

Observations from the HLPW05 Test Case 2 Configuration Build-Up using ADS Code Leo

Gregorio Robles Vega, Edward E. Bryan, Charles DeLorenzo, Michael Ni, Hsin-Hua Tsuei
AeroDynamic Solutions Inc., San Ramon, CA, 94583, U.S.A

As part of the 5th AIAA High Lift Prediction Workshop (HLPW05), wind-tunnel tests and CFD analyses were conducted to assess the aerodynamic performance of the High Lift Common Research Model (CRM-HL) on four progressively complex configurations. The present paper outlines observations obtained from steady RANS simulations using the ADS GPU accelerated flow solver, Code Leo. For test case 2.2, a mesh convergence study is performed using 3 mesh densities of the HeldenMesh mixed element unstructured mesh family provided by the committee. Simulated results were compared against 1D, loading and oil flow experimental data by ONERA. A second family of unstructured mesh, generated using Ennova, were also calculated and showed similar results. Lastly, a single mesh level from the committee provided ANSA unstructured meshes was also analyzed. At lower angles of attack, all selected results showed good agreement with the experimental data. However, discrepancies were noted at higher angles of attack, where the RANS simulations predicted lower lift and higher pitching moments than data.

Test case 2.4 was thoroughly investigated across 3 different mesh families provided by three different mesh generation software. It was shown that nearly identical results could be obtained with different mesh topologies from different users. This highlights the consistent results across mesh topologies and user experience. Several custom mixed element unstructured meshes of varying mesh densities were generated using Ennova. The computed results from the Ennova mesh family were compared against committee mesh simulations. It was found that the Ennova mesh family compared well to the committee provided meshes, but at a reduced cell count and more efficient computational time.

Lastly, a comparison against the HLPW04 was performed. The HLPW04 and HLPW05 models are based on the same geometry, the NASA CRM-HL, but have some small differences. The major geometric difference explored in the comparison was the tail-off vs tail-on prediction. This comparison demonstrated that with the same mesh topologies and solver the aerodynamic results are nearly identical, other than expected effects from geometrical differences, which enables the confidence level needed to know that small design changes or mesh topology decisions will not lead designers to inconsistent results.

Overall, the current research illustrated the capabilities and potential limitations of RANS simulations in predicting high-lift aerodynamic performance, particularly at higher angles of attack. The study also underscored the potential benefits of a different mesh generation topology for improving computational efficiency without compromising accuracy.

I. Nomenclature

<i>ADS</i>	=	<i>AeroDynamic Solutions Inc.</i>
<i>RANS</i>	=	Reynolds-Averaged Navier-Stokes
<i>CFD</i>	=	Computational Fluid Dynamics
<i>CPU</i>	=	Central Processing Unit
<i>GPU</i>	=	Graphics Processing Unit
<i>AoA</i>	=	Angle of Attack, degrees
<i>CP</i>	=	Pressure Coefficient
<i>CF</i>	=	Skin Friction Coefficient
<i>C_L</i>	=	Lift Coefficient
<i>C_D</i>	=	Drag Coefficient
<i>C_M</i>	=	Pitching Moment Coefficient

II. Introduction

Consistent and accurate predictions of aircraft in high-lift configurations have been and continue to be a challenge in the aeronautics industry. For this reason, workshops such as the AIAA High Lift Prediction Workshop (HLPW) [1] aim to establish collaboration among industry, government and academia to address this challenge. The latest of these efforts was the 5th High Lift Prediction Workshop (HLPW05) that took place in 2024 [2]. In test case 2 of HLPW05, components were incrementally added to determine at which configuration and conditions the prediction begins to struggle against measured performance data. This paper outlines the observations made with the ADS GPU-accelerated flow solver, Code Leo.

The results presented in this paper are all steady RANS solutions using the SA turbulence model. Steady RANS is still the fastest method of obtaining an initial design performance. This attribute is key for the design cycle when a significant number of geometric changes are analyzed in a fast-paced design cycle in the modern aviation industry. It is understood that steady RANS simulations are not ideal for highly separated unsteady flows. However, turbulence modeling and time-accurate analysis are not the only important factors in accurate design cycle predictions. Often, the most important item is the mesh size required for consistent predictions, and independence from user-to-user mesh variance. A smaller mesh can be run at greater speed and enable more thorough exploration of the design space. Similarly, independence from user-to-user mesh variance allows for a consistent comparison of conclusions across users of different experience levels and is a key for performance assessment and potential automation. From this perspective, the solutions from steady RANS can provide an accurate enough solution that captures trend directions from given design changes.

Recently a greater push has been made to accelerate simulations by switching from CPUs to GPUs. This not only brought significant improvement in per iteration turnaround time, but also a reduction in compute cost per simulation, resulting in over 100X improvement in compute efficiency. All the results presented on this paper were all run on GPUs ranging from consumer grade (laptop or workstation) to data-center grade.

The following sections will detail the steady RANS predictions made with Code Leo for two of the four configurations that are part of test case 2. The simulations were made with various committee [3] and custom generated meshes. Details regarding the 1D performance will be augmented with loading and contour comparisons. Finally, a comparison will be made between HLPW04 TC2 (test case 2) and HLPW05 TC2.4 results.

III. ADS CFD Code Leo

The present numerical simulations were carried out using Code Leo, a commercial GPU-accelerated flow solver from ADS CFD Inc. Code Leo is a density-based, compressible, explicit time-marching flow solver that uses the Ni-Scheme distribution formula [4] and is 2nd order accurate in both space and time. It has low numerical dissipation and uses a blend of 2nd and 4th order smoothing to reduce oscillations of the flowfield caused by shocks and transient flows. The solver is applicable to structured, unstructured, and mixed element meshes. Code Leo uses a finite volume cell-vertex discretization scheme and is able to run steady or time-accurate RANS, with either Dual-Time Stepping method (DTS) or Non-linear-Harmonic Balance Method (NHB). Higher fidelity simulations such as Delayed Detached Eddy Simulation (DDES) and Wall-Resolved Large Eddy Simulation (WRLES) are also available. Wall modeling and wall function approaches are supported. Parallel simulations scale efficiency to 1000s of cores through high speed MPI.

