research is sponsored by the NASA Transformational Tools and Technologies (TTT) Project of the Transformative Aeronautics Concepts Program (TACP) under the Aeronautics Research Mission Directorate (ARMD). Computational resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

We extend special thanks to: the 5th High-Lift Prediction Workshop leadership team for organizing and hosting the workshop; HLPW-5 WMLES TFG participants for their insights and technical guidance; Dr. Bram Rodgers and Keshav Sriram for their invaluable contributions to the development of the LAVA Voronoi meshing tool, a critical enabling technology behind this work; Leonardo Machado for his contributions to the meshing process; James Koch, Luis dos Santos Fernandes, and Gerrit-Daniel Stich for their review and essential feedback on this paper; and Johnny Chang and Kyle Rollin for their assistance with the NAS Cabeus GPU cluster.

References

[1] Goc, K. A., Lehmkuhl, O., Park, G. I., Bose, S. T., and Moin, P., "Large eddy simulation of aircraft at affordable cost: a milestone in computational fluid dynamics," *Flow*, Vol. 1, 2021, p. E14. https://doi.org/10.1017/flo.2021.17.

[2] Kiris, C. C., Ghate, A. S., Duensing, J. C., Browne, O. M., Housman, J. A., Stich, G.-D., Kenway, G., Fernandes, L. S., and Machado, L. M., "High-lift common research model: RANS, HRLES, and WMLES perspectives for CLmax prediction using LAVA," *AIAA SciTech 2022 Forum*, 2022, p. 1554. https://doi.org/10.2514/6.2022-1554, URL https://arc.aiaa.org/doi/10.2514/6.2022-1554.

[3] Wang, L., Anderson, W. K., Nielsen, E. J., Balakumar, P., Park, M. A., Carlson, J.-R., Iyer, P. S., and Diskin, B., "Wall-modeled large-eddy simulations for high-lift configurations using FUN3D," *AIAA SciTech 2022 Forum*, 2022, p. 1555. https://doi.org/10.2514/6.2022-1555, URL https://arc.aiaa.org/doi/abs/10.2514/6.2022-1555.

[4] Kiris, C. C., Ghate, A. S., Browne, O. M., Slotnick, J., and Larsson, J., "HLPW-4: Wall-modeled large-eddy simulation and lattice–Boltzmann technology focus group workshop summary," *Journal of Aircraft*, Vol. 60, No. 4, 2023, pp. 1118–1140. https://doi.org/10.2514/1.C037193, URL https://arc.aiaa.org/doi/10.2514/1.C037193.

[5] Slotnick, J. P., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D. J., "CFD vision 2030 study: a path to revolutionary computational aerosciences," 2014. URL https://ntrs.nasa.gov/citations/20140003093.

[6] Malik, M. R., "NASA's Revolutionary Computational Aerosciences (RCA) and CFD Vision 2030," *10th International Conference on Computational Fluid Dynamics (ICCFD10)*, 2018. URL https://ntrs.nasa.gov/citations/20200006466.

[7] Mauery, T., Alonso, J., Cary, A., Lee, V., Malecki, R., Mavriplis, D., Medic, G., Schaefer, J., and Slotnick, J., "A guide for aircraft certification by analysis," 2021. URL https://ntrs.nasa.gov/citations/20210015404.

[8] Rumsey, C. L., Slotnick, J. P., Long, M., Stuever, R. A., and Wayman, T., "Summary of the first AIAA CFD high-lift prediction workshop," *Journal of Aircraft*, Vol. 48, No. 6, 2011, pp. 2068–2079. https://doi.org/10.2514/1.C031447, URL https://arc.aiaa.org/doi/10.2514/1.C031447.

[9] Rumsey, C. L., and Slotnick, J. P., "Overview and summary of the second AIAA high-lift prediction workshop," *Journal of Aircraft*, Vol. 52, No. 4, 2015, pp. 1006–1025. https://doi.org/10.2514/1.C032864, URL https://arc.aiaa.org/doi/10.2514/1.C032864.

[10] Rumsey, C. L., Slotnick, J. P., and Sclafani, A. J., "Overview and summary of the third AIAA high lift prediction workshop," *Journal of Aircraft*, Vol. 56, No. 2, 2019, pp. 621–644. https://doi.org/10.2514/1.C034940, URL https://arc.aiaa.org/doi/10.2514/1.C034940.

