Seventh International Conference on Computational Fluid Dynamics (ICCFD7), Big Island, Hawaii, July 9-13, 2012

ICCFD7-2005

Conceptual Design of Low Sonic Boom Aircraft Using Adjoint-Based CFD

M. Wintzer*, I. Kroo*, M. Aftosmis** and M. Nemec***
Corresponding author: mwintzer@stanford.edu

* Stanford University, USA.
** NASA Ames Research Center, USA.
*** Science and Technology Corporation, USA.

Abstract: A multi-level framework for the design of low sonic boom aircraft is presented. An approach combining linear supersonic potential theory with atmospheric propagation and loudness prediction methods is used to generate low-boom near-field signals subject to lift and equivalent area constraints. This signal generator is then used to build a response surface fit that enables a conceptual-level design optimization to develop aircraft configurations with low boom. The imposition of a diverse range of multi-disciplinary, conceptual-level design constraints ensures a viable aircraft. Once a satisfactory baseline is obtained, the design is transferred to the CFD domain, where an adjoint-driven inverse design approach is used to determine the detailed aircraft shape. This optimization seeks to match the near-field pressure target associated with the conceptual baseline design, subject to geometric constraints that ensure conceptual-level performance predictions are preserved. An aircraft design example is presented, demonstrating the application of this approach to a Mach 1.5 supersonic business jet with 75,000 lb cruise weight.

Keywords: Multi-Disciplinary Optimization, Computational Fluid Dynamics, Conceptual Aircraft Design, Sonic Boom.

1 Introduction

The Shaped Sonic Boom Demonstrator (SSBD) program [1] confirmed in flight test that careful shaping of an aircraft can predictably alter the shape of the propagated ground signal of aircraft traveling at supersonic speeds. While indicating that the design of low-boom or even boom-less aircraft may be possible, the computational cost and complexity associated with current state-of-the-art tools makes effective conceptual-level design infeasible when using traditional optimization approaches. Ongoing efforts in low sonic boom design have focused on gradient-driven shaping of the signature in the near-field [2, 3] or on the ground [4, 5] to directly drive the high-fidelity shape optimization. Gradient-driven and gradient-free approaches incorporating an integrated objective in the form of a noise level are described by Alonso *et al.* [6]. Low-fidelity, linearized (area-rule) approaches are also described in Ref. [6]; however, these methods by definition are unable to accurately capture the non-linear features — namely, shocks — present in supersonic flows. While invaluable during design exploration, details of the shape produced by such approaches do not readily carry over to equivalent high-fidelity representations of the geometry.

This paper presents a methodology for low sonic boom aircraft design, integrating multiple levels of analysis fidelity in order to develop not only the detailed aircraft shape, but also to ensure the overall configuration satisfies multi-disciplinary constraints resulting in a viable design. The task of minimizing sonic boom loudness is solved using an approach based on linear theory, where a low-dimensionality parameterization based on the F-function [7] enables rapid discovery of waveform shapes with low sonic boom. The use of fast, linear methods enables the application of gradient-free optimization methods insensitive to the multiple minima and discontinuities that characterize the sonic boom design space. This boom-minimization

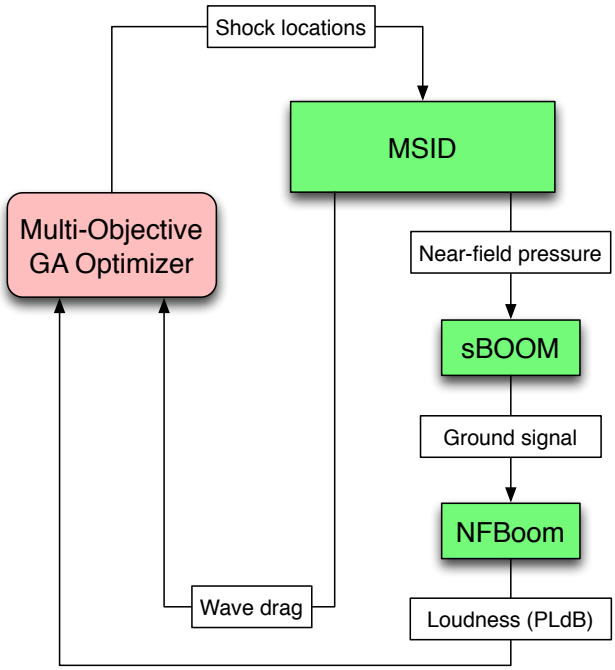

Figure 1: Boom minimization tool flow diagram

approach is used to build a response surface fit relating the equivalent area distribution (computed from the F-function), aircraft weight, and cruise altitude to a minimum loudness value. By integrating this fit into a conceptual design framework, an external aircraft configuration compatible with low boom is developed, along with a corresponding low-boom target waveform. Finally, an inverse-design method is used to solve for the detailed vehicle shape with near-field pressure disturbance matching the target waveform. The combined design approach is demonstrated on the design of a notional, low-boom, supersonic business jet.

2 Generation of Low-Boom Near-Field Signals

Near-field signals with low sonic boom are sought using the optimization approach depicted in Fig. 1. A multi-objective genetic algorithm (GA) based on NSGA-II [8] varies the distribution of shocks in the F-function representation of the near-field pressure signal. The "MSID" block encapsulates an inner optimization level, dedicated to solving for the near-field signal shape compatible with the provided shock distribution, subject to constraints on lift and enclosed volume. MSID refers to the Multi-Shock Inverse Design method of Haas and Kroo [9]; a brief summary of the approach is presented in §2.1. In addition, linear axisymmetric theory is used to estimate the wave drag of the equivalent body corresponding to the F-function shape. To obtain a loudness estimate, the near-field pressure signal output from the MSID block is propagated to the ground using the sBOOM augmented Burgers propagation code of Rallabhandi [10]; this ground signal is then processed using the NFBOOM [11, 12] implementation of the Stevens Mark VII procedure [13] to obtain the perceived loudness in decibels (PLdB). The loudness and wave drag are the objectives for the multi-objective optimization. The Pareto-optimal front developed through evolution of the GA population enables the designer to trade low-boom for low-drag; in this approach the focus is on low-boom, and use of a multi-objective GA ensures that those surviving low-boom population members have the lowest drag at their particular loudness level. The approach lends itself to parallelization, and the software implementation has been designed to make full use of high-performance computing clusters.

2.1 Multi-Shock Inverse Design

MSID extends the classic 2-shock Seebass-George-Darden (SGD) method [14, 15, 16] to handle an arbitrary number of shocks. In both approaches, the Whitham F-function [7] is used as a means of relating the equivalent axisymmetric area distribution A_e of a body

$$F(x) = \frac{1}{2\pi} \int_0^x \frac{A_e''(\xi)}{(x-\xi)} d\xi \tag{1}$$

to the near-field pressure disturbance

$$\frac{\Delta p}{p_\infty}(x) = \frac{\gamma M_\infty^2 F(x)}{\sqrt{2\beta r}} \tag{2}$$

where M_∞ is the freestream Mach number, γ is the ratio of specific heats, r is the vertical offset distance of the near-field from the body, and $\beta = \sqrt{M_\infty^2 - 1}$. For a lifting body terminating at a point, and taking as negligible the effects of wake and jet exhaust, the equivalent axisymmetric body has finite base area

$$A_e(l) = \frac{\beta L}{\rho_\infty U_\infty^2} \tag{3}$$

where l is the final body station, L the lift, and ρ_∞ and U_∞ the freestream density and speed, respectively.

