

Applications of a Cartesian Mesh Boundary-Layer Approach for Complex Configurations

Michael J. Aftosmis¹

NASA Ames Research Center, Moffett Field, CA 94035, USA

Marsha J. Berger²

Courant Institute, New York University, NY, NY 10012, USA

Juan J. Alonso³

Stanford University, Stanford, CA 94035, USA

This paper examines the performance of a coupled Euler boundary-layer approach in simulating viscous flows around a variety of aerospace configurations. The method combines an established multilevel Cartesian-mesh Euler solver with a transpiration boundary condition to account for the boundary-layer displacement thickness. This boundary condition is set via a strip-wise solution of the 2D boundary-layer equations which uses the inviscid solution as a driver. The implementation uses local flow topology to establish attachment and separation and an elliptic solve on the surface triangulation to couple surface transpiration velocities back to the inviscid solver. While interacting boundary-layer (IBL) approaches are not necessarily new, the current approach is strongly focused on complex configurations and the implementation includes some novel techniques for coping with geometric complexity, markedly improving its utility, and removing the need for additional viscous corrections. The use of IBL solvers is well established for transport aircraft configurations, and the current work examines the success of the technique for such cases and explores its utility outside this class of problems. The investigations demonstrate the technique's performance with both single-point and parametric studies on 2D supercritical airfoils, isolated wings, finned-missiles, and full-aircraft configurations. Results on the NACA RM-10 showed good agreement over a range of transonic and supersonic Mach numbers. Simulations on the DLR F-4 wing-body yielded aerodynamic force coefficients that agreed well with established results from the 1st AIAA Drag Prediction Workshop over a range of conditions. The discussion of these numerical results highlights regions of continued research.

1. Introduction

AUTOMATED Cartesian approaches for inviscid simulations of flows about complex air vehicles have matured substantially in the past decade.^[1-3] Despite the lack of viscous effects, these methods have shown a remarkable ability to accurately predict flight loads on broad classes of vehicles with diverse flight conditions.^[4-6] Moreover, the combination of rapid turn-around time and full automation has enabled leveraging of these methods to conduct massive parametric studies, enhancing the understanding of aerodynamic trends over a broad range of flight conditions.^{[7][8]} Despite this general success, the critical importance of viscous effects in particular aerodynamic situations means that practitioners of these approaches must be constantly wary of being misled by the simplified physical modeling used.

¹ Research Scientist, Senior Member AIAA

² Professor, Courant Institute, Senior Member AIAA

³ Associate Professor, Department of Aeronautics and Astronautics, Member AIAA

Copyright © 2006 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

High-lift and transonic simulations over high aspect-ratio configurations with supercritical airfoils are one area of particular difficulty for inviscid approaches. Although the recent series of Drag Prediction Workshops has highlighted outstanding deficiencies in this regime even for modern Navier-Stokes simulations, these methods clearly offer substantial improvements over purely inviscid approaches.^[9-11] In this particular class of problems, stiffness in the governing equations translates into high sensitivity to shape change. As a result, the small displacement effect of the boundary-layer may substantially modify the outer inviscid flow triggering erroneous shock locations and predictably poor values of both local flow properties and integrated aerodynamic coefficients.

Interacting boundary-layer (IBL) approaches offer an attractive way to improve these simulations without substantially increasing the cost of a solution. Such approaches typically increase the runtime of an inviscid simulation by only a few percent and the literature shows them to be very potent tools for this class of problems.^[11-14] Moreover, while full Navier-Stokes simulations offer a more general approach, runtime alone is typically 20-50 times that of Euler simulations.^{[11][15]} One goal of the present work is to determine the extent to which these results from the IBL literature carry over to real-world cases with truly complex geometry. Additionally we wish to investigate their ability to robustly predict accurate trend information through parametric studies of flow conditions.

2. Numerical Procedure

Two approaches are commonly used for simulating the boundary-layer displacement thickness with an inviscid solver. The most obvious is to simply modify the geometry, and apply the inviscid tangential flow condition on a displaced surface. In three dimensions, however, actually modifying the surface introduces the complexities of re-tessellating both the surface and volume meshes. This places even more burden on surface modeling and mesh generation, and with our current focus on geometry that is already complex, this is undesirable. The alternative is to introduce a surface transpiration velocity normal to the existing surface thereby deflecting the outer flow, mimicking the displaced surface. In the transpiration approach, a small error occurs since the boundary conditions are applied on the undisplaced surface, rather than the effective inviscid surface. The difference between these two surfaces means that the boundary-layer growth will be slightly incorrect. While this error may be appreciable in the vicinity of highly curved surfaces, it will obviously be small over the bulk of the surface.^[14] In these initial investigations, we primarily seek the gross effects of the boundary-layer's presence and therefore follow the transpiration approach. Note also that since our immediate focus is on high-Reynolds number flows, the boundary layer displacement is expected to be small.

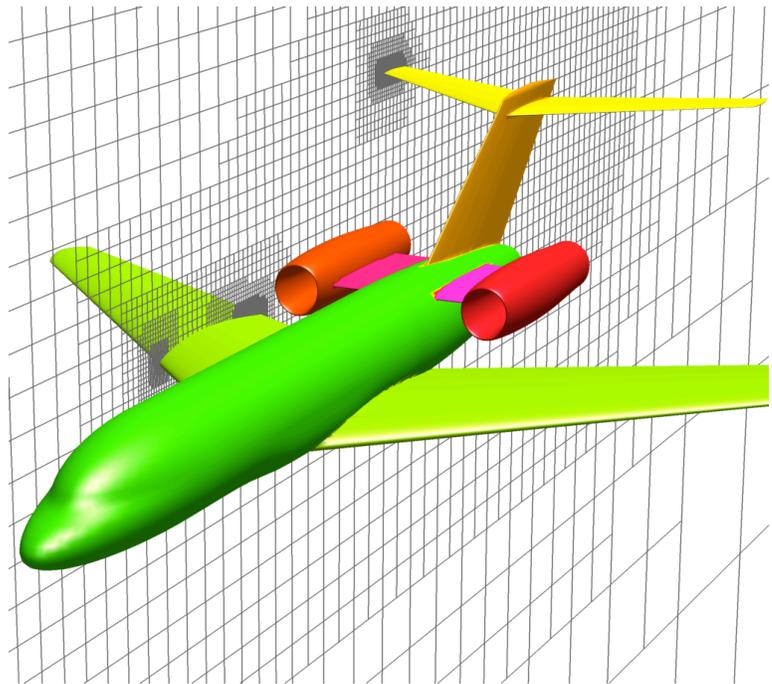


Figure 1: 2.3 Million cell multilevel Cartesian mesh around a business jet configuration using the parallel Cartesian method of reference [1].

The strip boundary-layer technique developed in this paper is coupled with the multilevel Cartesian mesh scheme of Ref. [1]. This parallel multigrid solver uses an adaptive Cartesian mesh generation scheme to refine the mesh to the geometry, and (optionally) the evolving solution.^[1-3] Figure 1 shows an example mesh (2.3M cells) produced by this package around a business jet. All geometry enters the package as water-tight surface triangulations. Cells in the mesh which intersect the body are cut against these surface triangulations, resulting in a layer of irregularly shaped polyhedra immediately adjacent to the surface. The finite-volume solver uses a second-order upwind-flux discretization for spatial differencing and an explicit modified Runge-Kutta method for time advance with multigrid acceleration.

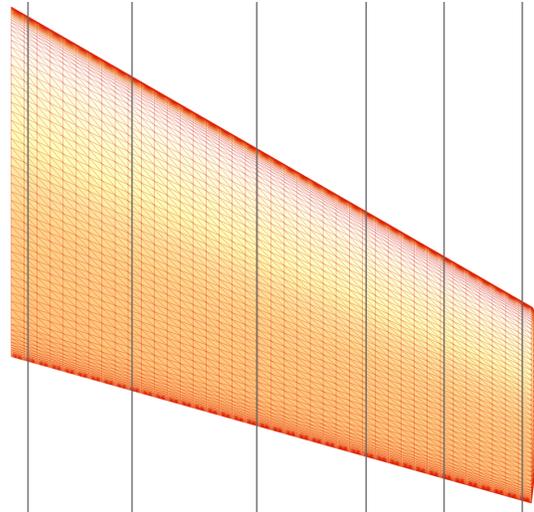


Figure 2: Sample distribution of boundary-layer stations spanwise along a simple wing.

2.1 Boundary-Layer Strips

The 2D boundary-layer equations are solved along strips of data generated by the intersection of coordinate-aligned cutting planes with the input geometry. The location of these strips is user-specified, and a typical setup is shown in Figure 2. Since the input geometry is a locally-manifold watertight triangulation, the intersection of these planes with the input geometry will always produce some number of closed loops. These loops are recorded as lists of line-segments. Each intersected triangle in the surface contributes one line-segment to these loops. Data from the inviscid solution is stored at the vertices of these loops which come from the intersection of edges in the surface triangulation with the cutting plane. This differs somewhat from the approach in reference [13] in which loops are formed by following edges in the triangulation nearest to the the cutting-planes.

To drive the boundary layer solution, flow data must be moved from the vertices of the triangulation to the vertices of the loops. Since these vertices are uniquely associated with edges in the triangulation, this data can simply be linearly interpolated from the endpoints of the corresponding edges in the triangulation.

