

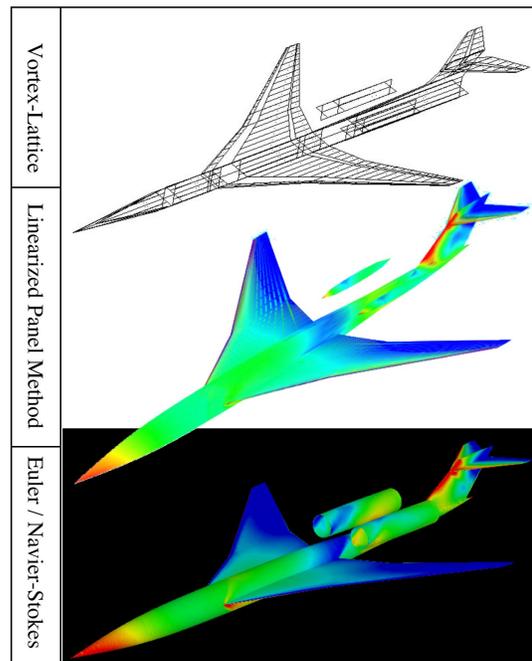


AIAA 2004-4371

**Multi-Fidelity Design Optimization of
Low-Boom Supersonic Business
Jets**

Seongim Choi, Juan J. Alonso, Ilan M. Kroo
Stanford University, Stanford, CA 94305

Mathias Wintzer
Desktop Aeronautics, Inc., Palo Alto, CA 94301



**10th AIAA/ISSMO Multidisciplinary Analysis
and Optimization Conference**

August 30 - September 1, 2004/Albany, NY

Multi-Fidelity Design Optimization of Low-Boom Supersonic Business Jets

Seongim Choi*, Juan J. Alonso†, Ilan M. Kroo‡
Stanford University, Stanford, CA 94305

Mathias Wintzer§
Desktop Aeronautics, Inc., Palo Alto, CA 94301

The practical use of high-fidelity multidisciplinary optimization techniques in low-boom supersonic business jet designs has been limited because of the high computational cost associated with CFD-based evaluations of both the performance and the loudness of the ground boom of the aircraft. This is particularly true of designs that involve the sonic boom loudness as either a cost function or a constraint because gradient-free optimization techniques may become necessary, leading to even larger numbers of function evaluations. If, in addition, the objective of the design method is to account for the performance of the aircraft throughout its mission (T/O and landing, climb, acceleration, etc.) while including important multi-disciplinary trade-offs between the relevant disciplines (performance, boom, structures, stability and control, propulsion, etc.) the situation only worsens. In order to overcome these limitations, we propose a hierarchical multi-fidelity design approach where high-fidelity models are only used where and when they are needed to correct the shortcomings of the low-fidelity models. Our design approach consists of two basic components: a multi-disciplinary aircraft synthesis tool (PASS) that uses highly-tuned low-fidelity models of all of the relevant disciplines and computes the complete mission profile of the aircraft, and a hierarchical, multi-fidelity environment for the creation of response surfaces for aerodynamic performance and sonic boom loudness (BOOM-UA) that attempts to achieve the accuracy of an Euler-based design strategy. This procedure is used to create three design alternatives for a Mach 1.6, 6-8 passenger supersonic business jet configuration with a range of 4,500 nmi and with a T/O field length that is shorter than 6,000 ft. Optimized results are obtained with much lower computational cost than the direct, high-fidelity design alternative. The validation of these design results using the high-fidelity model show very good agreement for the aircraft performance and highlights the need for improved response surface fitting techniques for the boom loudness approximations.

INTRODUCTION

SONIC booms have been the main obstacle for supersonic flight over inhabited areas. The minimization of the environmental impact from this kind of aircraft is one of the fundamental issues to be resolved to make this commercially viable supersonic flight a reality: recent studies have shown that such aircraft would have great market potential should they be allowed to fly supersonically over land.^{1,2} For these reasons, research efforts have been recently focused on various techniques for sonic boom mitigation³⁻⁷ and non-linear CFD has emerged as an essential tool owing to the requirement for accurate boom and performance information and the increasing availability of large computing resources. At the same time, it must be noted that single-minded efforts to reduce the

boom loudness typically result in aerodynamic performance shortcomings that translate into reduced range or payload, longer balanced field length, and stability & control problems that need to be resolved to make these aircraft viable.

Successful and realistic design methods must therefore carefully consider the balance between all of these performance measures, constraints, and requirements, while providing results that are sufficiently accurate to be believable. These two basic requirements embody the fundamental dilemma of high-fidelity, multi-disciplinary design: how can one produce results that are highly accurate in a reasonable time with limited resources?⁸⁻¹¹ Faced with these problems, the alternative of using approximation models of the actual analyses has received increasing attention in recent years.^{9,12} A number of different techniques such as polynomial response surfaces, radial basis functions, Kriging approximations, and piecewise polynomial methods have emerged.

It is widely accepted, however, that not every por-

*Doctoral Candidate, AIAA Member

†Associate Professor, AIAA Member

‡Professor, AIAA Fellow

§Engineer, AIAA Member

Copyright © 2004 by the authors. Published by the American Institute of Aeronautics and Astronautics, Inc. with permission.

tion of the multi-disciplinary analysis of these aircraft must be carried out with high-fidelity: simpler models can often provide very good approximations to such problems as stability & control, propulsion, and basic estimations of aerodynamic performance. In situations where high-fidelity methods are necessary, one may actually benefit from the combined use of low- and high-fidelity models in order to obtain identical solutions to those provided by the high-fidelity model, but with a lower computational cost.^{13,14} This is the approach that we have followed in the work in this paper: a combination of low-fidelity tools that are enhanced by multi-fidelity analyses in the areas where the increase in accuracy is needed. Our work, however, is based on the inexpensive construction of response surfaces using multi-fidelity tools, rather than on using the various fidelity models directly in the optimization process.

In this work, we combine ideas of multi-fidelity analysis and design into a hybrid concept that includes:

1. The Program for Aircraft Synthesis Studies (PASS): a multi-disciplinary design tool that incorporates carefully tuned fast models for the various disciplines in the design and is able to deal with all the major objective functions and constraints in typical aircraft synthesis problems.
2. A hierarchical, multi-fidelity response surface generation technique that uses results from classical supersonic aerodynamics, a linearized supersonic panel code (A502/Panair), and two different levels of mesh resolution in an unstructured adaptive Euler solver (AirplanePlus) to create models of the aerodynamic performance and boom loudness for use with PASS.
3. Automated tools based on a common geometry database to drive the analysis tools that are used in the generation of the response surfaces in this problem (BOOM-UA). This CAD-to-solution procedure is based on the CAPRI CAD-interface, the A502/Panair and AirplanePlus flow solvers, the Centaur mesh generation system, and the PC-BOOM software for propagation of the acoustic signatures.

Our previous work in high-fidelity supersonic design had focused on the cruise condition alone.^{3,15-17} Although recently we had included some basic constraints to handle other points in the mission,¹⁵ these were only surrogates for the real constraints that must be imposed for realistic designs to be produced. A major component of this work is the addition of PASS to formally include all of the necessary constraints.

Even when high-fidelity analyses are considered for the cruise condition alone, the design problem is made significantly more difficult since the design space for

the sonic boom objective has been shown to exhibit multiple local minima and to be rather noisy and even discontinuous.¹⁸ These characteristics of the design problem rule out the possibility of using gradient-based optimization (and the powerful adjoint method¹⁹⁻²²) for the boom portion of the problem and impose strict requirements on the response surface techniques used to model the high-fidelity behavior of the aircraft.

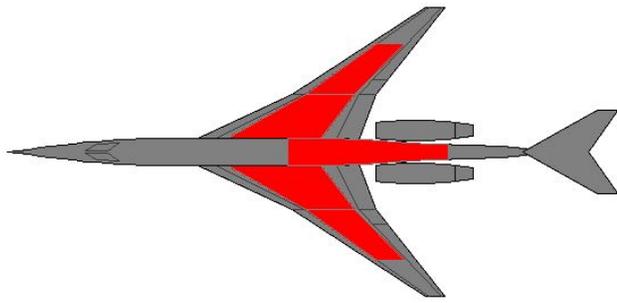
In the following sections we describe the various components of the design method that we have created. We start with the description of the PASS tool for conceptual multi-disciplinary design, its capabilities and the optimization algorithm used. We then provide details of the automated high-fidelity analyses for both the linearized supersonic panel code (A502/Panair) and the Euler / Navier-Stokes solver AirplanePlus, including all of the pre- and post-processing modules required to produce the necessary information. Following that description we detail the multi-fidelity approach that we have followed to generate response surfaces for the coefficient of drag of the aircraft, C_D , and the boom loudness (in dBA) using the hierarchy of flow solution techniques mentioned earlier. We finally conclude by presenting the results of several optimizations for a supersonic jet flying at $M_\infty = 1.6$, with a range of 4,500 nmi, and with a take-off field length no greater than 6,000 ft. High-fidelity validations of these results are also shown.

PASS CONCEPTUAL DESIGN TOOL

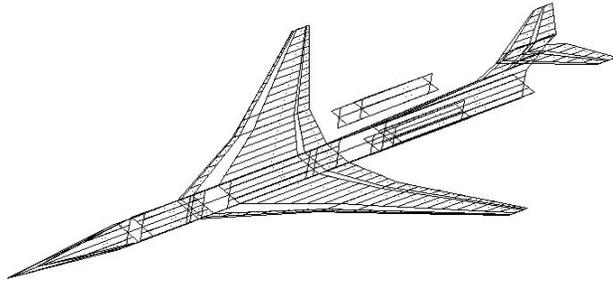
PASS (Program for Aircraft Synthesis Studies), an aircraft preliminary design tool created by Desktop Aeronautics, Inc., was used to generate all of the designs presented in this work. PASS uses a simplex method and an integrated set of predictive modules for all of the relevant disciplines in the design (including mission performance) to generate optimized designs that satisfy a number of imposed constraints.

