CREATE^{TM}-AV Kestrel Dual-mesh Simulations of the NASA Common Research Model

Timothy A. Eymann[*]

DoD HPCMP CREATE-AV Kestrel Team, Eglin AFB, FL 32542

Robert H. Nichols[†]

DoD HPCMP CREATE-AV Kestrel Team, Arnold AFB, TN 37389

We use NASA's Common Research Model and initial conditions from the sixth Drag Prediction Workshop (DPW VI) to investigate CREATE^{TM}-AV Kestrel's ability to accurately compute the drag coefficient for a transport aircraft. We leverage Kestrel's enhanced dual mesh capability to wrap the near-body, unstructured mesh with a higher-order Cartesian background. The dual-mesh simulation not only reduces the overall degrees of freedom in the system compared to the fully unstructured mesh but also enables feature-based mesh adaptation. This paper also explores the effects of subset distance and shock-based adaptation on the solution. Kestrel's drag coefficient values for the configurations with and without the nacelle compare very favorably to results from DPW VI and verify that Kestrel dual-mesh simulations produce accurate answers while minimizing user effort.

I. Introduction

THE Department of Defense (DoD) High Performance Computing Modernization Program started the Computational Research and Engineering Acquisition Tools and Environments (CREATE^{TM}) program in 2008 to develop a suite of powerful software tools for working-level acquisition engineers. The CREATE vision is to use physics-based simulation software to disrupt the design/build/test cycle that drives unsustainable cost increases and replace it with a cycle that includes pre-build virtual prototyping and post-test validation of the codes. CREATE is developing general-use tools for meshing and geometry as well as more specialized packages for ships, antennas, ground vehicles, and air vehicles (AV).

The CREATE-AV program is responsible for the rotorcraft package CREATE^{TM}-AV Helios and a fixed-wing product, CREATE^{TM}-AV Kestrel. Like all CREATE products, two of Kestrel's main requirements are ease of use and a product lifetime measured in decades. Our solution to meet these requirements is an event-based architecture known as the common scalable interface.[1] Under this architecture, all of the major functions of the code (mesh parsing, output, structural solver, etc.) are segregated into separate components. Components are not aware of each other and share data via pointers that are accessed through a data warehouse. Simulation control is handled by manager components that publish events such as "Iteration" that each component can subscribe and respond to. This event-based architecture and the use of components allows Kestrel to perform complex, multi-physics simulations without constructing a monolithic code base that is unmaintainable in the long-term.

A. Dual-mesh Components

Since version 5.0, Kestrel has included the ability to perform dual-mesh simulations containing an unstructured, near-body mesh communicating with a background, higher-order Cartesian mesh. The component containing the Cartesian solver and mesh framework in Kestrel v5 and v6 was SAMCart.[2] Kestrel v7 features a new core Cartesian solver based on the NXAIR code of Tramel and Nichols.[3] Version 7 retains the use of Lawrence Livermore National Laboratory's Structured Adaptive Mesh Refinement Application (SAMRAI)[4] to manage the parallel communication and mesh adaptation for the Cartesian solver, but we have upgraded to version 3.11.1. Due to the extensive changes in the component, the Cartesian solver in Kestrel v7 has been renamed SAMAir to reflect the SAMRAI library and the code's NXAIR/XAIR3D/AIR3D heritage.

---

[*]Research Engineer, Department of the Air Force, AIAA Senior Member

[†]Research Professor, University of Alabama at Birmingham, AIAA Associate Fellow

American Institute of Aeronautics and Astronautics

DISTRIBUTION A. Approved for public release; distribution unlimited.

In addition to SAMAir, the other critical components in a dual-mesh simulation are KCFD, the unstructured computational fluid dynamics (CFD) solver, and PUNDIT,[5] the domain connectivity package. Keeping with the larger goals of the CREATE program, we have automated the vast majority of the dual-mesh workflow. By default, the Cartesian mesh generation parameters are automatically set to provide an approximate match between the Cartesian and unstructured mesh spacing at the overset boundary. During domain connect events, PUNDIT identifies and maps donors and receptors for data transfer between KCFD and SAMAir without requiring any action from users. Once users set up their simulation and generate an input file using the Kestrel User Interface, the bulk of the work usually associated with overset simulations is automatically handled for them by Kestrel components.

Overset simulations reduce the mesh generation burden for users; however, this convenience does have costs associated with it. Increased numbers of meshes means more data transfer via various interpolation schemes. In Kestrel, these interpolation schemes are non-conservative and, at best, second-order accurate. The question we pose with this work is: given the known limitations of Kestrel dual-mesh simulations, can we still achieve accurate results?