Figure 3 details this process for three cuts through the wing of the business jet example. The illustration shows the loops in blue. Solid symbols mark the pierce-points of the edges with the cutting plane. Endpoints of these edges are shown with open green symbols and each pierce-point is uniquely associated with two of the green symbols.

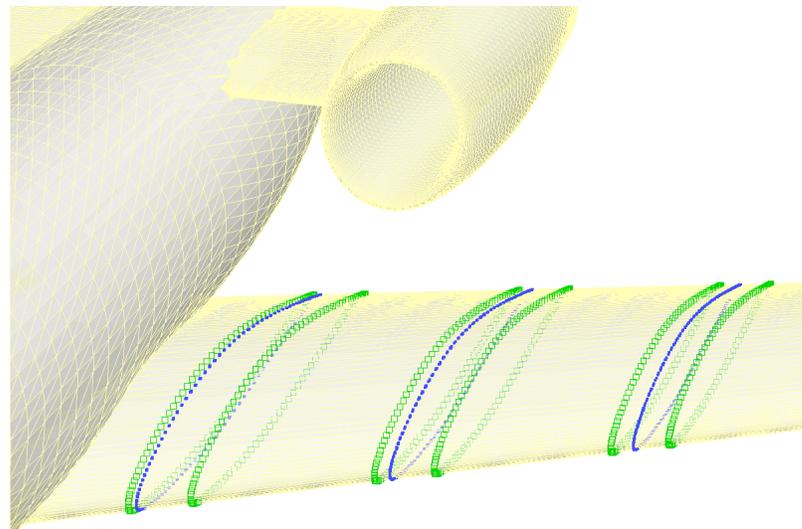


Figure 3: Loops formed by intersection of three boundary-layer stations with the wing of the business jet example. Solid blue symbols mark intersections of triangle edges with the cutting planes, and open green symbols mark the endpoints of these edges used for interpolating data to the loops.

In situations with complex geometry, the cutting-planes may form multiple loops when they intersect the geometry. Figure 4 shows a cut through a nacelle. While the triangulation is everywhere locally manifold, the surface has non-zero genus, and multiple loops are generated. In the particular case shown, the inboard loop is actually broken by the pylon attached to the nacelle, this breaks the loop where it is interrupted by the pylon. While obviously an approximation, such situations are automatically handled by simply closing the loop as if this gap didn't exist.

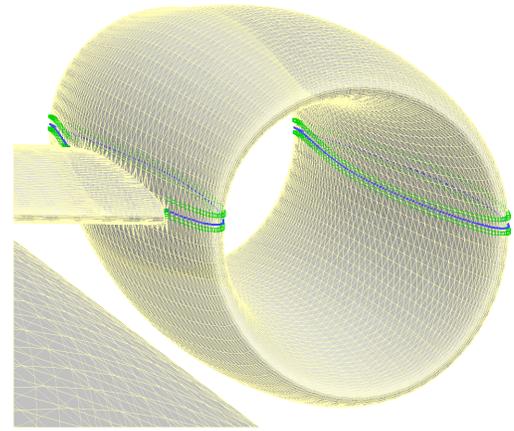


Figure 4: Multiple loops formed by intersection of horizontal cut through nacelle.

Figure 5 shows a close-up of the boundary-layer loops on a canard of a canard-controlled missile configuration. Pierce-points in this figure are marked with solid symbols and the illustration emphasizes the non-uniformity of the pierce-point distribution obtained when the cutting-planes arbitrarily slice through the triangulation. Since the boundary-layer routines are coded using differencing formulae that assume uniform mesh spacing, this irregular spacing can be problematic. To account for this, we spline the input data using standard cubic splines^[22] and resample using a constant number of stations on each closed loop. As one would expect, this uniform re-sampling dramatically improves the boundary-layer's ability to robustly produce consistent results regardless of the quality of the surface triangulation. Moreover, this uniform sampling means that smoothers applied to data along the cuts perform predictably irrespective of the actual number of pierce points in each loop.

2.2 Boundary Layer Solution

Solution along the loops is based upon two-dimensional solution techniques for compressible laminar and turbulent boundary layers.^[16-19] The laminar method is based upon Thwaites' correlation, and includes a free-transition option. The turbulent approach is an inverse method developed for compressible flow, and is capable of succeeding in regions of mild separation.^{[13][14]} This is an inverse approach where the boundary-layer edge velocity is matched iteratively, and Keller's box method is used to solve the finite-difference form of the equations. In separated regions, the convective terms are set to zero through the recirculation. While the laminar routines permit free transition, transition can also be specified either at a specific local Reynolds number, or at a specified location to simulate the use of trip-strips. Results with a similar boundary-layer model are also available in reference [11]. The two-dimensional boundary-layer solutions along all the strips proceed independently and are performed in parallel using loop-level parallel programming constructs.

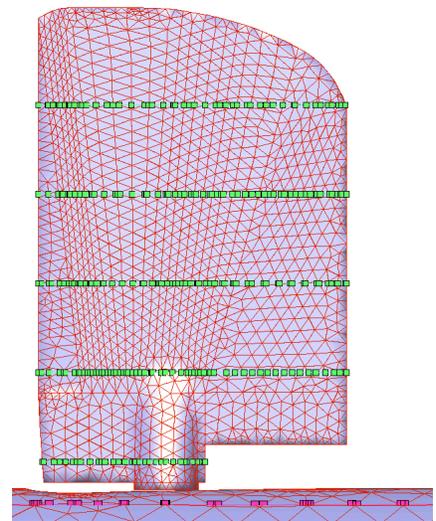


Figure 5: Close-up of boundary-layer cuts through surface triangulation showing non-uniform distribution of intersections.

2.3 Inviscid Trailing Edge Treatments and Smoothing

Without a wake model, the inviscid pressure profiles along the boundary-layer loops generally have very steep pressure gradients near the trailing-edge as the upper and lower surface pressures struggle to become single-valued when the profile closes. This issue

has been addressed extensively in the IBL literature and is particular concern in flows around wings with blunt trailing-edges.^[13,14] In these cases, inviscid pressure distributions typically show a sharp spike as the flow accelerates around the blunt trailing-edge. In a coupled-IBL approach with no wake model, these inviscid artifacts may trigger a trailing-edge flow separation that either hampers convergence of the method, or introduces unrealistic blowing-velocities. Approaches for circumventing these problems vary, but most involve some combination of smoothing of the inviscid data and “opening” the trailing-edge by modification of the pressure distribution within some small neighborhood of the trailing edge. In practice this combination of techniques can work quite well, but the approaches require user-specified parameters to control where and when to apply them and therefore increase the level of required user-expertise. In an effort to avoid these additional “knobs” in the current work, we settled on a treatment in which performed almost as well, and does not introduce free parameters. At the end of each boundary-layer strip, we reflect the last two points of the splined data from the inviscid solution (to simulate data from the beginning of a wake) and run an explicit smoother a fixed number of iterations on strip with a Neumann condition at the trailing edge. The amount of smoothing was “tuned” using several transonic-wing cases with incipient separation and then fixed. Each splined boundary-layer strip is discretized with (typically) 400 equally spaced stations. By fixing both the discretization and smoothing level on the strips, our hope is to get both predictable behavior while decreasing our reliance on user-specified inputs.

2.4 Boundary-Layer Interaction

Incorporating the boundary-layer solution into the inviscid outer flow simulation requires an iterative coupling procedure. An inviscid solution (perhaps only partially converged) provides an initial estimate of the flow over the geometry. At this point the pressure and Mach distributions along all the boundary-layer strips are extracted and the boundary-layer routines generate an estimate of the local displacement-thickness along each of these. This displacement thickness is then communicated back to the outer inviscid solver as a distribution of transpiration velocities over the surface. This feedback cycle is repeated periodically as the outer inviscid multigrid solver converges.

Transpiration velocities, V_n , are computed from the boundary-layer displacement-thickness, δ^* , through Lighthill’s relation,^[20]

$$V_n = \frac{1}{\rho_e} \frac{d}{ds} (\rho U_e \delta^*) \quad (1)$$

where ρ is the density, s is the running length from the stagnation point, and the subscript $()_e$ denotes quantities taken at the outer edge of the boundary-layer.

In the literature, unstructured IBL approaches have traditionally taken two approaches toward distributing the transpiration velocity over the full surface. While conceptually straightforward, this detail can have a major impact on the behavior of the coupled scheme, since it is these velocities which actually establish the boundary condition for the Euler method. In the first approach, the transpiration velocity along the cut is treated like a specified temperature field and Laplace’s equation is solved to distribute them over the remaining surface (*cf.* [13]). In the second, interpolation coefficients are pre-computed from the geometry of the cuts and the surface triangulation, and at every update the transpiration velocities over the triangulation are simply obtained through interpolation (*cf.* [14] and [11]).

