Adjoint-Based Low-Boom Design with Cart3D

Michael J. Aftosmis*
NASA Ames, Moffett Field, CA 94035

Marian Nemec†
Science & Technology Corp., Moffett Field, CA 94035

Susan E. Cliff‡
NASA Ames Research Center, Moffett Field, CA 94035

We consider analysis and design of low sonic-boom aircraft through the use of an inviscid, embedded-boundary Cartesian mesh method. Adjoint error estimation and adaptive meshing are used in the analysis portion of this study to determine resolution requirements of the computational domain, while adjoint-based gradients in conjunction with a parallel optimization framework are used for design. Two analysis examples are presented. The first is a verification study that shows excellent accuracy of mesh converged pressure signals when sampled both on-track and off-track at distances of 1.8 body lengths on meshes with ~13 million cells. The second is a validation study of a low-boom model that shows close agreement between wind tunnel experiments and off-body pressure signals computed via our Cartesian approach. The optimization example uses an inverse design approach to re-shape a low-boom Mach 1.6 aircraft traveling at 45,000 ft. A detailed parametric model is constructed in the RAGE modeler with approximately 180 parameters used as design variables. Optimization reduced the perceived ground noise level by ~10 dB to a rating of 76.7 PLdB while maintaining the lift and drag performance of the aircraft. Our discussion shows that for functionals with certain common forms, the adjoint solution provides a detailed bookkeeping that can highlight specific regions of the aircraft that contribute to the cost function. In such cases, examination of the adjoint solution directly highlights critical portions of the model, encouraging designers to efficiently select design variables.

I. Introduction

Analysis methods for sonic-boom prediction have improved dramatically in recent years as a result of both commercial and government interest in viable overland supersonic flight. Backed by programatic investment within NASA and elsewhere, a number of simulation tools have recently become available for predicting high-fidelity pressure signals several body lengths away from an aircraft.[1–11] At such distances, details of the three-dimensional aircraft geometry become less important and atmospheric propagation codes[12–14] can be employed to model wave propagation through the atmosphere and to the ground.

While several technologies have played roles in improving the effectiveness of CFD-based analysis, one of the keys has been the widespread use of adaptive mesh techniques.[1–4,6,7,15–17] Configurations designed for low sonic-boom necessarily send only weak pressure disturbances toward the ground. Accurate propagation of such weak waves over several body lengths is a challenge for any numerical simulation technique. Adaptive meshing techniques concentrate and orient mesh elements in the computational domain to more efficiently propagate these signals and have been far more successful than earlier efforts. Particularly noteworthy has been the contribution of adjoint equation or output-based meshing approaches which can prioritize mesh refinement specifically to annihilate error in the propagated signal.[3,4,6,16,18] The insight provided by these approaches has benefitted even fixed-mesh approaches since they aid in our understanding of flow sensitivities, meshing requirements and the role of discretization error in these simulations.

*Aerospace Engineer, *michael.aftosmis@nasa.gov* Associate Fellow AIAA.
†Senior Research Scientist, *marian.nemec@nasa.gov* Member AIAA.
‡Aerospace Engineer, *susan.e.cliff@nasa.gov* Associate Fellow AIAA.

This improved analysis capability has increased confidence in the use of numerical simulations in low-boom aircraft design. Moreover, the natural division between the near-body simulation and the long distance atmospheric signal propagation, sketched in Figure 1, lends itself to an appealing inverse design methodology. A target pressure distribution – with an acceptable ground signature – can be prescribed in the CFD domain. An optimizer can then use this prescribed target pressure distribution in shaping the aircraft.

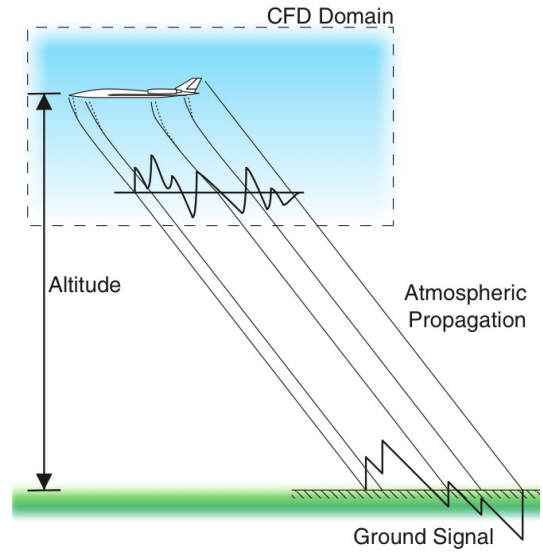

Figure 1. Sketch of basic methodology used for sonic-boom prediction with CFD.

In recent work, we presented an adjoint-based design framework for heterogeneous parallel computing environments.[8] Since the cost of the adjoint approach is almost independent of the number of design variables, it uses dramatically fewer resources than gradient-based methods based upon finite-differences.[19] The frameworks' flexibility stems from a core Cartesian meshing and analysis capability which is fast, robust and insensitive to geometric complexity.[20–22] This software was initially demonstrated on the low-boom shaping of a simplified aircraft geometry using a CAD-based geometry engine.[8] In the current work, we consider a more realistically detailed non-CAD model with empennage, pylons and nacelles. The optimization problem is also more realistic as it includes both on-track and off-track noise targets while seeking to maintain aerodynamic performance.

After first presenting relevant verification and validation studies, our presentation of the optimization example focuses on the role of the discrete adjoint solution in the optimization process. We show adjoints can provide a detailed mapping of offending regions of the pressure signal to specific parts of the vehicle. This aids not only in understanding the physics in the flow, but also in model layout, construction, and parameterization.

II. Shape Optimization with Adjoint-Based Gradients

The aerodynamic optimization problem we consider determines values of design variables, X[a], that minimize a given objective function

$$\min_{X} \; \mathcal{J}(X, Q) \tag{1}$$

where $Q = [\rho, \rho u, \rho v, \rho w, \rho E]^{\mathrm{T}}$ denotes the continuous flow variables and $\mathcal{J}$ represents a scalar objective function defined by pressure integrals either at the surface or at off-body locations in the domain. For sonic-boom and aircraft performance, this objective function frequently takes a form which combines off-body pressure sensors with metrics of aircraft performance (*i.e.* lift and drag) using a weighted sum,

$$\mathcal{J} = \frac{1}{p_\infty^2} \sum_{\Phi} \left[W_s \int (p - p_\star)^2 dS_\Phi \right] + W_L \left(1 - \frac{C_L}{C_{L\star}}\right)^2 + W_D \left(1 - \frac{C_D}{C_{D\star}}\right)^2 \tag{2}$$

where Φ is the sensor location and the subscript $(\star)$ indicates target values.

The flow variables satisfy the three-dimensional, steady-state Euler equations of a perfect gas within a feasible region of the design space Ω

$$\mathcal{F}(X, Q) = 0 \qquad \forall \; X \in \Omega \tag{3}$$

which implicitly defines $Q = f(X)$.

The optimization problem is solved through use of a gradient method. We use a discrete approach, where Eqs. 1- 3 are first discretized along with any constraints and then linearized to obtain the gradient $\mathrm{d}\mathcal{J}/\mathrm{d}X$. In the following sections, we first sketch a brief outline of the flow solution methodology for evaluating the objective function and then describe the gradient computation. Further detail of the theoretical approach is available in earlier work.[19,23,24]

[a]We assume X is a scalar to simplify notation of partial and total derivatives.

A. Objective Function Evaluation

At each step of the optimization procedure, we compute an approximation of the objective function $\mathcal{J}$ by solving the flow equations (3) on a multilevel Cartesian mesh with embedded boundaries. The mesh consists of regular Cartesian hexahedra everywhere, except for a layer of body-intersecting *cut-cells* adjacent to the boundaries as illustrated in Figure 2. The spatial discretization for the flow solver uses a cell-centered, second-order upwind finite-volume method with a weak imposition of boundary conditions. The resulting system of equations can be written in residual form

$$\mathbf{R}_H(\mathbf{Q}_H) = 0 \tag{4}$$

where H represents the average cell size and $\mathbf{Q} = [\bar{Q}_1, \bar{Q}_2, \ldots, \bar{Q}_N]^{\mathrm{T}}$ is the discrete solution vector of the cell-averaged values for all N cells of the mesh. Steady-state solutions are obtained using a five-stage Runge–Kutta scheme with local time stepping, multigrid and a domain decomposition scheme for parallel computing.[22,25,26]

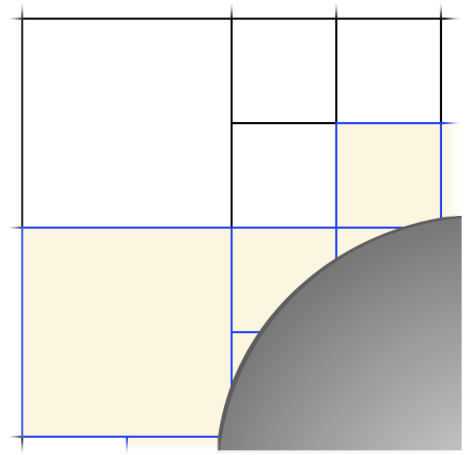

Figure 2. Cartesian mesh in two-dimensions with a cut-cell boundary.