II.   Problem description

To answer that question, we chose to compare Kestrel's simulation results to the well-studied NASA Common Research Model (CRM). The CRM was designed to be a non-sensitive, non-proprietary geometry representative of a modern transport aircraft.[6] The geometry has been extensively tested in wind tunnels around the world[7,8] and has been used for CFD code verification since the fourth AIAA Drag Prediction Workshop (DPW).[9,10]

In this work, we calculated the drag coefficient at a fixed lift value of $C_L = 0.5 \pm 0.0001$. For each mesh, we varied the angle of attack α until the average lift coefficient fell within the one-count tolerance. If the final lift value fell outside the tolerance, we varied α and began a new simulation rather than restarting from the converged result. Using the secant method, we typically arrived at the correct α value in three to four runs.

We simulated the CRM at its design cruise speed of Mach 0.85, Reynold's number of 5×10^6 based on a mean aerodynamic chord (MAC) of 275.8 in., and temperature of $100°$ F. These conditions match the free stream conditions specified for Case 2 of DPW VI. We studied both the wing-body (WB) and wing-body-nacelle-pylon (WBNP) configurations.

A.   Code Options

Robustness and ease-of-use are two of Kestrel's guiding design principles. Therefore the developers and the CREATE-AV quality assurance team have spent a lot of effort to determine the appropriate default values for inputs that are exposed to users. We left the majority of the code inputs at their default values. The primary exception is the turbulence model, where we employed the Spalart-Almaras Reynolds averaged Navier-Stokes model over the default delayed detached eddy simulation variant. Both the unstructured and Cartesian solver used the HLLE++ approximate Riemann solver to construct interface fluxes. KCFD ran with second-order spatial accuracy, while SAMAir used third-order and fifth-order face reconstructions. The majority of the runs used local time stepping (steady-state), although we also ran a few points time-accurately using a globally specified time step.
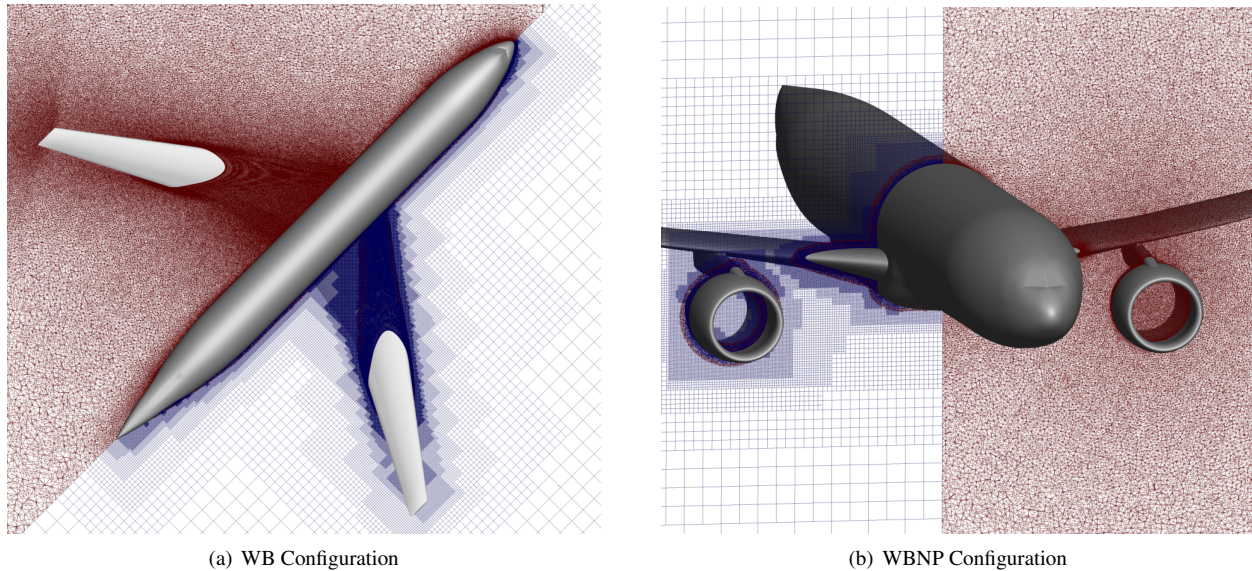
B.   Mesh Description

NASA's Geometry Laboratory (GeoLab) at Langley Research Center provided a series of unstructured meshes for DPW VI. The GeoLab supplied meshes with exclusively tetrahedral elements as well as mixed-element meshes with prismatic boundary layers. We ran Kestrel using the mixed-element meshes because Kestrel's gradient reconstruction techniques yield higher-quality solution gradients near the surface with prism elements compared to fully tetrahedral meshes. This study focused on the first three members of the WB and WBNP families provided for DPW VI.