2.1.1 MSID Background

The MSID methodology is briefly described to motivate a formal statement of the optimization problem. As it propagates through the atmosphere, the near-field form of the F-function distorts according to the atmospheric advance factor (α, in Darden's nomenclature), potentially giving rise to a non-physical, multi-valued signal. A single-valued signal is recovered by introducing shocks, located such that the areas of the regions enclosed by the signal are balanced about each shock. In MSID, treatment of each shock is analogous to that of the initial shock in the SGD approach. Figure 2 shows a portion of a notional F-function. The slope of the balancing line S is given by $1/\alpha$, x_{b_k} is the desired shock location, and S intercepts the F-function at x_{a_k}. The crossover point x' separates the enclosed areas A_1 and A_2. The optimizer varies the F-values to ensure $A_1 = A_2$ by solving

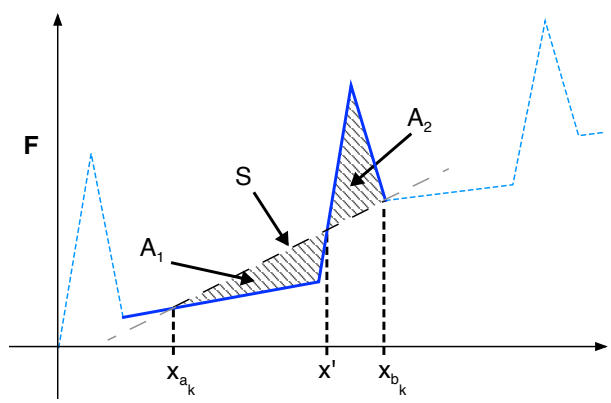

Figure 2: F-function shock-balancing

$$A_2 - A_1 = \int_{x_{a_k}}^{x_{b_k}} [F(x) - S(x)]\, dx = 0 \tag{4}$$

for each shock, where $F(x_{a_k}) = S(x_{a_k})$ for the signal to be valid. To suppress shock formation due to non-linear steepening of line segments not participating in area balancing, a slope limit of ε/α is imposed, with ε typically set between 0.7 and 0.8. For the propagated ground signal, constraints are imposed on the maximum allowable pressure jump, and minimum tolerated inter-shock interval. An integrated loudness metric for the ground signal arguably obviates the need for such ground-signal specific constraints. A recent study [17] on indoor perception of shaped sonic booms, however, indicates that ground signal features such as initial overpressure and shock frequency may lead to signals which are objectionable indoors despite being considered acceptable outdoors.

The MSID optimization problem for a signal with N shocks is stated as

$$
\begin{array}{lll}
\min & \mathcal{J} = \sum_{i=1}^{N} \left(\int_{x_{a_k}}^{x_{b_k}} [F(x) - S(x)]\, dx \right)^2 & \\
w.r.t. & F_i & \\
s.t. & A_e(l) - \frac{\beta L}{\rho_\infty U_\infty^2} & = 0 \\
 & F(x_{a_k}) - S(x_{a_k}) & = 0 \text{ for } k = 1, N \\
 & \Delta P_k - \Delta P_{\max_k} & \leq 0 \text{ for } k = 1, N \\
 & \tau_k - \Delta T_k & \leq 0 \text{ for } k = 2, N \\
 & \max(F'(x)\ \forall\ x \in [x_{b_{k-1}}, x_{a_k}]) - \varepsilon \frac{1}{\alpha} & \leq 0 \text{ for } k = 2, N \\
 & r_{\min_j} - r_j & \leq 0 \text{ for } j = 1, n_r
\end{array}
\tag{5}
$$

The additional terms introduced here are ΔP_k and ΔP_{max_k}, which are the actual and maximum allowable ground signal pressure jumps, respectively, τ_k and ΔT_k, the minimum allowed and actual ground signal inter-shock time interval, respectively, and r_{min_j} and r_j, the minimum allowed and actual radius at axial station j.

The optimizer seeks to minimize the sum of the squared area imbalance at each shock with respect to the F-function ordinates y_i. In addition to the MSID constraints, the additional radius constraint enables the designer to specify minimum radii at an arbitrary number of user-prescribed axial stations. Ref. [9] provides complete details of the approach along with a brief summary of the theoretical background.

2.1.2 Solving for Near-Field Signals with MSID

A gradient-based optimizer is used to solve for the F-function shape in Ref. [9]. Although this is an efficient approach to solving the optimization problem, the initial shape and point-distribution of the F-function required careful construction to ensure a well-posed problem. For sonic boom minimization, the solution of many different F-function shapes is required, and a more robust approach amenable to automation is needed.

To satisfy the requirement on robust optimization, a genetic algorithm (GA) is used to solve the inverse design optimization problem (Eq. 5), with constraints handled using linear exterior penalty methods. Despite an increase by several orders of magnitude in the number of required function evaluations, turnaround time to obtain and MSID solution is only a few seconds.

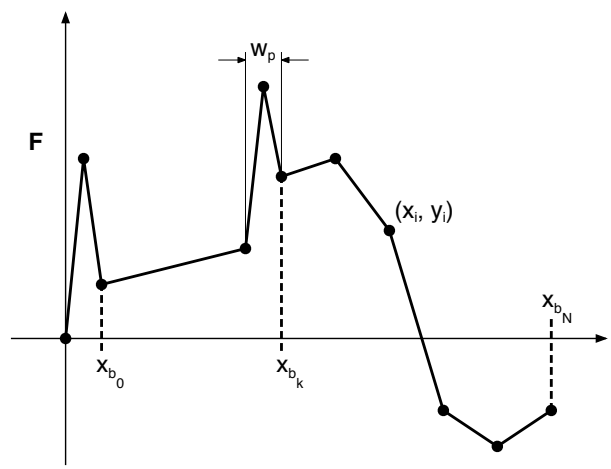

Figure 3: Automated F-function template construction based on balance point x_{b_k}. x_{b_0} and x_{b_N} correspond to the initial and final shocks, and w_p is the peak width

Proceeding towards full automation of the MSID process, consider that the character of the F-function is defined by the locations of the balance points, and that the initial and final balance points are always fixed. Based on this observation, a function template is constructed as follows: every interior shock except the last is assigned 3 vertices to form a peak of width w_p. The remaining region leading up to the final shock is assigned 5 equally spaced vertices. The abscissae specified, the ordinates are randomly assigned, and the GA is responsible for finding the F-values resulting in a valid F-function. Illustrated in Fig. 3, this spartan formulation of the F-function was heuristically determined over many trials, and found to offer sufficient degrees of freedom that a solution to the MSID problem was consistently found.

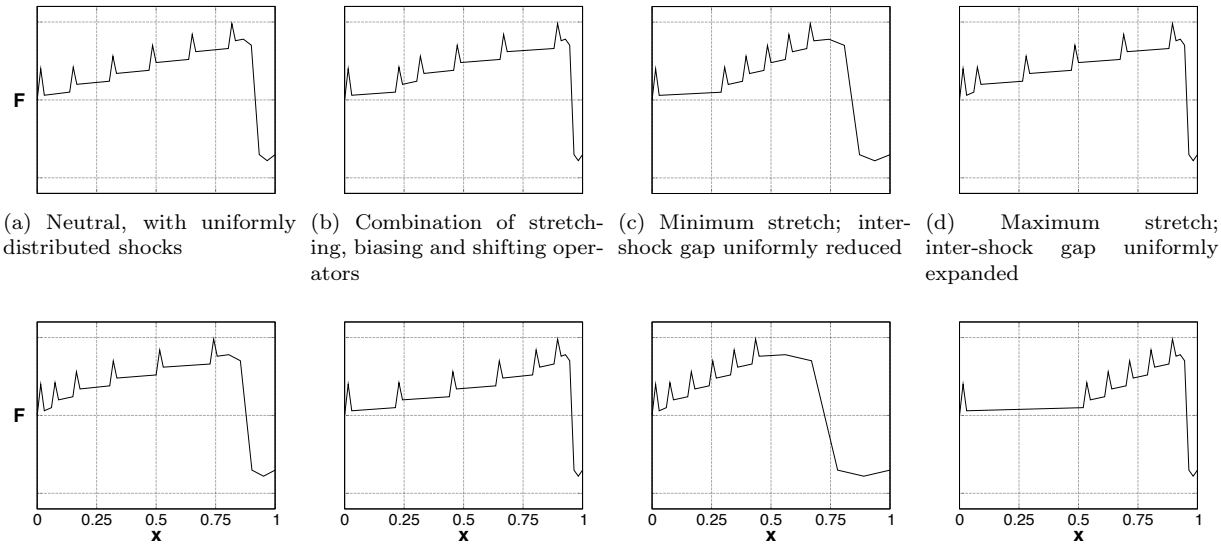

(a) Neutral, with uniformly distributed shocks

(b) Combination of stretching, biasing and shifting operators

(c) Minimum stretch; inter-shock gap uniformly reduced

(d) Maximum stretch; inter-shock gap uniformly expanded

(e) Power distribution with left bias

(f) Power distribution with right bias

(g) Uniform left shift of interior shocks

(h) Uniform right shift of interior shocks

Figure 4: Illustration of parametric balance point layout for an F-function with seven shocks. With three parameters controlling the uniform stretching, and left-right biasing and shifting of shocks, a wide variety of signal shapes is possible

2.2 Shock Parameterization for Efficient Optimization

A multi-objective GA is used to determine the balance point distribution minimizing sonic boom loudness and wave drag. When performing this optimization, a critical consideration is that the point distribution must remain in sequence for the signal to remain valid. Initial trials allowed unrestricted selection of the balance points by the GA, relying upon its inherent robustness to weed out invalid shock distributions. In a large proportion of cases, this resulted in a rapid and pervasive corruption of the population that in practice was unrecoverable. Pre-specifying fixed bounds for each balance point to explicitly forbid crossing of balance points prevented signal corruption, but such binning unduly constrained the design space, hiding many of the more interesting near-field signal shapes that propagated to quiet ground signals.

