43rd AIAA Aerospace Sciences Meeting & Exhibit
January 10–13, 2005/Reno, NV

Two-Level Multi-Fidelity Design Optimization Studies for Supersonic Jets

Seongim Choi[*], Juan J. Alonso[†], Sangho Kim[‡], Ilan M. Kroo[§]
Stanford University, Stanford, CA 94305

Mathias Wintzer[¶]
Desktop Aeronautics, Inc., Palo Alto, CA 94301

The *conceptual/preliminary* design of supersonic jet configurations requires a multi-disciplinary approach that provides the designer with information regarding the key trade-offs between the disciplines participating in the design. At the same time, at these stages of the design, the available tools must provide a level of flexibility that permits the exploration of large areas of the design space with significant changes to a baseline configuration. In order to achieve credible results one would like to use high-fidelity modeling tools for all of the components (and interactions) of the design. This can, however, be prohibitively expensive and in addition, it may significantly decrease the ability to make drastic modifications to the aircraft configuration in question. As our work has progressed in this area, we have come to realize that a truly hybrid, multi-fidelity approach that is properly managed is one of the answers to the supersonic design problem. In this paper we describe our approach to the two-level design of a supersonic business jet configuration where we combine a conceptual, SIMPLEX-based, low-fidelity optimization tool with a hierarchy of flow solvers of increasing fidelity (including simplified aerodynamic models, a linearized panel method and both structured and unstructured Euler solvers) and advanced adjoint-based Sequential Quadratic Programming (SQP) optimization approaches. Although this kind of aircraft has been studied in the past within the context of low supersonic boom, in this work we focus on the aerodynamic performance aspects alone: no attempt is made to reduce the acoustic signature. The results show that this particular combination of modeling and design techniques is quite effective to produce designs with optimum performance that meet or exceed all of the design constraints of the problem. In addition, we show that high-fidelity aerodynamic shape optimization techniques for complex configurations (such as the adjoint method) can be effectively used within the context of a truly multi-disciplinary design environment. Detailed configuration results of our optimizations are also presented.

[*]Doctoral Candidate, AIAA Member
[†]Associate Professor, AIAA Member
[‡]Research Associate, AIAA Member
[§]Professor, AIAA Fellow
[¶]Engineer, AIAA Member

I. Introduction

Previous work in the design of supersonic jet configurations has focused mostly on separate efforts at either the *conceptual* or the *preliminary* design stages. Conceptual design tools permit large variations in the design and use relative inexpensive multi-disciplinary analyses (MDAs). Preliminary design tools incorporate higher levels of fidelity (particularly in the aerodynamics) but can be quite costly. Traditionally it has been sufficient to follow a sequential process by which a rough configuration is developed during the conceptual design phase and it is later refined using preliminary design tools. The key question that we are revisiting in this paper is whether this sequential conceptual/preliminary design process is adequate for supersonic aircraft when the participating disciplines are closely coupled or whether higher-fidelity tools need to been included early on to ensure that the outcome of the *conceptual* designs can be believable and used for further design work. In some senses, we are proposing to merge the first two phases of the design, *conceptual* and *preliminary*, into a single one while ensuring that both the solution accuracy and turnaround time are acceptable to the aircraft designer.

There are, however, some fundamental problems in integrating conceptual and preliminary design tools. At the conceptual stage we have traditionally integrated a large number of disciplinary models of low-to-medium fidelity which can be coupled together into a single, all-encompassing multi-disciplinary analysis. This MDA procedure is then coupled to an optimizer and, after a suitable parameterization of the design space (including the shape of the aircraft, its structure, the flight mission profile, propulsion deck, etc.), the MDA is run repeatedly until the design converges to a user-defined optimum while satisfying all of the specified design constraints. The complexity of this *conceptual* design procedure is limited by the time of execution of a single MDA, the number of design variables used to parameterize the configuration, and the choice of optimizer which results on a certain number of analyses for the specified number of design variables.

The nature of the design spaces that result from these multi-disciplinary systems is such that they may include multiple local minima, may be discontinuous in some areas, may be noisy and ill-behaved, and therefore, in general, may only be amenable to the use of non-gradient based optimizers. Other optimizer choices (such as the more efficient gradient-based optimizers) may result in a lack of robustness of the design procedure that renders it unusable. Because of these reasons, in our previous work we have chosen to use either the SIMPLEX method or genetic algorithms for this phase of the design.[4,3,6,17,19,20] Although quite robust and able to navigate complex design spaces, these algorithms have very slow convergence properties and require very large numbers of MDAs in order to obtain a solution. This optimizer characteristic tends to limit both the number of design variables that can be used to parameterize the mission/configuration and the level of fidelity that can be used in the MDAs while completing design calculations in finite time.

An approach that has been often followed to introduce higher-fidelity analyses into these conceptual design tools is the use of pre-generated response surfaces that represent, in an accurate and inexpensive way, some key disciplines in the design: the computational cost is paid upfront when the response surfaces are generated by repeated evaluation of the MDA. As was mentioned above, the design spaces that must be represented by these response surfaces are quite complex, and therefore, the cost of response surface generation (using polynomial fits, radial basis functions, Kriging methods, etc.) grows very rapidly with increasing numbers of design variables. Unfortunately, the use of higher-fidelity models often requires higher-fidelity discretizations with much large numbers of design variables (this is indeed the case for aerodynamic shape optimization[22,23]), making the use of traditional response surface formulations impossible within the context of *conceptual* design and using large numbers of design variables.

On the other hand, much progress has been achieved with some specific high-fidelity design tools. This is true in both the structural design and aerodynamic shape optimization communities.[13,10,22,14] In these two areas both separately and in combination[12,7,27,28] novel methods such as the adjoint and direct methods have been used to carry out designs with high-fidelity and large numbers of design variables. Efforts to couple these advanced design methods to the large number of disciplines that are considered in traditional *conceptual* design are only at their beginning stages and the adjoint method for aerodynamic shape optimization, for example, has not been yet extended to treat the highly constrained design spaces present in supersonic

vehicle design.

In summary, successful and realistic design methods must therefore carefully consider the balance between all of the performance measures, constraints, and requirements of the problem, while providing results that are sufficiently accurate to be believable. In addition, they must allow the designer the necessary freedom to explore the design space using reasonably large numbers of design variables. These 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?[9,10,11,12]

It is widely accepted, however, that not every portion of the multi-disciplinary analysis of supersonic 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 form 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.[15,16,20]

This is one of the the fundamental ingredients of the approach that we have followed in this paper: a combination of low-fidelity tools that are enhanced by multi-fidelity analyses *only* in the areas where the increase in accuracy is needed, thus limiting the use of the expensive, high-fidelity tools to a minimum. Since we are trying to limit the computational expense required to construct these multi-fidelity response surface approximations, we are willing to work with approximations that may have errors around 5% in some areas of the design space where the sample points used for the response surface construction have not been sufficiently clustered. For this reason, the second component of our two-level optimization approach is to perform adjoint-based aerodynamic shape optimizations of the resulting multi-disciplinary designs to recover the performance that may have been lost due to inaccuracies in the response surfaces and to ensure that the performance of the resulting design is that predicted by the high-fidelity tools. In principle, a realistic design would iterate between these two optimization levels; in this work we have carried out a single iteration of the procedure for demonstration purposes. The resulting design, however, meets all of the design constraints and criteria at the end of this first design iteration.