One of the first things users must do to set up a dual-mesh case is to subset the mesh. While there are effective rules-of-thumb to use as guidance, choosing a subset distance continues to require experimentation and rely on user experience. Because the Cartesian solver uses higher-order reconstructions and it is more efficient per degree of freedom (DOF), it is generally better to minimize the number of unstructured cells present in a simulation. We selected a subset distance of 5% MAC, approximately 14 in. Table 1 shows that subsetting the mesh dramatically reduces the DOF count, dropping the number of solution points 20%-40%. Note that the point counts in the table do not include orphan or fringe points in the Cartesian mesh.

Figure 1 visually compares the meshes for the "coarse" WB and WBNP configurations. The figure clearly illustrates how the dual-mesh simulations achieve such a dramatic reduction in the mesh size by concentrating refinement to match the near-body spacing and using a 2:1 ratio between refinement levels.

American Institute of Aeronautics and Astronautics
DISTRIBUTION A. Approved for public release; distribution unlimited.

Table 1. Mesh sizes for WB and WBNP configurations

Name	WB Unstruc.	WB Dual	Reduction	WBNP Unstruc.	WBNP Dual	Reduction
Tiny	83,598,506	66,717,134	20.2%	120,990,279	78,742,253	34.9%
Coarse	122,878,088	91,420,871	25.6%	178,975,198	108,048,363	39.6%
Medium	182,037,523	129,056,885	29.1%	266,916,327	151,877,830	43.1%

(a) WB Configuration

(b) WBNP Configuration

Figure 1. Unstructured and dual-mesh systems for "coarse" refinement level

III.   Results

Our initial study consisted of a set of baseline Kestrel runs using purely unstructured meshes and overset Cartesian/unstructured meshes. The baseline runs all used a subset distance of 5% MAC and did not include any solution adaptation. The effects of the subset distance and solution adaptation were quantified during separate lines of investigation.
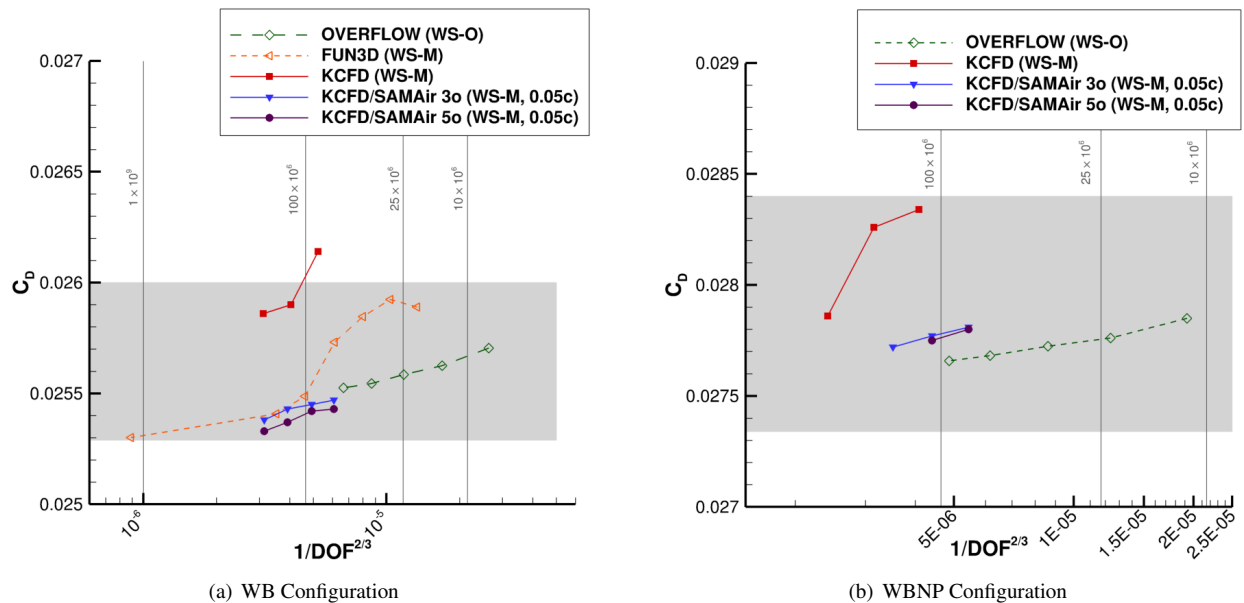
A.   Baseline Study

Figure 2 shows the mesh convergence history for the WB and WBNP configurations; the numerical values appear in Table 2. Kestrel's unstructured and dual-mesh simulations are plotted along with values from OVERFLOW[11] and FUN3D. FUN3D solutions were generated with second-order spatial accuracy while the OVERFLOW solutions plotted in Fig. 2 were computed with third-order methods. The legend indicates the source of the mesh with WS-O meaning the DPW VI overset meshes, and WS-M meaning the DPW VI mixed-element meshes. The gray, shaded region marks the range of answers from the other DPW VI participants,[12] excluding outliers.