B. Gradient Computation

To compute the discrete gradient, $\mathrm{d}\mathcal{J}_H/\mathrm{d}X$, we note that a variation in the design variables X influences both the computational mesh $\mathbf{M}$ and the flow solution $\mathbf{Q}$. We write the governing equations of the steady flow problem to explicitly include this dependency, resulting in a system of equations

$$\mathbf{R}(X, \mathbf{M}, \mathbf{Q}) = 0 \tag{5}$$

where the subscript H has been omitted to simplify the notation. The influence of shape design variables on the residuals is implicit through the computational mesh

$$\mathbf{M} = f[\mathbf{T}(X)] \tag{6}$$

where $\mathbf{T}$ denotes a triangulation of the wetted surface. The design variables that appear directly in Eq. 5 involve parameters that do not change the computational domain, such as the Mach number and angle of attack. The gradient is obtained by linearizing the objective function, $\mathcal{J}(X, \mathbf{M}, \mathbf{Q})$, and the residual equations, resulting in the following expression

$$\frac{\mathrm{d}\mathcal{J}}{\mathrm{d}X} = \frac{\partial \mathcal{J}}{\partial X} + \underbrace{\frac{\partial \mathcal{J}}{\partial \mathbf{M}} \frac{\partial \mathbf{M}}{\partial \mathbf{T}} \frac{\partial \mathbf{T}}{\partial X}}_{\mathrm{A}} - \psi^{\mathrm{T}} \left(\frac{\partial \mathbf{R}}{\partial X} + \underbrace{\frac{\partial \mathbf{R}}{\partial \mathbf{M}} \frac{\partial \mathbf{M}}{\partial \mathbf{T}} \frac{\partial \mathbf{T}}{\partial X}}_{\mathrm{B}} \right) \tag{7}$$

where the vector ψ represents the adjoint variables given by the following linear system

$$\left[\frac{\partial \mathbf{R}}{\partial \mathbf{Q}}\right]^{\mathrm{T}} \psi = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^{\mathrm{T}} \tag{8}$$

The solution algorithm for the adjoint equation uses the same time-marching scheme and parallel multigrid method of the flow solver.[23]

The most interesting part of the gradient computation is the evaluation of terms A and B in Eq. 7. Scanning these triple-product terms from left to right, the linearization with respect to $\mathbf{M}$ involves the flow solver, Eq. 5, while the middle term involves the mesh generator. In embedded-boundary Cartesian mesh methods, an infinitesimal perturbation of the boundary shape affects only the cut-cells. Unlike body-conforming approaches, there is no mesh perturbation scheme to smoothly map boundary shape changes into the interior of the volume mesh. Consequently, the mesh-sensitivities $\partial\mathbf{M}/\partial\mathbf{T}$, which contain the linearization of the Cartesian-face areas and centroids, volume centroids and the wall normals and areas with respect to the surface triangulation, are non-zero only in cut-cells. The last term of Eq. 7 is the inner product of the discrete adjoint solution with the residual sensitivities. This results in a fast and robust procedure for gradient computation, although, there is a reduction in the formal order of accuracy of the gradient.[24]

C. Design Framework and XDDM

The parallel Cartesian design framework[8] is responsible not only for control and execution of the flow and adjoint solvers, but also for coordinating construction of the geometry, the mesh and residual sensitivities for each design variable (A and B in Eq. 7). The shape sensitivities, $\partial \mathbf{T}/\partial X$, require access to and control of the geometric modeler. This is provided through the Extensible Design Description Markup protocol (XDDM).[8] XDDM's role is to provide modeler-neutral access to parameters and other information associated with geometry construction and queries. It can be used with CAD-based or non-CAD modelers. This XML-based protocol is used throughout the design framework to encapsulate meta-data (including design variables, analysis parameters, objective functions, constraints, etc.). It also provides access to a collection of standard services including computation of discrete surface sensitivities, symbolic function manipulation and model instantiation or regeneration.

While the framework is open to virtually any optimization package to solve the underlying minimization problem, interfaces for two optimizers are currently in place. An efficient quasi-Newton method with a backtracking line-search[24] is bundled with the framework. In addition, the popular SQP-based optimization package SNOPT(v.7)[27] is fully supported along with its elegant handling of constraints.

III. Numerical Investigations

Before shifting to shape optimization, we summarize two analysis studies useful both for illustrating aspects of the numerical approach and for providing insight into the meshing requirements of the boom-optimization study performed later.

A. Verification of Off-track Signal Convergence and Meshing Requirements

Our goal is to investigate designs with good performance off-track as well as on-track. This implies an understanding of the numerical scheme's behavior for sensors placed at various azimuths. To minimize computational expense, our meshing strategy aligns the Cartesian mesh with the characteristic wave propagation direction of the freestream by rotating the computational domain.[4,8] Since mesh cells are both aligned and stretched along the Mach angle of the oncoming supersonic flow, this meshing strategy clearly promotes propagation of signals from the geometry to sensors located on-track. The efficacy of this approach for sensors located at other azimuthal locations, however, is less obvious.

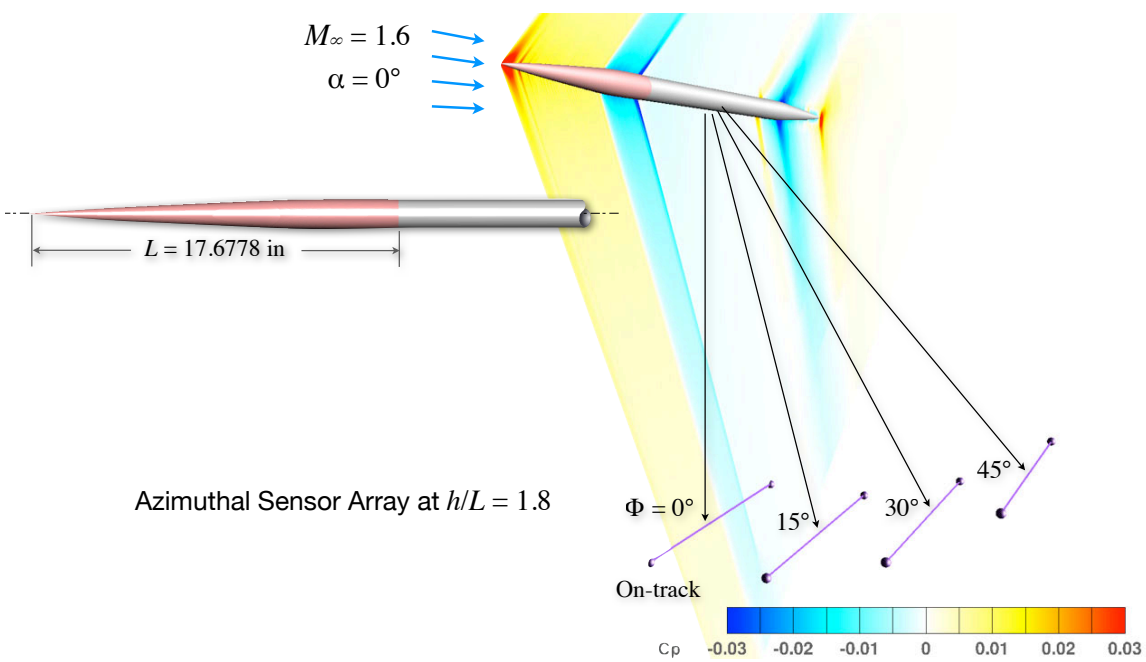

Figure 3. Axisymmetric geometry for verification study of signal propagation to off-track sensors. Pressure sensors are located at a distance $h/L = 1.8$ and at azimuthal locations $\Phi = \{0°, 15°, 30°, 45°\}$. $M_\infty = 1.6$, $\alpha = 0°$.

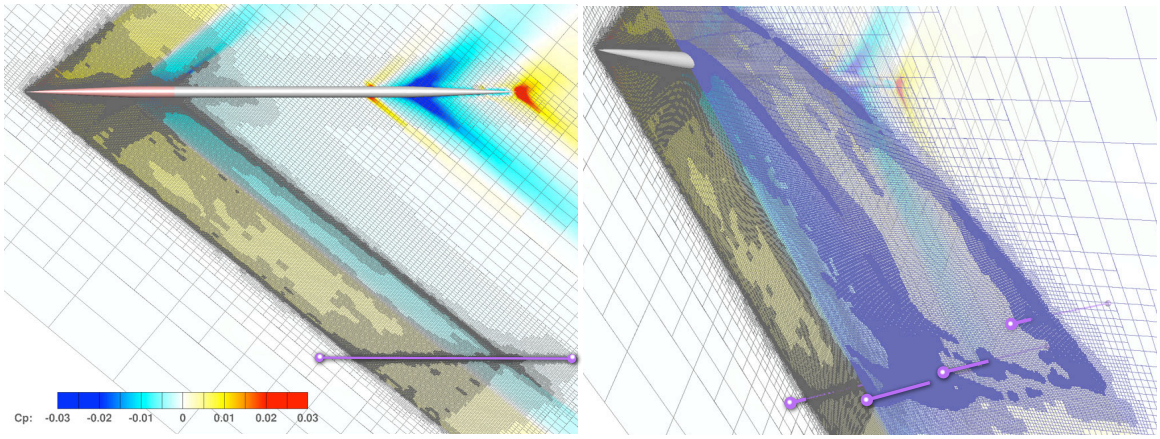

Figure 4. Symmetry plane (left) and perspective view (right) of cutting planes through the adaptively refined mesh used for the azimuthal performance study. Sensors are located at $\Phi = \{0°, 15°, 30°, 45°\}$ a distance $h/L = 1.8$ from the body. The final mesh contains ~13M cells. $M_\infty = 1.6$, $\alpha = 0°$.

To investigate numerical accuracy for such off-track signals, we perform a simple verification study using an axisymmetric model and examine signal propagation to an array of sensors located both on-track and off-track. Figure 3 shows the model and problem setup. The model is a body of revolution based upon the work of Darden, George and Seebass,[28,29] but incorporating lift relaxation near the aft portion of the shape.[b] The axisymmetric model is mounted on a cylindrical sting which tapers to a point downstream. Symmetry dictates that when the body is aligned with the freestream ($\alpha = 0°$), the flowfield will be axisymmetric and identical signals will be measured at sensors located at any azimuthal position. This symmetry provides an ideal basis for a verification exercise in which four sensors are placed in the simulation at azimuthal locations $\Phi = \{0°, 15°, 30°, 45°\}$, with the convention that on-track is at $\Phi = 0°$. The body length, L, is 17.678 inches and the sensors were all located a radial distance $h/L = 1.8$ away from the centerline.

Mesh adaptation is driven by a simplified form of the functional in Eq. 2 with the target values set to freestream pressure and without the lift and drag metrics

$$\mathcal{J} = \frac{1}{p_\infty^2} \sum_\Phi \int (p - p_\infty)^2 dS_\Phi \tag{9}$$

By comparing the off-track results with the on-track signal, we can quantify the scheme's performance at non-zero azimuths. Note that the output functional weighs errors along each of the four sensors equally. As a result, the adjoint-driven refinement scheme will attempt to equalize error contributions of each sensor by performing more adaptation near sensors that contribute higher error. Thus, the numerical experiment will also yield the resolution requirements to obtain equal error in off-track and on-track signals.

Figure 4 provides an overview of the mesh and flow solution at $M_\infty = 1.6$ and $\alpha = 0°$. The figure shows symmetry plane and perspective view with mesh cuts shaded by pressure coefficient. The initial mesh in this simulation contained ~10k cells. After 11 cycles of adaptive refinement this mesh had grown to nearly 13M cells (Fig. 4). The view in the symmetry plane shows the $\Phi = 0°$ sensor below the body and the expected pattern of mesh adaptation bounded in front by the cone-shock and in the rear by the last set of flow characteristics which impact the sensor. The flow over the forebody is essentially conical opening downstream, and the characteristics affecting the sensors are essentially conical opening upstream. In this supersonic flow, the adjoint-based refinement roughly concentrates cells within the intersections of these cones. Planes cutting in crossflow directions pick-out conic sections as well. The perspective view at the right of Fig. 4 has an inclined crossflow plane cut through the forebody flow capturing the parabolic expansion as it sweeps out toward the off-track sensor locations.