PASS uses a graphical user interface to explore the results of each of the participating disciplines in a design. In addition, the same interface can be used to define the design optimization problem. Design variables, objective functions and constraints can be setup using any of the relevant parameters and functions that are used in each of the disciplinary modules. A view of two of the various aircraft models (fuel tank arrangement and vortex lattice mesh) that are used by PASS can be seen in Figure 1 below.

Incorporating PASS into the analysis allowed for the evaluation of all aspects of mission performance, thus providing a balanced configuration not just limited to meeting some singular performance goal, but also capable of achieving field length, climb gradient, and cabin constraints (for example) required for a realistic aircraft design. Some of the most relevant capabilities of PASS for this work are briefly summarized below.



a) Fuel tank layout.



b) Vortex-lattice model.

Fig. 1 Different views of the configuration within PASS.

Drag Estimation

Lift- and volume-dependent wave drag, induced drag and viscous drag are evaluated at key mission points. Inviscid drag is estimated using linearized methods. The viscous drag computation is sensitive to Reynolds number and Mach number, and is based on an experimentally derived fit. Special attention is paid to transonic drag rise, with numerous points being sampled up to and through Mach 1. The analysis detail is of a level that allows configuration tailoring to minimize drag during supersonic cruise (i.e. the use of the area rule is contemplated.)

Weights and CG

Component weights are based on available data for modern business-jet class aircraft. Wing weight is estimated based on a bending index that is related to the fully stressed bending weight of the wing box, coupled with a statistical correlation. The weights of tail surfaces are similarly determined. Fuselage weight is based on both the gross fuselage wetted area and the use of a pressure-bending load parameter.

The CG location is computed based on typical placements and weights of the various aircraft components; CG movement during the mission due to fuel burn is also computed based on the fuel tank layout and the ability to transfer fuel between tanks is also used to aid in the trimming of the configuration throughout its mission.

Propulsion

Engines are typically modeled by sampling a manufacturer's deck at numerous Mach numbers and altitudes and constructing a fit. For this study, a generic deck was created and hand-tuned to give performance of a level achievable by available, mature technology, low-bypass turbofan engines.

Low-Speed Analysis

Low-speed stability and trim are computed using a discrete-vortex-lattice method. This data is then used to predict such things as the balanced field length (BFL) for the aircraft, stability derivatives and estimates for tail incidences at critical low-speed points (take-off rotation, for example.)

Mission Analysis

The mission analysis routine ties together all the various tools in PASS to run an aircraft through a typical flight and evaluate its overall performance. The key points analyzed are the takeoff run, takeoff rotation, 2nd segment climb, subsonic climb to acceleration altitude, subsonic-to-supersonic acceleration, supersonic climb to initial cruise altitude, cruise and landing. In our work in this paper, only the cruise condition benefits from enhanced computations for the aircraft performance.

Optimization

PASS provides a non-gradient based optimizer for configuration studies based on the Nelder-Mead simplex method. Given some variables, the optimizer will minimize an objective function subject to constraints. The variables, constraints and objective are all user-defined. Typically, the optimizer will be tied to the mission analysis computation. Constraints usually consist of performance goals such as range and balanced field length. Additional constraints to ensure a viable aircraft in the eyes of the FAA may also be imposed, to ensure, for instance, that the aircraft will climb out at the minimum 2.4% gradient stipulated by FAR regulations. Details of the optimization problem formulation are discussed in following sections.

Design Configurations

The baseline design, from which all work started, was created using the default PASS analysis modules which are based on classical supersonic aerodynamics and vortex-lattice methods. Subsequent designs are also created by PASS with two significant differences: 1) the inviscid aerodynamic drag prediction module was replaced by the response surface fits created using our multi-fidelity approach, and 2) a ground boom calculation module was added (also based on a multi-fidelity response surface fit) for addition of loudness-related objective functions or constraints.

It is worth noting that in the designs presented in this paper, no assumptions of future technology

have been made. All designs use models of existing propulsion plants, materials, and systems that can be incorporated into an actual design today.

AERODYNAMIC ANALYSIS : BOOM-UA

All of the necessary modules to carry out multi-fidelity aerodynamic analyses and ground boom signature computations are integrated into our multi-fidelity analysis tool, BOOM-UA. The current version of BOOM-UA is an evolution of our previous work¹⁵ that incorporates the ability to choose between two different aerodynamic solvers of different fidelities: the linearized panel code A502/Panair, and the Euler/Navier-Stokes flow solver AirplanePlus. Every other portion of the tool chain remains the same as before. The use of this integrated analysis tool, guarantees that both the geometry and the boom propagation tool used on multi-fidelity computations are the same. The differences in the results of the analysis of the same configuration using the alternate flow solver modules are solely due to the difference in the flow predictions between A502/Panair and AirplanePlus.

The complete procedure is as follows. Firstly, a parameterized geometry is represented directly with a CAD package (ProEngineer in this study) using a collection of surface patches. These surface patches can be used directly with A502/Panair or can serve as the geometric description for an unstructured tetrahedral mesh generated automatically by the Centaur software. Our geometry kernel, AEROSURF, generates multiple variations of this baseline configuration as required by the response surface construction tool. If the changes in the geometry are small enough (for the Euler solver) we can perturb the baseline mesh to conform to the deformed shape without problems with decreased mesh quality and/or edge crossings. If this is not the case, the mesh can be automatically regenerated to accommodate large changes in the geometry. Either our Euler solver, AirplanePlus or the linearized panel code A502/Panair calculates the surface pressure distributions and predicts both the C_L and C_D and the near-field pressures which can then be propagated to obtain ground boom signatures. A solution-adaptive mesh refinement procedure is used (when appropriate) to generate refined meshes adapted according to a criterion based on the pressure gradient information obtained by the flow solver. The boom prediction software PCBoom3 is used for the propagation portion of the solution procedure. Three-dimensional near-field pressure distributions are extracted on a cylindrical surface several body lengths beneath the configuration and are provided to the boom propagation tool. Figure 2 shows a brief schematic of all the processes that have been integrated into BOOM-UA. In this Figure, n refers to the number of design points. Each individual com-

ponent module is explained in detail in the following subsections.

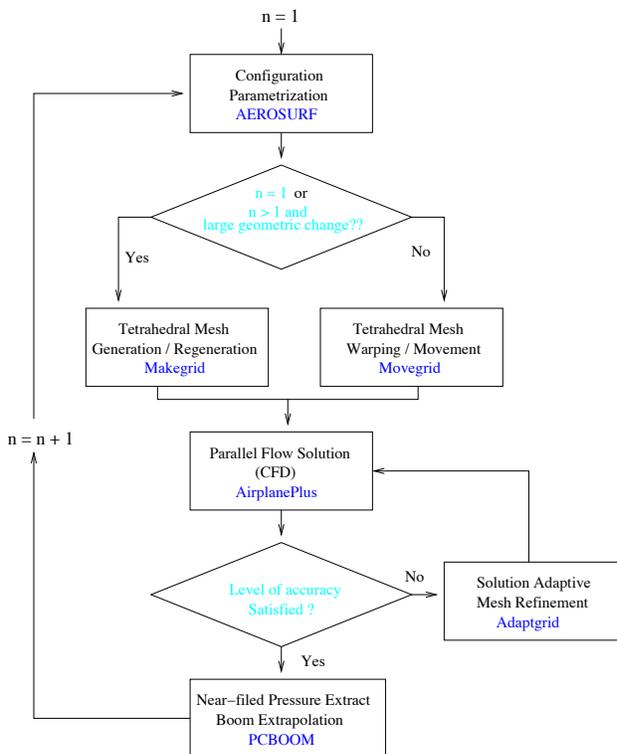


Fig. 2 Schematic of the aerodynamic analysis tool, BOOM-UA

CAD Geometry Representation

High-fidelity MDO requires a consistent high-fidelity geometry representation. In general, the geometric shape of an aircraft can be defined by an appropriate parameterization of the geometry. This parametric geometry kernel is available to all of the participating disciplines in the design so that both cost functions and constraints can be computed using the same geometry representation.

In our work, a CAD-based geometry kernel is used to provide this underlying geometry representation. Baseline shapes are developed in a CAD-package (ProE, in our case) and constitute our parametric master model: they uniquely define the parameterization of a configuration given a particular design intent.

Our geometry kernel (AEROSURF) is coupled with the parametric CAD description through the CAPRI interface of Haimes,^{23,24} such that it automatically generates watertight surface geometry patches. AEROSURF can be executed in parallel and uses a distributed geometry server to expedite the generation of a large number of different design alternatives, thus reducing the cost of running geometry re-generations in the CAD package. Using a master/slave approach and PVM for distributed computing, arbitrary numbers of slaves can be started simultaneously while a master program maintains a queue of geometry regen-

eration requests and keeps slaves busy doing CAD regenerations. Figure 3 shows a representative aircraft with 46 surface patches which is generated directly in ProEngineer by providing the values of 108 design variables.