The figure reveals that the fully unstructured KCFD solution does show similar mesh convergence trends as FUN3D and falls within the range of the DPW VI results. More interestingly, Fig. 2 shows that subsetting and wrapping the unstructured mesh with a Cartesian mesh dramatically lowers the drag count. The KCFD/SAMAir solutions are within one or two counts of the most refined OVERFLOW and FUN3D solutions even when using the coarser workshop meshes. For example, the dual-mesh "fine" solution at 180 million DOF predicts $C_D = 0.0254$ compared to the FUN3D value of $C_D = 0.0253$ using over *one billion* DOF. Figure 2 also shows that the fifth-order WENO solutions exhibit the same behavior as the third-order dual-mesh solutions and predict drag values within a half-count of the third-order solutions. Recall that these results were achieved primarily with default settings, and that domain connectivity is completely automatic. Although these results are very encouraging, runs with finer meshes need to be completed to fully assess Kestrel's dual-mesh behavior.

Table 2. Drag coefficients for WB and WBNP configurations

Mesh Name	WB Unstructured	WB Dual	WBNP Unstructured	WBNP Dual
Tiny	0.02614	0.02547	0.02834	0.02781
Coarse	0.02590	0.02545	0.02826	0.02777
Medium	0.02586	0.02543	0.02786	0.02772
Fine	-	0.02538	-	-

(a) WB Configuration

(b) WBNP Configuration

Figure 2. Drag coefficient with increasing mesh refinement

B. Subset Study

Table 2 and Fig. 2 shows that there is about a 7 count difference between the computed drag coefficients for the fully unstructured and dual-mesh runs using the "tiny" WB mesh. To characterize the source of the drag difference, we explored the flow solutions with increasing numbers of unstructured mesh points by varying the subset distance from 5% MAC to 200% MAC. Table 3 details the mesh sizes as the subset distance grew. Note that the 10% subset mesh actually contains fewer degrees of freedom than the 5% mesh because the number of Cartesian points increased less than the number of unstructured points decreased.

Table 3. Mesh degrees of freedom with increasing subset distance

Subset (% MAC)	Unstructured DOF	Cartesian Field DOF	Total
5	42,933,921	23,783,213	66,717,134
10	46,087,997	19,411,051	65,499,048
50	56,578,553	14,246,630	70,825,183
100	61,874,995	11,422,648	73,297,643
200	66,750,386	9,437,315	76,187,701

The drag coefficients plotted in Fig. 3 show a smooth progression and increasing values as more of the unstructured mesh is retained. Splitting the drag into its viscous and pressure components reveals the effect the off-body mesh has

American Institute of Aeronautics and Astronautics
DISTRIBUTION A. Approved for public release; distribution unlimited.

on the solution. The y-axes in Fig. 3b both have a range of 8 counts and the figure clearly shows that the pressure drag decreases with subset distance while the viscous drag due to skin friction remains fairly constant. This trend also shows up in many of the other DPW VI results[12] and indicates that the primary function of increasing the off-body accuracy – via mesh resolution, higher-order accuracy, or a combination of the two – is to improve the pressure drag calculation by better predicting the locations and strength of the wing shocks.

(a) Total Coefficient

(b) Split Coefficient

Figure 3. Drag coefficient with increasing subset distance

Plots of the pressure coefficient at various wing stations most clearly illustrate the differences in the computed pressure profiles using the 200% MAC subset distance and 5% MAC distance meshes. Figure 4 shows locations with significant differences between the single-mesh and dual-mesh runs. For the WB geometry, the most significant differences show up in the stations closest to the wing tip. Figure 5 shows examples of some of the more obvious differences, which occur at the $\eta = 0.8456$ and $\eta = 0.95$ stations. At $\eta = 0.8456$, the 200% mesh is predicting the shock too far aft and at $\eta = 0.95$ the same mesh misses the small compression feature at the quarter-chord location.
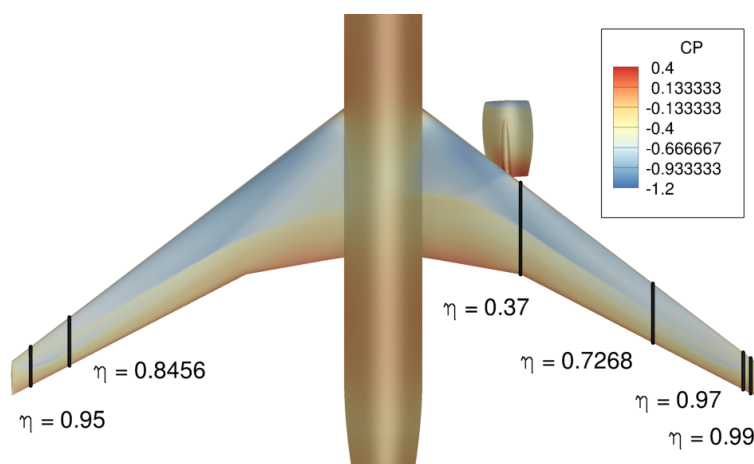