Code Leo has been applied extensively to turbomachinery applications, such as compressors [5], turbines [6], and fans [7, 8]. This includes full-wheel unsteady RANS simulations of multi-stage machines, as well as bypass/core flow domains. Conjugate Heat Transfer (CHT) [9] and flutter/aero-damping simulations [10] have also been conducted using Code Leo. This includes WRLES simulations for high-lift high-work low-pressure turbines [11, 12] as well as coupled CHT/WRLES simulations for cooled turbines. Code Leo is also used in various industries such as VTOL and external aerodynamics applications. For external aerodynamic analysis, Code Leo has been applied to airframe/engine coupling to analyze distortion effects at various flow conditions. The solver has been validated in cruise and high-lift conditions through many engineering projects.

As of 2018, Code LEO was completely ported over to GPUs without a full rewrite of the underlying numerical algorithm. Thus maintaining all validations from previous versions of the code whether a user is running on the CPU or the GPU. Regarding GPUs, Code Leo can run on both consumer-grade (laptop or workstation) to data-center grade GPUs. This flexibility enables a significant advantage to end-users, as consumer-grade GPUs can supply nearly 90% of the performance gains of data-center GPUs at only a fraction of the cost. This represents compute efficiency improvements of nearly 3 orders of magnitude over CPU based compute clusters.

IV. Mesh Generation

Simulations were conducted with both committee and custom generated Ennova meshes. Committee meshes were provided by the HLPW05 workshop and a summary for those used is detailed below. Details regarding the meshing approach and mesh settings for the custom generated meshes are also provided.

A. Mesh Description

Two committee provided mesh families were used in the simulations presented in this paper: HeldenMesh 2.R.01 provided by Helden Aerospace [3] and ANSA 2.L.01 provided by Beta-CAE. The mesh topology varied substantially between both mesh families and not all meshes in the mesh families were run. For the most part, at least three mesh sizes were used, except for TC2.2 for which only one ANSA mesh size was used. Based upon previous experience, custom meshes were generated using the Ennova

unstructured mesh generator. Table 1 lists the mesh sizes of all the HLPW05 meshes referenced in this paper. For the HeldenMesh and ANSA meshes, either the .cgns or .ugrid file types were used to convert the committee meshes into the Code Leo Restart file format used in the simulations. For Ennova meshes, the .reds file format is supported and can be imported into Code Leo directly.

Table. 1 Description of mesh counts and first layer height for the HLPW05 Meshes.

<i>Mesh Name</i>	<i>Number of Nodes (Million)</i>	<i>Number of Cells (Million)</i>	<i>First Layer Height (inches)</i>
TC2.2 HeldenMesh			
<i>Coarse</i>	3.6M	11.3M	0.0036444
<i>Medium</i>	9.3M	26.3M	0.0024296
<i>Fine</i>	56.8M	140.9M	0.0012148
TC2.2 ANSA			
<i>A</i>	45.0M	55.0M	0.00053
TC2.2 Ennova			
<i>Coarse</i>	20.7M	38.5M	0.0005
<i>Medium</i>	32.4M	59.4M	0.0005
<i>Fine</i>	55.8M	108.9M	0.0005
TC2.4 HeldenMesh			
<i>Coarse</i>	6.0M	20.2M	0.0036444
<i>Medium</i>	14.7M	43.8M	0.0024296
<i>Fine</i>	81.5M	205.1M	0.0012148
TC2.4 ANSA			
<i>B</i>	111.0M	130.0M	0.00053
<i>C</i>	220.0M	247.0M	0.00053
<i>C+</i>	249.0M	281.0M	0.00053
TC2.4 Ennova			
<i>Medium</i>	50.4M	90.7M	0.0005
<i>Medium+</i>	53.0M	96.3M	0.0005
<i>Fine+</i>	117.1M	246.3M	0.0005

The HeldenMesh 2.R.01 mesh family was generated using HeldenMesh v4.15 by Andrew Wick and Rick Hooker from Helden Aerospace. This mesh family was used for both TC2.2 and TC2.4. In both cases the Coarse, Medium and Fine meshes were used. Information regarding the mesh sizes and first layer heights can be found in Table 1. The mesh family had a first layer height that would decrease with increasing mesh density. Helden Aerospace reported that automated wake sources were added to the trailing edges and expected separation regions. This volume mesh refinement can be seen in Figure 1. The figure shows a y-cut at Y=420in of the volume mesh for the TC2.4 HeldenMesh Fine, where the mesh around the nacelle, slat, wing, and flap is visible.




Fig. 2 Image of volume mesh around the nacelle and wing at Y=420in for the TC2.4 HeldenMesh Fine.

The ANSA 2.L.01 mesh family was generated using ANSA v24.0.1 by Vangelis Skaperdas from Beta-CAE. This mesh family was used for both TC2.2 and TC2.4 with only the Level A mesh used for TC2.2 and for TC2.4, mesh Levels B through C were used along with level C+. Mesh level C+ was generated after the original mesh family and was guided by observations from the solutions from the original mesh family. It should be noted that the ANSA 2.L.01 meshes were part of the committee meshes for the HRLES Technical Focus Group (TFG). However, in this case the results presented are for steady RANS. Details regarding mesh sizes and first layer heights can be found in Table 1. The first layer height for the ANSA 2.L.01 mesh family (0.00053in) is maintained constant for all mesh levels at a value smaller than that used for the HeldenMesh 2.R.01 Fine mesh (0.0012148in). However, it should be noted that Helden Aerospace also provided a mesh family for the HRLES TFG named HeldenMesh 2.L.02 and this family is reported to have a constant initial first layer height that remains constant with mesh resolution. Thus, this difference in approach for the first layer height seems to be due to the TFG the mesh family was prepared for. Regarding the mesh topology, Figure 2 shows a similar cut for the TC2.4 ANSA C+ mesh as was provided for the TC2.4 HeldenMesh Fine in Figure 1. This volume y-cut is located at Y=420in and shows the mesh around the nacelle, slat, wing, and flaps. It can be observed that the mesh topology varies significantly between HeldenMesh and ANSA with the ANSA meshes using more hexahedral elements. Beta-CAE also reported that the ANSA mesh also used size functions to refine the wake of the wing, fuselage and tip vortex at an AoA range of 6 to 24 degrees.