At first glance, interpolation seems to be the method of choice. Interpolation produces an update with a single sweep over the triangulation and is obviously faster than any iterative approach. Nevertheless, the

technique has some distinct limitations. Since we're interpolating to a surface, two interpolation coefficients are required. Thinking of a simple trapezoidal wing, we need one spanwise coefficient, and one chordwise coefficient. In ref. [14] the spanwise coefficients are established as functions of the relative distance between adjacent strips, and the chordwise coefficients come from the local chordwise fraction. This approach has obvious shortcomings where the surface is not bounded by two spanwise strips, or the local chordwise profile differs substantially from that of the nearest spanwise station. To circumvent these shortcomings, Ref. [11] replaces this linear ruling with successive solutions of the Poisson equation to provide the spanwise and chordwise interpolation coefficients. Since the elliptic solves are done only at startup, this approach permits the interpolation to wrap around wing-tips, winglets, over pylons etc., while still reclaiming the speed of an interpolation scheme within the coupled solver iterations. Despite this advantage, both interpolation approaches rely upon having a template for interpolation which does not evolve with the solution. They both rely upon a one-time setup which essentially ties borders of rectangular patches of surface to the values of the transpiration velocity therein. Neither approach has an obvious extension if boundary-layers were being computed along surface streamlines, rather than along simple spanwise cuts, and neither adjusts the interpolation coefficients as the attachment and separation lines move with convergence of the solution.

In contrast to these interpolation methods, the approach in reference [13] is to solve Laplace's equation on the surface for the transpiration velocity (using point Jacobi iterations) after each boundary-layer update using the solution along the strips as a Dirichlet condition. While obviously slower than a single pass interpolation timing data in this report indicates that the cost is still inconsequential, even using a simple point Jacobi implicit scheme. The approach identifies leading and trailing edges and data along these is imposed with Dirichlet conditions. Unfortunately, in this earlier work, the leading and trailing edges are identified at startup from geometric (only) information, and this approach also holds these rigidly fixed through the entire simulation.

With its emphasis on complex geometry, and an eye toward future work with streamline-aligned boundary-layers, the restrictive nature of the interpolation approaches make them unattractive in the current work. However, the preceding paragraph suggests that issues still remain with the iterative approach. The choice of Dirichlet conditions along the boundary-layers is clear, but the treatment at leading and trailing edges is less clearly motivated. Figure 6 shows a section of a RAE 2822 airfoil that has been extruded to form an unswept 3D wing section with two boundary layer stations. The surface is mapped by contours of displacement-thickness after solution of Laplace's equation using *only* the displacement-thickness along the loops as a boundary condition. Examining these contours near the center of the trailing edge shows clear evidence that the upper and lower surfaces are contaminating each other around the trailing-edge.

Rather than simply search for leading and trailing-edges using geometry, a more general approach comes from examining the topology of the flow field.^{[21][22][24][26]} Flow attachment is indicated by the limiting streamlines diverging from a line on the surface, while converging streamlines are necessary (but not sufficient) to indicate sepa-

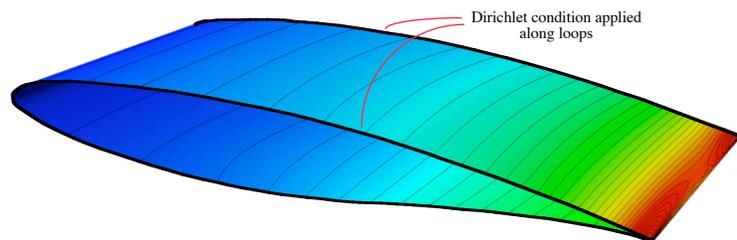


Figure 6: 3D RAE wing section mapped by contours of displacement-thickness after solution of Laplace equation on surface triangulation using only Dirichlet conditions along loops. With no treatment of attachment and separation lines, bleed from pressure surface is apparent at the trailing edge.

ration. At each boundary-layer update, we scan the edges in the triangulation and mark edges along either separations or attachments. Then, during solution of the heat equation, we impose a Neumann boundary condition on all the marked edges. The choice of the Neumann condition is clear since it prevents contamination across the edge, and ensures that isolines incident upon the edge are normal. Figure 7 illustrates attachment and separation lines found via this methods on the 3D RAE wing.

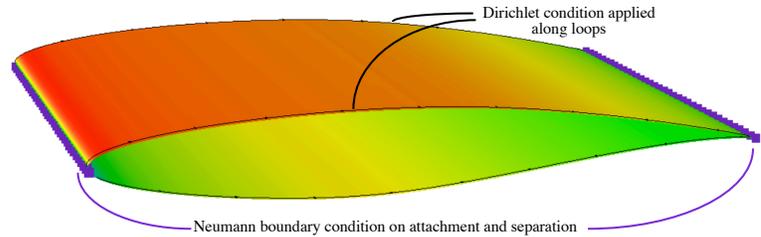


Figure 7: Edges along attachment and separation lines identified from local surface flow topology for imposition of Neumann boundary condition on surface.

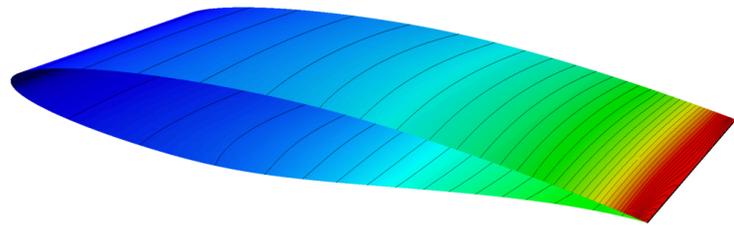


Figure 8: 3D RAE wing section mapped by contours of displacement-thickness including Neumann boundary conditions on edges along attachment and separation lines.

Rather than the edge-based, point-Jacobi iterative scheme of reference [13], we solve the heat equation using a triangle-based, point-implicit SOR scheme^[29] which has substantially

faster convergence. Figure 8 displays the distribution of displacement-thickness, δ^* , after the addition of the Neumann condition along the attachment and separation lines. Comparison with Figure 6 shows substantial improvement, and the Neumann condition prevents bleed around the trailing edge. While this example shows δ^* , the same method is used to distribute blowing-velocity, skin-friction, or any other boundary-layer parameter over the surface of the triangulation. Choosing SOR as the iterative method has the additional benefit that simple, loop-level parallelization can be used to speed up the elliptic solve without having to implement coloring techniques on the triangulation. In practice, this implementation yields parallel speedups on par with other loop-level approaches, and speedups of 4-8 are typical on 16 CPUs. While nowhere near as good as the near-ideal scalability of the Euler method used for the inviscid flow, this performance is sufficient to keep the cost of the boundary-layer updates small.

3. Numerical Results

This section presents an overview of several validation examples with the coupled-IBL method comparing results with data from both experimental and established computational sources. These cases examine the utility of the coupled-IBL method on a range of geometries in two- and three-dimensions. The majority of these cases focus on flows with strong inviscid-viscous interaction and therefore require coupled simulation. However since its incremental cost is negligibly small, and the method can be applied to general geometries, we also present cases showing use of the boundary-layer model to generate simple single-pass viscous corrections to inviscid simulations.

3.1 NACA RM-10: Decoupled Viscous Corrections

The underlying Euler simulation method was developed with a strong focus on complex geometry.^[1-3] This method has been proven to be very useful for generating accurate aerodynamic coefficients for a variety of complex vehicles under a wide variety of flight conditions.^[2-8] Integrated aerodynamic coefficients in these simulations are only weakly affected by viscous-inviscid interaction, but often contain

complex geometry. The need for a robust method to generate viscous corrections for general geometry was a main driver in our selection of the elliptic solve as a technique of distributing the transpiration velocities, skin-friction coefficients, etc. from the boundary-layer method over the wetted surface of the geometry.

The NACA RM-10 research model is a good example for demonstrating the utility of this approach for providing single-pass, *a posteriori*, viscous corrections to inviscid simulations around flight vehicles with weak viscous interactions. Reference [27] summarizes NACA's wind-tunnel and flight-test program for this vehicle which is a finned slender body designed to fly at transonic and supersonic Mach numbers. This research program was aimed at accurately measuring zero-lift drag, and is therefore a very good candidate for this validation exercise.

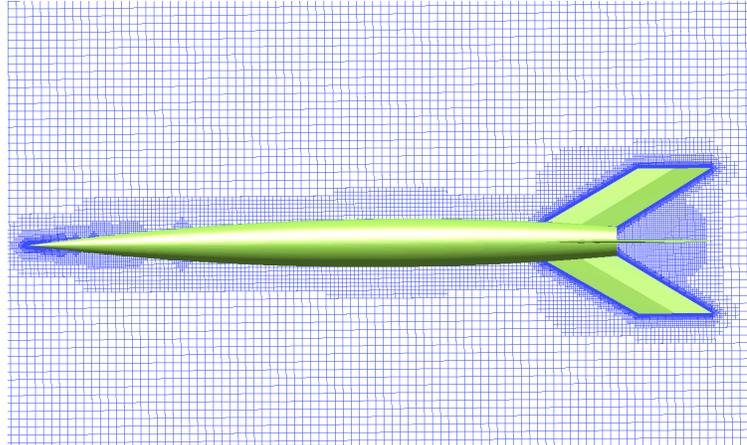


Figure 9: Cartesian mesh around RM-10 research model of reference [27] containing 1.8 M cells.



Figure 10: Computed flow around RM-10 research vehicle at several different Mach numbers and $\alpha = 0^\circ$. Density contours.

Figure 9 shows a side-view of the 146.5in flight-test model along with a cutting-plane through the Cartesian mesh used for all inviscid simulations. The geometry-adapted mesh contained 1.8 M cells and was used for computing flows from Mach 0.8-2.5. This parametric Mach-sweep included 20 simulations spanning this range of Mach numbers all run at 0° angle-of-attack. Figure 10 shows snapshots of the flow from a subset of these runs through density contours on the surface and symmetry plane. Snapshots at subsonic, transonic and two supersonic conditions are shown. The supersonic examples were chosen to show the flows when the Mach angle is both greater than and less than the fin-sweep angle.