Figure 4. WB and WBNP locations with significant differences in surface C_p between single and dual-mesh runs

The WBNP solutions show a similar trend where the dual-mesh simulations predict lower drag coefficients than the single-mesh, fully unstructured solutions. Here, the largest differences between the two Kestrel solutions are near the nacelle and, like the WB runs, toward the wing tip. Again, Fig. 4 shows the locations of the example cuts that are highlighted in Fig. 6.
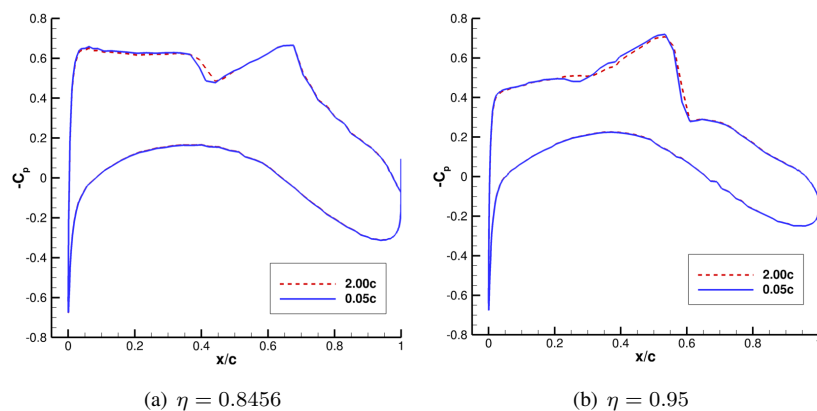
American Institute of Aeronautics and Astronautics
DISTRIBUTION A. Approved for public release; distribution unlimited.

(a) $\eta = 0.8456$

(b) $\eta = 0.95$

Figure 5. C_p **profiles for 5% and 200% subset meshes based on the WB "tiny" geometry**