Figure 5 examines convergence of the functional (Eq. 9), its adjoint-based correction and error-estimate, as described in our earlier work.[3,30] Recall in this case that the functional is the sum of the pressures along the sensors. Taken together these plots give a strong indication of mesh convergence in the simulation, since changes in these pressures are vanishing as the mesh is refined. The frame at the left shows convergence of both the functional and its correction (*i.e.* the functional value if the mesh were uniformly refined). Both

[b] A wind tunnel model of this geometry was provided by Lockheed Martin Corp.

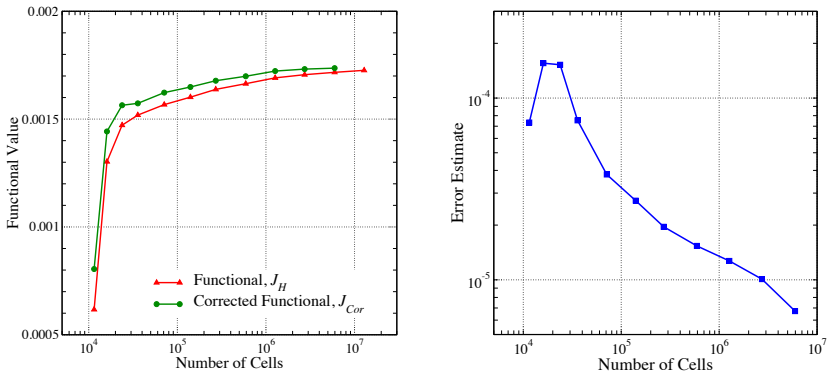

Figure 5. Mesh convergence of the functional (Eq. 9), its correction, and error-estimate.

asymptote with mesh refinement. Moreover, the fact that the correction systematically leads the functional is evidence that the problem is well behaved. The frame at the right shows convergence of the adjoint-based estimate of the error in $\mathcal{J}$. After about 30k cells, this estimate decreases steadily indicating that the mesh refinement is systematically eliminating discretization error in the computational domain.

The final plot in Figure 6 shows a comparison of the pressure signals ($\Delta p/p_\infty$) along all four sensors. Agreement between the sensors is very good and symbols have been added to distinguish among the various lines. Since the mesh was refined in response to estimates of the discretization error, the number of cells along each gives an indication of the meshing requirements to equalize error at the sensor locations. Cell counts along the sensors at 0°, 15°, 30° and 45° are 898, 908, 984 and 1209, respectively. This distribution is roughly cosine-like in the azimuth angle.

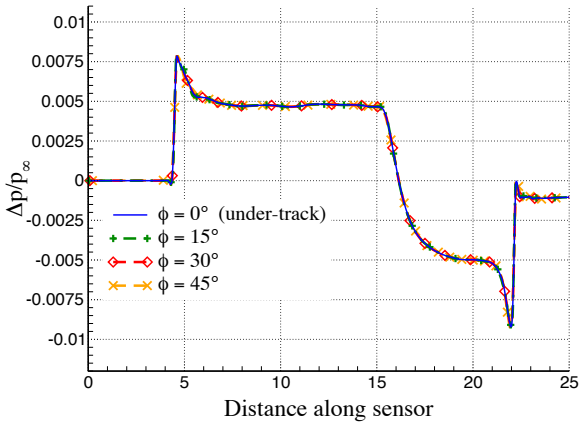

Φ	Num. Cells
0°	898
15°	908
30°	984
45°	1209

Figure 6. Comparison of computed normalized pressure signals ($\Delta p/p_\infty$) on azimuthal sensor array shown in Fig. 3. The table at the right gives the final number of cells along each sensor.

B. Experimental Comparison

The investigation in the preceding section gives insight into the meshing requirements for obtaining pressure signals at various azimuths of simple shapes. Before turning to shape optimization, we expand this understanding with a brief validation study of the on- and off-track signatures of a relevant low-boom model.

Figure 7 shows a perspective view of the Gulfstream Low Boom Model tested in the Langley Unitary Plan Wind Tunnel (UPWT)[31,32] under a cooperative agreement between Gulfstream and NASA through the Fundamental Aerodynamics Program's Supercruise Efficiency initiative. The model is a highly swept wing-body design which includes a vertical tail and a top-mounted blade shaped sting. The slender forebody incorporates a "Quiet Spike" nose and the model is designed for both low boom and low-drag.[32,33]

The simulations consider measurements taken at a free stream Mach number of 1.6 with pressure measurements taken both on- and off-track. The nominal fuselage length of 13.2 inches was used for establishing the h/L distances in the tests. Our comparison with data in this section considers two runs at $h/L = 1.2$, namely tunnel run 133 which was conducted at $\Phi = 0°$ and $\alpha = 0.256°$ and run 173 which was conducted at $\Phi = 48.2°$, $h/L = 1.32$ and $\alpha = 0.297°$.

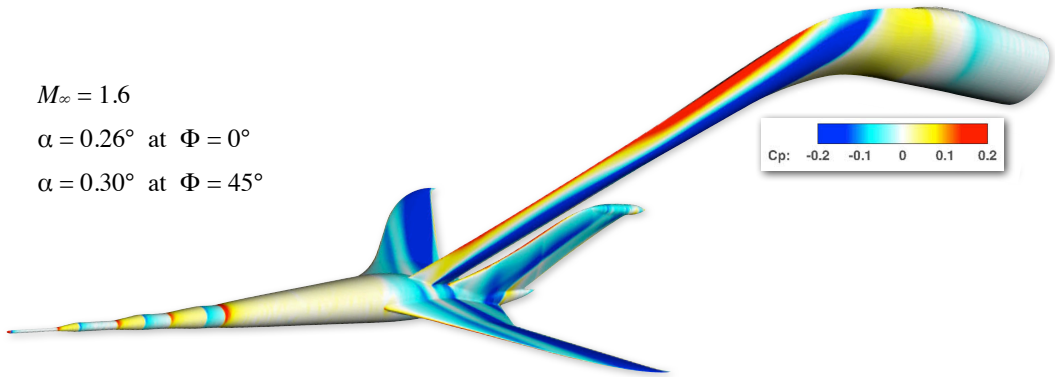

Figure 7. Perspective view of the Gulfstream Low Boom Model attached to the top-mounted sting. The surface triangulation used in the simulations had ~200k triangles. Test conditions were $M_\infty = 1.6$, $\alpha = 0.26°$ on-track and $\alpha = 0.3°$ at 45° azimuth. Surface mapped by pressure coefficient.

As in the earlier example, the Cartesian mesh was rotated to account for the Mach angle of the freestream and the incidence of the oncoming flow.[4,30] Separate adaptive-mesh simulations were performed for the two sensor positions. Since the angle-of-attack is quite similar between the two cases, a mesh rotation angle of 41.68° was used for both cases. The surface triangulation contained ~200k triangles, which is somewhat coarse for low-boom simulations requiring high fidelity. The background Cartesian mesh was stretched in the dominant propagation direction such that the cells are roughly square when viewed from the direction of the freestream flow. Both simulations began on coarse initial grids with ~4,100 cells in the computational domain. After 11 cycles of adaptive refinement, the final meshes in both cases contained roughly 14.7M cells.

As in the SEEB-like example presented earlier (§III.A), mesh adaptation was driven by an off-body functional measuring the pressure deviations along the sensor (Eq. 9). Figure 8 contains a view of the mesh in the symmetry plane and contours of pressure coefficient in the discrete solution for the on-track case $\Phi = 0°$. The series of compressions and expansions emanating from the nose are clearly visible in the C_p contours and have received significant attention by the error-driven refinement scheme. As in the axisymmetric SEEB case, the method has focused meshing between conical regions of the domain anchored in front by the tip of the geometry and in the rear by the aft end of the sensor. The mesh for the off-track ($\Phi = 45°$) case is similar, modulo the change in sensor position.

Pressure signals from the simulations are displayed in Figure 9. On the left, the blue line shows the computed signature ($\Delta p/p$) overplotted with data from tunnel run 133 taken on-track at $\Phi = 0°$ and compared with numerical results from inviscid simulations performed with both AIRPLANE[7] and FUN3D.[6,16] On the right, the off-track signal ($\Phi = 45°$) is compared with tunnel data from run 173. Off-track results for

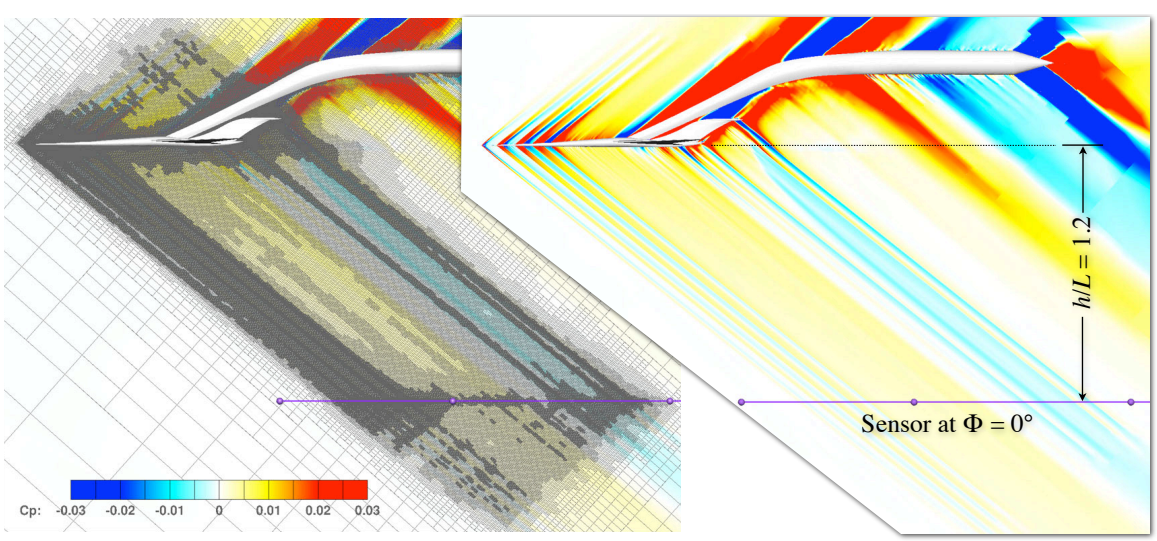

Figure 8. Final adapted mesh and pressure coefficient discrete solution for low-boom model with sensor located on-track at $\Phi = 0°$. $M_\infty = 1.6$, $\alpha = 0.26°$. The final mesh has ~14.7M cells after 11 levels of adaptive refinement.