In summary, in this work, we combine ideas of multi-fidelity analysis and design and a two-level optimization procedure 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 unstructured adaptive Euler solver (AirplanePlus) to create models of the aerodynamic performance.

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 of Haimes,[26,27] the A502/Panair and AirplanePlus flow solvers, and the Centaur mesh generation system.[29]

4. Adjoint aerodynamic shape optimization tools for both single-block wing-body configurations (SYN87-SB) and multiblock complete configurations (SYN107-MB) that use inexpensive gradient calculations with larger numbers of design variables to modify the twist and camber distributions of the wing (without changes to the wing planform) and to achieve the highest aerodynamic performance.

Our early work in high-fidelity supersonic design had focused on the cruise condition alone.[17,18,3,19] In September of 2004[20] we first introduced the PASS program to handle all of the major design constraints. Although we had already included some basic constraints to handle other points in the mission,[17] 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 simultaneous use of PASS and adjoint methods to formally include all of the necessary constraints in the design while allowing for high-fidelity results.

Even when high-fidelity analyses are considered for the cruise condition alone, the supersonic design problem is made significantly more difficult because of the large variations in the values of the design variables (and therefore in the geometry) and because the design space has been shown to contain multiple local minima and be rather noisy and even discontinuous.[21] These characteristics of the design problem rule out the possibility of using gradient-based optimization (and the powerful adjoint method[22,23,24,25]) directly. In this paper, we are therefore investigating the two-level hybrid method outlined above. First multi-fidelity response surfaces are constructed using both A502/Panair and the AirplanePlus Euler solver. The SIMPLEX method is used, together with these response surfaces, to find the locations of the optima in these large design spaces. Adjoint shape optimizations follow to modify the exact twist and camber distributions of a fixed wing planform so that maximum performance at the cruise condition can be recovered. Since the changes occur only in the wing sections (twist and camber, no thickness changes) it is assumed that the resulting configuration still meets all of the mission requirements imposed by PASS. Of course this may not be exactly true, but additional iterations of the design procedure are able to correct these inconsistencies.

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. These analysis tools have been presented previously[20] and the reader is referred to that publication for more details. 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. The result of a PASS optimization (enhanced by the multi-fidelity response surface) for a supersonic jet flying at $M_\infty = 1.6$, with a range of 4,000 nmi, and with a take-off field length no greater than 6,000 ft is presented. High-fidelity validations of these results (using both structured and unstructured Euler solvers) are also shown and compared to each other. Finally we show preliminary results of the adjoint wing redesign for drag minimization at the cruise condition to support the claim that an additional level of performance can be gained by using localized high-dimensionality parameterizations and high-fidelity flow solutions.

II. 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. 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.) It must be noted that the drag predictions in this module have been constructed so that they can be expected to be a reasonable upper bound in performance if sufficient detailed design work is carried out to properly re-twist and re-camber the wing at the conclusion of the design.

B. 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.

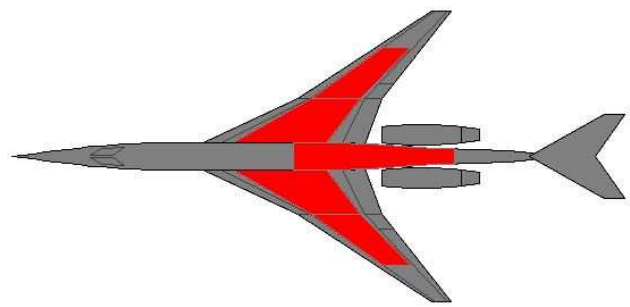

(a) Fuel tank layout.

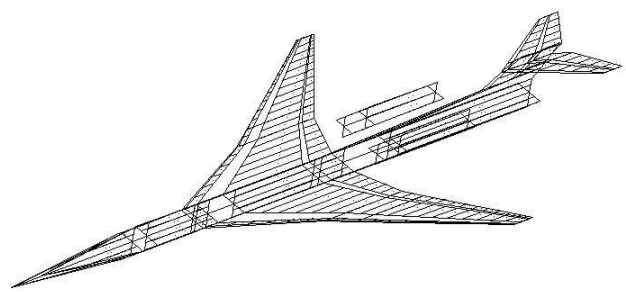

(b) Vortex-lattice model.

Figure 1. Two different views of the configuration within PASS.

C. 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.

D. 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.)

E. 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 in the form of multi-fidelity response surface approximations for the C_D of the complete aircraft.

F. 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.

G. 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 with two significant differences: 1) the inviscid aerodynamic drag prediction module in PASS was replaced by the response surface fits created using our multi-fidelity approach, and 2) once a candidate configuration was generated, an adjoint-based wing twist and camber optimization was run to further improve the performance at the cruise condition.

Unlike in our previous work[20] the ground boom calculation module was not used to guide the designs in any way. We will attempt to include sonic boom considerations into the design procedure at a later time.

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.

III. AERODYNAMIC ANALYSIS : A502 and AirplanePlus

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. As mentioned above, however, in this study we did not include the minimization of ground boom and therefore those portions of BOOM-UA were not used. The current version of BOOM-UA is an evolution of our previous work[17, 20] 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 the geometry definitions used on multi-fidelity computations are identical. 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 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 in the case of ground boom minimization problem. 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 component module is explained in detail in the following subsections. However in this study, the ground boom signature is not used as an objective function and, therefore, the full solution-adaptive procedure we have used in the past to obtain high-quality near-field pressures was not necessary. This greatly simplifies the calculations used to generate multi-fidelity response surfaces.

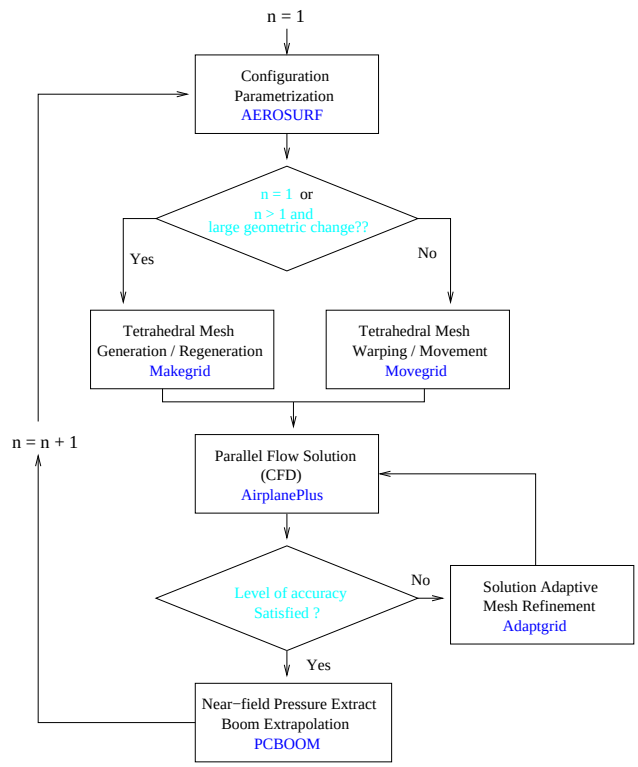

Figure 2. Schematic of the aerodynamic analysis tool, BOOM-UA

A. 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. Details of our CAD-based geometry engine (AEROSURF) have been presented earlier[20] and will not be discussed further here.