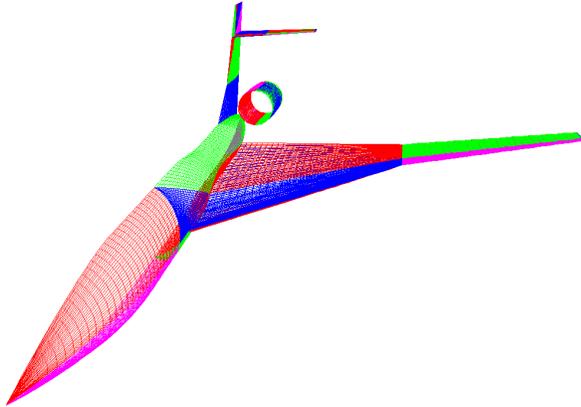


Fig. 3 Geometry representation by parametric, CAD-interfaced AEROSURF

Tetrahedral Unstructured Mesh Generation and Euler Flow Solution Approach

The high-fidelity portion of this work focuses on the use of unstructured tetrahedral meshes for the solution of the Euler equations around complete aircraft configurations. The Centaur²⁵ software is directly linked with the surface representation obtained from AEROSURF, and is used to construct meshes for aircraft configurations and to enhance grid quality through automatic post-processing. Only fine meshes need to be explicitly constructed since our multigrid algorithm is based on the concept of agglomeration and, therefore, coarser meshes are obtained automatically. Figure 4 shows a triangular mesh on the body surface and the symmetry plane of our configuration. For visualization purposes only, a coarsened grid is shown.

The three-dimensional, unstructured, tetrahedral AirplanePlus flow solver is used in this work. AirplanePlus is a C++ solver written by Van der Weide which uses an agglomeration multigrid strategy to speed up convergence. A modified Runge-Kutta time-stepping procedure with appropriately tailored coefficients is used to allow for high CFL numbers. Several options for artificial dissipation and the block-Jacobi preconditioning method are all available in the solver. The AirplanePlus solver is a tetrahedral unstructured flow solver loosely based on the ideas of the original AIRPLANE code.²⁶

The solution adaptation procedure that we use in BOOM-UA has been described previously and will not be repeated here. For more details, the reader is re-

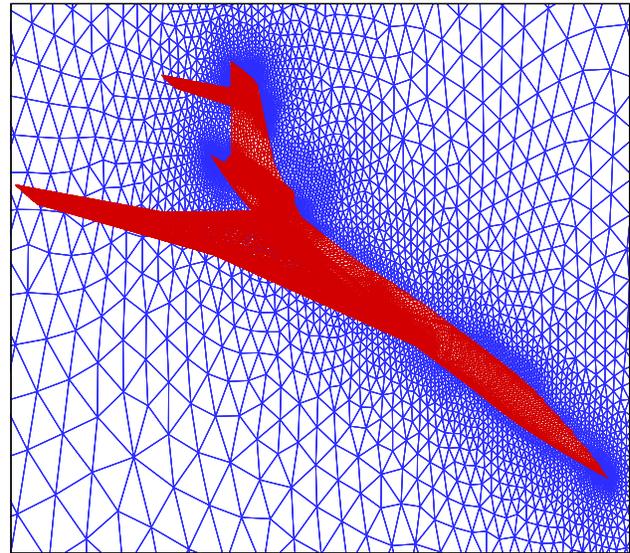


Fig. 4 Unstructured tetrahedral surface mesh around a low-boom aircraft

ferred to.¹⁵ Figure 5 shows a view of a typical mesh before and after two separate adaption cycles. The adaptive mesh capability allows for the efficient and accurate capturing of the near-field pressure distributions that are used as an input to the PCBoom3 for propagation. The alternative of a uniformly refined mesh all the way out to the near-field is far too expensive to be considered.

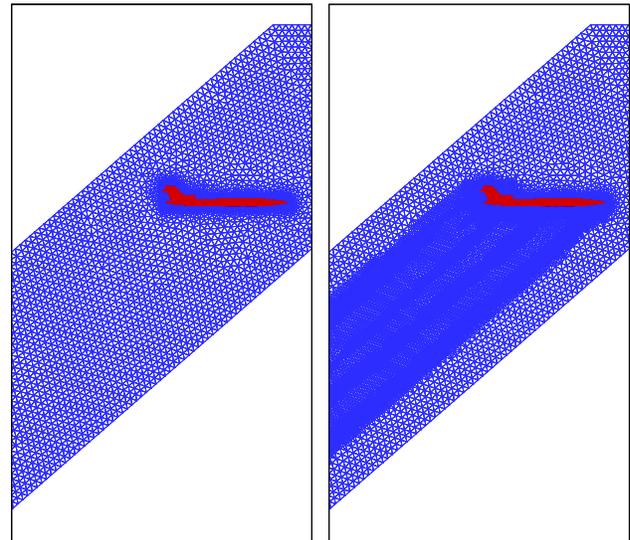


Fig. 5 Mesh before and after two adaption cycles

Three Dimensional Linearized Panel Method, A502/Panair

The A502 solver, also known as Panair,^{6,27} is a flow solver developed at Boeing to compute the aerodynamic properties of arbitrary aircraft configurations flying at either subsonic or supersonic speeds. This code uses a higher-order (quadratic doublet, linear source) panel method, based on the solution of the linearized potential flow boundary-value problem. Re-

sults are generally valid for cases that satisfy the assumptions of linearized potential flow theory - small disturbance, not transonic, irrotational flow and negligible viscous effects. Once the solution is found for the aerodynamic properties on the surface of the aircraft, A502 can then easily calculate the flow properties at any location in the flow field, hence obtaining the near field pressure signature needed for sonic boom prediction. In keeping with the axisymmetric assumption of sonic boom theory, the near field pressure can be obtained at arbitrary distances below the aircraft.²⁸

Since we are using two different flow solver modules for our multi-fidelity response surface fitting tool, it is important to be aware of the similarities and differences in the solutions provided by each of the flow solvers. A detailed study of these differences in the solution for a varying free stream Mach numbers (that put the normal Mach number in the transonic region) is presented in a later section and highlights the areas of the flow where A502/Panair is not usable.

Ground Boom Propagation

The basic strategy for the computation of ground boom signatures can be seen in Figure 6 below. At the near-field plane location, the pressure signature created by the aircraft is extracted and it is propagated down to the ground using extrapolation methods based on geometric acoustics.

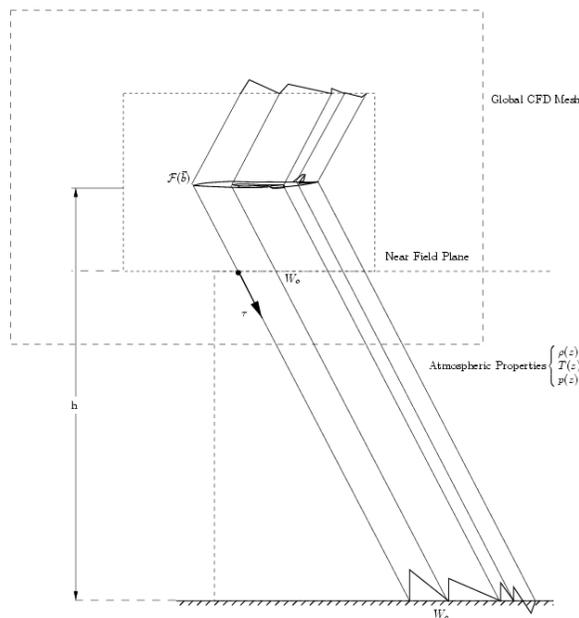


Fig. 6 Schematic of sonic boom propagation procedure.

The location of the near-field must be far enough from the aircraft so that the near-field flow field is nearly axisymmetric and there are no remaining diffraction effects which cannot be handled by the extrapolation scheme. Since A502/Panair only uses a surface mesh for all of its calculations, it is able to ob-

tain near-field pressures at arbitrary distances without changes in the computational cost. For Euler calculations, however, the larger the distance to the near-field, the larger the mesh required to compute the pressure distributions accurately. For these reasons, studies have been conducted in the past to establish the correct distances from the aircraft that are required to obtain accurate ground boom propagation. For the computations presented here, we extract near-field information at a distance of 1.2 body lengths beneath the aircraft.

In this work, we are using both the Sboom²⁹ and PCBoom3³⁰ extrapolation methods to propagate near-field signatures into ground booms. Although the PCBoom3 software is far more capable than Sboom (it is essentially a superset of it), we have only computed ground booms created by the aircraft in a steady-state cruise condition and, therefore, both codes are nearly equivalent. If ground booms caused by maneuvering aircraft were to be computed, the capabilities of the PCBoom3 software would have to be used.

The two sonic boom extrapolation methods account for vertical gradients of atmospheric properties and for stratified winds (the winds have been set to zero in this work.) Both methods essentially rely on results from geometric acoustics for the evolution of the wave amplitude, and both utilize isentropic wave theory to account for nonlinear waveform distortion due to atmospheric density gradients and stratified winds.

Our earlier research on low-boom aircraft design was mainly focused on the reduction of the magnitude of only the initial peak of the ground boom signature.^{3,5} This requirement, which had been suggested as the goal of the DARPA-sponsored Quiet Supersonic Platform (QSP) program ($\Delta p_0 < 0.3$ psf), hides the importance of the rest of the signature, which often arises from the more geometrically complex aft portion of the aircraft where empennage, engine nacelles and diverters create more complicated flow patterns. Moreover, such designs often have two shock waves very closely following each other in the front portion of the signature,^{6,18} a behavior that is not robust and is therefore undesirable.

For this reason, we have chosen to base our designs on the perceived loudness of the complete signature (dbA). Frequency weighting methods are used due to the unique property of the human hearing system which doesn't have an equal response to sounds of different frequencies. In these calculations, less weighting is given to the frequencies to which the ear is less sensitive. In addition, all signatures computed are post-processed to add a physical rise-time across the shock waves that yield loudness numbers that are more representative of those perceived in reality.

MULTI-FIDELITY RESPONSE SURFACE GENERATION

Now that we have described the two basic tools used to create supersonic designs, PASS and BOOM-UA, it is necessary to explain the procedure we have used to integrate them into a single analysis and optimization capability. The concept is straightforward: if the multi-fidelity analysis capability can be used to create response surfaces for the drag coefficient, C_D , and the ground boom loudness, the corresponding low-fidelity modules in PASS can be replaced by these response surface fits. This makes for a remarkably simple integration problem. PASS can then be used to generate optimized results and the outcome of the optimization can be analyzed using the high-fidelity tools to ensure that the response surface fits provide accurate representations of the true high-fidelity responses.