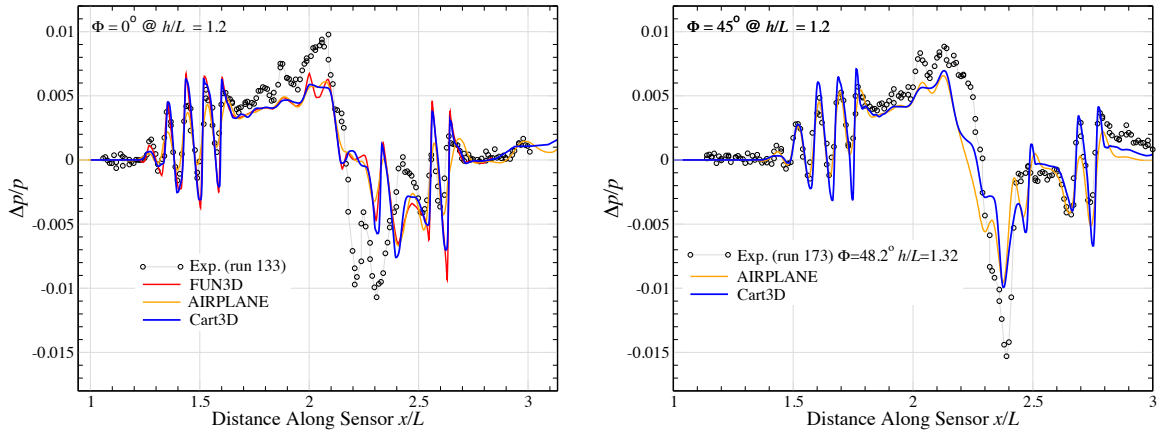

Figure 9. On-track and off-track pressure signals at $h/L = 1.2$ with comparison to experiment and other simulations. On-track: $\Phi = 0°$, $M_\infty = 1.6$, $\alpha = 0.26°$. Off-track: $\Phi = 45°$, $M_\infty = 1.6$, $\alpha = 0.30°$. Experimental data from UPWT tests.[31,32]

this case were not available from FUN3D, so only the AIRPLANE results are shown. Comparisons between experiment and simulation show good agreement, matching the experiment especially well over the forebody and the aft-portion of the signature. Correlation among the numerical simulations is particularly strong. Especially noteworthy is the close agreement with the FUN3D results. Both methods use adjoint-driven mesh adaptation with a similar choice of output functional, and both can document their approach to mesh convergence. Agreement with the AIRPLANE results is excellent for both sensor locations showing the efficacy of the meshing practices used in those simulations. All the simulations slightly under-predict the pressure rise associated with the main wing shock and differ through the expansion before the wing trailing edge. Similar trends are evident in a number of the numerical simulations and are being examined in other works.[6,7,31] The under-prediction of the maximum over-pressure suggests a possible mismatch in flow incidence between the simulations and experiment. Through the wing expansion, the simulations show a deviation at both sensor locations near x/L of ~2.2. Aside from this particular feature, agreement is quite good both on-track and 45° off-track, including the locations, sharpness, and magnitude of the signal at the sensors.

IV. Shape Optimization for Low Boom and Performance

The shape optimization example considers a vehicle based upon the wind tunnel model of the preceding section but with additional complexity. As shown in Figure 10, the wind-tunnel geometry was enhanced through the addition of pylons, nacelles and a horizontal tail. The vehicle has a semi-span reference area of 300 ft^2 and a length of 134.6 ft measured from the spike tip to the close of the tail-pod. Assuming a weight of 33,000 lbs and steady level flight at 45,000 ft in the standard atmosphere gives a design cruise lift coefficient of 0.0996.

To determine cruise angle-of-attack, we performed an adaptive mesh simulation taking lift coefficient as the output of merit and allowing the solver to adjust the incidence angle as the simulation advanced and the mesh was refined. This study yielded a mesh converged cruise α of 0.612°. This is roughly twice the incidence angle of that used in the preceding section with the tunnel model. At this lift coefficient, the baseline vehicle has a drag of $C_D = 0.0123$ giving an inviscid L/D ratio of ~8.1. In the shape optimization which follows, we seek to maintain these lift and drag values while reducing the boom-noise at the ground measured by sensors located both on-track and 15° off-track.

A. Parametric Model

For modeling support, we chose the analytic "Rapid Aerospace Geometry Engine" (RAGE). RAGE is lightweight, flexible, and fully parametric.[34–36] It is Java-based and provides high-quality loftings from a large catalog of geometric primitives. Since it is platform independent and fast, the modeler can run on the same hardware as the optimization through the threading options built into the design framework.[8]

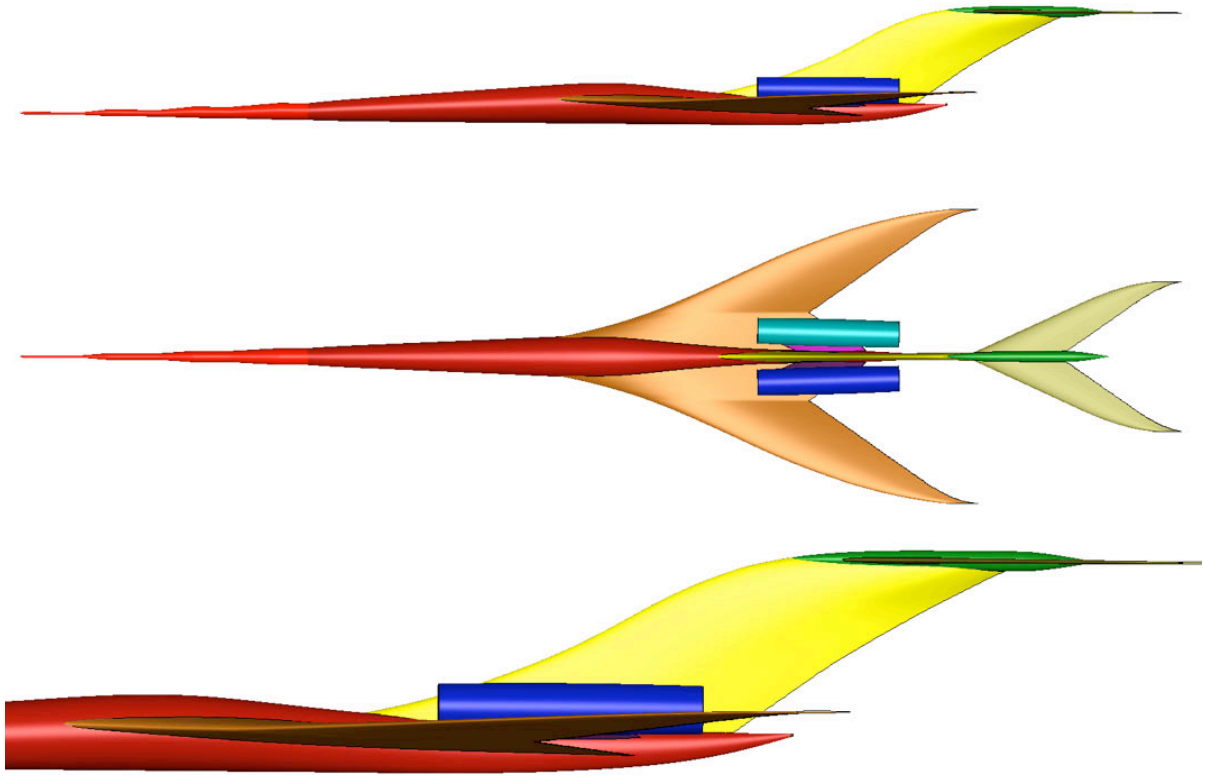

Figure 10. Baseline aircraft geometry used in shape optimization constructed using RAGE modeler and including empennage, pylons and nacelles. The surface triangulation contains ∼ 1M triangles.

Figure 10 shows views of the baseline model generated by RAGE with 12 watertight components. The fuselage is described by 24 circular or elliptical cross-sections lofted either linearly or via Akima splines.[34] The main wing planform is defined by 14 spanwise stations on each semi-span, while 7 are used for the the horizontal tail. In both cases, the semi-spans are lofted with Akima splines. The airfoils at each station use the Bernstein polynomial approach of Kulfan[37] with 5–11 parameters controlling the upper and lower surfaces of each section. The highly swept vertical tail uses a similar description, but with 8 symmetric airfoil sections and is capped with a circular cross-section tail-pod. The nacelles and pylons used construction similar to the wings and empennage and were included as "non-deformable" geometry with only their positions allowed to change parametrically.

In total, approximately 1150 parameters were used to describe the model and were available for control through the design framework. Access to these variables was provided through the framework's XDDM protocol, and a straightforward XDDM-to-RAGE wrapper was used to translate the XDDM meta data to native RAGE input files used for model regeneration. Shape sensitivities for the design variables were provided by the framework's finite-difference sensitivity service.

B. Forebody Design and Adjoint Analysis

Our design investigations open with consideration of a sub-problem narrowly focusing on inverse design of the aircraft forebody. This problem provides a more tractable starting point than optimization of the complete geometry and gives insight into the information offered by the adjoint solution in design. Reviewing the pressure signals shown in Fig. 9, the impact of the segmented nose on the pressure signal is obvious. Both on-track and off-track signatures commence with a series of four rapid oscillations as the flow negotiates the discrete compressions and expansions of the segmented nose geometry. These features are clearly discernible due to the hyperbolic nature of supersonic flow, and the fact that small disturbances travel with similar wavespeeds. This implies that in the near-field – before substantial wave coalescence – there exists a degree of stratification along the signal. Oscillations at the front of the signal are traceable to the forebody while those at the back come from the aft of the aircraft. While seemingly a statement of the obvious, there is an important subtlety. This stratification of the near-field permits a detailed accounting linking specific portions of the aircraft to different regions of the pressure sensor. Loosely speaking, this accounting is what

is being offered by the adjoint, and this information can aid directly in selection of design variables for use in the optimization.