B. 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 Centaur[29] 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. Figures 3 and 4 show the surface meshes for our baseline configurations (wing-body alone and complete configuration with wing-body-empennage-nacelles) that are used

to create design variations based on a number of design parameters. To verify the compatibility of the two aerodynamic analysis tools, A502/Panair and AirplanePlus, the drag polars for both the wing/body and the complete configurations will be compared to each other in the next section. The triangular meshes on the body surface are shown significantly coarsened for a visualization purposes: typical flow solutions are computed using around 2 million nodes of a tetrahedral mesh. This mesh size is significantly smaller (by over a factor of 2) when compared with the mesh resolution that is required to obtain solutions of sufficient accuracy for sonic boom computations. The drag polar comparisons are important as they provide useful information regarding the areas of the design space where significant disagreement between the two flow solvers is to be found.

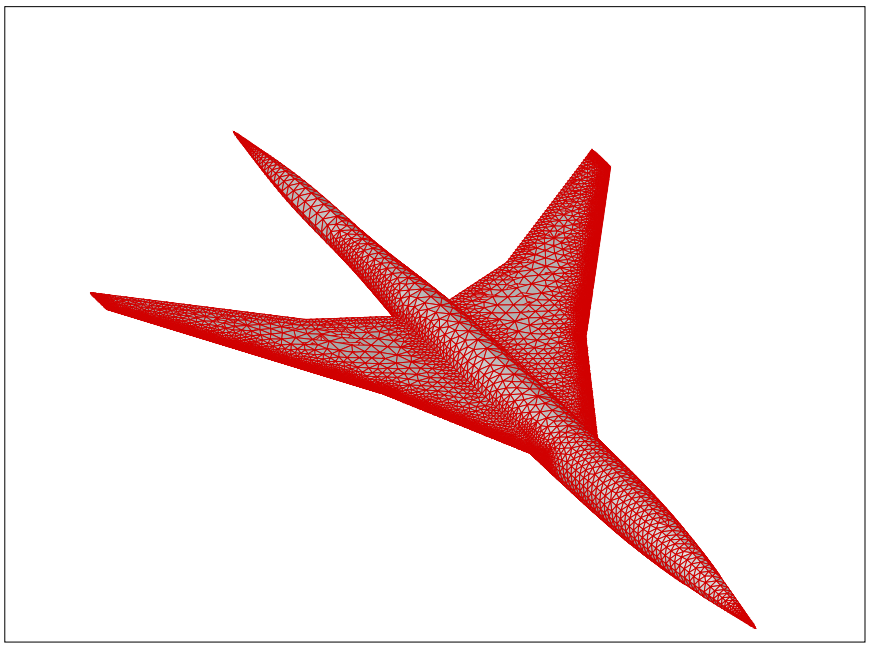

Figure 3. Unstructured tetrahedral surface mesh around wing/body configuration.

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.[30]

Figures 5 and 6 show typical surface pressure plots corresponding to the full configuration in the cruise condition. These figures are simply used to highlight the complexity of the wave patterns in the rear portion of the aircraft where the flow around the fuselage, wing, nacelles, and empennage interact in complex ways. The addition of pylons to support the nacelles and of power effects only serve to complicate the flow patterns further. This type of flow pattern can be hard to understand, particularly in the context of sonic boom computations.

C. Three Dimensional Linearized Panel Method, A502/Panair

The A502 solver, also known as Panair,[31,6] 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

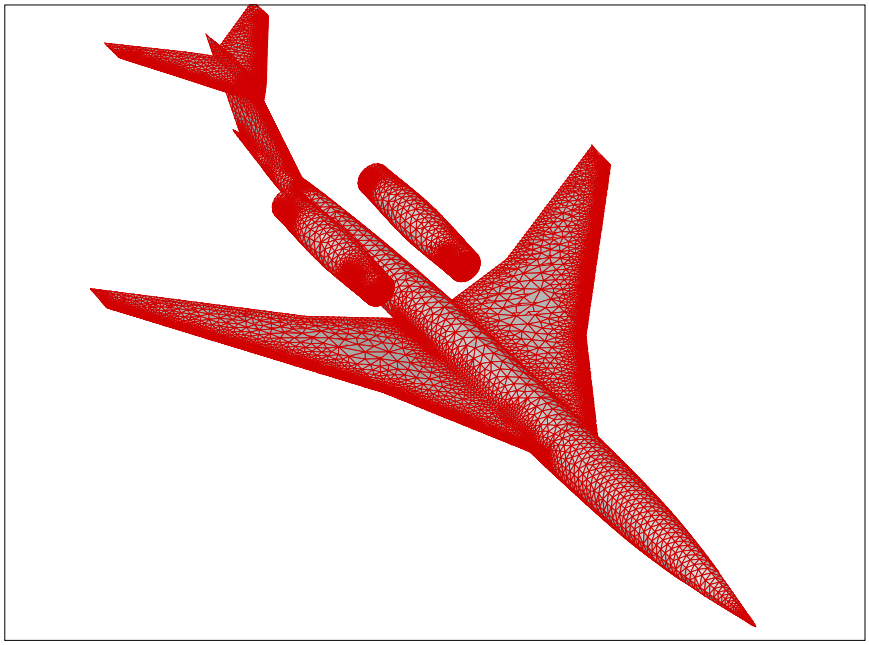

Figure 4. Unstructured tetrahedral surface mesh around full baseline configuration

uses a higher-order (quadratic doublet, linear source) panel method, based on the solution of the linearized potential flow boundary-value problem. Results 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 is required by the design (not in this work). In keeping with the axisymmetric assumption of sonic boom theory, the near-field pressure can be obtained at arbitrary distances below the aircraft.[34]

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. As mentioned in the previous section, the drag polars for the baseline configuration using the available flow solvers will be compared in a later section. The panel topologies for the same baseline configurations presented earlier are shown in Figures 7 and 8 below.

Typical pressure distributions corresponding to Figures 5 and 6 above can be found in Figures 9 and 10. For this particular case (Mach number and C_L) the validity of the panel code solutions is rather good and the flow patterns that can be observed are similar to those from the Euler results. It must be noted that the surface panel representations in Figures 7 and 8 have also been coarsened for presentation purposes. In addition, the nacelles are represented in A502 not by their true geometry, but instead, by an equivalent area representation that leads to differences in the predicted flow patterns. This is one of the shortcomings of the linearized panel code that can be overcome, at significant computational cost, by using the Euler solver.