To resolve this dilemma, rather than allow the optimizer to directly select the balance points, they are parametrically distributed using a combination of stretching, biasing and shifting operators, as illustrated in Fig. 4. The benefits of this approach are two-fold. First, with only three design variables, the optimizer can distribute an arbitrary number of shocks in a diverse range of ways. Second, even with this flexibility, shocks are guaranteed to remain in sequence, ensuring that every instance of near-field pressure in the GA population represents a physically realizable waveform.

3 Aircraft Conceptual Design with Minimum Loudness Objective

The Program for Aircraft Synthesis Studies [18] (PASS) provides a design environment encompassing all aspects of mission performance. Transonic acceleration and supersonic cruise performance are estimated using a linearized area rule method [19]; as a result an equivalent axisymmetric representation of the aircraft is generated for each instance of the aircraft configuration. Low-speed performance metrics such as balanced field length and second segment climb gradient are estimated with the aid of an integrated discrete-vortex method (a sample geometry is shown in Fig. 5(a)). This discrete-vortex model includes geometry features such as flaps, slats and control surfaces, allowing for the inclusion of constraints such as the maximum allowable elevator deflection for takeoff rotation in the design process. Aircraft component weights are based on available data for modern business-jets, while wing and empennage weights are estimated based on a statistically correlated bending index that is related to the fully stressed bending weight of the wing box.

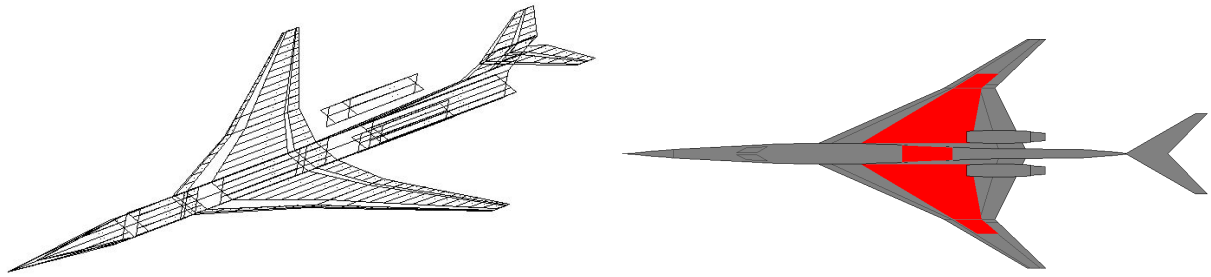

(a) Discrete-vortex model used to predict low-speed performance

(b) User-defined fuel tank layout used for weight and balance

Figure 5: Examples of PASS analysis models used during mission performance prediction

Fuselage weight is based on gross fuselage wetted area and a pressure-bending load parameter. The center of gravity (CG) location is computed based on weight and typical placement of the various aircraft components; CG movement during the mission due to fuel burn is also modeled based on the fuel tank layout, an example of which is shown in Fig. 5(b).

A typical optimization objective in the absence of boom considerations is the minimization of maximum takeoff weight. In this application, the goal is minimization of sonic boom loudness, a metric to which PASS is insensitive. Consider that the near-field target generator, due to its use of the F-function, can be used to derive the equivalent area distribution corresponding to the generated pressure signal; indeed this relation is what allows for the imposition of constraints on the axial development of cross-sectional area. At each mission evaluation, PASS produces an equivalent area distribution corresponding to the current aircraft configuration. By repeatedly calling the target generator with a variety of equivalent area constraints, cruise weights, and altitudes, a response surface relating these parameters to sonic boom loudness can be constructed. In this case, fitting is performed using an Ordinary Kriging implementation by Rajnarayan [20]. Introducing a call to this fit in the objective function following the mission analysis (at which point equivalent area, cruise weight and cruise altitude are available) enables the PASS optimization to obtain an estimate on loudness for the configuration state at any given iteration.

The PASS tool augmented with the sonic boom loudness fit is described graphically in Fig. 6. While the objective has changed, constraints ensuring a viable aircraft are still imposed, such as a minimum requirement on the 2nd-segment climb gradient, maximum allowable vertical tail C_L during engine-out conditions, and minimum allowable cruise-climb gradient. The fit relating aircraft parameters to loudness guides the optimized aircraft configuration towards one compatible — at least in a linear sense — with the target pressure signal. When simply specifying loudness as the objective, as is the case here, a concern is the unbounded increase in takeoff weight as the optimizer attempts to satisfy a range constraint, which may not readily align with configuration changes necessary for low sonic boom loudness. However, loudness tends to increase with aircraft weight, and as a result the latter parameter can be safely omitted from the objective.

Gross sizing parameters, such as the rough fuselage shape, wing and empennage planforms, and relative locations of each of these major components, are then transferred to the second stage of the design process, where the adjoint-driven inverse design framework performs detailed shaping of the aircraft. In this second stage, the PASS-derived gross sizing parameters are held fixed, and only parameters describing details such as airfoil section shapes, wing twist parameters and the fuselage radius distribution are allowed to vary. Additionally, geometric constraints from the PASS design definition are imposed to ensure features such as wing thickness and cabin size are preserved. By limiting the adjoint-design degrees of freedom to parameters to which the PASS analysis is largely insensitive, the performance predicted for the PASS-developed configuration is preserved. Experience has shown that with sufficient effort and design degrees of freedom, the inviscid L/D predicted by PASS can be reliably recovered in the CFD domain.

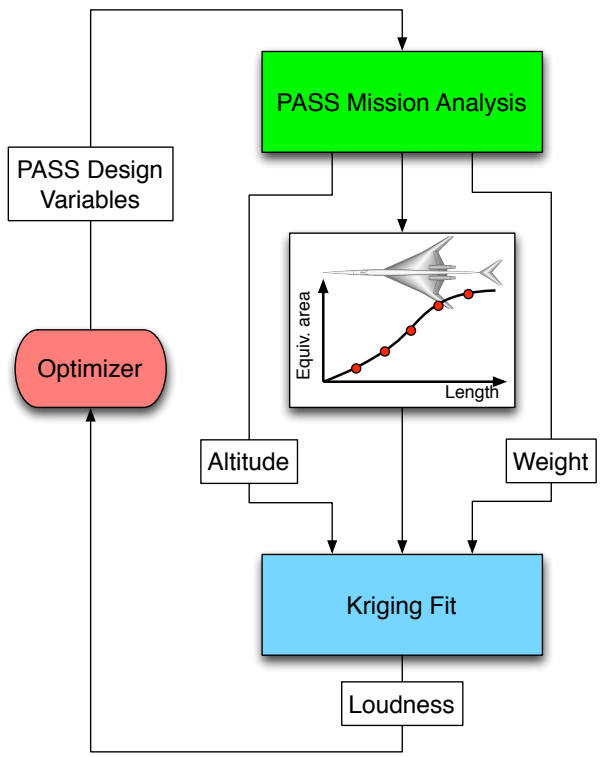

Figure 6: Development process for conceptual baseline with low-boom objective

4 Adjoint-Driven CFD for Conceptual Design

The high-fidelity component of this approach uses the NASA Cart3D [21] CFD design framework. It employs a multilevel, embedded boundary Cartesian mesh, and provides a robust and efficient parallel multigrid solver for the three-dimensional Euler equations. Cart3D also features adjoint-derived sensitivities [22] and adjoint-driven mesh refinement [23], both essential components for the design approach presented here.

The SNOPT [24] constrained optimization system is used to drive the high-fidelity optimization. SNOPT uses a Sequential Quadratic Programming (SQP) algorithm, solving a Quadratic Programming (QP) subproblem at each major iteration in order to determine the next search direction. Hessian updates are performed using the BFGS formula.