To illustrate, consider the inverse design functional

$$\mathcal{J} = \frac{1}{p_\infty^2} \int (p - p_{\text{target}})^2 dS \tag{10}$$

Optimization of this objective function seeks to minimize the differences between the actual signal produced and that of a prescribed target. Figure 11 contains a numerical experiment illustrating our discussion. In this example, the signal from the baseline parametric geometry was extracted at the design flight conditions ($M_\infty = 1.6$, $\alpha = 0.612°$) on-track at a distance of $h/L = 2$. We then specified a smooth, monotonic target signal over the forebody portion of the pressure signature ("active target" in Fig. 11). The objective function in Eq. 10 measures the deviation of this signal (black) from that of the target (blue).

Adjoint solutions show which regions of the flow influence the objective function. This is evident in the contours of the density adjoint for the baseline geometry shown at the left of Fig. 11. The parallel isoclines follow flowfield characteristics through the domain and trace differences between the functional and target back to specific portions of the surface geometry.

Put more precisely, solution of the adjoint equation (Eq. 8) shows the influence of residual perturbations on the objective, $\mathcal{J}$. When the gradient is computed for an off-body functional on a Cartesian mesh, the first two terms on the right side of Eq. 7 are zero because the objective function is not co-located with the design variables. Shape design variables affect only the surface thus the mesh sensitivities ($\partial\mathbf{M}/\partial\mathbf{T}$) are non-zero only in the cut-cells. The entire gradient expression simplifies to an inner product of the adjoint variables with the residual sensitivities at the body in the third term on the right of Eq. 7. Therefore, for fixed values of residual sensitivities, the adjoint variables give a direct measure of the sensitivity of the objective to the shape.

We can make an additional observation for inverse-design formulations, where pressure, lift, or drag are being driven to a specific target. The right side of the adjoint contains a term of the form $(p - p_\star)$ which is due to the linearization of the square in the objective function (see Eqs. 2, 9, or 10). Therefore as we approach optimality, *i.e.* $(p \to p_\star)$, we expect a weakening of the adjoint field.

The right side of Fig. 11 provides an excellent illustration and summarizes the situation after shape optimization. The shape was modified using 20 parameters controlling the vehicle forebody as design variables using the objective function in Eq. 10. After 16 design iterations, the optimizer had reduced the value of the objective by about 1.5 orders of magnitude. The resulting design is shown at the top of the right side of Fig. 11. Contours of the density adjoint are displayed along with the geometry to graphically illustrate the optimizer's progress. Isoclines on the right use the same color map as on the left, with white showing

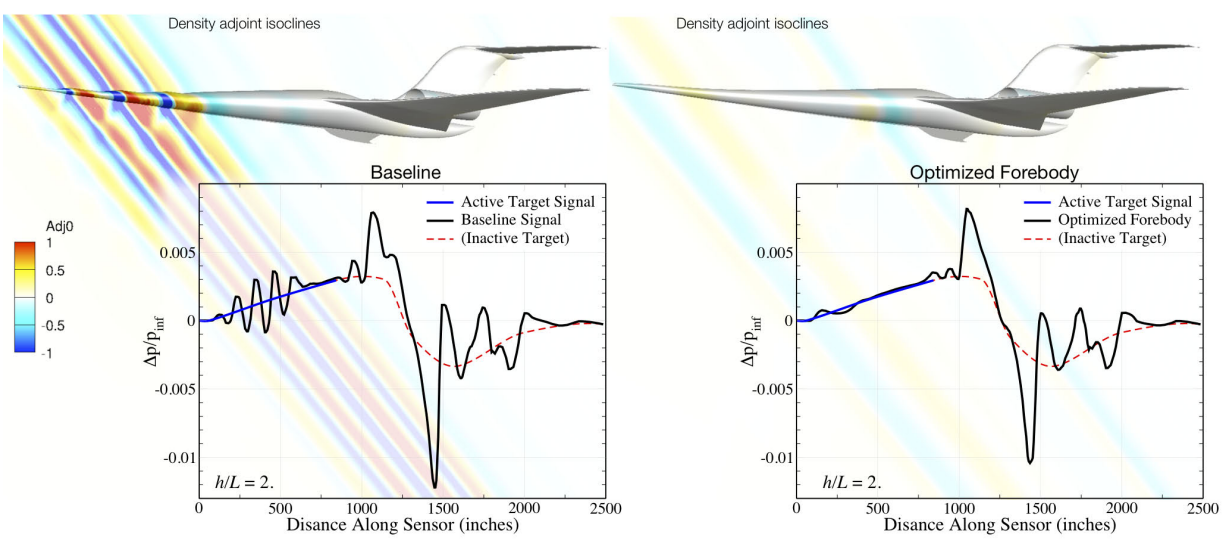

Figure 11. Forebody inverse-design example. Isoclines of the density adjoint are shown for an inverse design objective based on the difference between the forebody signal and the active target, Eq. 10. Pressure signals extracted undertrack at a distance of $h/L = 2$.

zero contribution to the functional. The line plot of the signal gives a quantitative view of the optimizer's success in achieving the target. We note the pronounced weakening of the adjoint field giving both a view of the optimizer's progress and a direct indication of regions where more work remains.

C. Baseline and Targets

The starting point for optimization is the RAGE model shown in Figure 10. Our optimization problem seeks to improve the ground signal on-track ($\Phi = 0°$) and off-track ($\Phi = 15°$), while maintaining lift and drag at fixed incidence angle. As described earlier, aerodynamic performance was established through error-driven adaptive simulations at fixed lift. To quantify the initial boom-levels, we ran adaptive simulations extracting pressure signatures at several distances on-track $h/L = \{1.2, 2.0, 3.0\}$. These signals were then propagated to sea level from the 45,000 foot cruise altitude using an atmospheric propagation code.[14] This investigation showed that signal extraction from $h/L = 2.0$ and $h/L = 3.0$ produced essentially the same ground signals and loudness metrics. Since the computational domain with the sensor at $h/L = 2.0$ was smaller, signal extraction from this position was used in optimization.

The computational domain used in optimization was meshed using a fixed grid containing pre-specified regions of mesh refinement.[11] Grid densities were guided by meshing requirements developed using a number of adaptive mesh simulations at the design conditions using the methodology outlined in the verification and validation studies presented earlier (§III). Figure 12 shows a view of the final mesh, the sensor location, and isobars in the discrete solution at the design conditions of $M_\infty = 1.6$, $\alpha = 0.612°$. This mesh has a total of 9.3 M cells with an aspect ratio of roughly 3:1. The expanded scale of C_p in this figure emphasizes the fact that the signals near the sensor are quite weak – even for the baseline model.

Figure 13 shows pressure signals on-track and off-track for the baseline configuration extracted at $h/L = 2.0$ for azimuths of $\Phi = 0°$ and $\Phi = 15°$. Also shown are the target signatures used to drive the shape optimization. These targets are approximately sinusoidal and were generated by simply splining selected points on the baseline signals and then evaluating their noise levels after propagation to the ground. More sophisticated methods exist for producing such near-field targets; however in the current work we simply used this "cut and try" approach. Since they were naively generated, there is no expectation of either attainability or optimality.

Figure 14 shows the resulting ground signatures both on- and off-track. These were generated taking the baseline and target signals in Fig. 13 and using the sBOOM augmented Burgers' equation propagation code of Rallabhandi.[14] Propagation through the standard atmosphere was computed from the cruise altitude of 45,000 ft to sea level for both azimuths. The graphs show overpressure with time (in milliseconds) on the abscissa. The peak-to-peak time is about 55 ms with a period of about 125 ms (8 hz). Legends in both plots include noise levels (in PLdB) for each of the signals. The baseline signal is 86.3 PLdB on-track, and 84 PLdB at 15° off-track. The noise ratings on the propagated target signals are 71.1 PLdB and 71.5 PLdB on-track and off-track, respectively. This is about a 15 decibel improvement over the baseline.

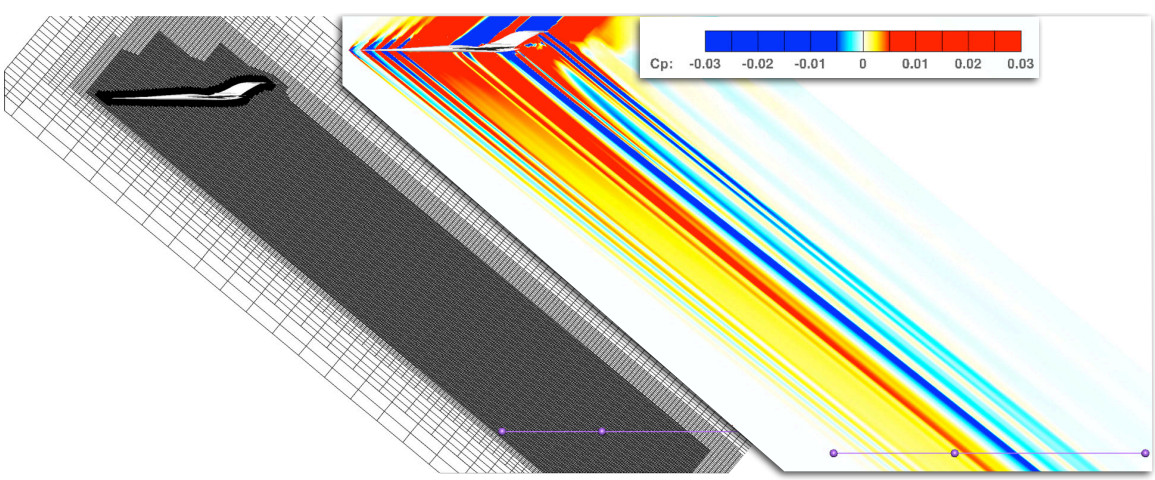

Figure 12. (left) Multilevel Cartesian mesh used for shape optimization containing 9.3 M cells with a pressure-sensor located at $h/L = 2.0$. (right) Isobars in the discrete solution of the baseline model shown in Fig 10. $M_\infty = 1.6$, $\alpha = 0.612°$.