To provide a viscous correction to the inviscid simulations, boundary-layer strips were placed on the geometry as illustrated in figure 11. Four loops were placed on each fin, and two (orthogonal) loops were placed on the fuselage running nose-to-base. In decoupled mode, a single call was placed to the boundary-layer after the inviscid flow was converged. Transition was fixed at $Re_x = 5.5 \times 10^5$.

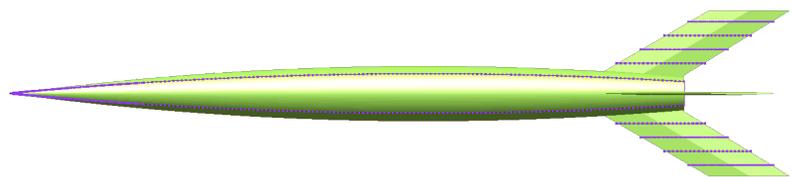


Figure 11: Top view of RM-10 vehicle showing boundary-layer stations.

Flight-test Reynolds numbers (based on maximum body diameter) varied nearly linearly from 35×10^6 at Mach 0.8 to 140×10^6 at Mach 2.5. Figure 12 compares the simulation results with the flight test data in figure 9a of reference [27]. Agreement for the corrected inviscid scheme is excellent, and the resulting axial force coefficient falls within the experimental uncertainty over the entire Mach range. The figure also shows variation of the computed value of integrated skin-friction (“ $C_x \text{ visc}$ ”) over the range of Mach numbers. These values compare very well with estimates of skin-friction drag in reference [27]. While the approach is simplistic, validation of the single-pass mode of the boundary-layer module is nevertheless important since its cost is negligible, and it can be applied to any new or previously computed inviscid simulation.

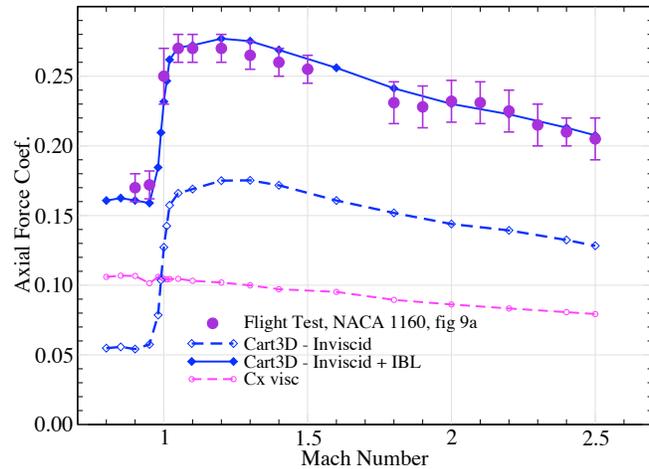


Figure 12: Comparison of computed axial force variation with Mach number with RM-10 flight test data from reference [27]. Reynolds number varies from 35×10^6 at Mach 0.8 to 140×10^6 at Mach 2.5. Boundary-layer correction computed decoupled (single-pass correction to inviscid solution).

3.2 Unit-Span RAE 2822: Coupled-IBL Simulations

Coupled solution techniques are called for whenever the boundary-layer’s displacement thickness substantially modifies the shape of the geometry perceived by the outer flow. Supercritical airfoils, high-lift systems, and transonic transports are only a few examples of important cases where viscous effects make purely inviscid simulations and single-pass viscous corrections poor predictive tools.

The abundance of experimental and simulation data for subsonic and transonic flow over an RAE 2822 make it an ideal first case for examining the performance of the current coupled-IBL technique. The airfoil section data was taken from Cook *et al*[23] and was used to generate the 3D unit-span unswept wing shown previously in the figures of §2. While the geometry is a simple airfoil, the solution was carried out using a full 3D mesh containing about 800k cells using the mesh generator described in reference [1].

The RAE 2822 was computed at $M_\infty = 0.74$, $\alpha = 2.643^\circ$ and $Re_C = 2.7 \times 10^6$ using the iterative coupling described in §2.4. Transition was fixed at 5% chord, and five boundary-layer stations were evenly spaced across the unit-span wing. The inviscid solver used 4-level multigrid acceleration and the boundary-layer was updated every 2 multigrid cycles. As a baseline for comparison, the same case was computed using pure inviscid modeling with the same mesh and solver setup.

Figure 13 displays convergence histories for both the inviscid and coupled-IBL simulations. Each graph shows convergence of the L1 norm of density residual as well as convergence of the integrated force vector resolved into its Cartesian components. In both simulations, forces are converged to engineering accuracy after about 50 multigrid cycles. Convergence of density residual is somewhat faster and deeper for the inviscid case than for that with IBL coupling, however it is still reasonable. From this single example its difficult to make firm statements about computational cost. However, with five boundary-layer strips, the SOR-driven elliptic solve on the surface triangulation converges the blowing velocities 6 orders of magnitude in about 800 sweeps over the surface triangulation. With the boundary-layer being re-evaluated every other multigrid cycle, this expense roughly doubled the cost of the computation over the

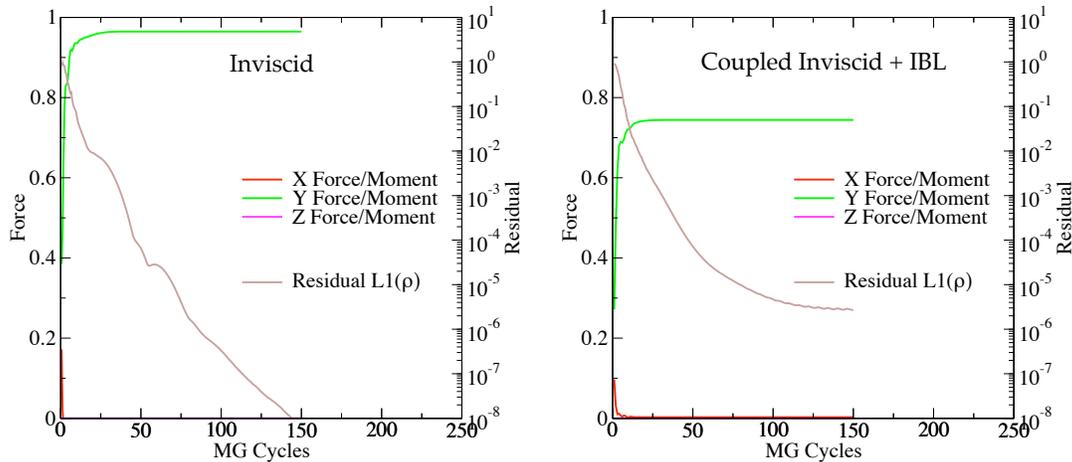


Figure 13: Comparison of residual and force convergence for inviscid (left) and coupled IBL simulations for flow over a unit span RAE 2822 at $M_\infty = 0.74$, $\alpha = 2.643^\circ$ and $Re = 2.7 \times 10^6$. Both cases were converged using 4 level multigrid for the inviscid scheme.

baseline inviscid simulation. This cost is obviously dependent upon many factors including number of CPUs, number of surface triangles, and desired depth of convergence. Thus far in the validation process, the additional cost has not been burdensome and many opportunities exist to reduce it further.

The large difference in normal force (“Y Force”) in fig. 13 gives an indication of the dramatic differences that viscous effects have in this flow. Figure 14 shows this more clearly by comparing pressure coefficient distributions for the inviscid and coupled-IBL simulations with results from both MSES and ARC2D. Drela’s MSES solver is a fully-coupled, implicit, inviscid-IBL code and is one of the most extensively validated 2D airfoil design codes in existence.^[28,25] Pulliam’s ARC2D code is one of NASA’s most widely disseminated 2D Reynolds Averaged Navier-Stokes (RANS) solvers, and has been extensively tested in over two decades of popular use and improvement.^[30,31] Simulations with both MSES and ARC2D held transition fixed at 5% chord, and ARC2D was run using the Spalart-Allmaras turbulence model. These codes were otherwise run using default input parameters and meshes.

The pressure distributions in figure 14 show the expected changes in shock-location due to the boundary layer as it has moved the shock forward from its inviscid location by about 12% chord and weakened it substantially. Beyond this, there are some minor differences between the three viscous profiles but most are due to shock-resolution in the discrete solutions. The current coupled-IBL approach does show slightly different behavior on the upper surface immediately post-shock. This is most likely due to a slightly delayed growth of the displacement thickness as a byproduct of over-smoothing the splined external flow data used to drive the boundary-layer.