Our multi-fidelity approach to the construction of the response surface fits relies on a hierarchy of four different aerodynamic analysis modules

1. PASS internal analysis based on classical aerodynamics.
2. A502/Panair supersonic linearized panel code.
3. Euler solutions using a relatively coarse mesh (around 250,000 nodes for the complete configuration) - we refer to these computations by the label "Coarse Euler (CE)."
4. Euler solutions of the highest fidelity using several cycles of unstructured mesh adaptation (with a total of around 3-5 million nodes for the complete configuration) - we refer to these computations by the label "Fine Euler (FE)."

In order to obtain response surface fits of the highest fidelity one could carry out a large number of FE solutions and fit the resulting data (for both C_D and boom loudness). Unfortunately, for large dimensional design spaces (we will be using 17 design variables later on), accurate fits require a large number of function evaluations. This is particularly true in our case since the ranges of variation of each of the design variables will be rather large.

The main objective in this section is to generate response surface fits of the same quality/accuracy that would be obtained by evaluating the FE solutions only, but at a much reduced cost. We accomplish this by relying on a fundamental hypothesis that will be tested later on: *the higher fidelity tools are only needed in small regions of the design space where the lower fidelity models have exhausted their range of applicability.* This is bound to be true as it is the premise upon which aerodynamic design has been predicated for the last 50 years: aerodynamicists and engineers use the fastest tools for a specific purpose (when they

are known to work well) and switch to more time-consuming, expensive tools only when they are needed. For example, in supersonic design, classical equivalent area concepts and linearized panel codes can provide very accurate results as long as non-linear effects (such as transonic flows in the direction normal to the leading edge of the wing) are not present and viscosity does not play a dominant role in the solution of the flow.

With this in mind, we have used the following five-step procedure to create the response surfaces used in this work. All databases of candidate designs are obtained by populating the design space using a Latin Hypercube Sampling (LHS) technique.

1. Run a large database of candidate designs (i 8,000) using the PASS aerodynamics module. Each evaluation takes roughly 1 second to compute on a modern workstation (Pentium 4, 3.2 GHz). This evaluation also flies each aircraft through the mission and returns a measure of the infeasibility of the design (an L-2 norm of the constraint violations.) Those designs that are found to significantly violate the requirements/constraints of the mission are removed from the database and are no longer considered in the response surface creation.
2. Run the remaining database of candidate designs using the A502/Panair solver. Each evaluation requires about 10 seconds of CPU time on the same modern workstation.
3. Select the design points whose relative error for C_D (based on the baseline design) is larger than a specified threshold, $\epsilon_{PASS-A502}$, and analyze only those designs using the Coarse Euler (CE) approach. In our work, we have set this threshold to about 30% resulting in a number of CE evaluations in the neighborhood of 1000. Each CE evaluation requires about 7 min on a modern workstation (5 minutes for the mesh generation and 2 minutes on 8 processors for the flow solution.)
4. We finally select the design points whose relative error for C_D (again based on the the baseline design) is larger than $\epsilon_{CE-A502}$ and analyze those designs using the Fine Euler (FE) approach. This threshold was set to 25% in our work, resulting in approximately 500 FE solutions. Each FE evaluation, from beginning to end, including geometry and mesh generation and adaptation (the bottlenecks in the process, since they are run serially) requires about 30 minutes of wall clock time. The flow solution portions (using AirplanePlus) are run in parallel using 16 Athlon AMD2100+ processors of a Linux Beowulf cluster.
5. Baseline quadratic response surface fits (using least squares regression) are created for the C_D

and boom results for A502/Panair. The error between the FE evaluations and the predictions of these quadratic fits is approximated with a Kriging method, and the resulting approximation is added to the baseline quadratic fits.

In sum, the response surfaces provided to PASS are the addition of the quadratic fits based on the A502/Panair results and the Kriging fits to the error between the FE solutions and those quadratic fits.

Figure 7 shows the result of the over 5,000 candidate designs evaluated using A502/Panair that are retained after the initial filtering of over 8,000 PASS results. The red dots in the Figure indicate those candidate designs for which the predicted values of C_D are off by more than $\epsilon_{PASS-A502} > 30\%$ between PASS and A502. Note that a number of these red dots have unreasonably large values of C_D since the geometries and design conditions are such that the limits of applicability of A502 are exceeded. These points for which the disagreement between PASS and A502 is large are taken for further evaluation using CE. Figure 8 shows in blue the results of the CE analyses for the red A502 results. Note that no attempt to correct the boom loudness (dBA) is made since the CE results do not have sufficient mesh resolution in the near field to yield accurate boom predictions. This same procedure is repeated once more but increasing the level of fidelity by one step. The points in the CE database that show large errors in comparison with the A502/Panair results are re-evaluated using the FE approach. The final result is a set of FE evaluations that are meant to be clustered around the areas where the lower fidelity models cannot accurately predict the flow physics. These results can be seen in Figure 9 where the blue and yellow dots represent the CE evaluations. The yellow dots represent the CE candidate designs that exhibit large errors ($> 30\%$) in C_D when compared with the results of A502/Panair. The designs corresponding to the yellow dots (around 500 in total) are evaluated using the FE procedure and their C_D and dBA values are corrected and represented with the purple dots. Notice that in this last step, the values of the boom loudness for the points representing the FE results have been corrected: the FE results are adaptively refined in the near field so that accurate near-field pressure distributions and ground booms can be obtained.

This multi-fidelity procedure has, to some extent, the flavor of Richardson's extrapolation in that it recursively uses results from different fidelities to arrive at a final answer/fit. It also has an adaptive nature to it, as results from the higher fidelity models are only evaluated in areas of the design space where the lower fidelity models are found to be insufficiently accurate. If the hierarchy of models is chosen in such a way that the areas where the lower fidelity models fail are small compared with the size of the design space,

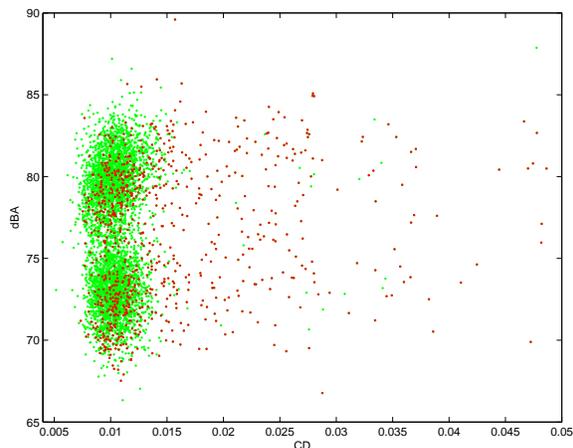


Fig. 7 Database of A502/Panair results.

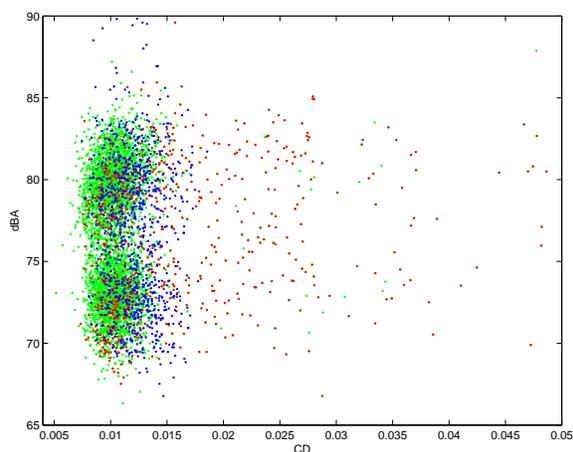


Fig. 8 A502/Panair and CE results.

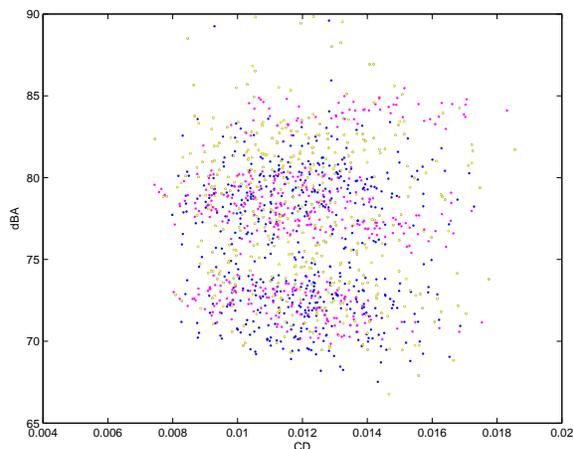


Fig. 9 CE and FE results.

then the procedure described above should be quite effective in producing results that are of nearly high-fidelity over the *entire* design space. Our experience shows that this is the case for aerodynamic performance: the PASS aerodynamic module is quite good at predicting the absolutely best wing (lower bound estimate on the C_D) that could be produced if considerable design work were done on the configuration (potentially using adjoint methods and a high degree of shape parameterization). However, it is unable to predict some of the finer details of aerodynamic performance and certainly fails when transonic effects are present. A502/Panair is also unable to deal with transonic flow effects (as was shown in a previous section) but produces more realistic results than the PASS analysis as the actual geometry of the configuration is truly accounted for. Finally, the Euler models (CE and FE) are quite good predictors of the aerodynamic performance of the complete aircraft as long as viscous effects are not dominant. It must be mentioned that, if sonic boom were not an issue in these designs, the CE evaluations would be sufficient as the differences in C_D between CE and FE are found to be insignificant (less than 5 counts) over the large range of variations pursued in this work.

RESULTS

In this section we first present some results related to the validation of the various solver components of our design framework, followed by the use of the multi-fidelity approximations for actual designs carried out with PASS.