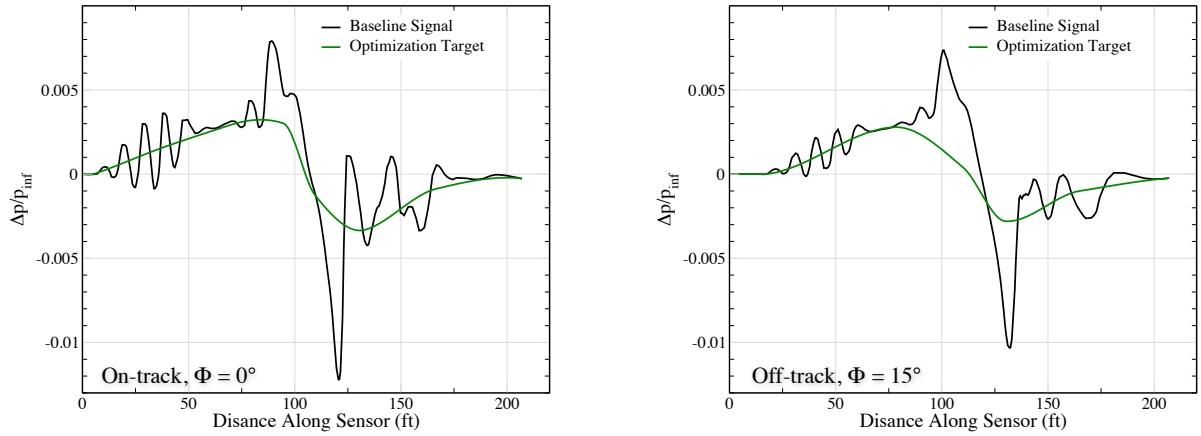

Figure 13. Pressure signals on-track and off-track for the baseline geometry at $\Phi = 0°$ and $\Phi = 15°$ at $h/L = 2.0$. Also shown are the signals used as optimization targets for each azimuth. $M_\infty = 1.6$, $\alpha = 0.612°$

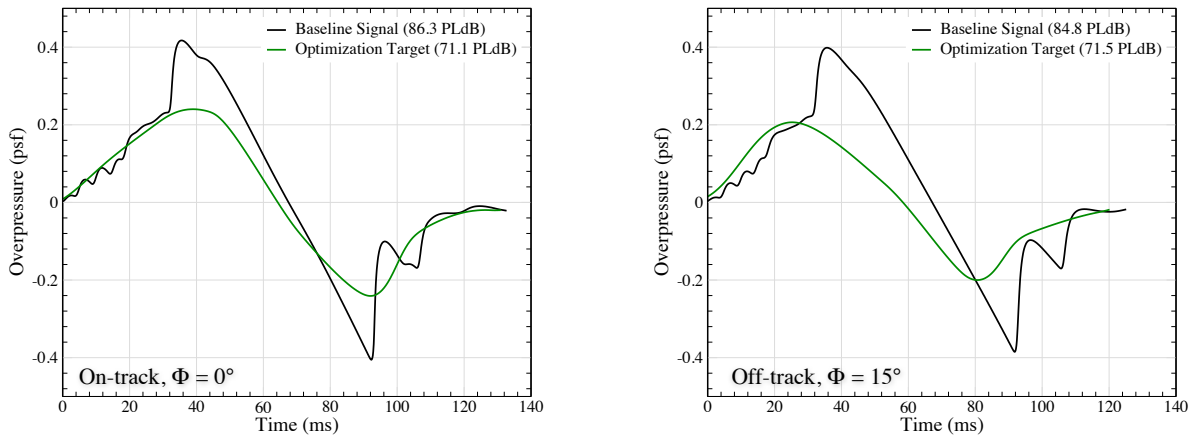

Figure 14. Ground signals on-track and off-track for the baseline geometry and optimization targets after propagation to the sea-level from a cruise altitude of 45,000 ft using sBOOM.[14]

The final objective function used to drive the optimization took the form of Eq. 2. Each sensor was weighted equally and the weights on the lift and drag penalty terms were both twice the sensor weight. With the targets and functional defined, optimization began with an evaluation of the baseline geometry as a starting point. Figure 15 shows the starting geometry, colored with isoclines of the density adjoint. The figure displays upper, lower and rear 3/4 views. White indicates regions of the geometry which do not contribute to the functional while the red and blue regions show where work is required. Perhaps unexpectedly, the mapping shows that residuals on the upper wing surface can affect the objective function even though it is dominated by the signal below the wing. Here the adjoint reminds us that the upper wing surface determines the flow over the nacelles and interacts with the empennage resulting in the design sensitivities shown.

D. Optimization Results

Optimization of the full geometry used design variables governing the shape of the fuselage, wing, and vertical and horizontal tails. Roughly 120 variables were dedicated to controlling the Bernstein polynomials on the main wing, while approximately 40 controlled the fuselage cross-section and lofting. Another 20 design variables were dedicated to control the thickness distribution of the vertical tail and the thickness and twist distribution of the horizontal tail. In total, 180 design variables participated in the optimization, although not all were active all the time. Numerous smaller optimization problems were performed on subproblems with ~20–40 design variables to help understand the signal's response to various shape parameters. Approximately 50 design iterations were used with the largest set of design variables. Optimization reduced the cost function by a factor of 16 before the design stopped progressing.

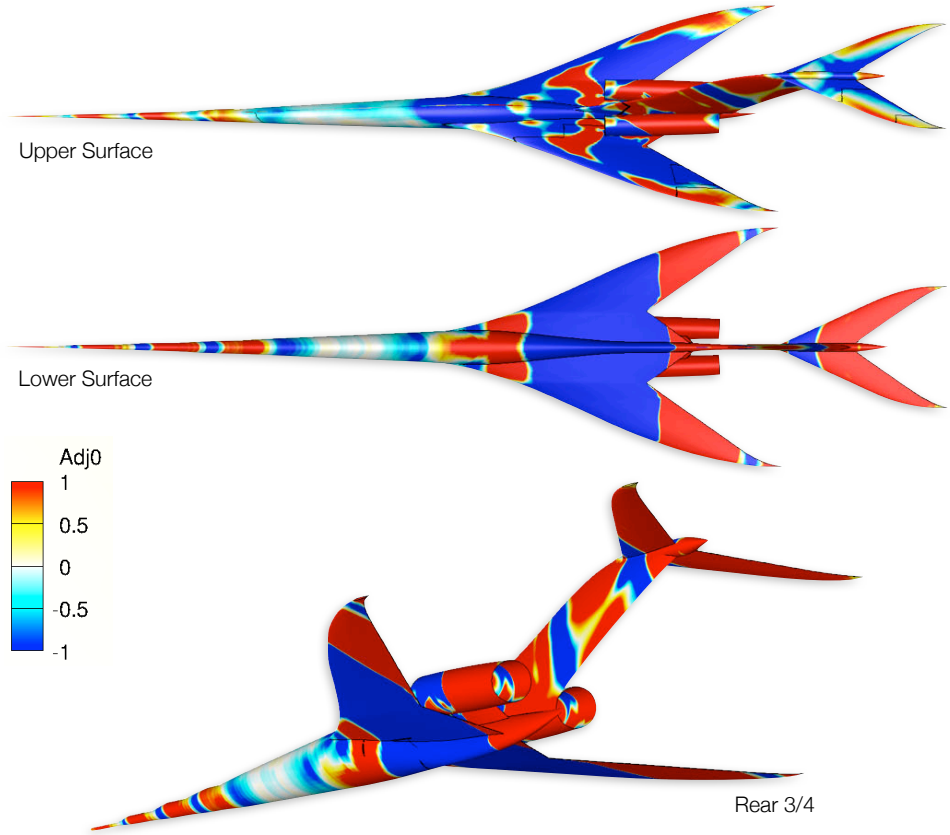

Figure 15. Views of baseline geometry mapped with isoclines of the density adjoint at cruise conditions for the composite functional given in Eq. 2.

Figure 16 shows the pressure signatures at $h/L = 2.0$ after optimization both on-track (left) and 15° off-track (right). For comparison, the optimization targets and the original signals of the baseline model are included on both plots. In both cases, the signatures have made substantial progress toward the target. The peak-to-peak values of the baseline signal are approximately $0.2p_\infty$ at both locations. The final design has reduced this peak-to-peak by approximately 60% on-track and 50% off-track. At both azimuths, the design was particularly successful in matching the front and aft-portions of the target signal. While the forebody was discussed earlier, the aft signals were particularly responsive to tuning the thickness of the airfoil sections on the vertical tail and the twist of the horizontal tail along with subtle reshaping of the aft fuselage.

Figure 17 shows the signatures at $\Phi = 0°$ and $\Phi = 15°$ after propagation from the cruise altitude of 45,000 ft to sea-level through the standard atmosphere using the sBOOM augmented Burgers' equation solver.[14] For reference, these graphs also include both the propagated signatures of the baseline and target.

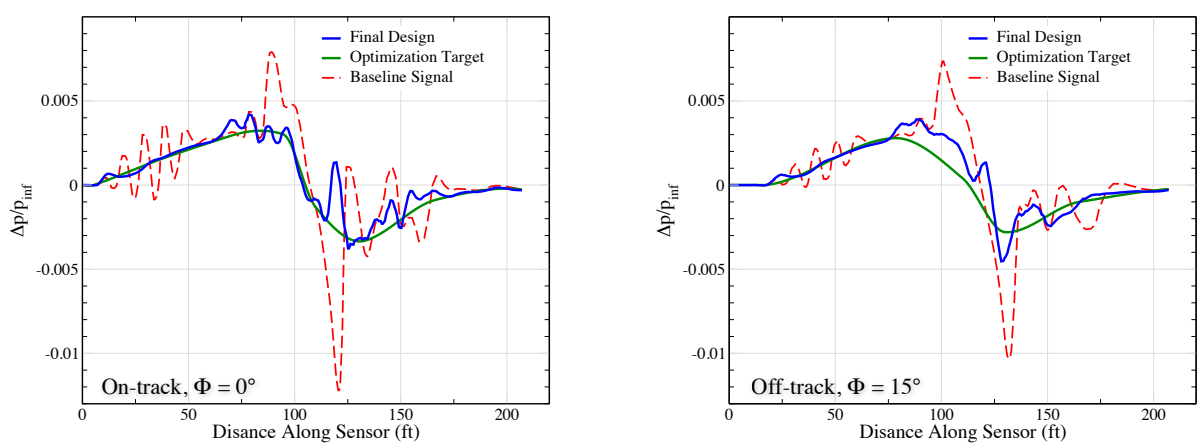

Figure 16. Pressure signals on-track and off-track for the baseline geometry at $\Phi = 0°$ and $\Phi = 15°$ at $h/L = 2.0$. Also shown are the signals used as optimization targets for each azimuth. $M_\infty = 1.6$, $\alpha = 0.612°$.