The high-fidelity objective function is a weighted sum of the form

$$\mathcal{J} = w_p \int \left(\frac{p - p_t}{p_\infty} \right)^2 ds + w_D(C_D) + w_L(C_L - C_{L_t})^2 \tag{6}$$

where the subscript (t) denotes a target value. This functional combines the near-field pressure target, drag, and lift constraint.

4.1 Adjoint-Driven Mesh Refinement

The accuracy of a CFD solution is dependent on the quality of the volume grid used. While hand-crafted grids are generally adequate at resolving body-based functionals such as lift and drag, accurate resolution of near-field pressure is more difficult. Adjoint-driven mesh refinement allows for the rapid and automated generation of volume grids that are appropriately tailored to the targeted functionals, while also providing an estimate of the functional error on the adapted grid.

In this refinement method, adjoint-derived cell-wise error estimates are used to flag cells in the volume grid that contribute the most error to the objective of interest. Flagged cells are then iteratively refined through

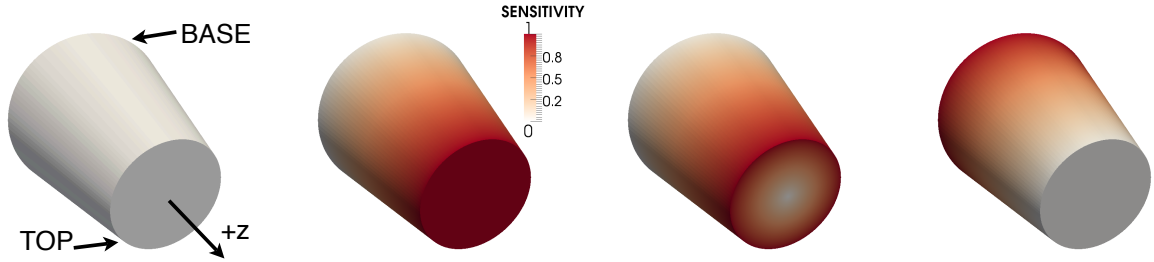

(a) Truncated cone geometry (b) Sensitivity to stretching along z-axis (c) Sensitivity to increase in top radius (d) Sensitivity to increase in base radius

Figure 7: Truncated cone providing a simple illustration of geometric sensitivities to the parametric variables defining its shape

successive flow and adjoint solutions until a preset error tolerance is satisfied. Details on the implementation of the adjoint solver and adaptive mesh refinement scheme are given in Refs. [22] and [25].

The validation of this method in supersonic flow is described in previous work [26], where mesh alignment and stretching techniques were used to dramatically reduce volume grid densities required for accurate near-field pressure prediction. In the current approach, adaptive refinement is used to determine meshing requirements for high-fidelity shape optimization, using a technique described in recent work by Aftosmis *et al.* [27].

4.2 Adjoint-Driven Design

The adjoint implementation of the Euler flow solver computes sensitivities of Eq. 6 with respect to shape parameters (design variables) of the model; details of the approach are given in Ref. [3]. This procedure requires computation of the sensitivities of the surface triangulation with respect to the design variables. These sensitivities are computed by central differencing each vertex $\mathbf{x}_i$ in the surface discretization with respect to each geometric design parameter v

$$\frac{\partial \mathbf{x}_i}{\partial v} = \frac{\mathbf{x}_i^+ - \mathbf{x}_i^-}{2\epsilon} \tag{7}$$

where ϵ is the perturbation of the design parameter, and $\mathbf{x}_i^+$ and $\mathbf{x}_i^-$ represent the locations of the ith vertex following positive and negative perturbations of the design parameter v, respectively. Cases where a perturbation is only feasible in one direction substitute one-sided differencing. The truncated cone shape depicted in Fig. 7(a) provides a simple illustration of surface sensitivities relative to axial stretching along the z-axis, and changes to top and base section radii (Figures 7(b) through (d), respectively). Red shading indicates the increasing magnitude of surface sensitivity with respect to the design variable.

The parametric geometry modeler used to produce the surface triangulations and associated surface sensitivities required for this work was initially conceived to produce specialized high-precision geometry for a prior study [28]. The capability and scope of this modeler has grown considerably since that time, and it now provides sufficient flexibility to produce full aircraft configurations, with surface sensitivities, as demonstrated in Fig. 8. Sensitivities associated with geometric features are also computed, allowing the imposition of constraints on quantities such as area, volume, and thickness in a way that can be efficiently handled by the optimizer. This approach retains much of the flexibility afforded by a CAD-based geometry generation system while offering significant time savings, especially as the number of design parameters is increased.

5 Aircraft Design Example

This example demonstrates the application of the presented design approach on a conceptual aircraft design problem. The baseline configuration was created using PASS, with the basic planform shape and layout of

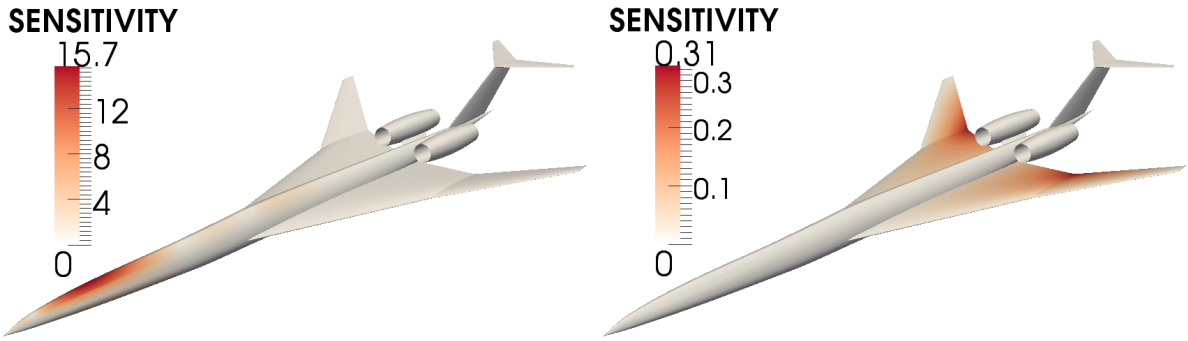

(a) Sensitivity to height change in fuselage station

(b) Sensitivity to wing break section twist

Figure 8: Illustration of surface sensitivities computed for a pair of typical geometry design parameters. Red shading indicates the increasing magnitude of surface sensitivity with respect to the design variable

major components established as described in §3. Salient characteristics of this vehicle are summarized in Fig. 9.

5.1 Geometry Parameterization

The geometry parameterization used during the inverse-design phase is defined using 250 shape design variables, and is illustrated in Fig. 10. The geometry consists of 6 components, which are constructed as follows:

- The fuselage is lofted using 27 defining stations; the resultant shaping authority, particularly near the nose, was found in previous work [29] to be necessary to match low-boom target signals of the form produced by the target generation approach described in §2. To minimize the structural weight penalty associated with pressurization of the forward fuselage, these stations are circular, with radius allowed to vary at every station. A 5 ft minimum radius is specified for the cockpit station (first red highlighted station), while a 6 ft minimum radius constraint is placed at the cabin constraint stations (remaining two red highlighted stations). Stations aft of the cabin use a significantly more flexible parameterization that allows for a wide range of cross-section shapes. The fuselage is lofted using a modification of the monotone cubic Hermite spline proposed by Rajnarayan [30, 31].

- The Kulfan "class function/shape function" (CST) parameterization [32] with 4th-order Bernstein polynomials is used to generate the airfoil section shapes at each of the defining spanwise stations, resulting in 10 design variables per section. The wing can be lofted either linearly or smoothly using a cubic spline interpolation. The wing incidence angle relative to the fuselage is adjustable, but not varied during optimization. The incidence angles of the all sections are allowed to vary, with the exception of the root section. The thickness-to-chord (t/c) ratio of each section is constrained to be equal to the thickness distribution established by the PASS optimization, which considers not only the effect of thickness on aerodynamic performance, but also on structural weight. In this case, the tip t/c ratio is 5%, steadily decreasing to 2% at the root section.

- The horizontal tail is all-moving, and its incidence angle is allowed to vary. Root and tip airfoil sections are defined using a similar CST parameterization as the wing, with 10 design variables per section. The t/c ratios of the defining sections are constrained to be equal to the PASS-optimized value of 3%.

- The vertical tail geometry is defined by four 4% thick CST sections with 10 design variables per section (with t/c constraint again from the PASS model). The vertical tail is lofted using the monotone spline interpolation.