[11] Rumsey, C. L., Slotnick, J. P., and Woeber, C., "HLPW-4/GMGW-3: overview and workshop summary," *AIAA Aviation 2022 Forum*, 2022, p. 3295. https://doi.org/10.2514/6.2022-3295, URL https://arc.aiaa.org/doi/abs/10.2514/6.2022-3295.

[12] Kiris, C. C., Housman, J. A., Barad, M. F., Brehm, C., Sozer, E., and Moini-Yekta, S., "Computational framework for launch, ascent, and vehicle aerodynamics (LAVA)," *Aerospace Science and Technology*, Vol. 55, 2016, pp. 189–219. https://doi.org/10.1016/j.ast.2016.05.008, URL https://www.sciencedirect.com/science/article/pii/S127096381630178X?via%3Dihub.

[13] Jensen, J. C., Stich, G.-D., Housman, J. A., Denison, M., and Kiris, C. C., "LAVA Simulations for the 3rd AIAA CFD High Lift Prediction Workshop Using Body Fitted Grids," *2018 AIAA Aerospace Sciences Meeting*, 2018, p. 2056. https://doi.org/10.2514/6.2018-2056, URL https://arc.aiaa.org/doi/10.2514/6.2018-2056.

[14] Duensing, J. C., Housman, J. A., Fernandes, L., Machado, L., and Kiris, C. C., "A Reynolds-Averaged Navier-Stokes Perspective for the High Lift Common Research Model Using the LAVA Framework," *AIAA Aviation 2022 Forum*, 2022, p. 3742. https://doi.org/10.2514/6.2022-3742, URL https://arc.aiaa.org/doi/10.2514/6.2022-3742.

[15] Browne, O. M., Housman, J. A., Kenway, G. K., Ghate, A. S., and Kiris, C. C., "A Hybrid RANS-LES Perspective for the High Lift Common Research Model Using LAVA," *AIAA Aviation Forum and Exposition 2022*, 2022. https://doi.org/10.2514/6.2022-3523, URL https://arc.aiaa.org/doi/10.2514/6.2022-3523.

[16] Ghate, A. S., Kenway, G. K., Stich, G.-D., Maldonado, D., and Kiris, C. C., "A Wall-Modeled LES Perspective for the High Lift Common Research Model Using LAVA," *AIAA Aviation Forum and Exposition 2022*, 2022. https://doi.org/10.2514/6.2022-3434, URL https://arc.aiaa.org/doi/10.2514/6.2022-3434.

[17] Sousa, V., Rodgers, A. K., Sriram, K., Sozer, E., Barad, M. F., Stich, G.-D., Cadieux, F., and Duensing, J., "LAVA Voronoi Mesher for Wall-Modeled Large-Eddy Simulations," *AIAA AVIATION FORUM AND ASCEND 2024*, 2024. https://doi.org/10.2514/6.2024-4306, URL https://arc.aiaa.org/doi/abs/10.2514/6.2024-4306.

[18] Rumsey, C. L., "High-Lift Prediction Workshops: Retrospective, Lessons Learned, and Future Prospects," *34th Congress of the International Council of the Aeronautical Sciences (ICAS)*, 2024. URL https://www.icas.org/icas_archive/icas2024/data/preview/icas2024_1217.htm.

[19] Mouton, S., Charpentier, G., and Lorenski, A., "Test Summary of the Full-Span High-Lift Common Research Model at the ONERA F1 Pressurized Low-Speed Wind Tunnel," *AIAA SCITECH 2023 Forum*, 2023, p. 0823. https://doi.org/10.2514/6.2023-0823, URL https://arc.aiaa.org/doi/10.2514/6.2023-0823.

[20] Wahls, R., "The National Transonic Facility-A Research Retrospective," *39th Aerospace Sciences Meeting and Exhibit*, 2001, p. 754. https://doi.org/10.2514/6.2001-754, URL https://arc.aiaa.org/doi/10.2514/6.2001-754.

[21] Larsson, J., Kawai, S., Bodart, J., and Bermejo-Moreno, I., "Large Eddy Simulation with Modeled Wall-Stress: Recent Progress and Future Directions," *Mechanical Engineering Reviews*, Vol. 3, 2015. https://doi.org/10.1299/mer.15-00418.