IV. 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

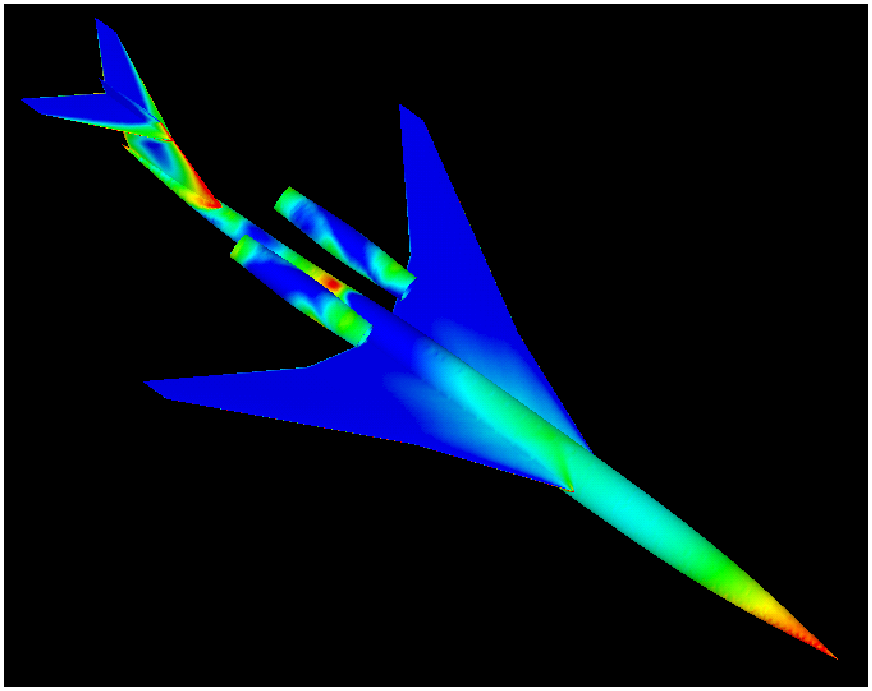

Figure 5. Upper surface pressure distribution for full baseline configuration using AirplanePlus Euler calculation.

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, the corresponding low-fidelity modules in PASS can be replaced by these response surface fits. This makes for a remarkably simple integration problem and also provides us with the ability to predict the changes in aerodynamic performance resulting from wing section changes. The baseline version of PASS is unaware of the actual wing sections used and assumes that, whatever the sections are, they have been adjusted in such a way that the camber and twist distributions are optimal (in the sense that they get close to elliptic load distributions in both the spanwise and streamwise directions). 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. The level of accuracy in the response surface representation depends greatly on the number of high-fidelity calculations that are used. Since we are trying to minimize this number, we will undoubtedly incur some errors in the fits. The validity of these fits will be assessed by direct analysis of the resulting optimized designs using two different Euler solvers.

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

1. PASS internal analysis based on classical aerodynamics.
2. A502/Panair supersonic linearized panel code.
3. Euler solutions of the highest fidelity using unstructured mesh (with a total of around 1-2 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. Unfortunately, for large dimensional design spaces (we will be using a

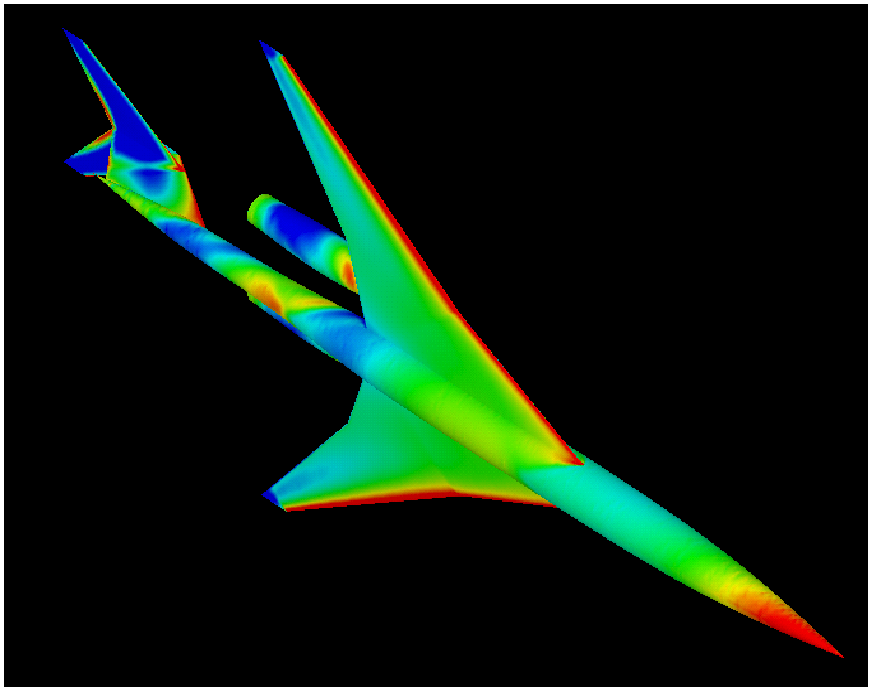

Figure 6. Lower surface pressure distribution for full baseline configuration using AirplanePlus Euler calculation.

total of 23 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 ($> 8,000$) using the aerodynamics module in PASS. 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 ($\approx 2,500$) using the A502/Panair solver. Each evaluation requires about 10 seconds of CPU time on the same modern workstation.

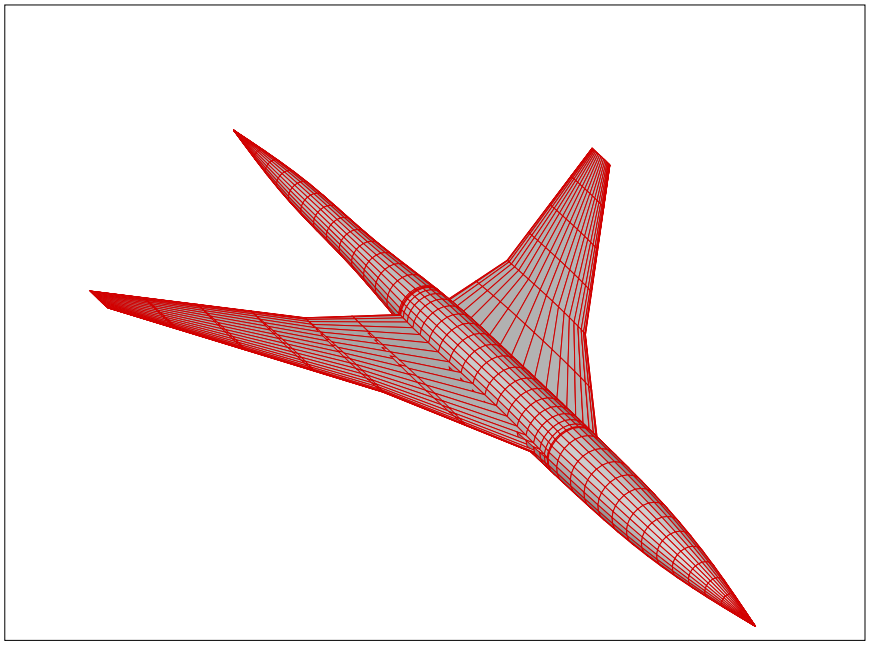

Figure 7. Panel representation around wing/body configuration

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 Fine Euler (FE) approach. In our work, we have set this threshold to about 45% resulting in a number of high-fidelity function evaluations in the neighborhood of 200. Each FE evaluation, from beginning to end, including geometry and mesh generation(the bottlenecks in the process, since they are run serially) requires about 40 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.