- The nacelle pylon geometry consists of a swept prismatic extrusion of a 10-parameter CST section constrained to be no less than 4% thick. Pylon incidence is allowed to vary.

Wing

S_{ref}	1,106 ft^2
AR	4.1
$\Lambda_{c/4}$	57.2°
λ	0.21
LEX	0.432
TEX	1.26
y_{break}/b	0.333

Empennage

S_h/S_{ref}	0.183
AR_h	3.0
$\Lambda_{c/4,\,h}$	57.0°
λ_h	0.28
S_v/S_{ref}	0.2
AR_v	0.75
$\Lambda_{c/4,\,v}$	59.0°
λ_v	0.55

Performance

Range	4,000 nmi
Cruise Mach	1.5
Cruise L/D	8.23
MTOW	83,250 lb
TOFL	6,160 ft
SLS thrust	34,300 lb
Initial/final cruise C_L	0.15 / 0.15
Initial/final cruise altitude	45,000 ft / 57,000 ft
Payload	5,640 lb

Fuselage

Length	140 ft
Max width	6.0 ft

Total design variables: 20

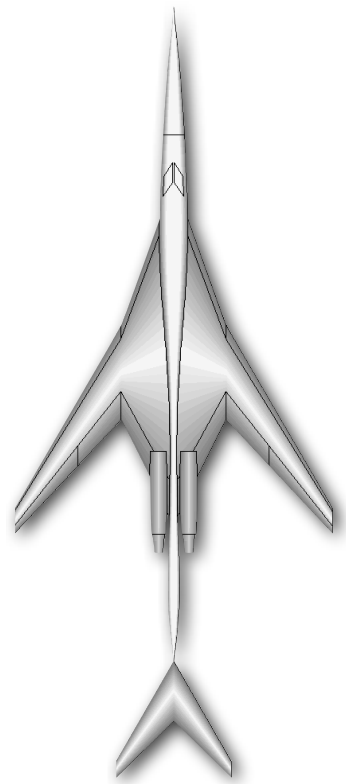

Figure 9: PASS baseline planform view and general characteristics

- The nacelle geometry is flow-through, with internal diameter corresponding to the engine inlet diameter predicted by PASS. The outer mold line is fixed; only pitch and yaw of the nacelle are allowed to vary.

The wing and empennage planform shapes, along with their positions relative to the fuselage, are specified by the PASS aircraft definition and not allowed to vary as part of the optimization. Longitudinal placement of the nacelles is similarly specified by PASS. The PASS baseline configuration makes certain assumptions about the spanwise and longitudinal lift distributions and does not include any airfoil twist parameters. In order to provide a reasonable starting point for the adjoint-driven shape optimization, a sequence of drag-minimization runs was performed to roughly characterize wing twist and angle of attack needed to meet the C_L requirement. This study presents untrimmed results; pitching moment in the final design would be easily corrected with fuel pumping (as was done, for example, on Concorde).

5.1.1 Evolution of Parameterization

The geometry parameterization shown in Fig. 10 corresponds to that of the final optimized configuration. Initially, this parameterization was less complex, with all circular fuselage sections, and a linearly-lofted wing defined by 3-sections (matching the PASS-parameterized planform). Figure 11 describes how the initial parameterization was expanded, enabling the optimizer to match the provided near-field signal with the allowed shaping authority. Mach contours, along with the corresponding best near-field pressure signal (relative to the target) achieved with the initial parameterization is shown in Fig. 11(a). The pronounced aft undershoot in the $\Delta p/p$ plot is due to the nacelle. Note that the initial spike present in the target signal is due to the linear theory approach used in its generation; due to the small distance over which the spike occurs, when the signal is area balanced the resulting initial shock is very weak. As a result, the discrepancy between the Euler-predicted signal and the target at this location does not manifest itself in the propagated ground signal. A more flexible aft fuselage parameterization was introduced, enabling the optimizer to hide the presence of the nacelle by flattening the aft fuselage, as seen in Fig. 11(b). The remaining discrepancy

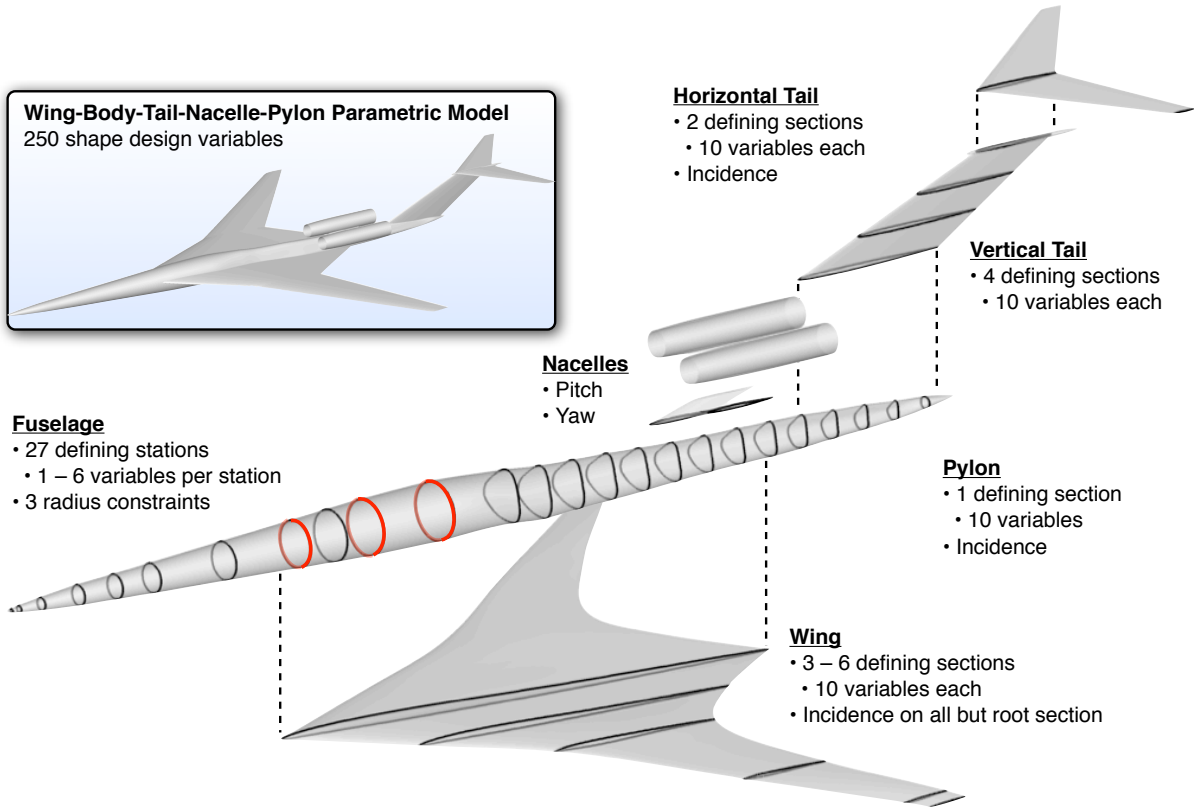

Figure 10: Parametric wing-body-tail model used for inverse design

in the near-field signal was associated with the hard kink in the wing trailing edge break. Transitioning to a cubic spline lofted wing planform and filling the hard break with a smooth fillet yielded the result shown in Fig. 11(c); at this point the optimizer was felt to have the shaping flexibility required to match the provided target. Note that the planform shape, other than the trailing edge fillet, is preserved.

5.2 Results

The volume mesh for this example was constructed using mesh stretching and rotation techniques described in previous work [26]. The mesh had a total of 5M cells with an aspect ratio of 4:1. Near-field pressure signals predicted for the starting and final geometries are compared against the target in Fig. 12(a). Agreement in general is very good. The propagated ground signals are shown in Fig. 12(b). The forward section of the ground signal shows good agreement with the target, while the slight increase in gradient of the aft portion of the signal in the near-field have formed a stronger aft shock, resulting in the increased loudness relative to the target. The final geometry PLdB is 79.4, 12 PLdB quieter than the initial design. The final geometry inviscid L/D is 14. Mach contours on the symmetry plane for the initial and final geometries are shown in Fig. 13(a) and (b), respectively. The reduction in shock strength is clearly visible between the two panels. C_p contours on the final geometry are shown in Fig. 14. Some features of note are the blunting of the nose, smoothly blended wing planform, and pronounced flaring out of the aft fuselage sections. This latter item is clearly visible in Fig. 14(b). The inverse design optimization from starting to final geometry used approximately 200 design iterations, consuming roughly 100 hours of wall-clock time on a desktop PC*.