[22] De Michele, C., and Coppola, G., "Asymptotically entropy-conservative and kinetic-energy preserving numerical fluxes for compressible Euler equations," *Journal of Computational Physics*, Vol. 492, 2023, p. 112439. https://doi.org/https://doi.org/10.1016/j.jcp.2023.112439, URL https://www.sciencedirect.com/science/article/pii/S002199912300534X.

[23] van Leer, B., "Towards the Ultimate Conservative Difference Scheme. II. Monotonicity and Conservation Combined in a Second-order Scheme," *Journal of Computational Physics*, Vol. 14, 1974, pp. 361–370. https://doi.org/10.1016/0021-9991(74)90019-9.

[24] Sozer, E., Brehm, C., and Kiris, C. C., "Gradient Calculation Methods on Arbitrary Polyhedral Unstructured Meshes for Cell-Centered CFD Solvers," *52nd Aerospace Sciences Meeting*, 2014. https://doi.org/10.2514/6.2014-1440, URL https://arc.aiaa.org/doi/abs/10.2514/6.2014-1440.

[25] Roe, P. L., "Characteristic-Based Schemes for the Euler Equations," *Annual Review of Fluid Mechanics*, Vol. 18, No. 1, 1986, pp. 337–365. https://doi.org/10.1146/annurev.fl.18.010186.002005, URL https://doi.org/10.1146/annurev.fl.18.010186.002005.

[26] Kim, K. H., Kim, C., and Rho, O.-H., "Methods for the Accurate Computations of Hypersonic Flows: I. AUSMPW+Scheme," *Journal of Computational Physics*, Vol. 174, No. 1, 2001, pp. 38–80. https://doi.org/https://doi.org/10.1006/jcph.2001.6873, URL https://www.sciencedirect.com/science/article/pii/S0021999101968731.

[27] Mittal, R., and Moin, P., "Suitability of upwind-biased finite difference schemes for large-eddy simulation of turbulent flows," *AIAA Journal*, Vol. 35, No. 8, 1997, pp. 1415–1417.

[28] Ducros, F., Ferrand, V., Nicoud, F., Weber, C., Darracq, D., Gacherieu, C., and Poinsot, T., "Large-Eddy Simulation of the Shock/Turbulence Interaction," *Journal of Computational Physics*, Vol. 152, No. 2, 1999, pp. 517–549. https://doi.org/10.1006/jcph.1999.6238.

[29] Jalali, A., Sharbatdar, M., and Ollivier-Gooch, C., "Accuracy analysis of unstructured finite volume discretization schemes for diffusive fluxes," *Computers & Fluids*, Vol. 101, 2014, pp. 220–232. https://doi.org/https://doi.org/10.1016/j.compfluid.2014.06.008, URL https://www.sciencedirect.com/science/article/pii/S0045793014002473.

[30] Gottlieb, S., Shu, C.-W., and Tadmor, E., "Strong Stability-Preserving High-Order Time Discretization Methods," *SIAM Review*, Vol. 43, No. 1, 2001, pp. 89–112. https://doi.org/10.2307/3649684, URL http://www.jstor.org/stable/3649684.

[31] Rizzi, A. W., and Inouye, M., "Time-Split Finite-Volume Method for Three-Dimensional Blunt-Body Flow," *AIAA Journal*, Vol. 11, No. 11, 1973, pp. 1478–1485. https://doi.org/10.2514/3.50614, URL https://doi.org/10.2514/3.50614.

[32] Musker, A. J., "Explicit Expression for the Smooth Wall Velocity Distribution in a Turbulent Boundary Layer," *AIAA Journal*, Vol. 17, No. 6, 1979, pp. 655–657. https://doi.org/10.2514/3.61193, URL https://doi.org/10.2514/3.61193.

[33] Vreman, A., "An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications," *Physics of Fluids*, Vol. 16, No. 10, 2004, pp. 3670–3681.

[34] Sozer, E., Ghate, A. S., Kenway, G. K., Barad, M. F., Sousa, V. C., and Kiris, C. C., "Evaluation of Voronoi Meshes for Large Eddy Simulations of High Lift Aerodynamics," *AIAA SCITECH 2023 Forum*, 2023. https://doi.org/10.2514/6.2023-0255, URL https://arc.aiaa.org/doi/abs/10.2514/6.2023-0255.

[35] Wang, L., "5th AIAA CFD High Lift Prediction Workshop (HLPW-5)," , 2024. URL https://hiliftpw.larc.nasa.gov.