Comparison of Results: A502/Panair vs. AirplanePlus

The accuracy and efficiency of our design procedure is predicated on the fact that A502/Panair can be assumed to provide accurate information in large regions of the design space. It is expected to fail in regions where non-linearities (such as transonic flows and shock waves) are present. To assess the validity of this claim, we carried out the following computational experiments. For the baseline configuration from our previous work,¹⁵ with an inboard wing leading edge sweep angle of 61° degrees, we ran both A502/Panair and AirplanePlus at the following free stream Mach numbers; 1.76, 1.863, 1.9665, and 2.07, corresponding to normal Mach numbers of 0.8533, 0.9032, 0.9534, and 1.0 respectively. All computations were carried out at a fixed $C_L = 0.10$. Figure 10 shows the computed near-field pressure signatures from both A502/Panair and Euler computations for these four cases. As can be seen, as the normal Mach number enters the region where transonic flow effects are significant, the predictions of A502/Panair start to deteriorate to the point where they become unusable. However, one must note that even for normal Mach numbers of 0.90, the predictions from A502 are quite accurate. A similar

Cruise Mach	1.6
Range	4,500 nmi
BFL	6,000 ft
Minimum static margin	0.0
Alpha limit	15°
MTOW	96,876 lbs

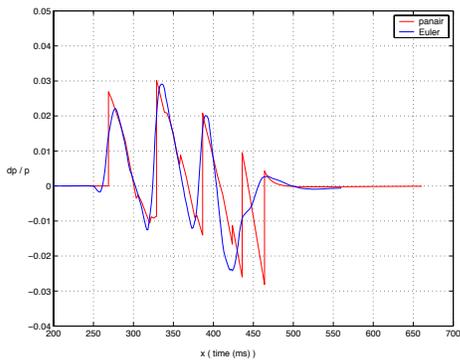
Table 1 Performance requirements for optimized baseline configuration.

trend is seen for the C_D of the configuration: the agreement is fairly good at the first two flow conditions (with a maximum relative error of only 5.5 %) but it quickly worsens (the relative error is 20 % at the last flow condition). This experiment, together with additional experience not reported here shows that it is only in this very narrow region of the design space that A502/Panair does not deliver useful results. This is important to ensure that our multi-fidelity approach can create high-fidelity information at fairly low cost. Initially one may think that if the transonic flow regime is avoided altogether, then A502/Panair should be sufficient to accomplish all design tasks related to aerodynamic performance. Unfortunately, the results in later sections point out that the optimizer appears to want to arrive at wing geometries that operate in this regime and, therefore, the higher fidelity models are necessary to obtain believable designs.

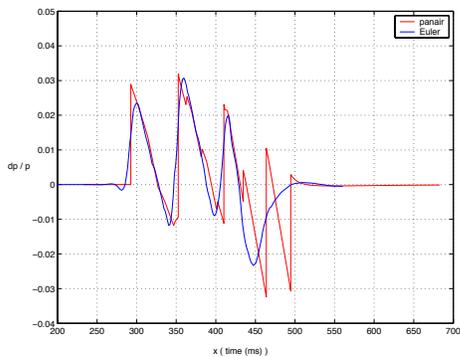
Baseline Configuration: Standard PASS Optimization

For subsequent comparisons, an optimized baseline geometry was generated by running the standard version of PASS for a mission with the performance objectives summarized in Table 1. Mission requirements and geometric constraints for the baseline configuration were based on numbers that were felt to be representative of current industry interest. As mentioned before, in an effort to generate an aircraft achievable using current levels of technology, advanced technology assumptions were kept to a minimum.

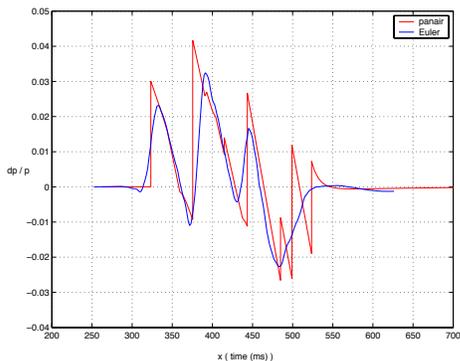
The values of the design variables of the resulting baseline configuration (which are also used in subsequent designs) are provided in Table 2. Note that the values highlighted in red were not allowed to vary during the optimizations. In addition to these variables, 6 variables representing the radii of fuselage stations located at 5%, 10%, 15%, 62.5%, 75%, and 87.5% of the fuselage length were added to allow for both performance and boom tailoring, and to maintain cabin and cockpit compartment constraints. The outer mold line of the fuselage was generated by fitting an Akima spline to the specified radius distribution. Note that the allowable ranges for all of the design variables (for this baseline configuration and all subsequent designs) were rather large, being at least $\pm 30 - 40\%$ of their baseline values. This large range of variations allows for a more complete design space to be searched but



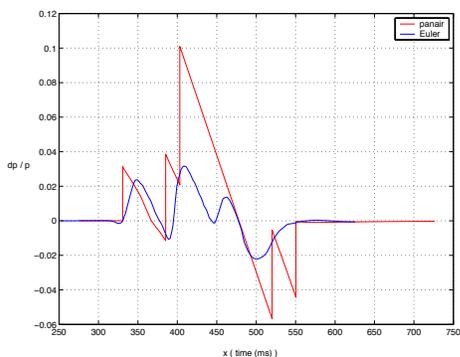
a) $M_\infty = 1.76, M_n = 0.8533$



b) $M_\infty = 1.86, M_n = 0.9032$



c) $M_\infty = 1.97, M_n = 0.9534$



d) $M_\infty = 2.07, M_n = 1.0036$

Fig. 10 Near-field pressure comparisons between A502/Panair and Euler computations for changing free stream Mach numbers and corresponding normal Mach numbers.

Wing and Tail Geometry

Wing reference area (S_{ref})	1,078 ft ²
Wing aspect ratio (AR)	4.0
Wing quarter-chord sweep (Λ)	53.35°
Wing taper	0.15
Wing dihedral	3°
Leading edge extension	0.278
Trailing edge extension	0.197
Break location	0.4
Location of wing root LE	0.294
Root section t/c	2.5%
Break section t/c	3.0%
Tip section t/c	2.5%
Vertical tail area (% S_{ref})	0.125
Vertical tail AR	0.65
Vertical tail Λ	56°
Vertical tail λ	0.6
Horizontal tail area (% S_{ref})	0.6
Horizontal tail AR	2.0
Horizontal tail Λ	56°
Horizontal tail λ	0.3

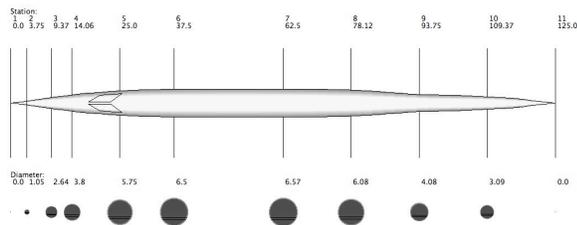
Fuselage Geometry

Maximum fuselage length	125 ft
Minimum cockpit diameter	60 inches
Minimum cabin diameter	78 inches
Cabin length	25 ft

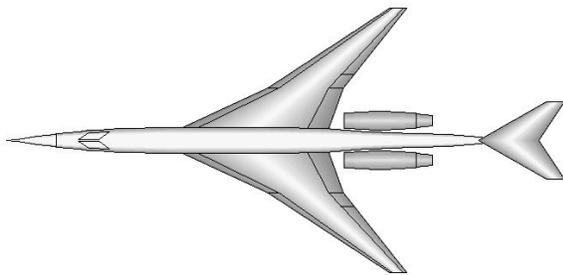
Table 2 Geometric design variables for design optimization and values for baseline design.

also makes the job of both the optimization algorithm and the response surface fitting techniques more complicated. The values of the leading and trailing edge extensions are normalized by the trapezoidal wing root chord. The location of the wing root leading edge is normalized by the fuselage length and is measured from the leading edge of the fuselage. Both the vertical and horizontal tail areas are normalized by S_{ref} .

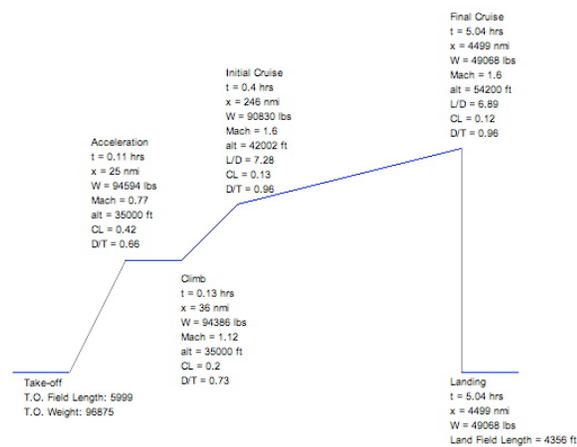
This baseline geometry, according to PASS, meets all of the requirements of the mission specified in Table 1 as can be seen in the mission profile in Figure 11. Note that no attempt to minimize or tailor the ground boom signature has been made in this design. Furthermore, this design can be considered a *best-case* scenario since it was calculated using the aerodynamic prediction module in PASS, which computes coefficients of drag that can be assumed to represent a lowest bound estimate. In some senses, two of the main points of this paper are to, first, see how close we can get to the performance of this baseline design when using the higher-fidelity models and, second, how much additional work is required to obtain such a design. In subsequent designs, unless otherwise specified, the boom loudness is also taken into account, which complicates the job of the optimizer even further.



a) Fuselage layout and stations



b) Top view of configuration



c) Mission profile

Fig. 11 Summary of baseline configuration.

PASS Optimizations Using Response Surface Fits

Changes were made as necessary to the PASS code to incorporate the C_D and boom loudness (dBA for ground signature with rise time added) fits.