*A dual-hex (12 cores) 2.93 GHz Intel Xeon PC

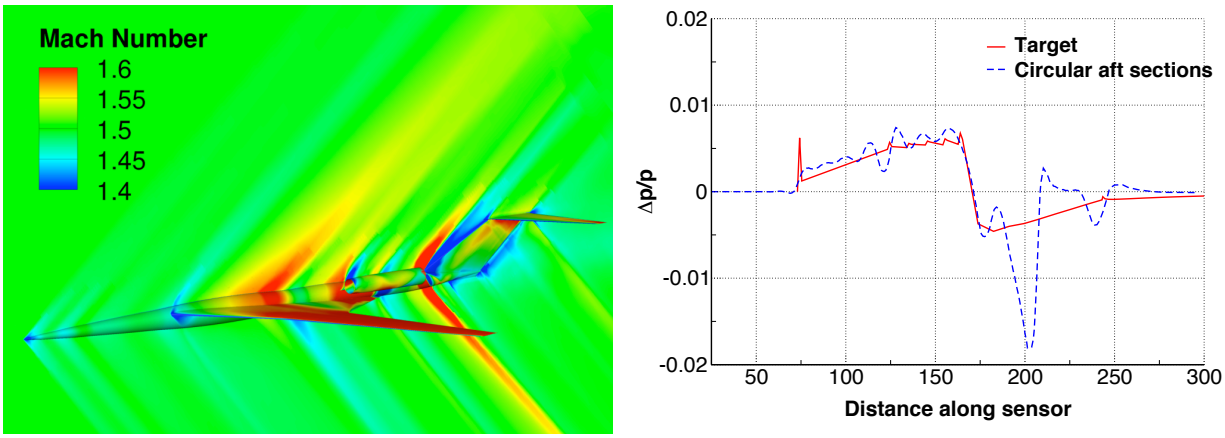

(a) Initial parameterization. Circular fuselage sections, linearly lofted wing with 3 defining stations

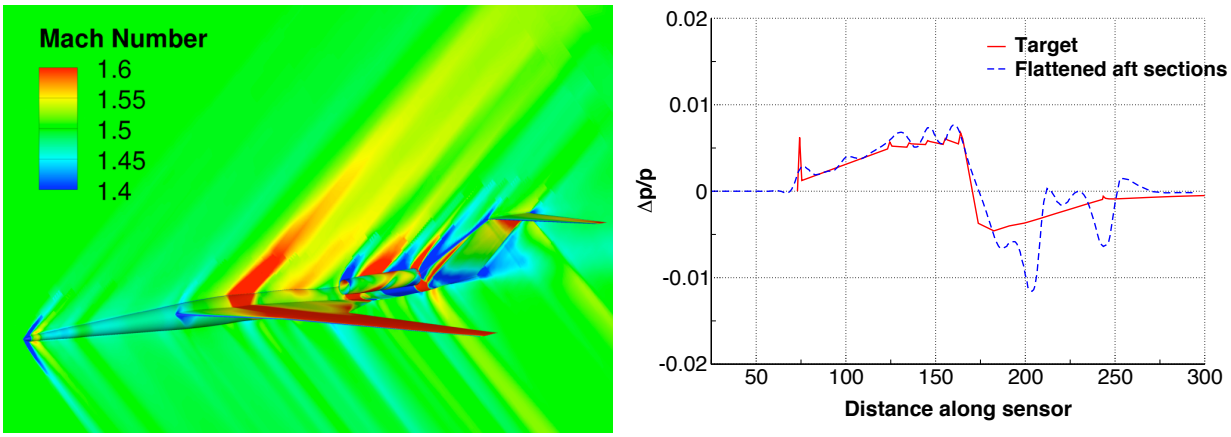

(b) More flexible aft fuselage sections introduced

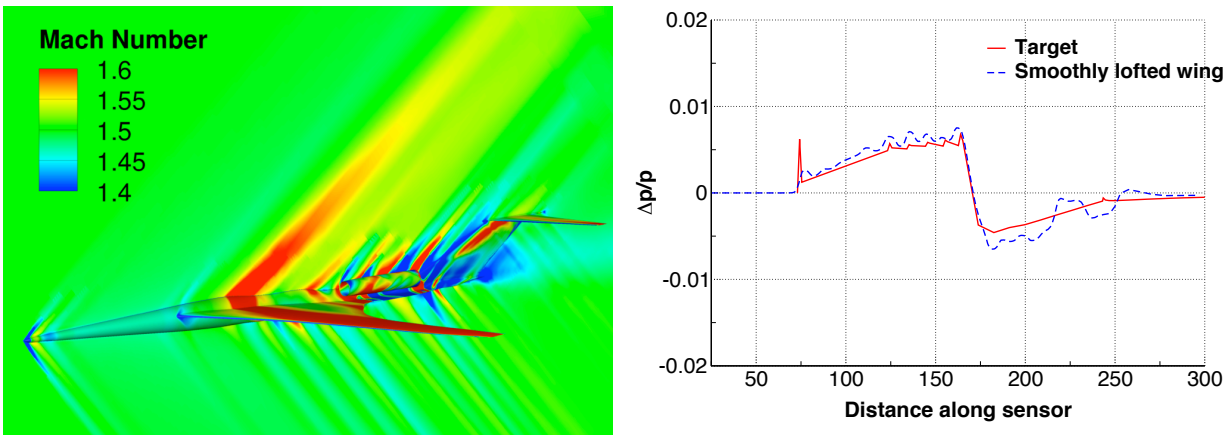

(c) Wing converted from linear to cubic spline lofting; with the exception of the fillet at the break section trailing edge, planform shape is preserved relative to 3-section wing

Figure 11: Evolution of geometry parameterization leading to that used to develop final configuration. Left panel in each figure shows Mach contours on geometry and symmetry plane, while right panel shows best achieved near-field signal, relative to target, with the noted parameterization. Near-field signals are sampled at 2.4 body lengths beneath the aircraft

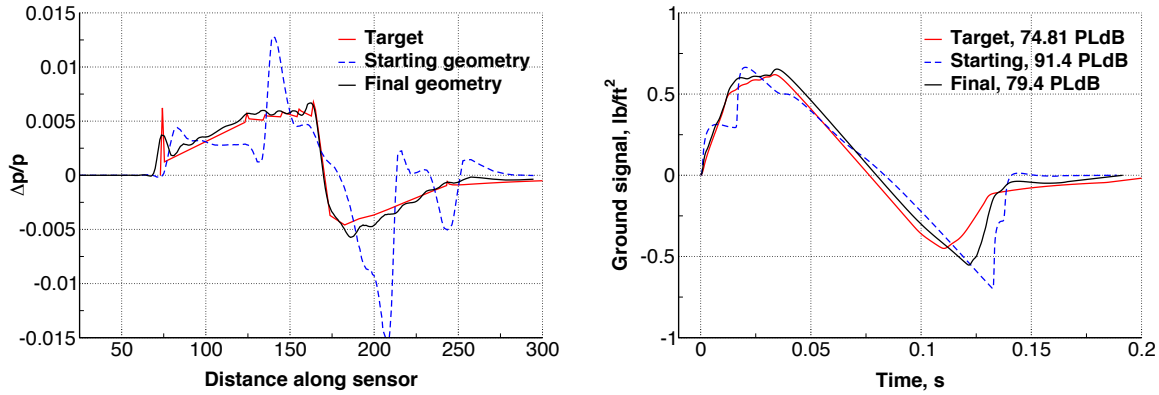

(a) Near-field pressure signals

(b) Propagated ground signals

Figure 12: Comparison of near-field and propagated ground signals for target, starting and final geometries. Near-field pressures sampled at $h/L = 2.4$. Cruise $M_\infty = 1.5$, altitude 45,000 ft, and $C_L = 0.15$ (this corresponds to a cruise weight of roughly 75,000 lb)