Fig. 2 Image of volume mesh around the nacelle and wing at Y=420in for the TC2.4 ANSA level C+.

Lastly, a set of custom mixed-element unstructured meshes were generated using Ennova v1.10. The custom TC2.2 Ennova meshes have been provided to the workshop and are listed as 2.R.05. All custom Ennova meshes have no wake refinement added to the volume mesh. Meshes were generated automatically after setting a few global and local mesh size and meshing parameters. A first layer height of 0.0005in was used across all meshes after it was found that values around the order of 0.001in would yield higher than desired y^+ in critical areas around the slat and flap. It should be noted that for TC2.4 the “+” in the mesh name denotes a change in prism layer meshing approach. In this case more prism layers were added to the boundary layer. Figure 3 shows a view of the TC2.4 Ennova Fine+ volume mesh. This volume cut is at $y=420$ in and shows a similar view as the other mesh types used in the analysis.

An aspect of the three mesh types that should be considered is the significant differences in meshing approach across all three types. The volume mesh differences can be observed across Figures 1-3. These show differences in the dominant element types as well as a different approach to volume refinement in expected wake regions. Similarly, Figure 4 shows a comparison of the surface mesh around a slat support bracket and the pylon region. This figure highlights the differences between element types and surface feature refinements preferred across meshes. The selection of such varying mesh types was done intentionally to assess how the flow solver performs across a wide range of meshing techniques.




Fig. 3 Image of volume mesh around the nacelle and wing at Y=420in for the TC2.4 Ennova Fine+.




Fig. 4 Surface meshing approach across the three mesh types explored.

V. Simulation Approach

The simulation setup and execution approach are detailed in this section. The simulations are initialized based on the conditions specified by the workshop. The viscosity was scaled to achieve the target Chord Reynolds Number of 5.9×10^6 . The simulations were initialized with a uniform flowfield throughout the domain. All simulations presented here used the SA model and wall-integration though additional work was done using a 2 equation KW-98 turbulence model. When running across multiple GPUs,

the domain was broken up across the number of GPUs available using MeTIS breakup methods. AoA sweeps were then executed in a serial manner using a warm-start approach. In other words, each consecutive AoA was ran using the converged solutions from the previous AoA.

Generally speaking, simulations reached a stable C_L , C_D and C_M by 10,000 iterations. Most cases were carried out to 30,000 iterations to ensure full convergence. At higher AoA periodic oscillations appeared in some of the convergence history. If the simulations did not reach a stable value or a periodic oscillation, the simulation was run for another 30,000 iterations. For the simulations with periodic oscillations, the final 10,000 iterations were averaged and reported as the 1D value.

VI. HLPW05 TC2.2 Results

A. Test Case 2.2 Description

The first test case studied in this paper is TC2.2 from HLPW05. The geometry was named the ONERA Large Reference Model (LRM). The configuration for TC2.2 consists of the fuselage, horizontal stabilizer, vertical stabilizer, wing, and slats. The ONERA LRM is based on the NASA Common Research Model-High Lift (CRM-HL), but with variations in slat support placement, nacelle geometry, among others. The geometry used in the simulations is full-scale and half-span. Simulations were run free-air with no wind tunnel walls or support structures. The target freestream mach number is 0.2 with Reference Static Temperature equal to 518.67 °R and Reference Static Pressure of 14.696 psi. The target Chord Reynolds Number is 5.9×10^6 and is achieved by scaling the viscosity in the simulation.

The CFD results are compared against the test data collected on a 1/19.5 full-span model in the ONERA F1 wind tunnel and provided to the workshop. The data consisted of four angle of attack (AoA) sweeps: an upsweep and downsweep for both a mono-strut and tri-strut configuration. In the figures below, the four AoA sweeps are combined into one dataset and error bars were added based on those reported [13].

B. Test Case 2.2 1D Results

Two of the three mesh types being evaluated for TC2.2 had a mesh convergence check performed across three mesh levels. The mesh convergence checks were conducted for the HeldenMesh and Ennova mesh types. Three mesh densities were run for both types. Details regarding mesh sizes can be found in Table 1. The HeldenMesh family of meshes range from 9.3M to 140M cells from Coarse to Fine levels. The HeldenMesh results can be seen in Figure 5 and show reasonable mesh convergence up to AoA=10.0°. Higher AoA showed an increasingly larger C_L increase between the Medium and Fine mesh levels. The HeldenMesh Fine mesh level was used for comparison against experimental data. The Ennova mesh family has cell ranges between 38M and 100M cells from Coarse to Fine with simulated results seen in Figure 6. The Ennova results also showed reasonable mesh convergence up to AoA=10.0° for all three mesh sizes, but at higher AoA, C_L converges after the medium level of 59M cells. This suggests that overall mesh convergence can be achieved with either mesh family at an overall cell count well under the 140M cells of the HeldenMesh fine mesh.

The selected two mesh levels from the mesh convergence study discussed above are compared to ANSA A and the ONERA experimental data. Figure 7 shows the C_L , C_D and C_M results from the AoA sweeps. All three mesh types showed similar predictions, with HeldenMesh Fine having only a slightly lower C_L than the other two at higher AoA. In general, the percent differences in C_L were less than 0.9% at all points prior to peak lift, except for AoA=17.7° where the largest percent difference rose to 1.6%. This consistency is also observed in the AoA when peak lift was achieved. HeldenMesh Fine and Ennova Fine both predicted the peak lift at AoA=23.8° while ANSA A is at AoA=23.5°. When compared against the experimental data, the predicted results compare very well to data at the lower angles up to AoA=19.7°. At the higher AoA, the predicted C_L is lower than the data. The HeldenMesh Fine, Ennova Fine and ANSA A have a 3.9%, 3.1 and 5.1% percent error, respectively when comparing to the $C_{L,Max}$ of the experimental data. Overall, the predictions compare well to both C_L data and each other.