4. A baseline quadratic response surface fit (using least squares regression) is created for the C_D obtained with A502/Panair. The error between the values of 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 of the error between the FE solutions and those quadratic fits.

Figure 11 shows the result of the over 2,500 candidate designs (green dots) 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 $\epsilon_{PASS-A502} > 45\%$ 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 FE. Figure 12 shows in blue the results of the FE analyses for a subset of about 200 of the red A502 results. 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.

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

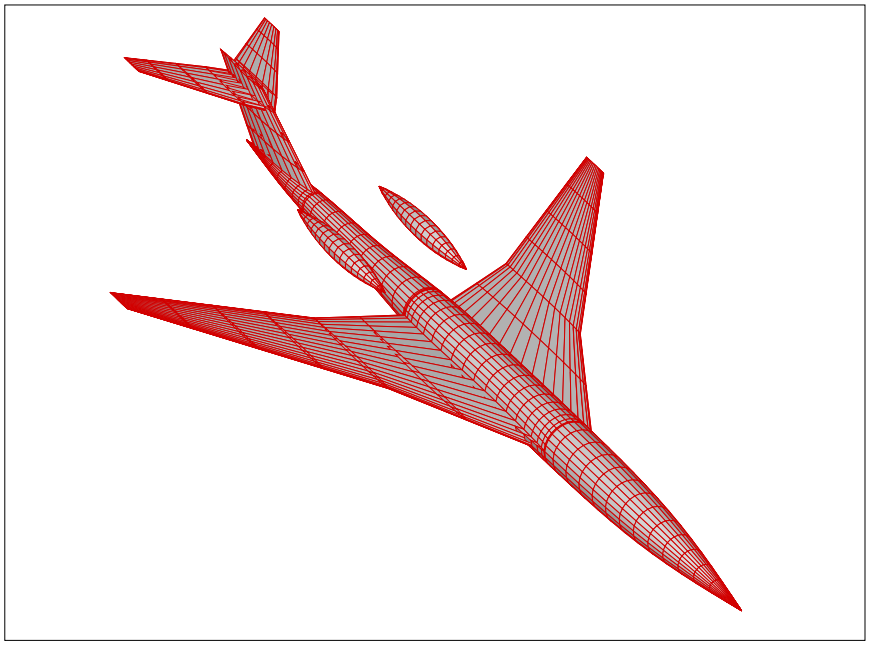

Figure 8. Panel representation around complete baseline configuration

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, 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-dimensional 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 but produces more realistic results than the PASS analysis as the actual geometry of the configuration is truly accounted for. Finally, the Euler models 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, since sonic boom has not been considered in these designs, the Coarse Euler (CE) evaluations which we have used in previous work[20] would be sufficient as the differences in C_D between CE and FE calculations were found to be small (less than 5 counts) over the large range of variations pursued in this work.[20]

V. DETAILED OPTIMIZATION : Adjoint Method

Non-gradient based methods (direct searches, simplex method, genetic algorithms) have been shown to work with a wide range of problem types, and the additional complexity required to compute gradients, approximate Hessians, perform line searches and determine the optimal steps, is not needed. These properties have a tendency to make this search procedures simple and robust.

In addition, these search procedures are able to handle rather noisy and large design spaces, making them a reasonably good match for conceptual design procedures, as was mentioned in the Introduction. However all of these algorithms tend to require a very large number of function evaluations for convergence

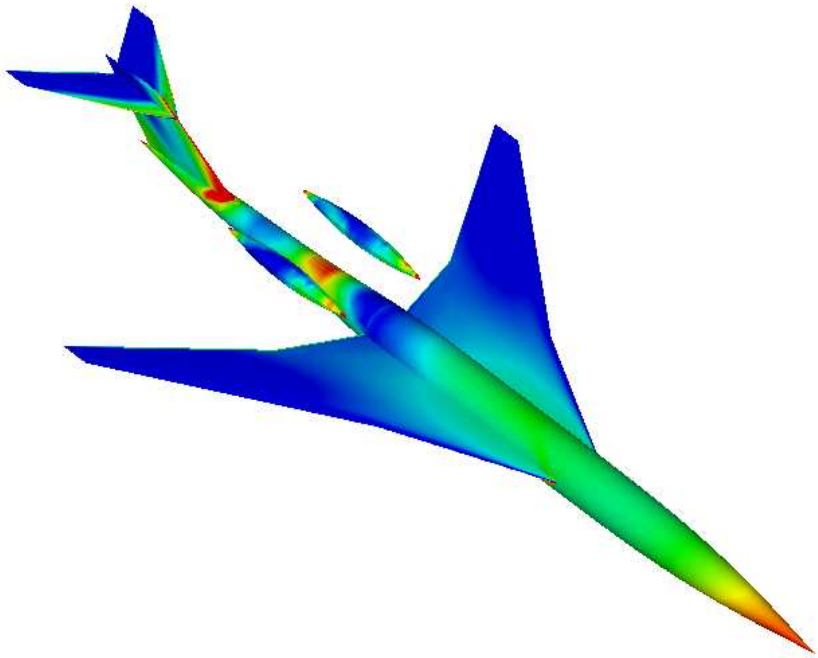

Figure 9. Upper surface pressure distribution for full baseline configuration using A502/Panair