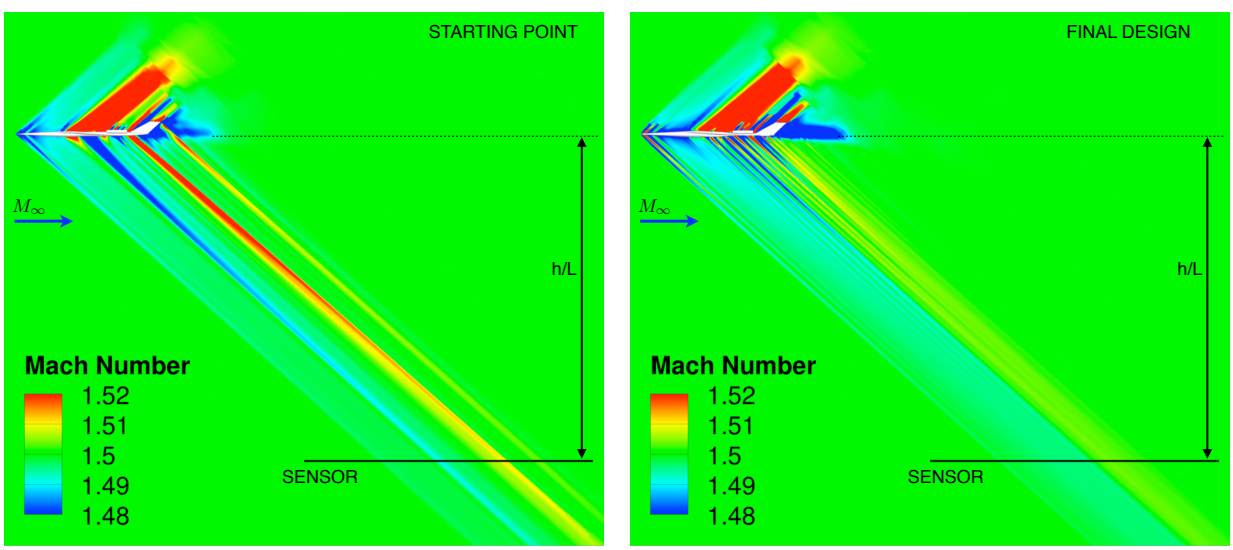

(a) Starting geometry

(b) Final geometry

Figure 13: Comparison of Mach contours plotted on the symmetry plane for starting and final geometries. Cruise $M_\infty = 1.5$, $C_L = 0.15$

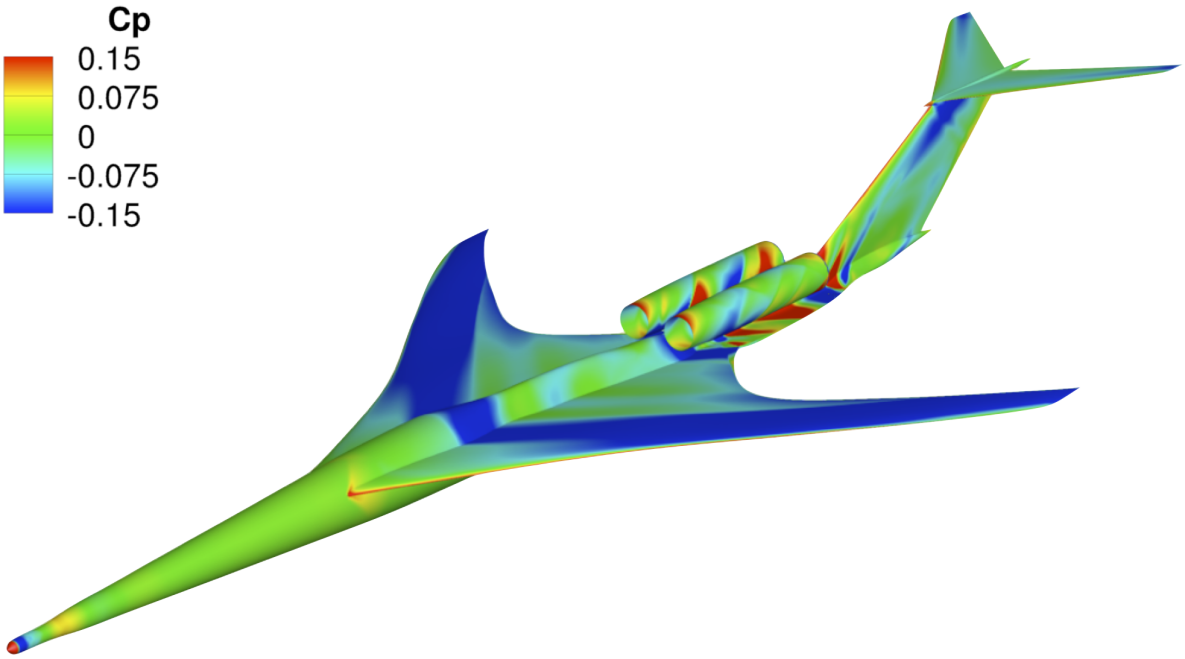

(a) Front ¾ view

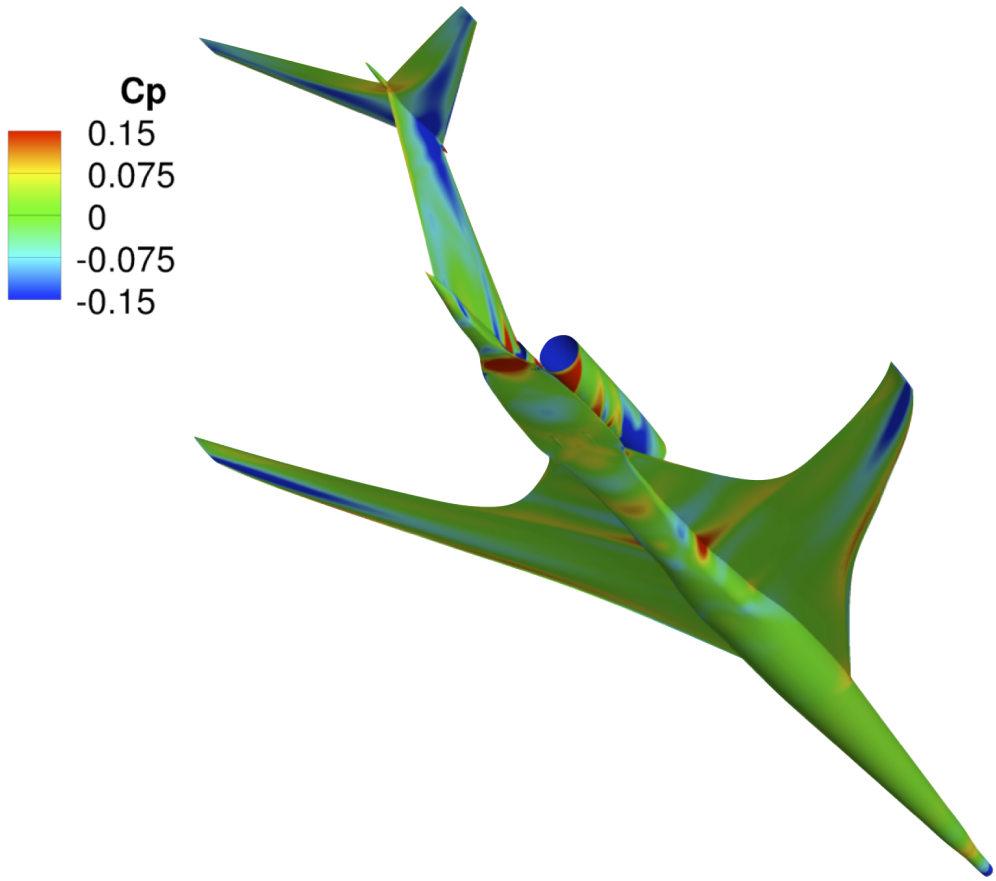

(b) Rear ¾ view

Figure 14: C_p contours on final geometry. $M_\infty = 1.5$, $C_L = 0.15$

6 Conclusions and Future Work

This work presented a multi-level framework for the conceptual design of aircraft with low sonic boom. Fast, linear methods enabled the rapid discovery of low-boom near-field pressure signals. The incorporation of this capability into a sophisticated, multi-disciplinary conceptual design tool through a response surface fit enabled the development of aircraft configurations compatible with low-boom. A full mission analysis, including constraints on range, balanced field length, climb performance, and transonic ac celeration ensured a viable aircraft. Detailed shaping of the outer mold line was then performed in the CFD domain using an adjoint-driven, inverse design approach, subject to geometric constraints ensuring that the CFD parameterization only altered the geometry in ways that the mission analysis and associated constraints were insensitive.