Before proceeding to the presentation of design calculations, it is worth re-evaluating (analysis only) the baseline configuration just presented. If the C_D information provided by the response surface fits is identical to the information provided by the PASS model, then the same aircraft ought to be able to fly the mission and complete it while satisfying all of the constraints that were imposed on the original baseline design.

PASS, however, makes certain theoretically realistic assumptions related to the shape of the spanwise and longitudinal lift distributions when computing inviscid drag. Though achievable in principle, these estimates for best possible performance can in practice only be attained with a significant CFD design effort. Given

Range	3,650 n.mi.
BFL	5,482 ft
MTOGW	92,018 lbs
Ground signature (dBA)	86.15

Table 3 Results of PASS + response surface analysis of the baseline configuration.

that the baseline geometry and all of its perturbations were evaluated as is, without any additional design work, it perhaps should not be surprising that the drag reported by the response surface fit to the CFD evaluations fit was substantially higher than that reported by PASS for the baseline case. The aerodynamic module in PASS reports an inviscid drag coefficient at the beginning of the cruise segment of $C_{D_{PASS}} = 0.0101$, while that of the response surface is, as expected, substantially higher, $C_{D_{RS}} = 0.0121$.

The baseline case was evaluated with the modified version of PASS. Even with optimized values for TOGW and initial and final cruise altitudes, it was unable to make the revised range requirement, far less the original goal. The results of this analysis using the information from the response surface fits is presented in Table 3. Notice also that the predicted loudness (from the FE analysis) is 86.15 dBA.

Although the C_D predictions are rather far off, both the near-field pressure distribution and the ground boom signatures produced by A502/Panair and Euler are rather similar, as can be seen in Figure 12. A graphical view of the mission and its details can be seen in Figure 13.

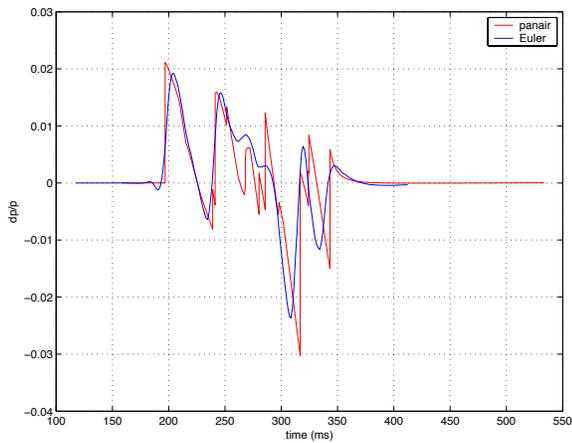
As a result of the significant drag increase, performance goals were reduced such that meaningful optimizations could be run. The changes consisted of a reduction of the range goal from 4,500 n.mi. to 4,000 n.mi. and an increase in the TOFL/LFL from 6,000 ft to 6,500 ft.

For subsequent optimization runs, the optimizer was configured such that the input variables and associated bounds mirrored the perturbation values used in generating the fits. These inputs and bounds are summarized in Table 7. The design constraints presented in Table 8 were also imposed.

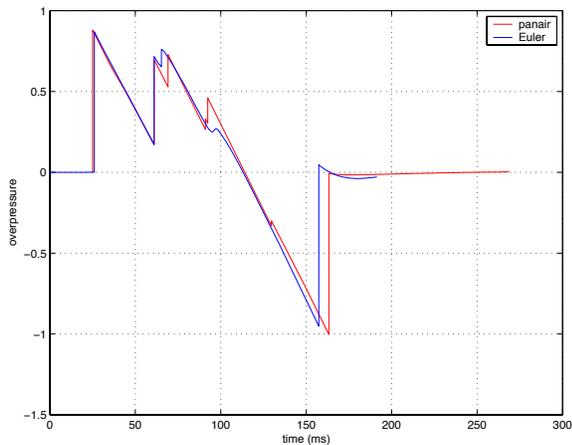
Minimum Boom Loudness (dBA) Optimization

The minimum boom optimization used the response surface fit value for sonic boom strength (measured as dBA with finite rise time) as the objective. Although a seemingly viable design was found, as shown in the summary Table 4, a glance at the planform view in Figure 14 indicates that the optimizer is exploiting a weakness in the wing weight routine in an effort to drive the objective value down. It is worth calling attention to the interesting fuselage shaping that has been converged upon. A plot of the iteration history for this case is shown in the Figure 15.

Figure 16 shows the change (from the baseline de-



a) Near-field pressure distribution



b) Ground boom signature

Fig. 12 A502/Panair vs. Euler comparison for baseline design.

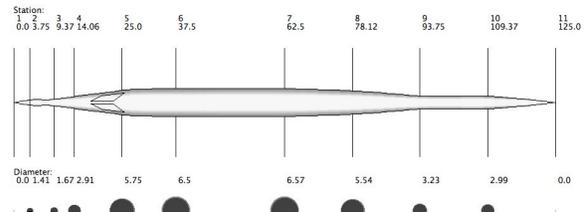
Performance

Range	4,000 n.mi.
BFL	6,500 ft
MTOGW	97,130 lbs
Ground signature (dBA)	78.87

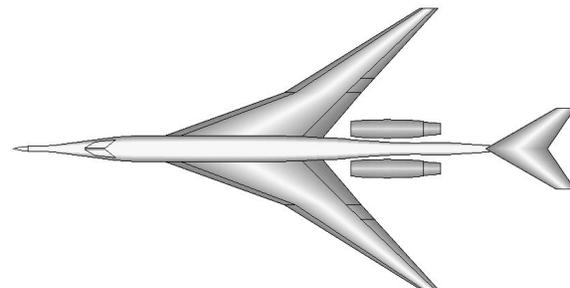
Wing and Tail Geometry

Wing reference area	1,155 ft^2
Wing aspect ratio	4.2
Wing quarter-chord sweep	55°
Leading edge extension	0.55
Trailing edge extension	0.026
Wing root leading edge	0.23
Root section t/c	3.0%
Break section t/c	2.56%
Tip section t/c	2.0%

Table 4 Results of PASS + response surface optimization for boom minimization.



a) Fuselage layout and stations



b) Top view of configuration

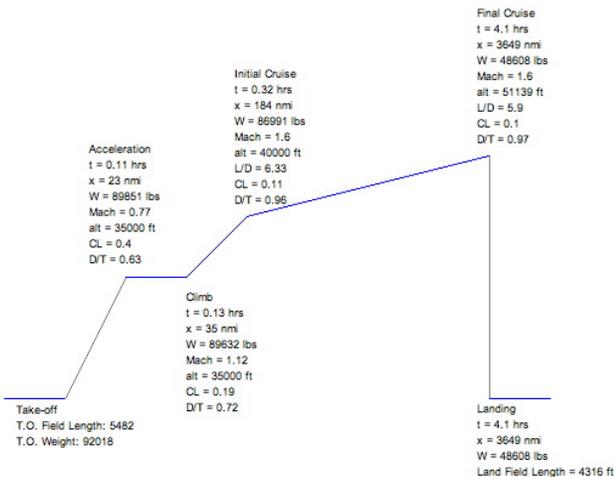
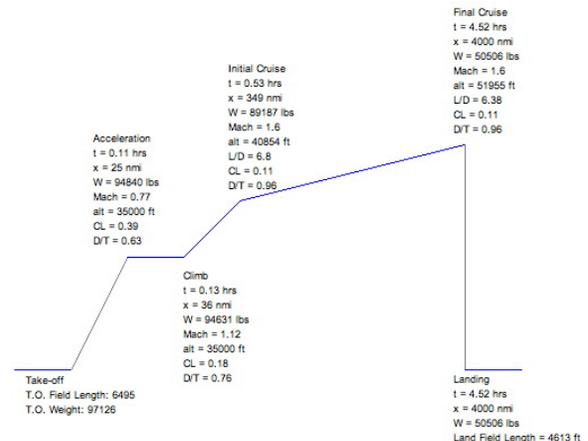


Fig. 13 Mission profile for the analysis of the baseline configuration using PASS+RS.



c) Mission profile

Fig. 14 Summary of boom optimized configuration.

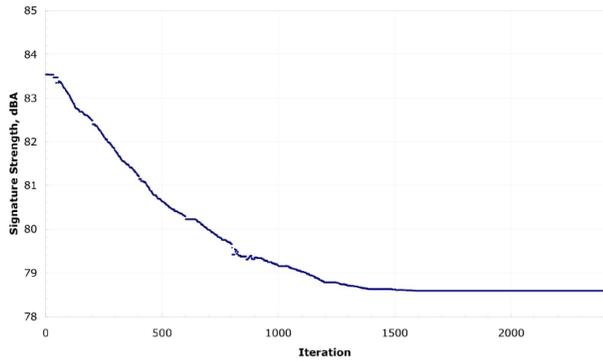
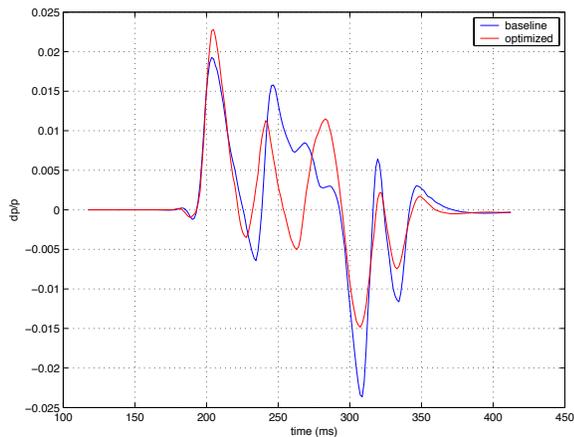
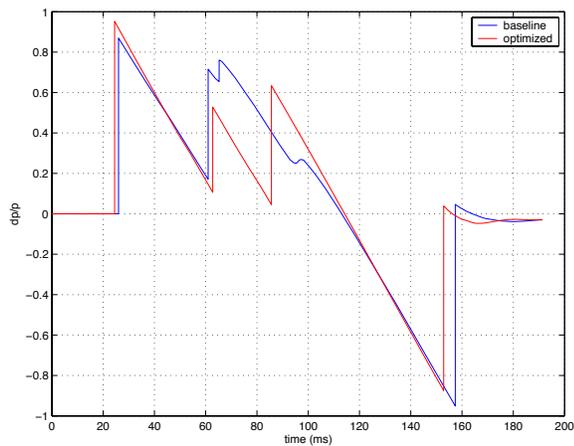