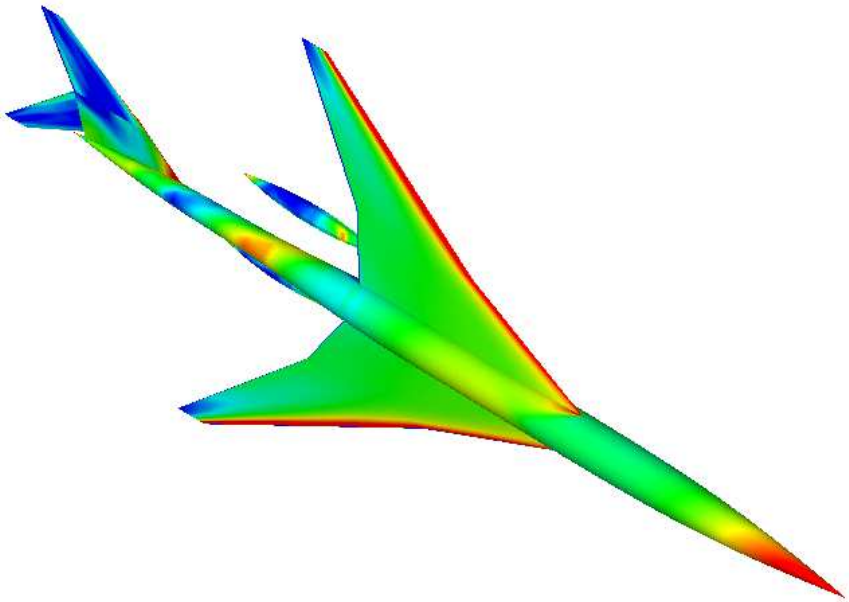

Figure 10. Lower surface pressure distribution for full baseline configuration using A502/Panair

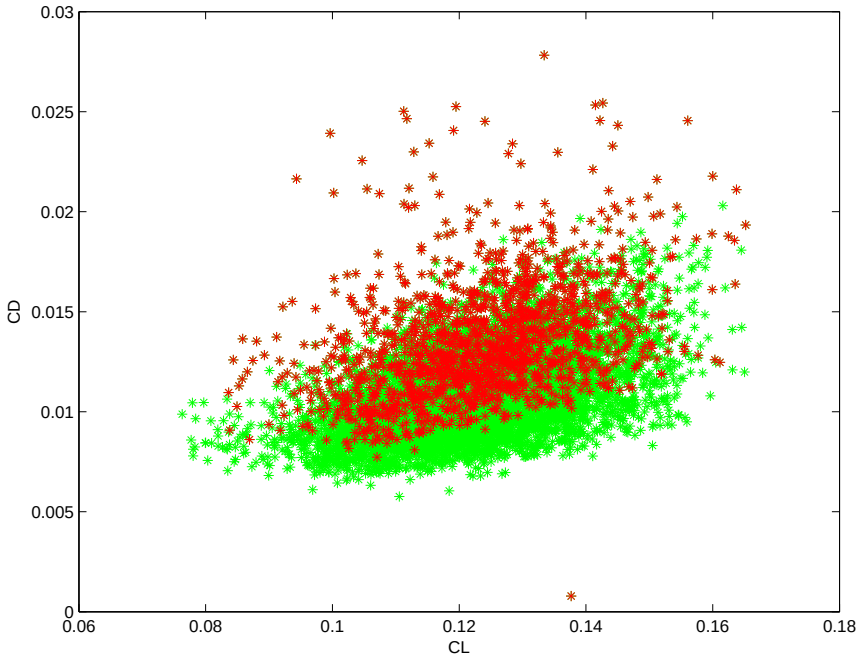

Figure 11. Database of PASS (green) and A502/Panair (red) results.

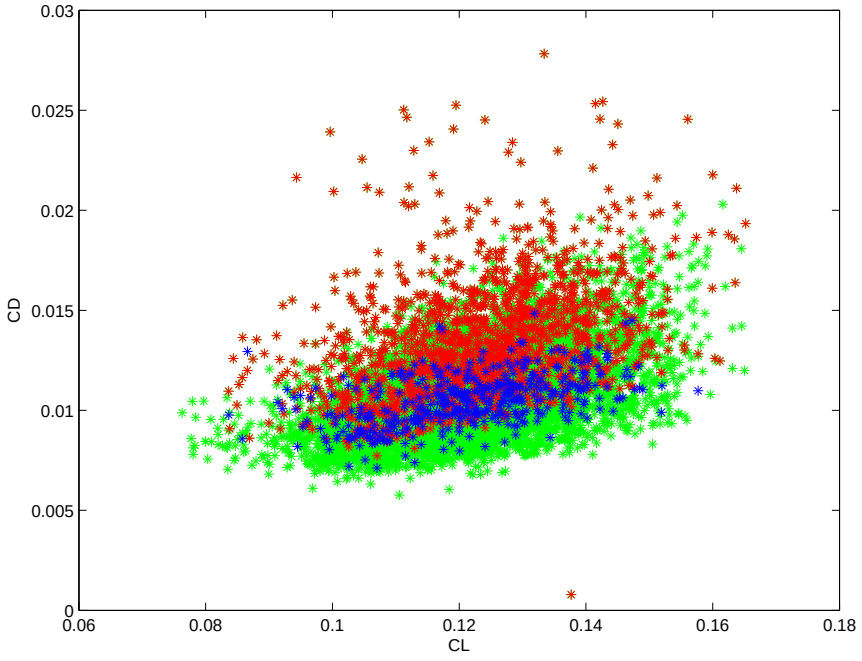

Figure 12. Database of PASS (green), A502/Panair (red) and FE (blue) results.

and therefore their computational cost can be very high.

On the other hand, if a design problem shows a smooth response to the variation in the design variables and gradient information is readily available and can be obtained inexpensively (as is the case with the adjoint method) gradient-based optimization techniques can be shown to have significant advantages over non-gradient search procedures.

In this work we are seeking to combine the advantages of both gradient- and non-gradient-based optimization procedures. As mentioned earlier, PASS uses a simplex method and is able to produce reasonable designs (using a maximum of around 20-25 design variables) even with very large variations of the design variables. Once the simplex method has converged to an optimum (local or global) we may limit ourselves to smaller changes in the configuration. These changes are more likely to result on well-behaved design spaces that can be tackled with an adjoint procedure and a gradient-based optimization algorithm.

In this second level of our optimizations we limit ourselves to modifications in the twist and camber of the wing, while maintaining the same wing planform, fuselage and relative positioning of the nacelles and empennage.

Two different tools are available for this portion of the overall optimization:

1. *SYN87-SB.* A single-block, wing-body Euler adjoint optimization code that uses the NPSOL SQP algorithm for the optimization with or without constraints. SYN87-SB allows for arbitrary changes to the shape of the fuselage and wing and is able to enforce thickness, curvature, and fuel volume constraints.

2. *SYN107-MB.* A multi-block, complete configuration, RANS adjoint optimization code that also uses the NPSOL SQP algorithm for optimization and that allows similar geometry controls, cost functions, and constraints as SYN87-SB, but that can be made to treat arbitrarily complex geometries such as the complete aircraft configurations that are the subject of this work.

Because of time limitations, this second level of adjoint-based optimization was carried out using the SYN87-SB code only, although the flow solver portion of SYN107-MB was used to carry out validation runs using a 5.9 million node multiblock mesh that was constructed for the PASS/response surface optimized configuration.

The basic theory and methodology of the adjoint method has been presented numerous times before and will not be repeated here. The reader is referred to[32, 33] for more details of the actual formulation and for representative design calculations using this method. By now, the adjoint method has been used for a very large number of design calculations in all flow regimes: low-speed high-lift configurations, transonic designs, and supersonic configurations. This method has been applied in both academic and industrial environments is by now a fairly established procedure.