Driven using a multi-level optimization scheme and coupled with the novel parameterization and automation methods developed as part of this work, MSID was found to be an effective way of developing near-field signal shapes incorporating lift and volume constraints. The coupling of this method through a response surface fit with the conceptual design component enables the development of baseline configurations compatible with low boom. Baseline designs are then transferred to the CFD domain, where an inverse-design method re-shapes selected aspects of the aircraft outer mold line in order to match the near-field waveform target. By limiting the shaping authority of the inverse-design optimization to parameters to which the conceptual-level analyses are insensitive, performance predictions made by these analyses — such as cruise range, takeoff field length and structural weight — are preserved in the final, CFD-based design. Finally, the airplane design example showed that even with the optimizer lacking any planform shaping and positioning authority, the shape of the near-field pressure on the optimized aircraft geometry was in good agreement with the target signal.

Future work will explore whether alternate aircraft configurations, such as delta-wing and oblique wing designs, can deliver similarly low sonic boom levels while offering even higher aerodynamic performance. The inclusion of power effects, including inlet and exhaust flows, and their effect on shape and performance will be investigated. In terms of the design approach, the potential for performance improvements and lower sonic boom loudness resulting from the integration of the PASS analysis into the shape optimization will be studied.

7 Acknowledgements

This work was supported by NASA grant NNX09AJ21G, and the NASA Education Associates Program. The authors wish to thank Sriram Rallabhandi (National Institute of Aerospace) for kindly providing access to the sBOOM atmospheric propagation code used in this work. The authors are also grateful to Alex Haas for many insightful discussions and sharing his expertise with the multi-shock inverse design method that formed a core part of this design approach.

References

[1] Joseph Pawlowski, David Graham, Charles Boccadoro, Peter Coen, and Domenic Maglieri. Origins and overview of the shaped sonic boom demonstration program. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, number 2005-0005, Reno, NV, January 2005. AIAA.

[2] Siva K. Nadarajah, Antony Jameson, and Juan J. Alonso. An adjoint method for the calculation of remote sensitivities in supersonic flow. In *40th AIAA Aerospace Sciences Meeting and Exhibit*, number 2002-0261, Reno, NV, January 2002. AIAA.

[3] Marian Nemec and Michael J. Aftosmis. Parallel adjoint framework for aerodynamic shape optimization of component-based geometry. In *49th AIAA Aerospace Sciences Meeting*, number 2011-1249, Orlando, Florida, January 2011. AIAA.

[4] Charbel Farhat, Brian Argrow, Melike Nikbay, and Kurt Maute. A shape optimization methodology with f-function lobe balancing for mitigating the sonic boom. In *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, number 2002-5551, Atlanta, GA, September 2002. AIAA/ISSMO.

[5] David H. Graham, John H. Dahlin, Keith B. Meredith, and J. Vadnais. Aerodynamic design of shaped sonic boom demonstration aircraft. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, number 2005-0008, Reno, NV, January 2005. AIAA.

[6] Juan J. Alonso, Ilan M. Kroo, and Antony Jameson. Advanced algorithms for design and optimization of quiet supersonic platforms. In *40th AIAA Aerospace Sciences Meeting and Exhibit*, number 2002-0144, Reno, NV, January 2002. AIAA.

[7] G. B. Whitham. The flow pattern of a supersonic projectile. *Communications on Pure and Applied Mathematics*, 5(3):301–348, 1952.

[8] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, April 2002.

[9] Alex Haas and Ilan Kroo. A multi-shock inverse design method for low-boom supersonic aircraft. In *48th AIAA Aerospace Sciences Meeting*, number 2010-0843, Orlando, FL, January 4–7 2010. AIAA.

[10] Sriram K. Rallabhandi. Advanced sonic boom prediction using augmented burger's equation. In *49th AIAA Aerospace Sciences Meeting*, number 2011-1278, Orlando, Florida, January 2011. AIAA.

[11] Donald A. Durston. A preliminary evaluation of sonic boom extrapolation and loudness calculation methods. In *High-Speed Research: Sonic Boom*, number NASA CP 10133, Moffett Field, CA, May 12-14 1993. NASA Ames Research Center.

[12] Donald A. Durston. *NFBOOM User's Guide*. NASA Ames Research Center, Moffett Field, CA 94035, june 1992 edition, November 2000.

[13] S. S. Stevens. Perceived level of noise by mark vii and decibels (e). *The Journal of the Acoustical Society of America*, 51(2B):575–601, February 1972.

[14] R. Seebass and A. R. George. Sonic-boom minimization. *Journal of the Acoustical Society of America*, 51(2C):686–694, January 1972.

[15] A. R. George and Kenneth J. Plotkin. Sonic boom waveforms and amplitudes in a real atmosphere. *AIAA Journal*, 7(10):1978–1981, October 1969.

[16] Christine M. Darden. Minimization of sonic boom parameters in real and isothermal atmospheres. Technical Note D-7842, NASA Langley Research Center, Hampton, VA, March 1975.

[17] Natalia V. Sizov, Kenneth J. Plotkin, and Christopher M. Hobbs. Predicting transmission of shaped sonic booms into a residential house structure. *The Journal of the Acoustical Society of America*, 127(6):3347–3355, 2010.

[18] Ilan Kroo. *PASS: Program for Aircraft Synthesis Studies*. Desktop Aeronautics, Inc., Palo Alto, CA, 2.1 edition, 2011.

[19] Robert T. Jones. Theory of wing-body drag at supersonic speeds. Report 1284, National Advisory Committee for Aeronautics, 1953.

[20] Dev Gorur Rajnarayan. *Trading Risk and Performance for Engineering Design Optimization Using Multifidelity Analyses*. Ph.D. dissertation, Stanford University, Stanford, CA, June 2009.

[21] Michael J. Aftosmis, Marsha J. Berger, and John E. Melton. Robust and efficient cartesian mesh generation for component-based geometry. *AIAA Journal*, 36(6):952–960, June 1998.

[22] Marian Nemec and Michael J. Aftosmis. Adjoint sensitivity computations for an embedded-boundary cartesian mesh method. *Journal of Computational Physics*, 227:2724–2742, 2008.

[23] Marian Nemec, Michael J. Aftosmis, and Mathias Wintzer. Adjoint-based adaptive mesh refinement for complex geometries. In *46th AIAA Aerospace Sciences Meeting*, number 2008-0725, Reno, NV, January 2008. AIAA.

[24] Philip E. Gill, Walter Murray, and Michael A. Saunders. *User's Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming*. Department of Mathematics, University of California, San Diego, La Jolla, CA, February 2006.

[25] Marian Nemec and Michael J. Aftosmis. Adjoint error estimation and adaptive refinement for embedded-boundary cartesian meshes. In *18th AIAA Computational Fluid Dynamics Conference*, number 2007-4187, Miami, FL, June 2007. AIAA.

[26] Mathias Wintzer, Marian Nemec, and Michael J. Aftosmis. Adjoint-based adaptive mesh refinement for sonic boom prediction. In *26th AIAA Applied Aerodynamics Conference*, number 2008-6593, Honolulu, HI, August 2008. AIAA.

[27] Michael J. Aftosmis, Marian Nemec, and Susan E. Cliff. Adjoint-based low-boom design with Cart3D.

In *29th AIAA Applied Aerodynamics Conference*, number AIAA 2011-3500, Honolulu, HI, June 27–30 2011. AIAA.

[28] Mathias Wintzer. Span efficiency prediction using adjoint-driven mesh refinement. *Journal of Aircraft*, 47(4):1468–1471, July-August 2010.

[29] Mathias Wintzer and Ilan M. Kroo. Optimization and adjoint-based CFD for the conceptual design of low sonic boom aircraft. In *50th AIAA Aerospace Sciences Meeting*, number AIAA-2012-0963, Nashville, TN, January 9–12 2012. AIAA.

[30] Dev Rajnarayan. Monotonic cubic spline interpolation. Unpublished note, Stanford University, Stanford, CA, July 2009.

[31] Mathias Wintzer. *Optimization and Adjoint-Based CFD for the Conceptual Design of Low Sonic Boom Aircraft*. Ph.D. dissertation, Stanford University, Stanford, CA, July 2012.

[32] Brenda M. Kulfan. Universal parametric geometry representation method. *Journal of Aircraft*, 45(1):142–158, January-February 2008.