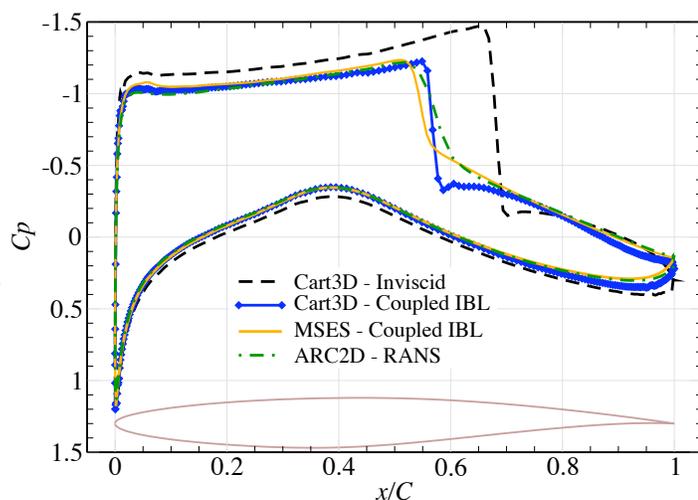


Figure 14: Comparison of pressure distribution for inviscid and coupled IBL solutions with simulations performed using MSES (inviscid+IBL)^[28,25] and ARC2D (RANS)^[30,31]. RAE 2822, $M_\infty = 0.74$, $\alpha = 2.643^\circ$ and $Re = 2.7 \times 10^6$.

Table 1 summarizes the force coefficients in the aerodynamic frame for all four simulations. In this table, the lift coefficients for ARC2D and MSES agree to within about 2%, and their drag predictions differ by 2.4 counts. Lift with the current coupled-IBL method agrees with ARC2D to 3 digits. The predicted drag falls in between the two reference simulations. This case emphasizes the improvement in aerodynamic force prediction provided by the coupled solver. The pure inviscid baseline case in figure 14 overpredicts ARC2D's lift by over 28% with a similar discrepancy in drag.

	C_L	C_D
ARC2D	0.748	0.0202
MSES	0.733	0.0178
Cart3D Inviscid	0.960	0.0248
Cart3D Coupled-IBL	0.748	0.0196

The RAE 2822 case in figures 13 and 14 examines the behavior of the coupled-IBL method under conditions when flow separation is incipient. To provide a broader picture of the method's performance, an angle-of-attack sweep was performed using all four methods listed in Table 1. Figure 15 highlights results of this investigation showing both the lift-curve (left) and drag polar (right). Mach number and Reynolds number stayed fixed at 0.74 and 2.7×10^6 while incidence angle was varied from -3° to 4° .

Table 1: Comparison of integrated aerodynamic forces for RAE 2822 at $M_\infty = 0.74$, $\alpha = 2.643^\circ$ and $Re = 2.7 \times 10^6$

The results in figure 15 show excellent agreement between the present coupled-IBL approach with those of both ARC2D and MSES over nearly the entire range of incidence angles. In examining the lift-curves, the primary differences are in prediction of the airfoil stall somewhere between 3.5° and 4° . Closer examination of these cases reveals that ARC2D, MSES and the current coupled-IBL approach failed to converge at 4° and above (the circled datapoints in both frames of fig. 15). Examination of the partially-converged flow at these conditions shows massive shock-induced separation on the upper surface, resulting in a flow that appears unsteady. Results from the baseline inviscid solver are included on both plots. Its interesting to note that, for this case, the inviscid drag polar can be corrected reasonably well by simply adding in the drag from the viscous predictions at zero-degrees incidence (dotted line in drag polar at right of fig. 15).

3.3 ONERA M6 Wing

The transonic, turbulent flow over the ONERA M6 wing provides the first example with fully three-dimensional geometry. Experiments for CFD code validation using this case are detailed in reference [32],

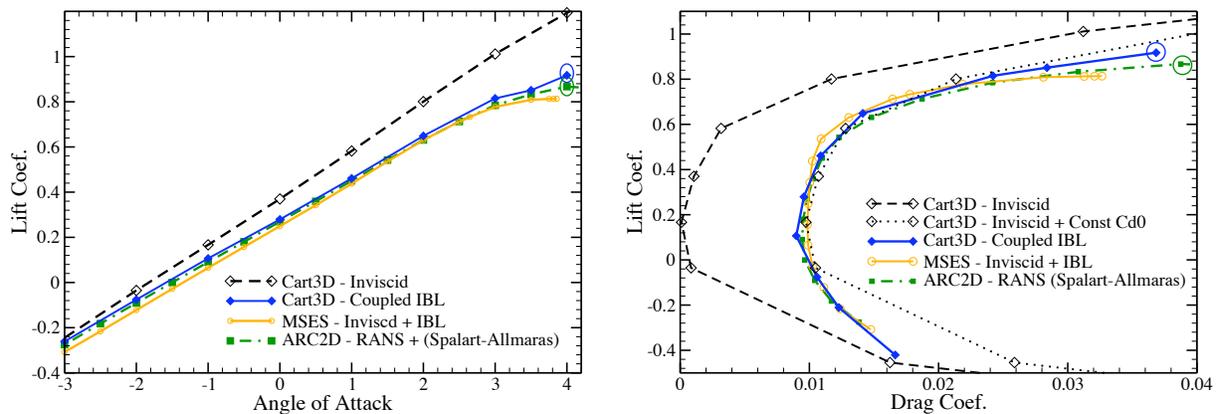


Figure 15: Comparison of lift coefficient variation with angle-of-attack (left) and accompanying drag polar (right) for inviscid and coupled IBL solutions with simulations performed using MSES (inviscid + IBL)^[28,25] and ARC2D (RANS)^[30,31] for RAE 2822 at $M_\infty = 0.74$, and $Re = 2.7 \times 10^6$. Results labeled "Inviscid + Const C_{d0} " use the 0° drag from the coupled IBL method to correct the pure inviscid polar. Circled data indicates cases with poor convergence.

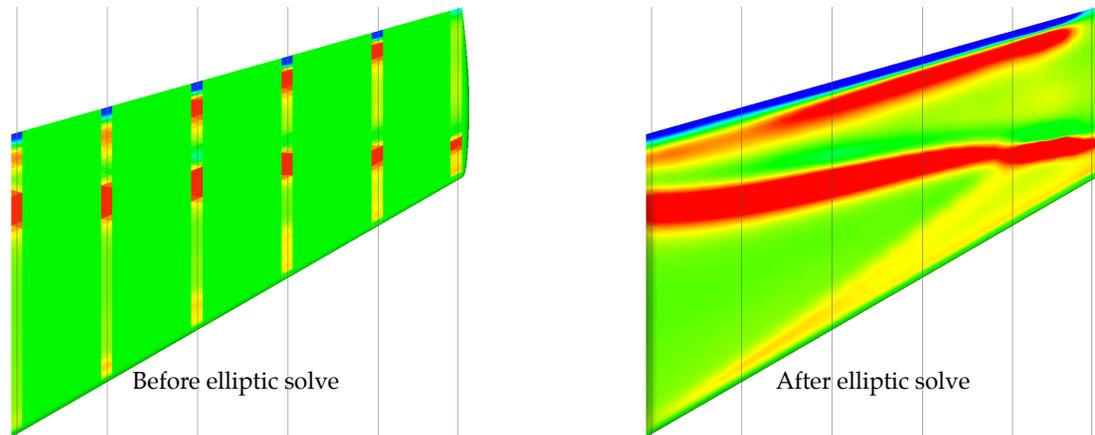
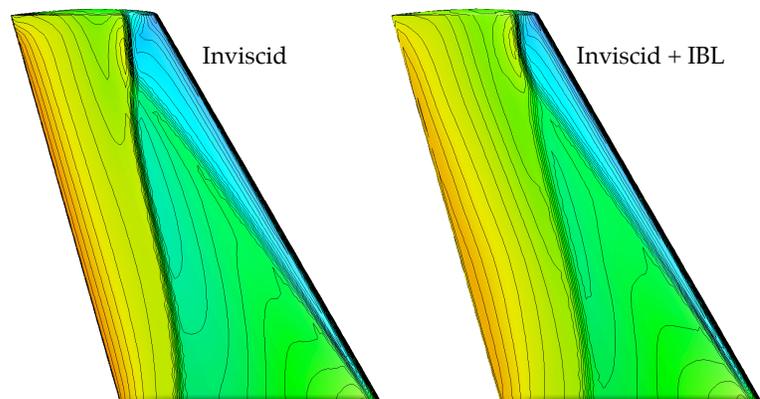


Figure 16: Blowing velocities on upper surface of ONERA M6 wing before and after elliptic solve. *Left:* Boundary-layer cuts and blowing-velocities mapped to surface triangulation strips. *Right:* Blowing velocities on surface triangulation after elliptic solve to distribute over surface. ONERA M6 wing $M_\infty = 0.8395$, $\alpha = 3.06^\circ$, and $Re_C = 11.72 \times 10^6$ from reference [32].

and the case is a popular validation case for RANS solvers (*cf.* ref. [33]). Reference [32] details the geometry of the wing which has an extruded and tapered airfoil section with a taper ratio of 0.562, an aspect ratio of 3.8 and a leading edge sweep of 30° . Simulations were performed at $M_\infty = 0.8395$, $\alpha = 3.06^\circ$, and $Re_C = 11.72 \times 10^6$ (based on mean aerodynamic chord). These conditions correspond to test 2308 in reference [32]. The computational mesh for this example contained 1.1M cells and is a standard regression test for the Cart3D software package.

In addition to examining results of the flow simulation, this geometry also provides an opportunity to examine performance of the surface elliptic solver in accurately painting the geometry's surface with boundary-layer scalars as described in §2.4. Figure 16 shows top views of the wing before and after performing the elliptic solve. The frame on the left shows blowing-velocities computed with Lighthill's relation (eq.(1)) along the six splined boundary-layer stations and transferred to triangles on the surface that the stations pass through. The SOR elliptic solver converged the blowing-velocity distribution 6 orders-of-magnitude using about 400 iterations over the 13,200 triangles describing the surface of the wing. The resulting blowing-velocity distribution is shown in the frame at the right of figure 16.