VI. RESULTS

A. Comparison of Results: A502/Panair vs. AirplanePlus

Before we generate the multi-fidelity response surface for C_D, a comparison of the quality of the results produced by different fidelity analyses can be helpful.

If the errors between the two approximations are large over wide areas of the design space, the fundamental premise of the multi-fidelity approach breaks down. In order to do a limited test of this assumption, we show in Figure 13 drag polars computed using A502 and AirplanePlus for the baseline configuration in its wing-body and complete configuration versions. The corresponding drag polars show fairly similar trends and absolute values with typical differences around 3 counts for the wing-body configuration and slightly larger differences for the complete configuration. These slightly larger differences come partly from the fact that both solvers use different nacelle geometries. A502/Panair uses a representation of the nacelle based on an equivalent body of revolution which differs from the actual geometry of the real nacelles used in the

Euler code. However, the relative errors are still reasonable and around 5-6%. It must be mentioned that for the wing-body configuration, the results of the drag polar show somewhat larger errors when compared to the classical aerodynamics performance predictions in the baseline version of PASS. Since PASS assumes an *optimized* (in the load distribution sense) we carried out a number of adjoint-based twist and camber optimizations that were able to reduce the values of the Euler C_D so that the agreement with PASS was rather good (around 1-2 counts). These results are not presented here for brevity.

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 wing/body configuration, and for a given wing inboard leading edge sweep angle (around 65°), we changed the freestream Mach number to the following values: 1.4, 1.6, 1.8, and 2.03, so that the corresponding normal Mach numbers were 0.5949, 0.68, 0.7649, and 0.85 respectively. The errors calculated for two methods are -2.2%, -1.58% ,8.6694% and 28.22% respectively. All computations were carried out at a fixed $C_L = 0.13$. The drag coefficients for the A502/Panair and Euler calculations are compared in Figure 14. For the fully subsonic (normal to the leading edge) Mach numbers, the two aerodynamic analyses shows fairly good agreement. However as the normal Mach number enters the transonic regime (where the value of the normal Mach number is typically larger than 0.75) the differences in drag coefficient between become very large. This explains that there are areas in the design space where the high-fidelity tool should be applied to provide corrections to the objective function value predicted by the lower fidelity tools.

B. Baseline Configuration: Standard PASS Optimization

For subsequent design work, 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. The value of the MTOW is the result of the optimization as this was the objective function of the design. As mentioned before, in an effort to generate an aircraft achievable using current levels of technology, advanced technology assumptions were kept to a minimum.

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

Table 1. Performance requirements for optimized baseline configuration.

The values of the design variables for the resulting baseline configuration (which are also used as starting points for subsequent designs) are provided in Table 2. Note that the values highlighted in red were not allowed to vary during this initial optimization. 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 performance improvements and to maintain cabin and cockpit compartment constraints. Finally wing section changes were allowed at three defining stations. The twist at the root/symmetry plane section, the leading edge crank section and the tip section were allowed to vary. Furthermore, the value of the maximum camber and the location of maximum camber were also allowed to change at the first two wing defining stations. This makes up for an addition 7 design variables for the wing.

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

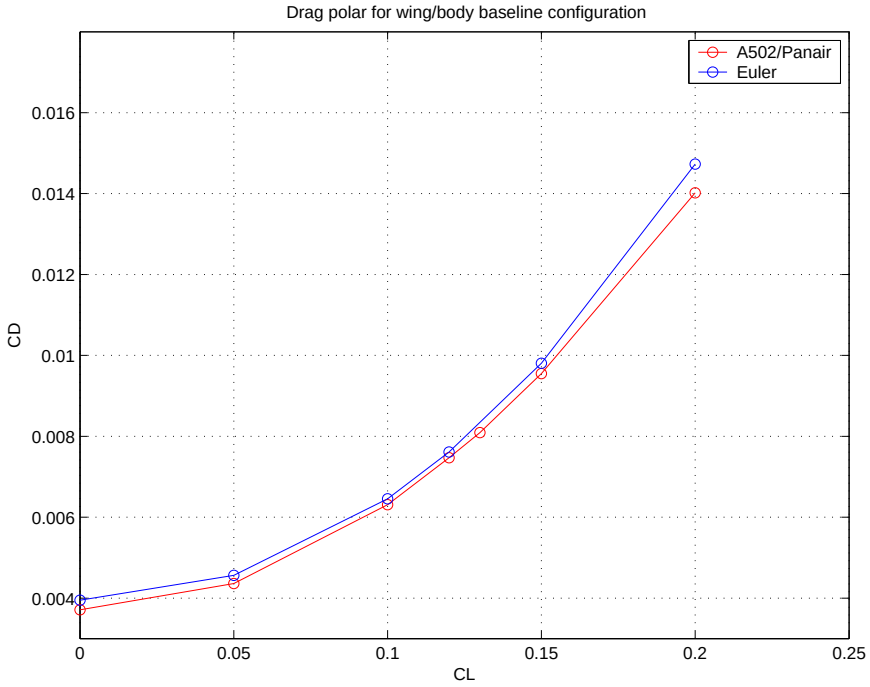

(a) Drag polar of wing/body baseline configuration.

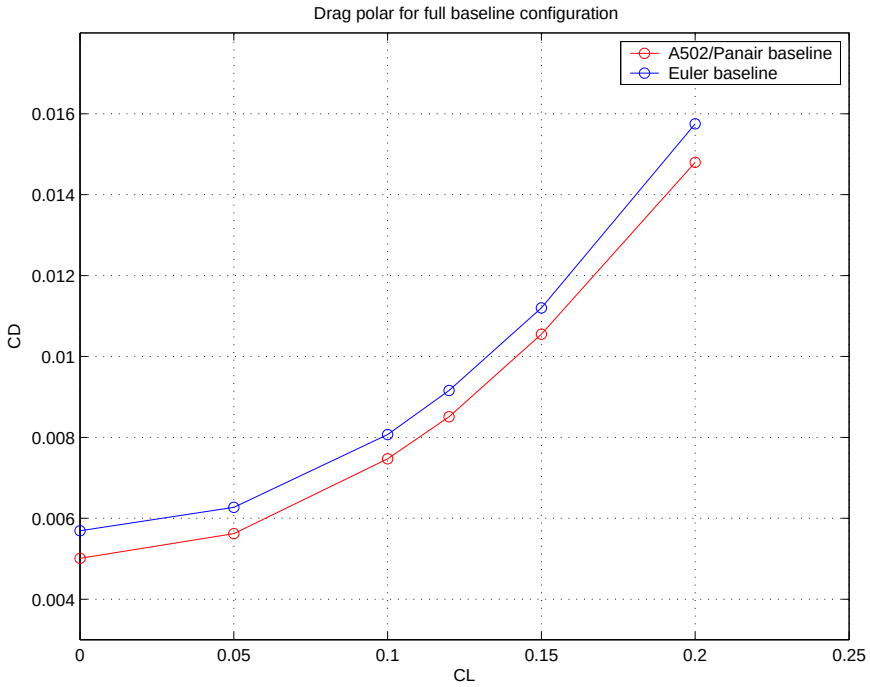

(b) Drag polar of complete baseline configuration.