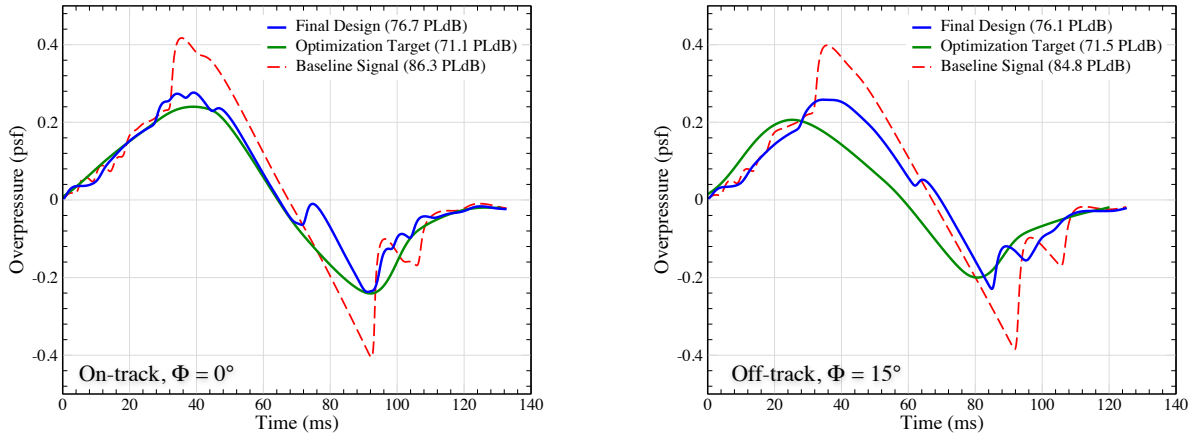

Figure 17. Ground signals on-track and off-track for the baseline geometry and optimization targets after propagation to the sea-level from a cruise altitude of 45,000 ft using sBOOM.[14]

As before, the graphs show overpressure with time (in milliseconds) on the abscissa. At both azimuths, the peak-to-peak time is about 55 ms with a period of about 125 ms (8 hz). Legends in both plots include perceived noise levels for each of the signals. On-track, the baseline signal is rated at 86.3 PLdB. The final design achieved a 9.6 dB improvement to a noise rating of 76.7 PLdB. At 15° off-track, the final design achieves 76.1 PLdB as compared to the baseline aircraft's rating of 84.8 PLdB. Using the acoustic "A-weighting" filter, the final design's acoustic noise ratings are 61.3 dB(A) on-track and 60.1 dB(A) 15° off track.

Figure 18 contains isobars in the discrete solution of the final model in the symmetry plane from the aircraft to the sensor location at $h/L = 2.0$. The colormap for the C_p distribution is same as that shown in Fig. 12 for the baseline design. Comparing these two figures we see a marked weakening of the waves arriving at the sensor. The strong over- and under-pressures of the baseline vehicle have been completely eliminated and the field is noticeably smoother.

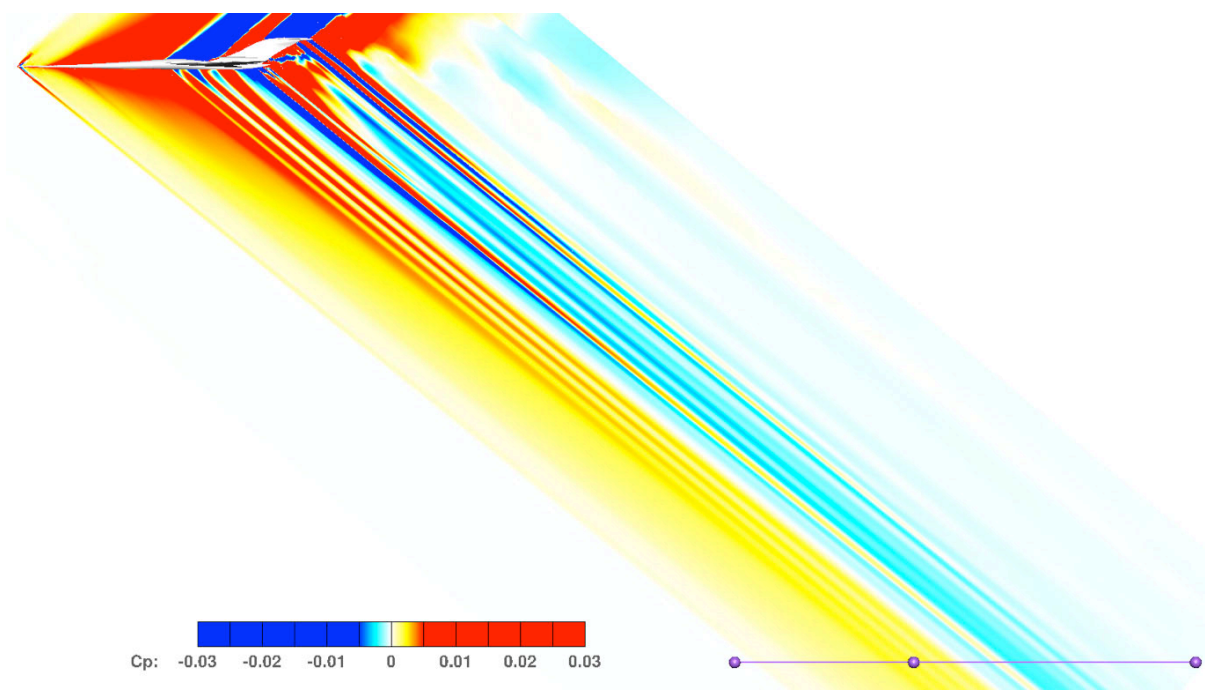

Figure 18. Isobars in the discrete solution of the final model after design, sensor at $h/L = 2.0$ shown. Colormap for C_p distribution is same as in Fig. 12. $M_\infty = 1.6$, $\alpha = 0.612°$.

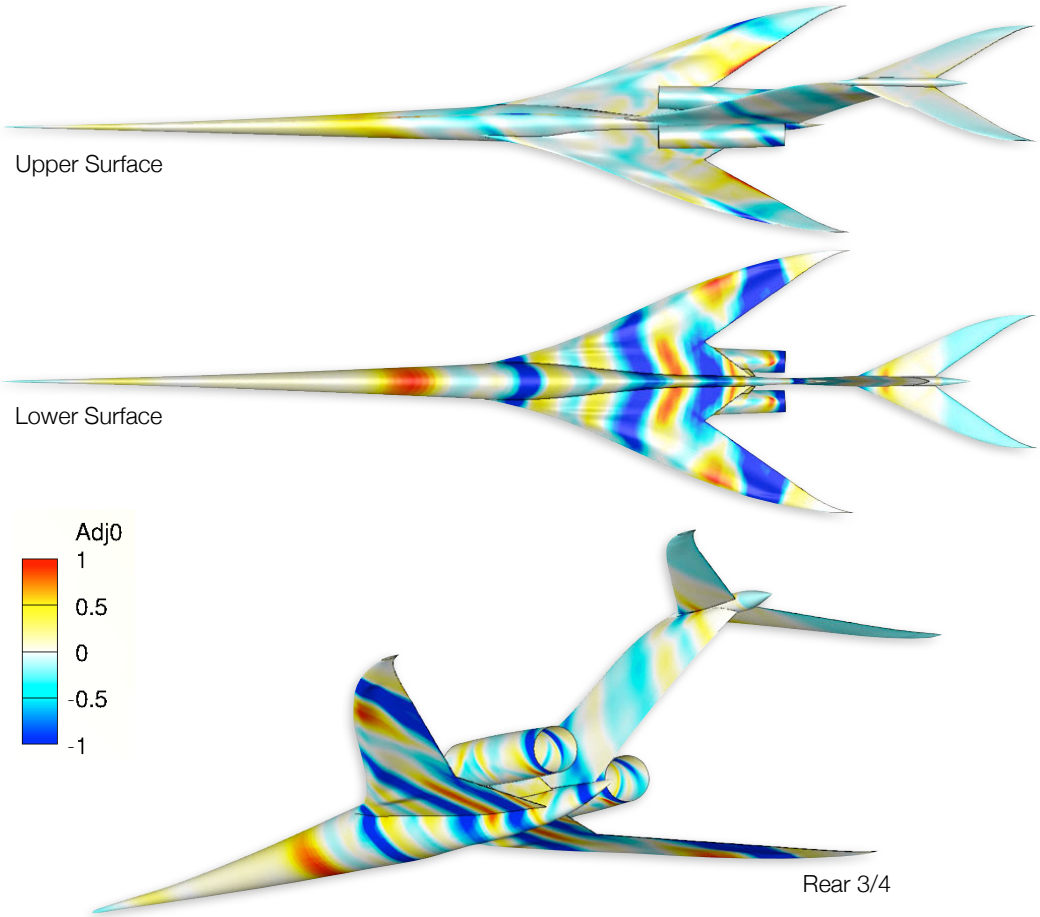

Figure 19. Views of final design mapped with isoclines of the density adjoint at cruise conditions for the composite functional given in the text. Colormap identical to that in Fig. 15.

Lastly, Figure 19 shows views of the design mapped with isoclines of the density adjoint. The colormap is identical to that used in Fig. 15 to facilitate direct comparison. Optimization has reduced the magnitude of the composite cost function (Eq. 2) by a factor of 16 from the baseline design, and comparing the adjoints before and after reshaping, this improvement is readily apparent. Only the underside of the wing shows regions that still contribute meaningfully to the cost function.

The aerodynamic performance of the final design was verified using adjoint-based adaptive meshing to obtain mesh converged values of the primary aerodynamic coefficients. At the design conditions of $M_\infty = 1.6$ and $\alpha = 0.612°$, the aircraft has $C_L = 0.0960$ and $C_D = 0.0119$. This is a loss of ~3.5% in lift with a drag improvement of 4 counts. The resulting L/D value is 8.1, which is the same as the baseline configuration at these flight conditions. Additional simulations showed that if warranted, increasing the incidence angle by 0.065° restores the lost lift while maintaining a slight (3 count) drag advantage over the baseline vehicle. At this flight condition, the inviscid L/D of the final design is 8.3.

V. Summary

This work explored the analysis and design of low sonic-boom aircraft through the use of an inviscid, embedded-boundary Cartesian mesh method. Adjoint error estimation and adaptive meshing were used in the analysis portion of this study to determine resolution requirements of the computational domain, and the design examples were conducted using adjoint-based gradients in conjunction with a parallel optimization framework. The work demonstrates rapid shape optimization of highly parametric, complex models to simultaneously meet both noise and performance targets.