Figure 17 presents a comparison of surface pressures on the upper surface of the wing between the baseline inviscid solver and the coupled-IBL method using flooded isobars. Results with both simulations display the the lambda shock structure which is so characteristic of this case. As seen in the RAE case, the shock system in the coupled solution has migrated forward somewhat from its location in the inviscid simulation.



Reference [32] contains experimental pressure distributions at several stations along the wing's span. These data have been widely used for validation of

Figure 17: Comparison of surface pressure distributions for inviscid (left) and coupled (inviscid + IBL) for ONERA M6 wing at $M_\infty = 0.8395$, $\alpha = 3.06^\circ$, and $Re_C = 11.72 \times 10^6$.

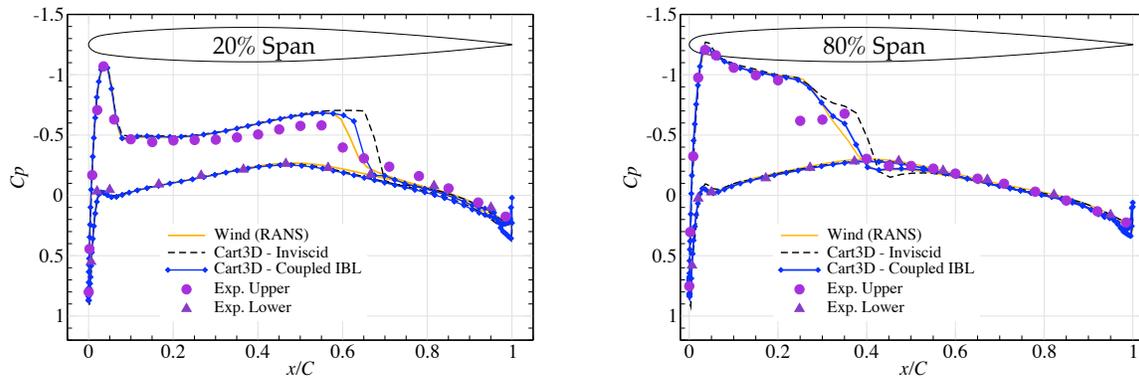


Figure 18: Comparison of surface pressure distributions with data from experiment in reference [32] and RANS simulations in [33] at 20% span (left) and 80% span (right) for ONERA M6 wing at $M_\infty = 0.8395$, $\alpha = 3.06^\circ$, and $Re_C = 11.72 \times 10^6$.

RANS solvers, and reference [33] contains detailed comparisons for the NPARC Alliance's WIND solver. Figure 18 contains comparisons to both these experimental and simulation data for both the baseline inviscid code as well as the coupled-IBL solver. Pressure distributions are shown at 20 and 80 percent span and these results are typical of those found at other spanwise locations. Overall, agreement is very good with the RANS results and reasonable with the experiment. These pressure distributions make clear the forward migration of the wing shocks. Both plots show a slight pressure glitch due to the trailing-edge treatment in the boundary-layer solve.

3.4 DLR F4

The recent series of AIAA sponsored Drag Prediction Workshops (DPW) have been conducted to establish the capability of current simulation techniques to predict aerodynamic coefficients and detailed flow structure around transonic transport geometries.^[9,10,11,35] These workshops have produced a wealth of simulation data and an invaluable database for solver validation. The first workshop (DPW-I) examined the flow about the DLR-F4 geometry. This transonic wing-body configuration has a supercritical wing-section and a blunt trailing-edge. Reference [11] estimates that inviscid simulations of this configuration can be performed for about 1/50th the cost of full RANS simulations. Unfortunately, as with the RAE case presented in §3.2, inviscid simulations of this geometry can dramatically over-predict lift and shock strength, making them of little use as predictive tools.

Figure 19 provides an overview of the DLR-F4 geometry and shows surface pressure contours for simulations performed using both the coupled-IBL and baseline inviscid solvers. These simulations were performed at $M_\infty = 0.75$, $\alpha = 1.0^\circ$, and $Re_C = 3 \times 10^6$ (based on mean aerodynamic chord) and correspond to one of the points on the drag polar required at the workshop. Both simulations were performed on the same Cartesian mesh which used approximately 2.5M cells for the full-span configuration. The supercritical airfoil section on this model's wing has a very flat top which makes the shock location quite sensitive to boundary-layer development. The top-views of pressure in figure 19 clearly illustrate the disparity in shock-location between the inviscid and viscous simulations. As a result, the inviscid simulation's lift coefficient of 0.85 is 40% higher than the experimental value for these flow conditions.

The inset wing-planform at the left of Figure 19 shows the location of boundary layer strips placed on each wing of the model for the full-span simulations. In addition to these, four additional boundary-layer strips were placed down the fuselage of the model for the coupled-IBL solution. Figure 20 shows conver-

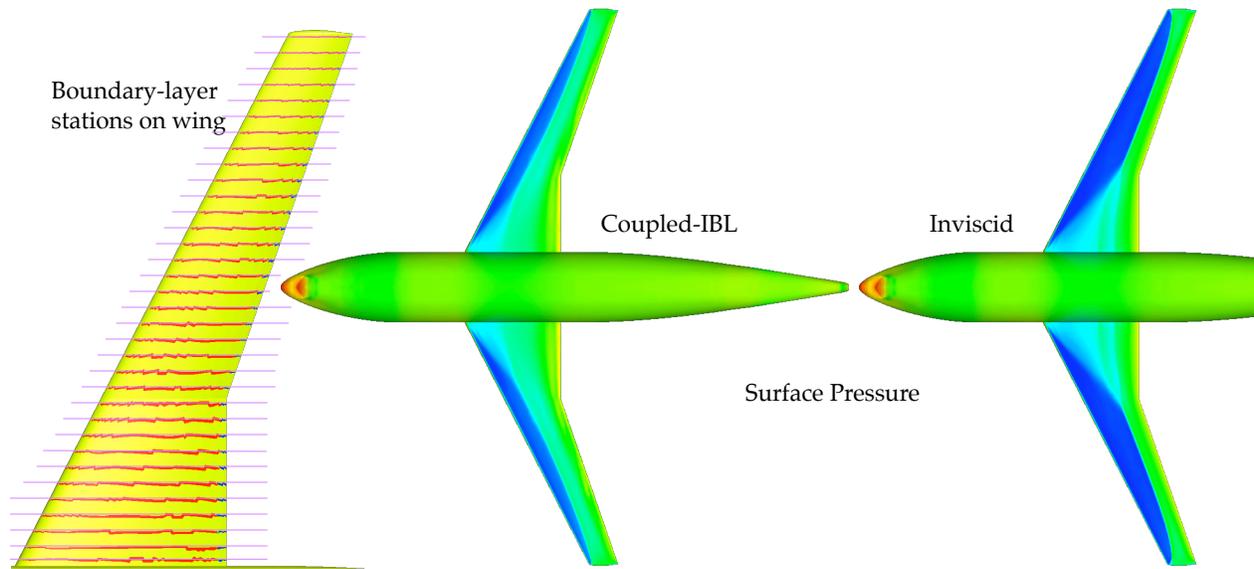


Figure 19: Surface pressure distribution for coupled-IBL (left) and inviscid (right) simulations on DLR-F4 wing-body geometry at $M_\infty = 0.75$, $\alpha = 1.0^\circ$, and $Re_C = 3 \times 10^6$. In addition to the boundary layer stations on the wing (far left), the coupled-IBL simulations used four additional stations on the fuselage (not shown).

gence of the coupled-IBL solution using 4-levels of multigrid. Boundary-layer updates were performed every other multigrid cycle. Convergence of integrated forces is somewhat slower than in the previous example on the RAE wing, but is still reasonable. While initial convergence of the L1 density norm is quite good, the residual stalls abruptly after converging just over four orders of magnitude. Detailed inspection of the developing boundary-layer profiles indicates that this may be due to a slight separation at some of the inboard stations. The right frame in figure 20 compares the surface pressure profiles of the inviscid and coupled-IBL method with RANS results at 40.9% span. The RANS results included for comparison were reproduced from reference [11] and are based on the NSU3D solver. The C_p distributions in Figure 20 show the shock moving forward by about 20% of the wing-chord at this spanwise location, and the profile of the coupled method agrees very well with that produced by NSU3D. As in earlier examples, there is a slight irregularity of the C_p profile near the wing trailing-edge. Reference [11] noted a similar irregularity when computing the same flow using a coupled-IBL approach based on essentially the same boundary-layer routines.