(a) $\eta = 0.37$

(b) $\eta = 0.7268$
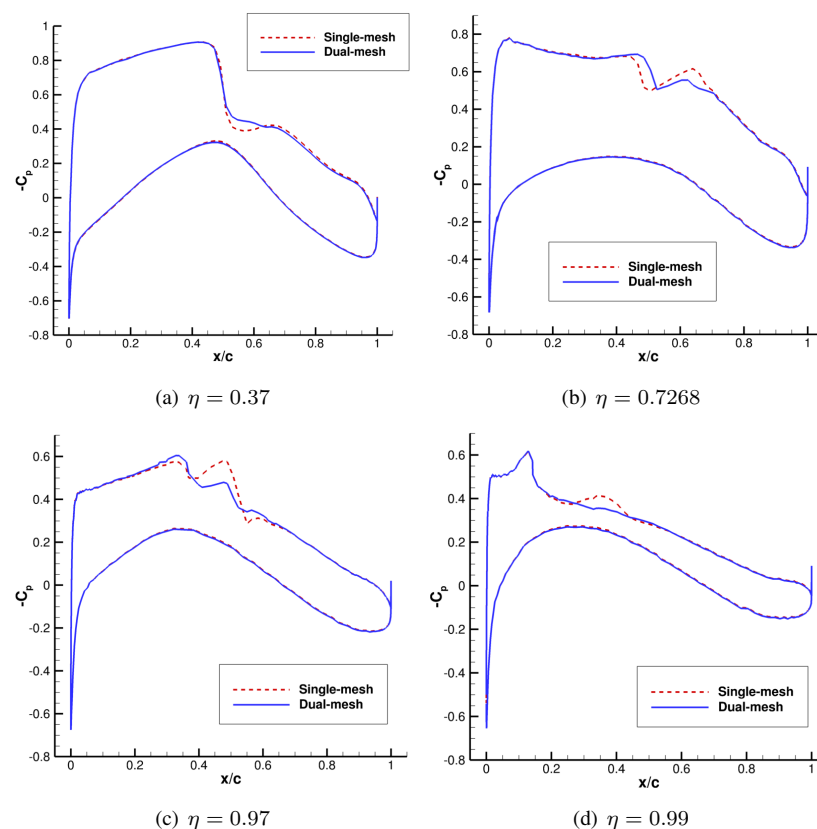
(c) $\eta = 0.97$

(d) $\eta = 0.99$

Figure 6. C_p **profiles for single and dual-mesh WBNP "coarse" mesh**

C.   Refinement Study

One of SAMAir's primary advantages is its ability to perform feature-based, adaptive mesh refinement (AMR). The pressure coefficient profiles generated using various subset distances demonstrated the importance of accurately predicting the flow and shock system around the nacelle and toward the wing tip. To further guide the choice of adaptation variables, we compared the Kestrel solutions with and without SAMAir. Because the unstructured and dual-mesh solutions use identical near-body meshes, differencing the two solutions in this region can highlight where the Cartesian mesh is having the greatest impact on the solution. In other words, we can attribute any solution differences in the near-body mesh solution to the presence of the overset boundary and Cartesian mesh. To visualize the difference,

we computed the ΔC_p quantity defined in Eq. (1) using the pressure coefficient from the fully unstructured solution and the pressure coefficient computed from the dual-mesh solution.

$$\Delta C_p = |C_{p_{\text{unstruc.}}} - C_{p_{\text{dual}}}| \tag{1}$$

Figure 7 shows an isosurface of this pressure coefficient difference. The isosurface indicates regions where the fully unstructured and overset solutions differ. The figure clearly shows that the area around the shocks and the overset boundary near the leading edge have small, but significant pressure differences between the two solutions. The features shown in Fig. 7 reinforce the conclusion that the shocks are a leading contributor to the drag coefficient differences shown in Figs. 2-3. This suggests that refining on the shock sensor value[2] that contributes to the Kestrel limiter should affect the solution behavior.
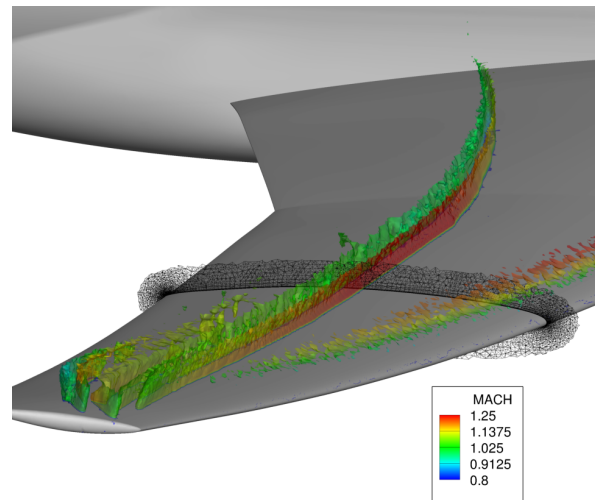

Figure 7. Isosurface of $\Delta C_p = 0.02$ indicating differences between the fully unstructured and dual-mesh simulations

We refined the mesh every 500 iterations from iteration 500 to 8000 where the shock sensor value exceeded a threshold of 1.0. The unstructured mesh and the adapted dual-mesh system are shown in Fig. 9. The slice that is shown is representative of the refinement along the entire span of the wing and illustrates the improved shock resolution. While the AMR did sharpen the shock as expected, Fig. 8 shows it did not have a major impact on the converged loads. Using solution refinement during the "fine" WB simulation shifted the drag coefficient by 0.8 counts. The difference with AMR is larger for the "tiny" WBNP mesh, but runs with finer meshes are needed to establish whether this is due to the effects of AMR or the relatively coarse near-body mesh.