Fig. 15 Convergence history of boom objective function.



a) Near-field pressure distribution



b) Ground boom signature

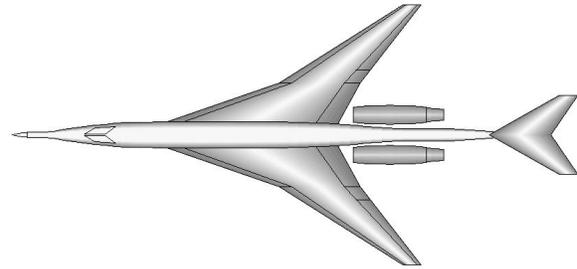
Fig. 16 Euler comparison for optimized and baseline designs.

sign) in the near-field pressures and the predicted ground boom signatures. Although it appears that the optimized design has a slightly higher initial pressure rise, the rest of the signature is such that it leads to a lower noise level.

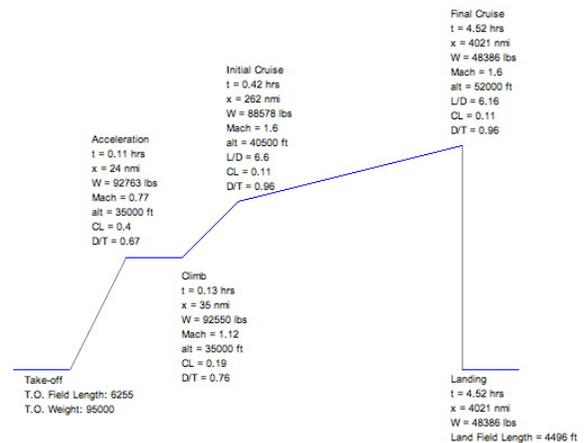
Performance	
Range	4,021 n.mi.
BFL	6,255 ft
MTOGW	95,000 lbs
Ground signature (dBA)	80.3

Wing and Tail Geometry	
Wing reference area	1,100 ft^2
Wing aspect ratio	3.7
Wing quarter-chord sweep	54°
Wing root leading edge	0.23

Table 5 Results of PASS + response surface optimization for boom minimization with small manual modifications.



a) Top view of configuration



b) Mission profile

Fig. 17 Summary of boom optimized configuration with small manual modifications.

Variation on Result of Minimum Boom Loudness Optimization

Minor reductions in reference area (from 1,155 ft^2 to 1,100 ft^2), aspect ratio (from 4.2 to 3.7) and sweep (55 degrees to 54 degrees) were manually applied to the above design to return the wing to a more reasonable shape. The resulting configuration suffered only a modest gain in boom strength, while still managing to meet all other constraints, as can be seen in Table 5 and Figure 17.

Minimum TOGW Optimization

Finally, an optimization was run with TOGW as an objective function, and boom removed entirely from consideration. Given the results presented in Table 6

Performance	
Range	4,000 n.mi.
BFL	6,300 ft
MTOGW	94,747 lbs
Ground signature (dBA)	84.9
Wing and Tail Geometry	
Wing reference area	1,000 ft^2
Wing aspect ratio	4.2
Wing quarter-chord sweep	54.1°
Leading edge extension	0.34
Trailing edge extension	0.14
Wing root leading edge	0.29
Root section t/c	2.6%
Break section t/c	3.06%
Tip section t/c	2.55%

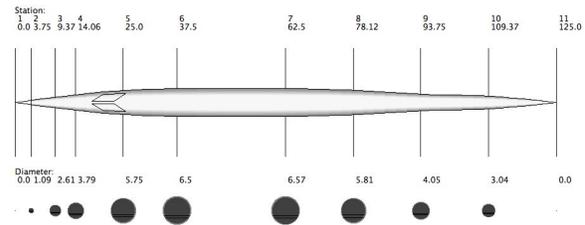
Table 6 Results of PASS + response surface optimization for TOGW optimization.

and Figure 18, it can be surmised that the outlandish low-boom configuration was largely due to the boom requirement, since this pure-performance optimization results in a decidedly more believable design. The resulting design, however has an increased boom loudness which is lower than the baseline, but only by less than 2 dBA. Considering that the minimum boom design had achieved a 7.5 dBA reduction, this improvement in performance comes at a substantial cost.

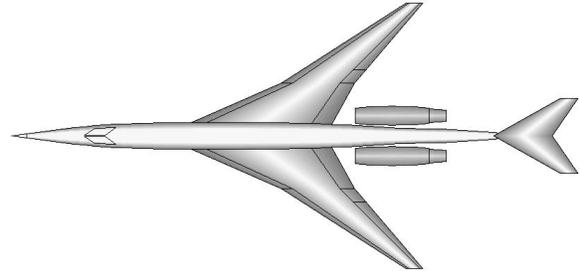
High-Fidelity Validation of Optimization Results

In this section we present the results of the validation of the predicted boom and performance for the designs carried out using the PASS+RS approach. We present validations for the boom-optimized configuration and for the TOGW optimization only. These results serve two main purposes. Firstly, they assess, indirectly, the accuracy of the response surface fits, at least in the areas where the optima are found. Secondly, they provide the necessary confidence in the outcome of the designs from the PASS+RS design methodology.

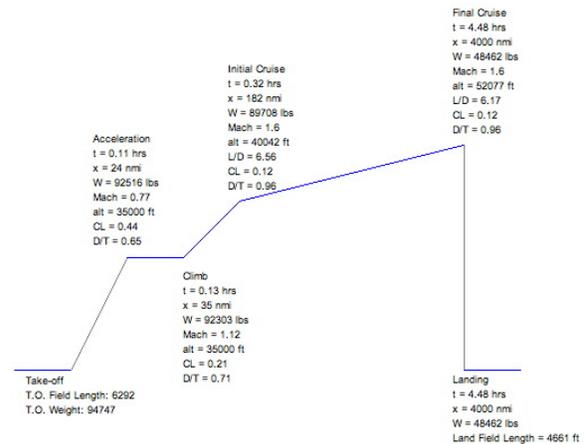
Table 9 shows the results of this validation study. As can be seen the errors in the prediction of the inviscid C_D of the aircraft are rather small: for the boom-optimized design, the response surface fit predicted a $C_D = 0.00911$, while the actual Euler re-analysis produced $C_D = 0.00949$. The same is true, or even better of the TOGW-optimized result: the response surface fit had predicted a $C_D = 0.01154$, while the actual Euler re-analysis comes in very close at $C_D = 0.01174$, a very small error indeed. The comparisons for the predicted boom levels are not as good, but still quite acceptable. For the boom-optimized configuration, the response surface fit had predicted a value of dBA = 79.24, while the Euler re-analysis yielded dBA = 77.65. In other words, the response surface had over-predicted the noise level of the configuration by almost 2 dBA. For the TOGW-optimized configuration, the



a) Fuselage layout and stations



b) Top view of configuration



c) Mission profile

Fig. 18 Summary of TOGW optimized configuration.

response surface fit predicted a dBA = 83.31, while the actual Euler analysis achieved dBA = 80.47.

CONCLUSIONS AND FUTURE WORK

In this paper we have developed a methodology for the design of supersonic jets using a multi-fidelity approximation to the response of the vehicle (for both boom and performance) in the cruise condition. The method incorporates an aircraft synthesis tool, PASS, which is able to account for a large number of realistic constraints throughout a specified mission, and the BOOM-UA analysis environment which can be used for rapid generation of multi-fidelity fits for the C_D and boom loudness of the aircraft. The advantage is that a tool such as PASS can be leveraged while providing results that are of high-fidelity with a reasonable additional cost. Response surfaces for C_D

and boom dBA are constructed throughout the design space by targeting for high-fidelity analyses only those areas that are computed to have large errors. For Euler and linearized supersonic panel codes, these areas appear to be confined to the transonic flow region (in the direction normal to the leading edge of the wing of the configuration.) Low-fidelity analyses are mainly used to construct the approximations away from these areas. Optimizations were carried out using both the baseline PASS approach and the one with the multi-fidelity response surfaces and, for the designs considered here, it is clear that the addition of the high-fidelity information has a substantial impact in the quality of the solutions obtained. Validations of the optimized designs show that the overall design strategy carries an error of about 1-4% in the C_D predictions, while the accuracy of the boom loudness predictions is quite similar (typical errors are in the neighborhood of 2-3.5%).

The work presented in this paper represents our first effort to validate this design procedure. Much has been learned during the process and we intend to carry out similar design studies with improvements to the accuracy and efficiency of the process. We are mainly concerned with two issues: the ability to reconcile the lower-bound estimates of performance produced by PASS with the outcome of Euler analysis/optimizations, and the ability to construct response surfaces for the performance and boom loudness that are more accurate than the quadratic+Kriging approach that we have currently followed. Finally, we would like to make improvements to the underlying simplex optimization techniques inside of PASS so that larger design spaces and multi-fidelity models can be more easily incorporated.