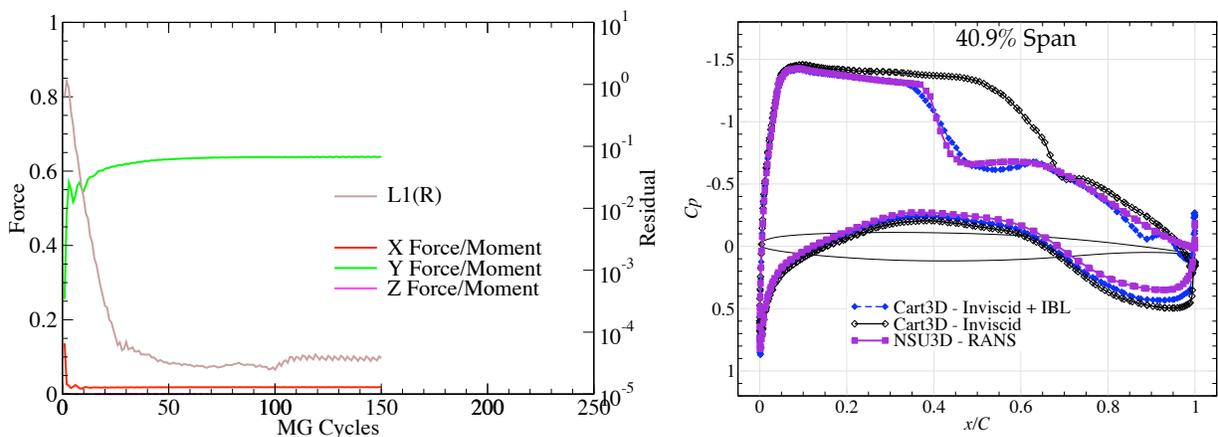


Figure 20: DLR-F4 convergence history (left) and computed surface pressure coefficients at 40.9% span at $M_\infty = 0.75$, $\alpha = 1.0^\circ$, and $Re_C = 3 \times 10^6$. Pressures with the present coupled-IBL approach are compared with results from the pure inviscid solver and published data using the NSU3D RANS solver from reference [11].

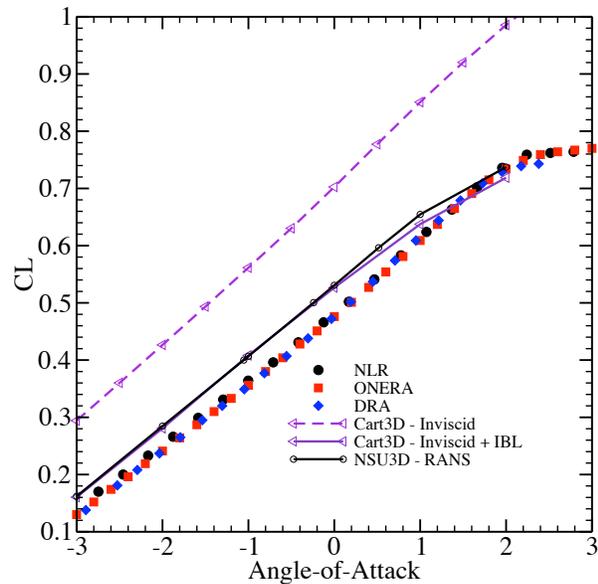


Figure 21: Comparison of lift coefficient variation for DLR-F4 wing-body using the coupled IBL approach with inviscid flow, experimental data and RANS simulation using NSU3D. $M_\infty = 0.75$, $Re_C = 3 \times 10^6$.

The coupled-IBL solution in figures 19 and 20 produced an integrated C_L of 0.637 which compares well to the RANS result of 0.655. Both of these simulations slightly over-predict the experimentally determined value of C_L , of around 0.615 ± 0.01 .^[34]

The required elements of DPW-I included computation the drag polar for the DLR-F4 wing-body at $M_\infty = 0.75$, and $Re_C = 3 \times 10^6$ by varying incidence angle from -3° to 2° . This study was reproduced using the current coupled-IBL solver as a broader examination of its predictive ability. Figure 21 shows the resulting lift-curve for these simulations compared against those of the baseline inviscid solver, the NSU3D RANS solver, and experimental data from three sources.^{[11][34]} In this plot, the current results nearly overlay the predictions from NSU3D. Both the coupled-IBL and RANS results display a slight vertical shift of the lift-curve from the experimental data. This behavior is representative of that reported by other RANS solvers (OVERFLOW, CFD++, FUN3D, TAU and others) that participated in the workshop.^[35]

Figure 22 presents the drag polar resulting from this parametric study. This plot shows the data from the coupled method agree extremely well with both the experimental values and the predictions from NSU3D. Unlike the RAE 2822 case, correcting the inviscid polar by a constant viscous drag increment from the $\alpha = 0^\circ$ case does not accurately track the drag polar. Considering the simplistic nature of the viscous modeling used in this strip boundary-layer approach, results for the coupled-IBL approach are very encouraging.

4. Discussion

The examples in the preceding section show that the coupled-IBL method can predict accurate values of integrated aerodynamic coefficients for a wide variety of cases and flow conditions. Nevertheless, these same examples also highlight some shortcomings of both the technique and its current implementation that merit further discussion.

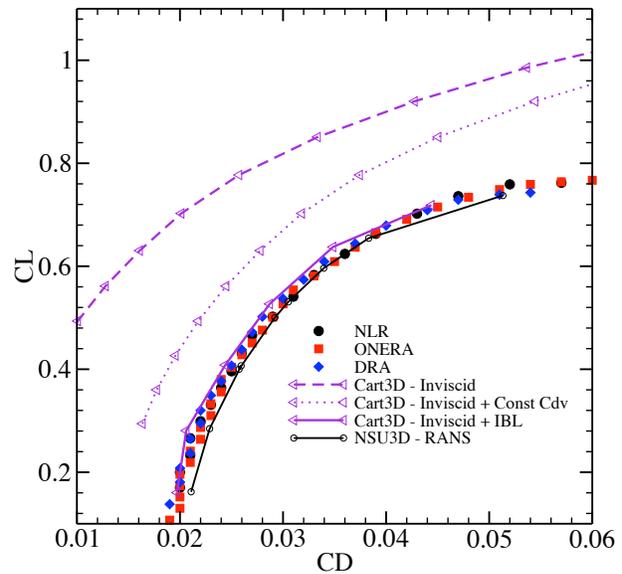


Figure 22: Comparison of computed drag polar for DLR-F4 wing-body using the coupled IBL approach with inviscid flow, experimental data and RANS simulation using NSU3D. Results labeled "Inviscid + Const C_{dv} " use the 0° drag from the coupled IBL method to correct the pure inviscid polar. $M_\infty = 0.75$, $Re_C = 3 \times 10^6$.

Numerical results were presented for coupled simulations of wing-sections with cusped, sharp and blunt trailing-edges. Reviewing the C_p profiles for each of these cases, one observes that while the trailing-edge treatment from §2.4 performed reasonably well for geometry with cusped and sharp trailing edges, it was less successful for the blunt trailing-edge of the DLR-F4 wing. Thus, despite generally good aerodynamic predictions for this case, there are some discrepancies in the flow near the wing trailing-edge. Given the importance of flow in this region to the circulation around the entire wing, issues in this region still need to be addressed. A variety of approaches have been examined, but no unified treatment has yet been developed.

The examples in Figures 14, 16 and 20 were all performed at flow conditions either near or containing limited regions of separated flow. In simulations of both the RAE 2822 and the DLR-F4, the number of points in the splined boundary-layer inputs needed to be adjusted from its default of 400 to a more modest 250. Without this reduction, these cases converged poorly and some of the boundary-layer solutions tended to show premature shock-induced or trailing-edge separation. Detailed inspection of the boundary-layer suggested that under some circumstances, it can prematurely predict flow separation from sharp adverse pressure gradients. Model verification cases performed using ARC2D and MSES seem to indicate that this particular boundary-layer model may be somewhat less robust than that used in MSES.^[28,25] Despite these misgivings, the results in §3 show that with some care near C_{Lmax} , the current model is very capable of accurately predicting aerodynamic trends.

Several of the results in the previous section were accompanied by estimates of the computational overhead required by the boundary-layer solution. Since the strips are processed in parallel, the time required for boundary-layer integration even using large numbers of strips is extremely small. Moreover, using more strips has the advantage of accelerating the elliptic solve on the surface triangulation since it strengthens the boundary condition enforcement. Nevertheless, the elliptic solve does remain the most expensive part of the boundary-layer update. On small numbers of processors, the time for this solve is vanishingly small. It is only performed on the $\mathcal{O}(N^2)$ triangles of the surface geometry as compared with the $\mathcal{O}(N^3)$ cells in the volume mesh integrated by the Euler solver. On larger numbers of processors, however, the near-ideal parallel scalability of the Euler solver makes the elliptic solve comparatively more expensive. Thus, while the coupled-IBL solver may be only 10% more expensive than the baseline inviscid code on 8 CPUs, it may be twice as expensive as the inviscid code on 128 CPUs. Domain-decomposition, multigrid and other powerful techniques offer many simple remedies for dramatically improving performance of this elliptic solve and this is not an area of concern.

5. Summary

This paper presented an overview of the development and selected validation of a coupled-IBL approach for simulating viscous flows around complex aerospace configurations. The method combined an established multilevel Cartesian-mesh Euler solver with a transpiration boundary condition to account for the boundary-layer displacement thickness. This transpiration condition is set via a strip-wise solution of the 2D boundary-layer equations which uses the inviscid solution as a driver. The implementation uses local flow topology to establish attachment and separation and solution of Laplace's equation on the surface triangulation to couple the transpiration velocities back to the inviscid solver. This elliptic system is solved using an implicit SOR approach which is implemented using loop-level parallel programming constructs to decrease its expense on parallel computing hardware. This approach allows boundary-layer properties to be distributed over the entire surface of complex vehicles removing the need for additional viscous corrections for fuselages, nacelles, or other components of a complex configuration.