For C_D , the predictions compared well to each other and the data. At AoA below 14.0° all three predictions were within the error bars and experimental variation between mono-strut and tri-strut. However, it was observed that the variations of the predictions do start to increase at higher AoA. Though all of them followed the general trend and magnitudes, Ennova Fine does have higher C_D values at higher AoA relative to the other predictions.

Regarding C_M , the predictions compare well to each other and data at angles below 17.7°. The HeldenMesh Fine and Ennova Fine have different paths but both are within the variation observed between the mono-strut and tri-strut in the experiment. Pitch break occurs near the same AoA of 17.7° when all three predictions trend towards higher C_M than data. The HeldenMesh also starts to deviate away from the other two meshes between 19.7° and 21.5°. At angles above 21.5° all three meshes trend towards lower C_M but are still above the data.




Fig. 5 AoA sweep of TC2.2 HeldenMesh across three mesh levels.




Fig. 6 AoA sweep of TC2.2 Ennova across three mesh levels.




Fig. 7 Comparison of HeldenMesh, Ennova and ANSA mesh types against the ONERA experimental data.

C. Test Case 2.2 Loading Results

To get a more detailed understanding of the differences of the three mesh types, loading profiles were evaluated at two AoA: 17.7° and 21.5°. The loading profiles were also compared against the ONERA experimental data. The profiles are taken at Belt A, Belt E and Belt I which were defined by HLPW05. The loading profiles can be seen in Figure 8.

For AoA=17.7°, shown to the left of Figure 8, all three solutions agreed well with each other and data at Belt A. At Belt E, a divergence in the mesh types was observed. Ennova Fine and ANSA A both agreed well with data and each other. However, the HeldenMesh Fine showed flow separation about mid-chord of the wing, as well as slightly lower lift on the slat. Interesting to note is the difference in incidence that started to appear at this station. Finally, at Belt I all three predictions started to align together again though some differences were observed for the slat incidence and at the wing mid-chord and trailing edge. At this station, the ANSA A had the best comparison against data with the Ennova Fine mesh following a similar prediction except for the wing trailing edge.

Shifting to the AoA=21.5°, a good comparison among all three predictions and data was observed for Belt A and Belt E. At Belt I the mesh types diverged from each other. In this case the Ennova Fine mesh showed lower lift on the slat and wing compared to both other predictions and data. Upon further analysis, it was found that the Ennova Fine mesh experienced the wedge-shaped flow separation commonly referred to as a “pizza-slice separation” thus leading to the lower lift in this section of the wing.

Compared to the ANSA A and Helden Fine predictions, the HeldenMesh Fine had slightly lower lift on the slat and leading edge of the main wing. All three predicted lower lift than experimental data at the main wing leading edge.




Fig. 8 CP loading profiles HeldenMesh Fine, Ennova Fine and ANSA A compared to ONERA experimental data at AoA=17.7° (left) and AoA=21.5° (right) for Belt A, E, and I.

D. Test Case 2.2 Surface Contour Results

The last comparison performed was to evaluate how the skin friction (CF) surface contours compare across all three mesh types. The first comparison is at an AoA=17.7° as seen in Figure 9. The comparison showed the observed flow separation around Belt E (midspan) for the HeldenMesh Fine. Another interesting perspective is how the Ennova Fine showed sharper signatures of the two vortices coming off each slat support, while ANSA A and HeldenMesh Fine had similar but more dissipated features. Similarly, ANSA A showed a larger area of lower CF values at the trailing edge of the wing. Regarding the fuselage, the flow features between ANSA A and HeldenMesh Fine were similar. On the horizontal stabilizer, the flow features between Ennova Fine and ANSA A compared more favorably to each other than the HeldenMesh Fine.

A similar comparison at AoA=21.5° is shown in Figure 10. The CF contours show the appearance of a pizza-slice separation on the Ennova Fine. This led to further exploration of which meshes experienced this phenomenon and at which point. It was found that the HeldenMesh Fine started observing pizza-slice separation at AoA=22.5° and higher. The Ennova Fine mesh experienced it at AoA=19.7° and higher. The ANSA A did not experience the pizza-slice separations at any point that was modeled. It should also be noted that both HeldenMesh Fine and Ennova Fine did not experience the pizza-slice separation along the same slat support bracket. The HeldenMesh Fine experienced it along the 2nd bracket from the wing tip while the Ennova Fine experienced it along the 3rd bracket from the wing tip. An observation is that there was no clear signature of a change in behavior in the 1D parameters at the point when the pizza-slice separations formed. Regarding the slat support vortices, all three exhibited similar features but Ennova Fine had a sharper resolution. Again, the ANSA A results do show lower skin friction that was more washed-out at the trailing edge. The observations for the fuselage also hold for this AoA, with the HeldenMesh Fine and ANSA A having more resolved flow features. On the horizontal stabilizer, all three meshes show slight variation on the features.




Fig. 9 Skin Friction surface contours at $AoA=17.7^\circ$.




Fig. 10 Skin Friction surface contours at $AoA=21.5^\circ$.

E. Section Summary

In summary, three mesh types with significantly different topologies were compared for TC2.2. The 1D results showed good agreement both among the predicted 1D values and the data. The 1D comparison was followed by exploring the loading profiles and Skin Friction surface contours. This portion of the comparison revealed good agreement between data and predicted loading at Belt A (inboard station). However, at $AoA=17.7^\circ$ on Belt E (midspan) the HeldenMesh Fine simulation experienced a sudden flow separation on the wing that was not seen in the other two meshes. This feature resolved itself by $AoA=21.5^\circ$ and all three mesh types again compared well against data. The last loading profiles compared were at Belt I, where the Ennova Fine simulation experienced a pizza-slice separation which caused its CP to deviate from the other solutions and data. The use of CF surface contours further shed light on the features predicted by the different meshes. Variations across the mesh types were observed regarding the sharpness of the slat bracket vortices traces on the wing and the onset/location of the pizza-slices.

VII. HLPW05 TC2.4 Results

A. Test Case 2.4 Description

The second test case that was explored was TC2.4 from HLPW05. The geometry is again the full-span ONERA LRM with the addition of the flaps, nacelle, and associated parts as compared to TC2.2. The CFD model is full-scale and half-symmetry model. The freestream mach number was 0.2 with a Reference Static Temperature of 518.67 °R and Reference Static Pressure of 14.696 psi. The Chord Reynolds Number of 5.9×10^6 was achieved by scaling the viscosity. Like TC2.2, the results are compared against a combination of the four AoA sweeps of wind tunnel data provided by ONERA with the reported error bars incorporated into the plots.