As a prelude to the optimization work, two analysis experiments were conducted to provide V&V for specific aspects of the optimization problem. A verification study using an axisymmetric body and an array

of on- and off-track pressure sensors was conducted to demonstrate the ability of the Cartesian method to accurately predict pressure signatures at various azimuth angles. This study was conducted at a distance of 1.8 body lengths and provided mesh convergence and error information. The mesh converged result (~13 M cells) showed that the signals collapsed to plotting accuracy with no slowdown in the rate of error convergence. This analysis was instructive for providing relative meshing requirements necessary for controlling discretization error at off-track sensor locations. A validation study showed comparison of on-track and off-track wind-tunnel data for a low-boom model similar to that used in the optimization example. This study showed close agreement between computed pressure signals and the wind tunnel experiments.

The design problem performed shape optimization of a 33,000 lb aircraft traveling at Mach 1.6 at an altitude of 45,000 ft to meet both on-track ($\Phi = 0°$) and off-track ($\Phi = 15°$) noise targets while maintaining lift and drag performance. Low-noise target pressure distributions were prescribed in the computational domain at a distance of two body-lengths away from the model. A fully parametric low-boom model was constructed using the RAGE modeler, and a total of approximately 180 parameters were used as design variables. The design framework used an inverse design formulation to re-shape the model to achieve ground noise levels of approximately 76.7 PLdB which was a 9.6 dB improvement over the baseline design while maintaining the lift and drag performance of the original aircraft.

Our discussion focused on the role of the discrete adjoint solution in the design process. We showed that for common functional definitions, the adjoint solution provides a detailed bookkeeping which highlights specific regions of the aircraft that contribute to the cost function being minimized through optimization. For appropriate objective functions, examination of the adjoint solution highlights portions of the geometry with significant influence on the objective function, and therefore offers insight for model construction, parameterization and design variable selection.

Acknowledgments

This work was supported by the NASA Fundamental Aeronautics Program Supersonics project and by NASA Ames Research Center Contract NNA10DF26C. We gratefully acknowledge the aid and support of David L. Rodriguez and Dev Rajnarayan at DesktopAeronautics for their work and responsiveness in providing both the flexible RAGE aircraft model used in the optimization example as well as the XDDM interface to the RAGE modeler. We wish to thank Sriram K. Rallabhandi at the National Institute of Aerospace in Hampton, VA, for graciously providing access to the sBOOM signal propagation code which was used to generate ground signals in this work. We are also grateful to Mathias Wintzer for many useful discussions and sharing his insights into supersonic flight and boom prediction.

References

[1]Choi, S., Alonso, J. J., and Van der Weide, E., "Numerical and Mesh Resolution Requirements for Accurate Sonic Boom Prediction of Complete Aircraft Configurations," *AIAA Paper 2004-1060*, January 2004.

[2]Ozcer, I. A. and Kandil, O., "FUN3D/OptiGRID Coupling for Unstructrred Grid Adaptation for Sonic Boom Problems," *AIAA Paper 2008-0061*, January 2008.

[3]Nemec, M., Aftosmis, M. J., and Wintzer, M., "Adjoint-based adaptive mesh refinement for complex geometries," *AIAA Paper 2008-0725*, January 2008.

[4]Wintzer, M., Nemec, M., and Aftosmis, M. J., "Adjoint-Based Adaptive Mesh Renement for Sonic Boom Prediction," *AIAA Paper 2008-6593*, August 2008.

[5]Howe, D. C., "Hybrid Cart3D/OVERFLOW Near-Field Analysis of a Low Boom Configuration with Wind Tunnel Comparisons," *AIAA 29th Applied Aerodynamics Conference (in publication)*, June 2011.

[6]Park, M. A., "Low Boom Configuration Analysis with FUN3D Adjoint Simulation Framework," *AIAA 29th Applied Aerodynamics Conference (in publication)*, June 2011.

[7]Cliff, S. E., Elmiligui, A. A., Campbell, R., and Thomas, S. D., "Evaluation of Refined Tetrahedral Meshes with Projected, Stretched, and Sheared Prism Layers for Sonic Boom Analysis," *AIAA 29th Applied Aerodynamics Conference (in publication)*, June 2011.

[8]Nemec, M. and Aftosmis, M. J., "Parallel Adjoint Framework for Aerodynamic Shape Optimization of Component-Based Geometry," *AIAA Paper 2011-1249*, January 2011.

[9]Carter, M. B., Campbell, R., and Nayani, S., "USM3D Analysis of a Low Boom Configuration," *AIAA 29th Applied Aerodynamics Conference (in publication)*, June 2011.

[10]Alauzet, F. and Loseille, A., "High-Order Sonic Boom Modeling Based on Adaptive Methods," *Journal of Computational Physics*, Vol. 229, No. 3, 2010, pp. 561–593.

[11]Cliff, S. E., Thomas, S. D., McMullen, M., Melton, J. E., and Durston, D. A., "Assessment of Unstructured Euler Methods

for Sonic Boom Pressure Signatures Using Grid Refinement and Domain Rotation Methods," NASA/TM 2008-214568, NASA, September 2008.

[12]Thomas, C., "Extrapolation of Sonic Boom Pressure Signatures by the Waveform Parameter Method," NASA TN-D-6832, June 1972.

[13]Durston, D. A., "A Preliminary Evaluation of Sonic Boom Extrapolation and Loudness Calculation Methods," *High-Speed Research: Sonic Boom*, NASA CP 10133, NASA Ames Research Center, May 12-14 1993, pp. 301–323.

[14]Rallabhandi, S. K., "Advanced Sonic Boom Prediction Using Augmented Burger's Equation," *AIAA Paper 2011-1278*, January 2011.

[15]Campbell, R., Carter, M. B., Deere, M., and Waithe, K. A., "Efficient Unstructured Grid Adaptation Methods for Sonic Boom Prediction," *AIAA Paper 2008-7327*, June 2008.

[16]Jones, W. T., Nielsen, E. J., and Park, M. A., "Validation of 3D Adjoint Based Error Estimation and Mesh Adaptation for Sonic Boom Prediction," AIAA Paper 2006–1150, Jan 2006.

[17]Kandil, O. and Ozcer, I. A., "Sonic Boom Computations for Double-Cone Configuration using CFL3D, FUN3D and Full-Potential Codes," *AIAA Paper 2006-0414*, January 2006.

[18]Alauzet, F., Dervieux, A., and Loseille, A., "Fully Anisotropic Goal-Oriented Mesh Adaptation for 3D Steady Euler Equations," *Journal of Computational Physics*, Vol. 229, No. 8, 2010, pp. 2866–2897.

[19]Jameson, A., "Aerodynamic Shape Optimization Using the Adjoint Method," Lecture notes, von Karman Institute for Fluid Dynamics, Brussels, Belgium, Feb. 2003.

[20]Aftosmis, M. J., Berger, M. J., and Adomavicius, G. D., "A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries," AIAA-Paper 2000-0808.

[21]Aftosmis, M., Berger, M. J., and Murman, S. M., "Applications of Space-Filling-Curves to Cartesian methods in CFD," AIAA-Paper 2004-1232.

[22]Berger, M. J., Aftosmis, M. J., Marshall, D. D., and Murman, S. M., "Performance of a new CFD flow solver using a hybrid programming paradigm," *J. Parallel Distrib. Comput.*, Vol. 65, 2005, pp. 414–423.

[23]Nemec, M., Aftosmis, M. J., Murman, S. M., and Pulliam, T. H., "Adjoint Formulation for an Embedded-Boundary Cartesian Method," AIAA Paper 2005–0877, Reno, NV, Jan. 2005.

[24]Nemec, M. and Aftosmis, M. J., "Adjoint Sensitivity Computations for an Embedded-Boundary Cartesian Mesh Method," *Journal of Computational Physics*, Vol. 227, 2008, pp. 2724–2742.

[25]Aftosmis, M. J., Berger, M. J., and Adomavicius, G., "A Parallel Multilevel Method for Adaptively Refined Cartesian Grids with Embedded Boundaries," *AIAA 2000-0808*, Jan. 2000.

[26]Aftosmis, M. J., Berger, M. J., and Murman, S. M., "Applications of Space-Filling-Curves to Cartesian Methods for CFD," *AIAA Paper 2004–1232*, Jan. 2004.

[27]Gill, P. E., Murray, W., Michael, and Saunders, A., "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM Journal on Optimization*, Vol. 12, 1997, pp. 979–1006.

[28]Darden, C. M., "Sonic boom minimization with nose-bluntness," NASA-TP 1348, National Aeronautics and Space Administration, 1979.

[29]George, A. R. and Seabass, R., "Sonic Boom minimization including both front and rear shocks," *AIAA Journal*, Vol. 9, No. 10, 1971, pp. 2091–2093.

[30]Nemec, M. and Aftosmis, M. J., "Adjoint Error-Estimation and Adaptive Refinement for Embedded-Boundary Cartesian Meshes," *AIAA Paper 2007–4187*, June 2007.

[31]Elmiligui, A. A., Cliff, S. E., Wilcox, F., Bangert, L., Nemec, M., Aftosmis, M. J., and Parlette, E., "Sonic Boom Computations for a Mach 1.6 Cruise Low Boom Configuration and Comparisons with Wind Tunnel Data," *AIAA 29th Applied Aerodynamics Conference (in publication)*, June 2011.

[32]Wayman, T., Howe, D. C., Waithe, K. A., Bangert, L., and Wilcox, F., "Near Field Acoustic Test on a Low Boom Configuration in Langley's 4x4 Wind Tunnel," *AIAA 29th Applied Aerodynamics Conference (in publication)*, June 2011.

[33]Howe, D. C., III, F. S., and Freund, D., "Developmenet of the Gulfstream Quiet Spike for Sonic Boom Minimization," *AIAA Paper 2008-0124*, 2008.

[34]Rodriguez, D. L. and Sturdza, P., "A Rapid Geometry Engine for Preliminary Aircraft Design," *44th AIAA Aerospace Sciences Meeting*, No. 2006-0929, AIAA, January 2006.

[35]Suwaratana, D. L. and Rodriguez, D. L., "A More Efficient Conceptual Design Process Using the RAGE Geometry Modeler," *49th AIAA Aerospace Sciences Meeting*, No. 2011-0159, AIAA, Janiuary 2011.

[36]Rodriguez, D. L., "*http://www.desktop.aero/docs/rage/*," .

[37]Kulfan, B., "Universal Parametric Geometry Representation Method," *Journal of Aircraft*, Vol. 45, No. 1, Jan.-Feb. 2008, pp. 142–158.