Numerical investigations were carried out on a variety of discriminating test cases. Results were presented for wings with cusped, sharp and blunt trailing edges and included cases with both strong and weak viscous coupling. Drag polars and lift-curves were presented for both an RAE 2822 wing and the DLR-F4 wing-body used in the 1st AIAA Drag Prediction Workshop. Both investigations yielded aerodynamic force coefficients that agreed well with established results from RANS solvers and experiment. In all the coupled examples, examination of the surface pressure distributions revealed accurate predictions of the shock location on transonic wings. Results for zero-lift drag of the NACA RM-10 finned missile geometry showed good agreement with flight-test measurements over a range of trans- and supersonic Mach numbers.

The discussion outlined some areas of outstanding interest. These include the development of a unified trailing-edge treatment without a wake model and the premature onset of flow separation near regions of steep pressure gradients. Despite these outstanding issues, the low computational cost and demonstrated predictive ability of the coupled-IBL approach make it a very attractive simulation tool to both widen the applicability of the underlying inviscid solver and to provide preliminary insight into flows that would otherwise require much more expensive RANS simulations.

Acknowledgements

The authors wish to thank John Vassberg of The Boeing Company for providing the basic boundary-layer solver used in this work. In addition we thank Dimitri Mavriplis, Tom Pulliam, Marian Nemec and Scott Murman both for providing simulation data and for many insightful discussions. Marsha Berger was supported in part by DOE grants DE-FG02-88ER25053 and DE-FC02-01ER25472 and AFOSR Grant F49620-03-1-0122.

Reference

- [1] Aftosmis, M.J., Berger, M.J., Melton, J.E., "Robust and efficient Cartesian mesh generation for component-based geometry." *AIAA J.* **36**(6):952-960, Jun. 1998.
- [2] Aftosmis, M. J., Berger, M.J., and Adomavicius, G., "A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries." *AIAA Paper 2000-0808*, Jan. 2000.
- [3] Aftosmis M.J., and Berger M.J., "Multilevel error estimation and adaptive h-refinement for Cartesian meshes with embedded boundaries." *AIAA Paper 2002-0863*, Jan. 2002.
- [4] Chaderjian, N. M., Rogers, S. E., Aftosmis, M. J., Pandya, S.A., Tejnil, E., and Ahmad, J., "Automated CFD database generation for a 2nd generation glide-back booster." *AIAA Paper 2003-3788*, Jun. 2003.
- [5] Murman, S. M., Aftosmis, M.J., and Berger M.J., "Numerical simulation of rolling-airframes using a multi-level Cartesian method." *AIAA Paper 2002-2798*, 20th AIAA Applied Aerodynamics Conference, St. Louis, MO, Jun. 2002. Also *Jol. Spacecraft and Rockets*, **41**(3):426-435. May 2004.
- [6] Cliff, S., Thomas, S., Baker, T., Jameson, A., and Hicks, R., "Aerodynamic shape optimization using unstructured grid methods." *AIAA Paper 2002-5550*, Sep. 2002.
- [7] Rogers, S. E., Aftosmis, M. J., Pandya, S.A., and Chaderjian, N. M., Tejnil, E., and Ahmad, J., "Automated CFD parameter studies on distributed parallel computers." *AIAA Paper 2003-4229*. 16th AIAA Computational Fluid Dynamics Conference, Jun. 2003.
- [8] Murman, S.M., Aftosmis, M.J., and Nemec, M., "Automated parameter studies using a Cartesian method." *AIAA Paper 2004-5076*, Aug. 2004.

- [9] Lee-Rausch, E.M., Buning, P.G., Mavriplis, D.J., Morrison, J.H., Park, M.A., Rivers, S.M., and Rumsey, C.L., "CFD sensitivity analysis of a Drag Prediction Workshop wing/body transport configuration." *AIAA 2003-3400*, Jun. 2003.
- [10] Rumsey, C.L., Rivers, S., and Morrison, J., "Study of CFD variation on transport configurations from the second Drag-Prediction Workshop." *AIAA-2004-0394*, Jan. 2004.
- [11] Mavriplis, D.J., "Drag Prediction Using Unstructured Mesh Solvers, D. J. Mavriplis VKI Lecture Notes, CFD-Based Aircraft Drag Prediction and Reduction." Von Karman Institute for Fluid Dynamics, Rhode-Saint-Genese, Belgium, *VKI LS 2003-02*, Feb. 2003.
- [12] Cebeci, T., Sedlock, D., Chang, K.C., and Clark, R.W., "Analysis of wings with flow separation." *J. of Aircraft* **26**(3):214-220, Mar. 1989.
- [13] Potsdam, M.A., "An unstructured mesh Euler and interactive boundary layer method for complex geometries." *AIAA Paper 94-1844*, In Proceedings of the 12th AIAA Aerodynamics Conference, Colorado Springs, CO, Jun. 1994.
- [14] Al-Saadi, J.A., Smith, W.D., "Analysis of a supercritical wing with a blunt trailing edge using an unstructured-grid Euler method." *NASA TM 4688*, Oct. 1995.
- [15] Buning, P.G., Gomez, R., Scallion, W., "CFD approaches for simulation of wing-body stage separation." *AIAA-2004-4838*, Aug. 2004
- [16] McDonald A.G., and Nash J.F., "A turbulent skin-friction law for use at subsonic and transonic speeds." *Aeronautical Research Council Report No. ARC-CP-948*, Jan. 1967.
- [17] Cebeci, T., and Cousteaux, J., *Modeling and Computation of Boundary Layer Flows*. Springer, Horizon Publishing Inc., Long Beach, CA, 2001.
- [18] Clarence B. Cohen, Eli Reshotko. "The compressible laminar boundary layer with heat transfer and arbitrary pressure gradient." *NACA Report 1294*, Jan. 1956.
- [19] Mc Nally, W.D., "FORTRAN program for calculating compressible laminar and turbulent boundary layers in arbitrary pressure gradients," *NASA TN D-5681*, 1970.
- [20] Lighthill, M.J., "On displacement thickness." *Jol. of Fluid Mechanics*, 4(4):383-392), 1958.
- [21] Buning P.G., and Steger, J., "Graphics and Flow Visualization in Computational Fluid Dynamics," *Proc. AIAA 7th Computational Fluid Dynamics Conf.* American Institute of Aeronautics and Astronautics, New York, pp.162-170, 1985.
- [22] Press, W.H., Flannery, B.P., Teukolsky, S.A., and Vetterling, W.T., *Numerical Recipes in C: The Art of Scientific Computing*. 2nd Edition. Cambridge University Press, 1993.
- [23] Cook, P.H., M.A. McDonald, M.C.P. Firmin, "Aerofoil RAE 2822 - Pressure Distributions, and Boundary Layer and Wake Measurements," *Experimental Data Base for Computer Program Assessment, AGARD Report AR 138*, 1979.
- [24] Tobak M., and Peake, D.J., "Topology of 3D Separated Flows," *Ann. Rev. Fluid Mechanics*, 14:61-85, 1982.
- [25] Giles, M.B., and Drela, M., "Two-Dimensional Transonic Aerodynamic Design Method," *AIAA J.*, **25**(9)1199-1206, Sep. 1987.
- [26] Dallmann, U., "Topological Structures of 3D Flow Separation," *Tech. Report DFVLR-IB 221-82 A 07*. German Research Institute for Aerospace. Göttingen, Germany, Apr. 1983.

- [27] Evans, A. J., "The zero-lift drag of a slender body of revolution (NACA RM-10 research model) as determined from tests in several wind tunnels and in flight at supersonic speeds." *NACA Report 1160*, 1954.
- [28] Drela, M., and Giles, M. B., "ISES: A Two-Dimensional Viscous Aerodynamic Design and Analysis Code," *AIAA Paper 87-0424*, Jan. 1087.
- [29] Saad, Y., *Iterative methods for sparse linear systems*, PWS Publishing Company, Boston, MA, 1996.
- [30] Pulliam, T.H., "Efficient Solution Methods for The Navier-Stokes Equations", Lecture Notes for the von Kármán Institute For Fluid Dynamics Lecture Series :Numerical Techniques for Viscous Flow Computation In Turbomachinery Bladings, von Kármán Institute, Rhode-St-Genese, Belgium , 1985.
- [31] Pulliam T. H. and Steger, J. L. , "Recent Improvements in Efficiency, Accuracy, and Convergence for Implicit Approximate Factorization Algorithms", *AIAA 85-0360*, Jan. 1985.
- [32] Schmitt, V. and F. Charpin, "Pressure Distributions on the ONERA-M6-Wing at Transonic Mach Numbers," Experimental Data Base for Computer Program Assessment. Report of the Fluid Dynamics Panel Working Group 04, *AGARD AR 138*, May 1979.
- [33] Slater, J. W., " NPARC Alliance Validation Archive: ONERA M6 Wing: Study #1," <http://www.grc.nasa.gov/WWW/wind/valid/m6wing/m6wing01/m6wing01.html>. Aug. 30, 2002.
- [34] Redeker, G., "A Selection of Experimental Test Cases for the Validation of CFD Codes," *AGARD AR 303 Volume 2*, ISBN 92-836-1003-2, Aug. 1994.
- [35] Levy, D.W., Zickuhr, T., Vassberg, J., Agrawal, S., Wahls, R.A., Pirzadeh, S., and Hensch, M.J., "Data summary from the first AIAA computational fluid dynamics drag prediction workshop," *Jol. of Aircraft*, **40**(5):875-882, Sep.-Oct. 2003.