ACKNOWLEDGMENTS

This work has been carried out under the support of the NASA Langley Research Center, under contract NAG-1-03046. We gratefully acknowledge the support of our point of contact, Dr. Peter Coen. The second author would like to acknowledge the support of the Air Force Office of Scientific Research under grant No. AF F49620-91-1-2004. We also appreciate the invaluable help of Dr. Edwin van der Weide for allowing us to use the AirplanePlus flow solver for this study, of Dr. Yoshikazu Makino for his assistance with the running of the A502/Panair solver and his contributions to the computation of the boom loudness using finite rise times, and of Mr. Curran C. Crawford for creating the CAD parametric model of the aircraft that was used for all computations.

References

¹National Research Council, High Speed Research Aeronautics and Space Engineering Board *U.S. Supersonic Commercial Aircraft : Assessing NASA's High Speed Research Program* , National Academy Press, Washington, D.C., 1997.

²National Research Council, High Speed Research Aeronautics and Space Engineering Board *Commercial Supersonic Technology : The Way Ahead* , National Academy Press, Washington, D.C., 2002.

³H. S. Chung, S. Choi, J. J. Alonso, "Supersonic Business Jet Design Using Knowledge-based Genetic Algorithm with Adaptive, Unstructured Grid Methodology," *AIAA 2003-3791*, 21st Applied Aerodynamic Conference, Orlando, FL, June 2003.

⁴H. Chung and J. J. Alonso "Design of a Low-Boom Supersonic Business Jet Using Cokriging Approximative Models," *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, AIAA 2002-5598, Atlanta, GA, September 2002.

⁵J. J. Alonso, I. M. Kroo and A. Jameson. "Advanced Algorithms for Design and Optimization of Quiet Supersonic Platforms," *AIAA Paper 02-0114*, 40th AIAA Aerospace Sciences Meeting and Exhibit, January 2002, Reno, NV.

⁶M. Chan, "Supersonic Aircraft Optimization for Minimizing Drag and Sonic Boom," PhD thesis, Stanford University, Stanford, CA 94305, 2003.

⁷D. Aronstein and K. Schueler "Conceptual Design of a Sonic Boom Constrained Supersonic Business Aircraft," *42nd AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, Jan 2004*, AIAA 2004-0697, Reno, NV, January 2004.

⁸Thomas A. Zang and Lawrence L. Green. "Multidisciplinary Design Optimization Technique: Implications and Opportunities for Fluid Dynamics" *30th AIAA Fluid Dynamics Conference*, Norfolk, Virginia, AIAA 99-3798, June 1999

⁹Jaroslav Sobieszczanski-Sobieski and Raphael T. Haftka "Multidisciplinary Aerospace Design Optimization: Survey of Recent Developments," *34th AIAA Aerospace Sciences Meeting and Exhibit* , Reno, Nevada, AIAA 96-0711, January 1996.

¹⁰J. Martin, J. J. Alonso and J. Reuther "High-Fidelity Aero-Structural Design Optimization of a Supersonic Business Jet," *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA 2002-1483, Denver, CO, April 2002.

¹¹J. Martin and J. J. Alonso "Complete Configuration Aero-Structural Optimization Using a Coupled Sensitivity Analysis Method," *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, AIAA 2002-5402 Atlanta, GA, September 2002.

¹²Resit Unal, Roger A. Lepsch, Jr. and Mark L. McMillin "Response Surface Model Building and Multidisciplinary Optimization Using Overdetermined D-Optimal Designs," *7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, AIAA 98-4759, September 1998

¹³N. M. Alexandrov, R. M. Lewis, C. R. Gumbert, L. L. Green, and P. A. Newman "Optimization with Variable-Fidelity Models Applied to Wing Design," *Accepted for Publication, AIAA Journal of Aircraft*

¹⁴S. E. Gano, V. M. Perez, J. E. Renaud, S. M. Batill "Multilevel Variable Fidelity Optimization of a Morphing Unmanned Aerial Vehicle," *45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, AIAA 2004-1763 Palm Springs, CA, April 2004.

¹⁵S. Choi, J. J. Alonso and H. S. Chung "Design of a Low-Boom Supersonic Business Jet Using Evolutionary Algorithms and an Adaptive Unstructured Mesh Method," *45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA 2004-1758, Palm Springs, CA, April 2004.

¹⁶S. Choi, J. J. Alonso and E. Weide. "Numerical and Mesh Resolution Requirements for Accurate Sonic Boom Prediction of Complete Aircraft Configurations". *AIAA Paper 2004-1060* 42nd AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, Jan 2004.

¹⁷H. S. Chung and J. J. Alonso "Multiobjective Optimization using Approximation Model-Based Genetic Algorithms," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA 2004-4325 Albany, NY, September 2004.

¹⁸H. S. Chung. *Multidisciplinary Design Optimization of Supersonic Business Jets Using Approximation Model-Based Genetic Algorithms*. PhD thesis, Stanford University, Stanford, CA 94305, 2004. To be published.

¹⁹J. Reuther and J. J. Alonso and A. Jameson and M. Rimplinger and D. Saunders "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers: Part I," *Journal of Aircraft*, 36:51-60. 1999

²⁰J. Reuther and J. J. Alonso and A. Jameson and M. Rimplinger and D. Saunders "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers: Part II," *Journal of Aircraft*, 36:61-74. 1999

²¹A. Jameson "Aerodynamic Design via Control Theory," *Journal of Scientific Computing*, 3:233-260 1988

²²Nadarajah, S. K., Kim, S., Jameson, A., and Alonso, J. J. "Sonic Boom Reduction Using an Adjoint Method for Supersonic Transport Aircraft Configuration," *Symposium Transonicum IV, International Union of Theoretical and Applied Mechanics*, September 2-6, 2002, DLR Gottingen, Germany

²³Hames, R. and Follen, G. "Computational Analysis Programming Interface," *Proceedings of the 6th International Conference on Numerical Grid Generation in Computational Field Simulations*, Greenwich, 1998

²⁴J. J. Alonso, J. R. R. A. Martins, J. J. Reuther, R. Hames, and C. A. Crawford "High-Fidelity Aero-Structural Design Using a Parametric CAD-Based Model," *16th AIAA Computational Fluid Dynamics Conference*, AIAA 2003-3429, Orlando, FL., 2003.

²⁵ *Centaur System Users' Manual*, <http://www.centaursoft.com>

²⁶A. Jameson, T. J. Baker and N. P. Weatherill. Calculation of Inviscid Transonic Flow over a Complete Aircraft. *AIAA Paper 86-0103*, 24th Aerospace Sciences Meeting, January 6-9, 1986, Reno, Nevada

²⁷Carmichael, R.I., Erickson, L.I., "A Higher Order Panel Method for Predicting Subsonic or Supersonic Linear potential Flow about Arbitrary Configurations," *AIAA Paper 81-1255*, June 1981.

²⁸Ashley, H., Landahl, M. "Aerodynamics of Wings and Bodies", Dover, New York, republished 1985, p.104 and p.178

²⁹Thomas, Charles L. "Extrapolation of Wind-Tunnel Sonic Boom Signatures without Use of a Whitham F-function" *NASA SP-255*, pp.205-217, 1970

³⁰Kenneth J. Plotkin "PCBoom3 Sonic Boom Prediction Model-Version 1.0e," *Wyle Research Report WR 95-22E*, October. 1998.

Variable	Min. Value	Max. Value
Max TOGW (lbs)	80,000	100,000
Initial cruise altitude (ft)	40,000	42,000
Final cruise altitude (ft)	40,000	56,900
S_{ref} (ft ²)	1,000	1,300
Wing aspect ratio (AR)	3.0	5.0
Wing quarter-chord sweep (Λ)	45°	60°
Location of wing root LE	0.2	0.4
Root section t/c	2.0%	3.0%
Break section t/c	2.5%	3.5%
Tip section t/c	2.0%	3.0%
Fuselage station 2 radius / fus. length (x = 3.75 ft)	0.005	0.020
Fuselage station 3 radius / fus. length (x = 9.37 ft)	0.005	0.020
Fuselage station 4 radius / fus. length (x = 14.06 ft)	0.005	0.020
Fuselage station 8 radius / fus. length (x = 78.12 ft)	0.015	0.035
Fuselage station 9 radius / fus. length (x = 93.75 ft)	0.010	0.030
Fuselage station 10 radius / fus. length (x = 109.37 ft)	0.005	0.020

Table 7 Bounds for geometric design variables and for samples for response surface generation.

Constraint	Min. Value	Max. Value
Cruise range (n.mi.)	4,000	N/A
TOFL (ft)	N/A	6,500
LFL (ft)	N/A	6,500
Minimum stability	0.0	N/A
Vertical tail CL for engine out condition	-1.0	1.0
Second segment climb gradient	0.024	N/A
Main gear location as fraction of chord length	0.6	0.95
Initial cruise-climb gradient	0.005	N/A
Final cruise-climb gradient	0.005	N/A
Horizontal tail CL margin, takeoff	0.0	N/A
Horizontal tail CL margin, takeoff rotation	0.0	N/A
Horizontal tail CL margin, climb	0.0	N/A
Horizontal tail CL margin, initial cruise	0.0	N/A
Horizontal tail CL margin, final cruise	0.0	N/A
Horizontal tail CL margin, landing	0.0	N/A
Wing CL margin, climb	0.0	N/A
Wing CL margin, initial cruise	0.0	N/A
Wing CL margin, final cruise	0.0	N/A
Elevator deflection, takeoff rotation	-25.0	N/A
Elevator deflection, takeoff	-10.0	N/A
Elevator deflection, climb	-10.0	N/A
Elevator deflection, landing	-15.0	N/A

Table 8 Mission constraints for optimization runs.

	C_D RS	C_D Euler	dBA RS	dBA Euler	C_D Rel. Error	dBA Rel. Error
Min. Boom Design	0.00911	0.00949	79.24	77.65	4%	2.0 %
Min. TOGW Design	0.01154	0.01174	83.31	80.47	1.7 %	3.5%

Table 9 Results of validation of minimum boom and minimum TOGW designs.