The biggest effect of the shock-based AMR was that it accelerated the lift and drag convergence of the time-accurate Kestrel runs. There was a less pronounced effect on the steady-state runs. Figure 10 shows that, for the case with shock sensor AMR, both C_L and C_D settle into their mean values shortly after the refinement ends at iteration 8000. The case with no solution refinement converges several thousand additional iterations later, taking over 20,000 iterations to reach a stable mean. The increased shock resolution from the AMR likely reduces limiter chatter around the shock, leading to the improved loads convergence. Although feature-based adaptation increases the mesh size, the increased convergence it can unlock translates into significant wall-clock savings for users. In general, using the Cartesian mesh in the off-body, even without AMR, improved the loads convergence and accelerated the time to a solution up to a factor of two compared to the single-mesh runs.
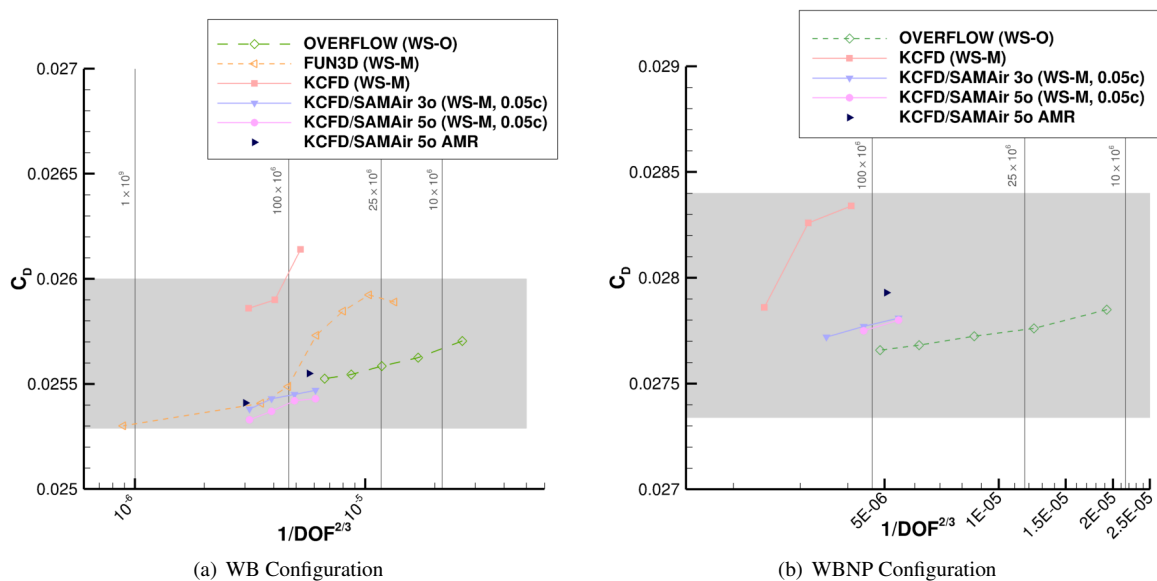
(a) WB Configuration

(b) WBNP Configuration

Figure 8. Drag coefficients with shock-sensor based AMR

(a) Fully unstructured

(b) Dual-mesh with AMR

Figure 9. Shock sensor based AMR; plane location at $\eta = 0.3415$

(a) Lift coefficient
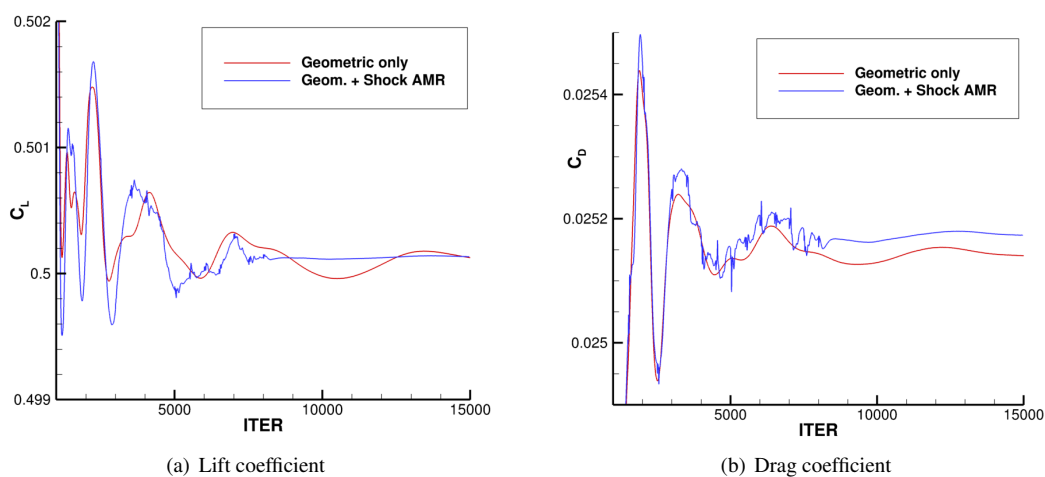
(b) Drag coefficient

Figure 10. Time-accurate loads convergence with and without mesh refinement

IV.   Summary and Conclusions

These results show that Kestrel's default settings provide an excellent starting point for dual-mesh simulations of cases like the CRM. Kestrel's predicted drag coefficient values are in-line with results from DPW VI. The majority of the error in the drag coefficient is due to the pressure drag and Kestrel's dual-mesh solutions using SAMAir in the off-body provide an effective way for users to improve this calculation and their overall simulation accuracy. SAMAir's higher-order reconstructions are much more efficient at reducing error than second-order methods, leading to solutions that converge in nearly half the wall-clock time and with a quarter of the error of the single-mesh runs. Notably, Kestrel's automation allows users to compute these answers with far less effort than other unstructured and overset solvers. While AMR is critical for off-body feature tracking and multi-body simulations with complex aerodynamic interactions, it had little effect on the loads prediction for the CRM. Returning to the question posed in the introduction, the answer – based on this exploration of the CRM WB and WBNP drag prediction – is yes; Kestrel's dual-mesh simulations are certainly capable of producing accurate results for this class of problems.