B. Test Case 2.4 1D Results

A mesh convergence check was performed for all three mesh types in TC2.4. Each mesh convergence check consisted of three mesh densities and an entire AoA sweep. Details regarding the mesh counts were provided in Table 1.

The first mesh convergence check is the one performed for the HeldenMesh is shown in Figure 11. The HeldenMesh family of meshes ranged from 20.2M to 200M cells with large jumps between medium and fine mesh levels which resulted in significant changes in behavior across mesh levels. Regarding C_L , the Coarse and Medium meshes predicted similar values up to $AoA=14^\circ$. The slope of this region then changed for the Fine mesh. At the higher AoAs, the $C_{L,Max}$ and slope also varied significantly between the three mesh levels. The behavior for C_D , was more consistent at lower AoA with variation only at AoA past 17.7° . The C_M variation is apparent at low AoA. The Fine mesh level was the one selected for the HeldenMesh family and used in the additional comparisons.

The second mesh convergence check was for the Ennova mesh family. It should be noted that only two mesh levels, Medium+ and Fine+, followed the same meshing approach. The Medium level was similar but had fewer prism layers than the other two and cell counts varied from approximately 90M for the Medium and Medium+ meshes to 249M for the Fine+ mesh. The results can be seen in Figure 12. The C_L values for all three mesh levels for Ennova compared very well up to $AoA=10^\circ$. After that point the Medium+ and Fine+ mesh levels diverged from the Medium for $AoA=14^\circ$. Then for $AoA=16^\circ$ the C_L aligned again before again diverging at higher AoA. The Medium+ and Fine+ agreed very well in C_L along the entire AoA sweep even though their mesh sizes differed by 150M cells. Regarding C_D , all three mesh levels agreed well except for $AoA=14^\circ$ and post peak C_L . Similar to the other parameters, C_M agreed well until $AoA=14^\circ$. The Fine+ mesh level was the one selected for the Ennova family and used in the additional comparisons.

The last mesh convergence check was for the ANSA mesh family. This level again had two mesh levels which used a similar meshing approach, Level B and Level C, and one that deviated from that approach, Level C+. Mesh cell counts ranged from 130M to 281M cells from B to C+. The results for this mesh convergence check can be seen in Figure 13. A good agreement was observed between the Level C and Level C+ mesh across all properties. This is as expected since Level C+ was intended to be similar to Level C but with some improvements based on observations from the solutions. For C_L , the Level B mesh showed good agreement up to $AoA=10^\circ$ with the other two meshes. At higher AoA its C_L was lower by a consistent amount across various AoAs. All three mesh levels showed a similar C_D except at $AoA=14^\circ$. Similarly for C_M , there is a difference in slope between Level B and the other two mesh levels. The Level C+ mesh level was the one selected for the ANSA family and used in the additional comparisons.

The selected meshes from each mesh type were compared against each other and the ONERA experimental data in Figure 14. The comparison showed reasonable agreement between Ennova Fine+ and ANSA C+ up to peak lift. The two mesh types had a good prediction of C_L relative to data up to $AoA=14^\circ$ in the linear regime of the lift curve. On the other hand, the HeldenMesh Fine deviated from data after $AoA=7.6^\circ$. Past 14° all mesh types predicted lower C_L than data, though Ennova Fine+ and ANSA C+ had higher C_L than HeldenMesh Fine in that range. Regarding peak lift, all three mesh types predicted peak lift to occur at $AoA=19.7^\circ$. The HeldenMesh Fine, Ennova Fine + and ANSA C+ had a $C_{L,Max}$ percent error of -4.6%, -3.7% and -2.5% compared to experimental data, respectively.

Regarding C_D , all three mesh types compared well against data. At lower AoA the HeldenMesh Fine and ANSA C+ had slightly lower values than Ennova Fine+ and data. This trend reversed for AoA above 16° when Ennova Fine+ had slightly higher values than the other mesh types and data.

Lastly, C_M compared well for all three mesh types and data at very low AoA up to 10° . All three selected meshes also predicted the point of pitch break around the same $AoA=14^\circ$, but far from experimental data of 16° . At $AoA=14^\circ$ the Ennova Fine+ and ANSA C+ had lower C_M than HeldenMesh Fine but were still above data. All three meshes deviated from data for higher AoAs.




Fig. 11 AoA sweep of TC2.4 HeldenMesh across three mesh levels.




Fig. 12 AoA sweep of TC2.4 Ennova across three mesh levels.




Fig. 13 AoA sweep of TC2.4 ANSA across three mesh levels.




Fig. 14 Comparison of HeldenMesh, Ennova and ANSA mesh types against the ONERA experimental data.

C. Test Case 2.4 Loading Results

The three mesh types were further investigated by extracting the loading profiles for Belt A, E and I at $AoA=10^\circ$ and 19.7° . These results can be seen in Figure 15 where the left set of loading profiles belongs to $AoA=10^\circ$ and the right set is for $AoA=19.7^\circ$. Focusing on the $AoA=10^\circ$ loading profiles, it was observed that overall, there was good agreement between all three mesh types and data. Throughout all three stations the Ennova Fine+ mesh predicted slightly higher lift on the slat, wing leading edge and flaps than the other two mesh types. This was especially noticeable in the flaps. The calculated incidence for Belt E and I also showed variation between mesh types. Nonetheless, the loading distributions compare well with each other and the experimental data. This agrees with the 1D lift for this AoA .

Shifting to $AoA=19.7^\circ$, which is the peak lift for all three mesh types, more variation from one another is observed. Belt A showed reasonable agreement between all three mesh types for the wing and flaps but not for the slat. On the slat HeldenMesh Fine had a noticeably lower lift than the other two mesh types. All three meshes compared well to data on the wing loading. However, on the flap all three meshes predicted higher lift than data. For the slat, HeldenMesh Fine compared better to data in portions but failed to get the correct incidence. For Belt B, Ennova Fine+ and ANSA C+ compared well to each other while HeldenMesh Fine deviating from the two. Ennova Fine+ and ANSA C+ compared well to the slat and wing data, but had lower lift on the flap. Though the HeldenMesh Fine had lower lift on the slat and wing it compared reasonably well to data on the flap. Finally, for Belt I all three mesh types predicted lower lift to data. This was because all three meshes experiencing pizza-slice separations.