Acknowledgements

References

[1]McDaniel, D. R., Tuckey, T. R., and Morton, S. A., "Multiple Bodies, Motion, and Mash-Ups: Handling Complex Use-Cases with Kestrel," *52nd AIAA Aerospace Sciences Meeting*, National Harbor, MD, January 2014, AIAA Paper 2014-0415.

[2]Eymann, T. A., Nichols, R. H., McDaniel, D. R., and Tuckey, T. R., "Cartesian Adaptive Mesh Refinement with the HPCMP CREATE™-AV Kestrel Solver," *53rd AIAA Aerospace Sciences Meeting*, 2015, AIAA Paper 2015-0040.

[3]Tramel, R. W. and Nichols, R. H., "A Highly Efficient Numerical Method For Overset-Mesh Moving-Body Problems," *13th AIAA Computational Fluid Dynamics Conference*, Snowmass Village, CO, June 1997, AIAA Paper 1997-2040.

[4]Anderson, R., Arrighi, W., Elliott, N., Gunney, B., and Hornung, R., "SAMRAI Concepts and Software Design," LLNL-SM-617092-DRAFT, Lawrence Livermore National Laboratory, February 2013.

[5]Sitaraman, J., Floros, M., Wissink, A., and Potsdam, M., "Parallel Domain Connectivity Algorithm for Unsteady Flow Computations Using Overlapping and Adaptive Grids," *Journal of Computational Physics*, Vol. 229, 2010, pp. 4703–4723.

[6]Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," *26th AIAA Applied Aerodynamics Conference*, Honolulu, HI, August 2008, AIAA Paper 2008-6919.

[7]Rivers, M. B. and Dittberner, A., "Experimental investigations of the NASA Common Research Model in the NASA Langley National Transonic Facility and NASA Ames 11-Ft Transonic Wind Tunnel," *49th AIAA Aerospace Sciences Meeting*, Orlando, FL, January 2011, AIAA Paper 2011-1126.

[8]Koga, S., Kohzai, M., Ueno, M., Nakakita, K., and Sudani, N., "Analysis of NASA Common Research Model Dynamic Data in JAXA Wind Tunnel Tests," *51st AIAA Aerospace Sciences Meeting*, Grapevine, TX, January 2013, AIAA Paper 2013-0495.

[9]Vassberg, J., Tinoco, E., Mani, M., Rider, B., Zickuhr, T., Levy, D., Brodersen, O., Eisfeld, B., Crippa, S., Wahls, R., Morrison, J., Mavriplis, D., and Murayama, M., "Summary of the Fourth AIAA CFD Drag Prediction Workshop," *28th AIAA Applied Aerodynamics Conference*, Chicago, IL, June 2010, AIAA Paper 2010-4547.

[10]Levy, D. W., Laflin, K. R., Tinoco, E. N., Vassberg, J. C., Mani, M., Rider, B., Rumsey, C. L., Wahls, R. A., Morrison, J. H., Brodersen, O. P., Crippa, S., Mavriplis, D. J., and Murayama, M., "Summary of Data from the Fifth Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 51, No. 4, July 2014, pp. 1194–1213.

[11]Coder, J. G., Hue, D., Kenway, G., Pulliam, T. H., and Sclafani, A. J., "Contributions to the 6th AIAA CFD Drag Prediction Workshop Using Structured, Overset Grid Methods," *55th AIAA Aerospace Sciences Meeting*, Grapevine, TX, January 2017, AIAA Paper 2017-0960.

[12]Tinoco, E. N., Brodersen, O. P., Keye, S., Laflin, K. R., Feltrop, E., Vassberg, J. C., Mani, M., Rider, B., Wahls, R. A., Morrison, J. H., Hue, D., Gariepy, M., Roy, C. J., Mavriplis, D. J., and Murayama, M., "Summary of Data from the Sixth AIAA CFD Drag Prediction Workshop: CRM Cases 2 to 5," *55th AIAA Aerospace Sciences Meeting*, Grapevine, TX, January 2017, AIAA Paper 2017-1208.