Fig. 15 CP loading profiles at $AoA=10^\circ$ (left) and $AoA=19.7^\circ$ (right) for Belt A, E, and I.

D. Test Case 2.4 Surface Contour Results

The last comparison performed for TC2.4 was the surface CF contours at 10° and 19.7° , which can be seen in Figure 16 and 17, respectively. At $AoA=10^\circ$, all three mesh types showed the same flow features on the fuselage, horizontal stabilizer, and outboard wing. However, there were some variations downstream of the nacelle. At $AoA=19.7^\circ$, all three mesh types had pizza-slice separations along the same slat support bracket, 3rd from outboard. Further investigation into the inception of the pizza-slice separation found that HeldenMesh Fine experienced the separation as early as $AoA=14^\circ$, while the other two mesh types developed it at $AoA=16^\circ$. Focusing on the fuselage and horizontal stabilizer, all three mesh types observed similar flow features. On the inboard wing, the flow features were similar except for the flow features resolutions coming off the nacelle. Note that the nacelle flow features were similar among all three mesh types.

Figure 18 shows a comparison of an oil flow image from the ONERA wind tunnel test to the three mesh types at $AoA=19.7^\circ$. This figure shows that there were four flow separation zones observed in the predicted surface contours:

1. Inboard section of the inboard flap
2. Near Fairing Support 1
3. Intersection of inboard and outboard flap
4. Near Fairing Support 2

An interesting aspect of this observation was how the flap flow features variation is dependent on the mesh type. The HeldenMesh Fine exhibited all four separation zones. Ennova Fine+ exhibited only zones 2 and 3, and ANSA C+ exhibited zones 2, 3 and 4. The oil flow images from the experiment show separation only at zones 2 and 3. Notably, the oil flow image showed an additional zone that was not seen in any of the mesh types and was near the Fairing Support 3 but was much smaller than the other two.




Fig. 16 Skin Friction surface contours at AoA=10°.




Fig. 17 Skin Friction surface contours at AoA=19.7°.




Fig. 18 Streaklines comparison to ONERA oil flow at AoA=19.7°.

E. Section Summary

In summary, a total of nine meshes from three mesh types were analyzed for TC2.4. Each mesh type had a mesh convergence check performed using three mesh levels. Two of the mesh levels were able to pass the mesh convergence check. A critical observation was how the Ennova Medium+ and Ennova Fine+ predicted similar performance throughout the AoA sweep even though the mesh size varied by 150M cells. The 1D results showed how the Ennova Fine+ and ANSA C+ compared more favorably to data in both the linear and higher AoA portions. The loading profiles also showed good agreement to data as well in regions not affected by the pizza-slice separations. Lastly, the surface contours showed similar flow features in most regions except for the flaps.

VIII. Comparison of HLPW04 vs HLPW05

TC2.4 stood out as a case of interest due to the drop in C_L at higher AoA once the pizza-slice separations occurred. To check if these same observations were present for a similar geometry, the HLPW04 test case 2 was examined. Previous experience with this case had shown good agreement with experimental data though it had been run with a different turbulence model. Therefore, HLPW04 TC2 was explored again with the same turbulence model as HLPW05 and meshes that were similar to those used in HLPW05.

A. HLPW04 Test Case Description

TC2 from HLPW04 was a $C_{L,Max}$ study with the NASA CRM-HL. As previously stated, the ONERA LRM from HLPW05 is the same geometry as NASA CRM-HL with variations in the slat support bracket shape and placement, nacelle geometry, and an axial shift in the inboard flap. All differences were expected to be minor and have a minimal impact on the aerodynamic performance. The major difference between HLPW05 TC2.4 and HLPW04 TC2 was that the latter did not include the horizontal or vertical stabilizers. Thus, the two solutions should be similar but with the expected contribution of the tail missing.

The HLPW04 TC2 case was also at a freestream Mach number of 0.2 though at a slightly different Chord Reynolds Number of 5.49×10^6 . Also, the Reference Static Temperature was 521 °R and the Reference Static Pressure was 24.67 psi. The HLPW04 TC2 was compared against data from the QinetiQ 5-metre wind tunnel which recorded data for a half-span model. This is opposed to the full-span model used by ONERA for HLPW05.

B. Mesh Generation

A mesh was generated for HLPW04 that used the same meshing approach, including a first layer height of 0.0005in, used for HLPW05 TC2.4 Ennova meshes. The HLPW04 Ennova mesh had 67.9M nodes and 157.8M cells. An image of the volume mesh at a similar y-plane as seen for the HLPW05 meshes can be seen in Figure 19. This shows the meshing topology is very similar to what had previously been used.




Fig. 19 Image of volume mesh around the nacelle and wing at Y=420in for the HLPW04 Ennova Fine.

C. 1D Results

The AoA sweep of the HLPW04 mesh compared to data is shown in Figure 20. These simulations were executed using the same simulation approach and settings as the HLPW05 results. The 1D results show good comparison with data. The predictions were above data at some AoAs. The drop in C_L at AoA=18.05° was because of a separation region behind the nacelle that forms only for that AoA. In contrast to HLPW05 results discussed in earlier sections, pizza-slice separations appear along the same slat bracket as HLPW05 but at a weaker strength. Moreover, there is no sharp deviation in C_L or for C_M observed for the HLPW04 solution. This shows that steady RANS resolves the flow physics reasonably well in this type of configuration. Though there is high sensitivity to the slat bracket support geometry.




Fig. 20 Comparison of HLPW04 Ennova mesh to HLPW04 experimental data.

D. Comparison of HLPW04 vs HLPW05

To finalize the HLPW04 and HLPW05 comparison, the experimental data are compared against each other as well as HLPW05 TC2.4 Ennova Fine+ and HLPW04 Ennova. The major geometric difference between them is the absence of the tail in the HLPW04 geometry and mesh. These results are shown in Figure 21. Focusing just on the CFD results, it was noted how both solutions are very consistent to each other, with differences attributed to the presence of the tail. The experimental data showed a much larger C_L difference between HLPW04 and HLPW05 data and is much larger than the differences observed in CFD. Furthermore, a comparison of CF surface contours at AoA=17° shown in Figure 22 shows very similar flow features between both CFD solutions. The outboard slat for HLPW04 is noted to have higher CF than HLPW05.




Fig. 21 Comparison of HLPW04 Ennova to HLPW05 TC2.4 Ennova Fine+ along with corresponding data.




Fig. 22 Comparison of CF surface contours between HLPW05 TC2.4 Ennova Fine+ and HLPW04 Ennova at approximately $AoA=17^\circ$.

IX. Simulation Times and Memory Requirements

To Demonstrate the value of GPUs, run time, memory requirements and hardware cost are recorded for meshes from TC2.2 and TC2.4 using version 11.24 of ADS CFD software suite. Run time is stated in increments of 10,000 iterations as that is typically enough for low AoA simulations to converge. Memory requirements are also a very important factor as it determines the size of simulation that can be run before needing additional GPUs and thus more hardware costs. Hardware costs are taken as the cost to rent an 8 A100 p4d.24xlarge system from Amazon Web Services. As of 2024, the cost per hour is \$32 per 8 A100 GPUs and though this is not necessarily representative of on-premise cost for GPU hardware, it is still applicable. For consumer grade cards, such as the NVIDIA RTX4090, the cost/hour is calculated to be 1/10 the cost of an NVIDIA A100 GPU as of 2024. This results in \$0.4 per hour.

Simulations were conducted on four types of GPUs that consisted of both consumer-grade GPUs as well as data-center grade GPUs. In the consumer-grade category, simulations were conducted on a laptop with a single Nvidia Mobile RTX4090 GPU with 16GB of on-chip memory. This single GPU was capable of running meshes up to approximately 50M cells. In the same consumer-grade category was a workstation version of the Nvidia RTX4090 GPU with 24GB of memory which could run meshes of up to approximately 75M cells. In the data-center grade category was an 8 Nvidia A100 GPU system with 80GB of on-chip memory per GPU. The A100s could run meshes up to approximately 250M cells per GPU. Details such as timing comparison between these systems across a set of cases can be found in Table 2. The table shows not only the speed gains from GPUs but also the reduction in cost. Though not shown in Table 2, similar to data-center GPUs, simulations can be run across multiple consumer-grade GPUs using MPI communication making them ideal for return on investment.

Of importance is the memory utilization per GPU. This value indicates how much GPU memory capacity was required for a given mesh. The more efficient a GPU-accelerated code is in memory utilization, the larger the meshes that can be run on a given set of GPUs without requiring more resources which would increase cost. As is observed from Table 2, Code Leo has very efficient memory utilization. This allows the flow solver to be executed on consumer-grade GPU laptop. This is important since the data-center grade GPUs are more expensive and difficult to purchase due to their low supply and high demand. However, computer gaming laptops have consumer-grade GPUs that would allow for reasonably sized meshes to be executed at significantly faster turnaround times than a typical CPU solution. This reduces the cost to entry for smaller design teams to gain the advantages of GPU-accelerated flow solvers. Thus, providing a win-win scenario.

Compared against an equivalent CPU system, the Ennova TC2.4 Medium+ mesh was run on a 32 core CPU system equivalent to the x2idn.32xlarge and required 17 hours to run 10,000 iterations. The x2idn.32xlarge instance cost is \$13.338/hr and it was used because the system memory of 1024GB compares well to the 1152GB of system memory for the p4d.24xlarge. At 17 hours, the cost for the simulation is \$226.75 compared to \$17.06, an improvement in compute efficiency of 414x. If a compute optimized system is used instead, such as the c8g.24xlarge, where the cost per hour is \$3.81, this 17 hour calculation would cost \$64.77 compared to \$17.06, an improvement in compute efficiency is 120X.

Table. 2 Description simulation times and cost across various GPUs.

Mesh Name	Nodes/Cells (Million)	GPU Type	# of GPUs	Cells/ GPU	Memory Utilization per GPU (%)	Running Time for 10,000 iterations (minutes)	Hardware Cost (\$)
Consumer-Grade Mobile GPU (Laptop)							
HeldenMesh TC2.2 Medium	9.3M/26.3M	Nvidia RTX4090 Mobile (16GB)	1	26.3M	52.6%	56 min	\$0.37
Helden Mesh TC2.4 Medium	14.7M/43.8M	Nvidia RTX4090 Mobile (16GB)	1	43.8M	87.6%	94 min	\$0.62
Consumer-Grade GPU (Workstation)							
HeldenMesh TC2.2 Medium	9.3M/26.3M	Nvidia RTX4090 (24GB)	1	26.3M	35.1%	38min	\$0.25
Helden Mesh TC2.4 Medium	14.7M/43.8M	Nvidia RTX4090 (24GB)	1	43.8M	58.4%	64 min	\$0.42
Data-Center Grade GPU							
HeldenMesh TC2.2 Medium	9.3M/26.3M	Nvidia A100 (80GB)	8	3.29M	1.3%	10 min	\$5.33
HeldenMesh TC2.4 Medium	14.7M/43.8M	Nvidia A100 (80GB)	8	5.48M	2.2%	11 min	\$5.87
Ennova TC2.4 Medium+	53M/96.3M	Nvidia A100 (80GB)	8	12.0M	4.8%	32 min	\$17.06
HeldenMesh TC2.4 Fine	81.5M/205.1M	Nvidia A100 (80GB)	8	25.6M	10.2%	52 min	\$27.70
Ennova TC2.4 Fine +	117.1M/246.3M	Nvidia A100 (80GB)	8	30.8M	12.3%	77 min	\$41.07
ANSA TC2.4 C+	249M/281M	Nvidia A100 (80GB)	8	35.13M	14.1%	120 min	\$64.00

X. Conclusions

The investigation presented in this paper focused on the use of steady RANS on the modeling of high-lift configurations. The investigation centered on test case 2 from the 5th High Lift Prediction Workshop. Two configurations from the build-up study were focused on, TC2.2 and TC2.4. Three drastically different mesh topologies were used for each test case with the goal of assessing how the GPU-accelerated flow solver, Code Leo, would perform. The results were evaluated using 1D performance values, CP loading at various locations along the wing and skin friction contours. Results were compared against wind tunnel data provided by ONERA to the workshop. To further ground the investigation, results from the HLPW05 were compared to the HLPW04 where the same geometry, with minor differences, was modeled without a tail.

The conclusion from these investigations demonstrated how Code Leo proved to be relatively agnostic to mesh topology. If given a mesh with enough degrees of freedom, the solutions showed similar solutions. This quality is important as meshing best practices differ not only between organizations but also individual engineers. In addition, it was demonstrated that mesh convergence could be achieved across all AoA at a significantly lower cell count. These two conclusions prove to be key for using steady RANS in a design-cycle for the aviation industry.

Nonetheless, some limitations were also observed from the steady RANS model. Wedge-shaped flow separations referred to as pizza-slice separations on the outboard of the wing were observed from RANS model results, which were theorized to inhibit higher CL prediction when compared to data. This was of interest to the authors since this observation was different in HLPW04 with a different turbulence model. Thus, a new set of meshes were created for HLPW04 that used the identical meshing approach used for HLPW05 TC2.4 and executed using the same approach. These simulations did show a pizza-slice separation along the same bracket as HLPW05 but at a weaker state. Moreover, there was no sharp change in slope in the performance parameters due to the appearance of the pizza-slice separations. The lift curve of HLPW04 did show a consistent solution relative to the HLPW05 TC2.4 results with the differences primarily coming from the lack of the tail in the geometry. This consistency between HLPW04 and HLPW05 provided confidence that steady RANS can be used to predict the effects of small geometric changes.

Though a precise prediction of peak lift still alludes steady RANS, the value of this modeling methodology and the consistency of results prove to be highly valuable in the design-cycle. The coupling of consumer-grade and data-center grade GPUs to this modeling methodology further amplifies this point, allowing both small and large aeronautic companies to take on similar challenges through the use of a GPU laptop or multi-node GPU clusters.

References

- [1] Rumsey, C. L., Slotnick, J. P. and Woeber, C., "HLPW-4/GMGW-3: Overview and Workshop Summary", AIAA 2022-3295
- [2] Diskin, B., Rumsey, C. L., Wang, L. and Iyer, P. S., "High-Lift Prediction Workshop 5: Fixed-Grid Reynolds-Averaged Navier-Stokes Technology Focus Group Summary", AIAA SciTech Forum, January 2025
- [3] Wick, A. T. and Hooker, R., "HeldenMesh Contributions and Grid Generation Lessons Learned From the 5th AIAA High-Lift Prediction Workshop", AIAA SciTech Forum, January 2025
- [4] Ni, R. H., "A Multiple Grid Scheme for Solving the Euler Equations", AIAA 81-1025, A Collection of Technical Papers, AIAA 5th Computational Fluid Dynamics Conference, 1981, pp. 257-264
- [5] Robles Vega, G., Bosshart, A., Ni, M., Ni, R. H., 2024, "NASA HECC Geometry and Performance Review Part I: Validation of a Computational Model for the Vaneless Diffuser Configuration with As-Manufactured Impeller Geometry", Proceedings of ASME Turbo Expo 2024, GT2024-125128
- [6] Schmitz, J. T., Perez, E., Morris, S. C., Corke, T. C., Clark, J. P., Koch, P. J., Puterbaugh, S. L., 2016, "Highly Loaded Low-Pressure Turbine: Design, Numerical, and Experimental Analysis", Journal of Propulsion and Power, January 2016, Vol. 32 Number 1
- [7] Robles Vega, G., Lee, D. Y., Ni, M., Ni, R. H., 2023, "Part I: Validation and Verification of CFD Analysis for NASA SDT Transonic Fan Stage Test Rig", Proceedings of ASME Turbo Expo 2023, GT2023-103856
- [8] Lee, D. Y., Robles Vega, G., Ni, M., Ni, R. H., 2023, "Part II: Validation and Verification of CFD Analysis for NASA SDT Transonic Fan Stage Coupled with Nacelle", Proceedings of ASME Turbo Expo 2023, GT2023-103871
- [9] Ni, R. H., Humber, W., Fan, G., Clark, J. P., Anthony, R. J. and Johnson, J. J., 2013, "Comparison of Predictions from Conjugate Heat Transfer Analysis of a Film-Cooled Turbine Vane to Experimental Data" GT2013-94716, ASME J Turbomach
- [10] Ni, R. H., Humber, W., Ni, M., Capece, V. R., Ooten, M., and Clark, J. P., 2016, "Aerodynamic Damping Predictions for Oscillating Airfoils in Cascades Using Moving Meshes", Proceedings of ASME Turbo Expo 2016, ASME Paper No. GT2016-57017
- [11] Ni, M., Ni, R., & Clark, J. P. "LES Modeling of High-Lift High-Work LP Turbine Profiles: Part I: Approach." *Proceedings of the ASME Turbo Expo 2023: Turbomachinery Technical Conference and Exposition. Volume 13B: Turbomachinery — Axial Flow Turbine Aerodynamics*. Boston, Massachusetts, USA. June 26–30, 2023. V13BT30A019. ASME. <https://doi.org/10.1115/GT2023-102625>
- [12] Kerestes, J., Marks, C., Clark, J., Wolff, J. M., Ni, R. (., and Fletcher, N. (September 27, 2023). "LES MODELING OF HIGH-LIFT HIGH-WORK LPT BLADES: PART II—VALIDATION AND APPLICATION." ASME. *J. Turbomach*. doi: <https://doi.org/10.1115/1.4063509>
- [13] Mouton, S., Charpentier, G., & Lorenski, A. (2024). Testing the full-span high-lift common research model at the ONERA F1 pressurized low-speed wind tunnel. *AIAA AVIATION FORUM AND ASCEND 2024*. <https://doi.org/10.2514/6.2024-3512>