Figure 13. Drag polars for A502/Panair and AirplanePlus

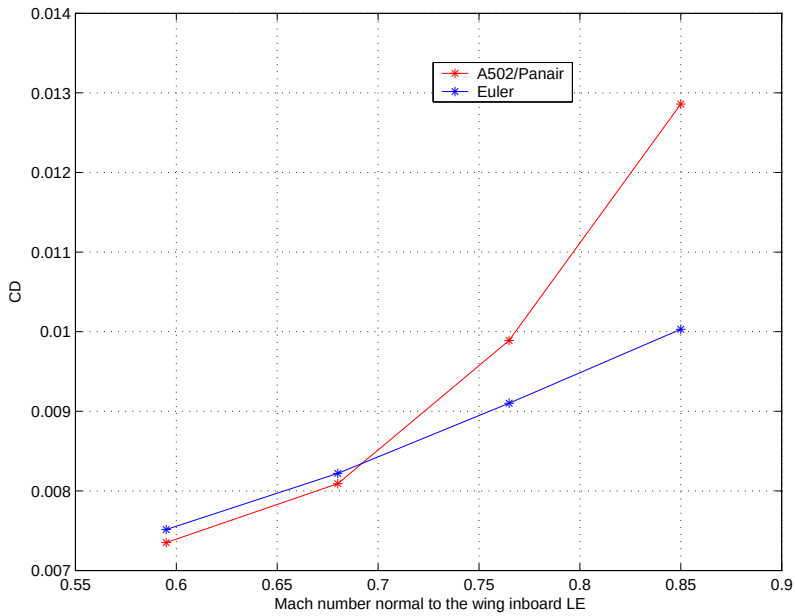

Figure 14. Drag coefficient at different normal Mach numbers to the wing inboard section

of variations allows for a more complete design space to be searched but 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 in the Tables 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}.

Figure 16 also shows some interesting results. In the Figure we show the drag polars of the baseline configuration (designed using the baseline version of PASS and without allowing for changes in the wing twist and camber) analyzed using both A502/Panair and AirplanePlus. The results represented by the dashed lines) are in fairly good agreement with each other. If we carry out an optimization where only the fuselage shape and the wing twist and camber are allowed to vary and use A502/Panair (within PASS) as the aerodynamic prediction module, we see how the drag (for both A502/Panair or AirplanePlus analyses) is significantly reduced: a decrease of nearly 20 drag counts is found. However, since the optimization was carried out using a linearized potential flow model, there are discrepancies of about 10 drag counts between the predictions of A502 and AirplanePlus. This is one of the main reasons why we will follow our PASS/simplex optimizations with an adjoint-based twist/camber design: if it were possible to reduce the drag of the Euler designs to the levels predicted by A502 (a difference of about 10 drag counts uniformly across the C_L range) then one could use A502 as an aerodynamic model with relative confidence, knowing that, at the end of the design, a high-fidelity wing re-twist and re-camber will be necessary. The issue of the differences between the representation of the nacelles in both codes still remains and needs to be accounted for in A502/PASS to provide realistic and attainable goals.

In fact, we will later show that Euler optimizations recover some of the lost drag, but not all of it. Our preliminary computations show that about half (5 counts) of the drag difference can be recovered by using a large number of design variables and the Euler adjoint method implementations.

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.

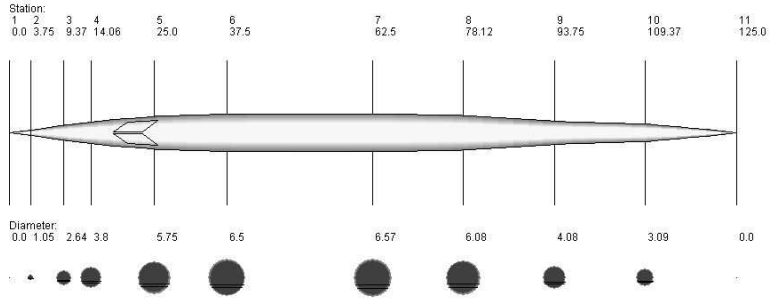

(a) Fuselage layout and stations

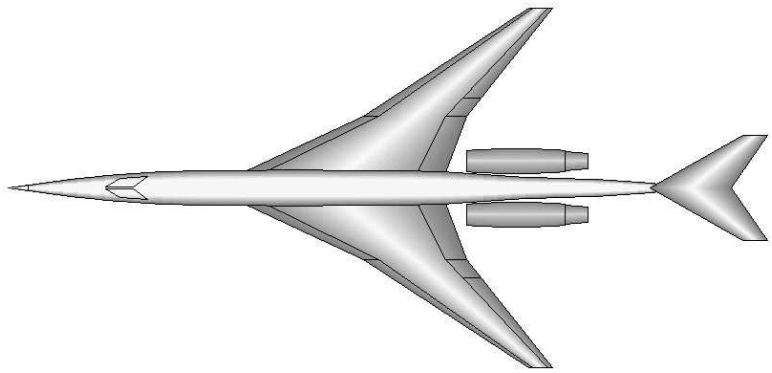

(b) Top view of configuration

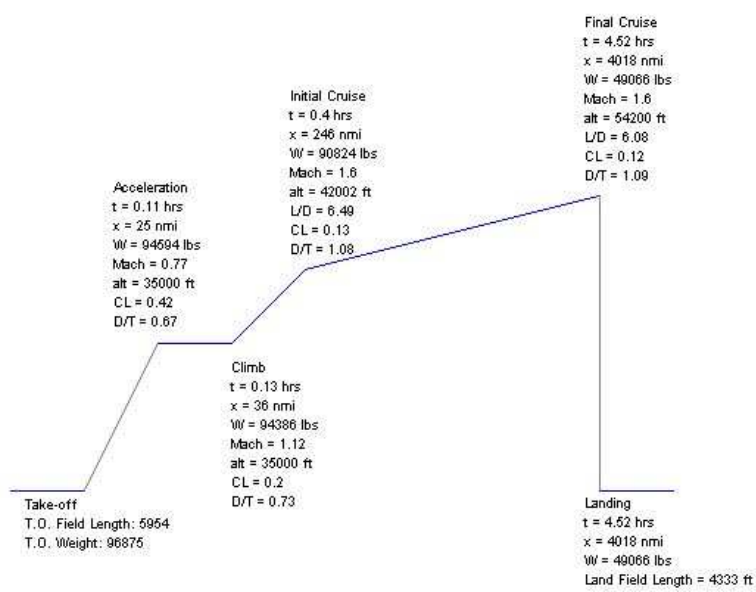

(c) Mission profile

Figure 15. Summary of baseline configuration.

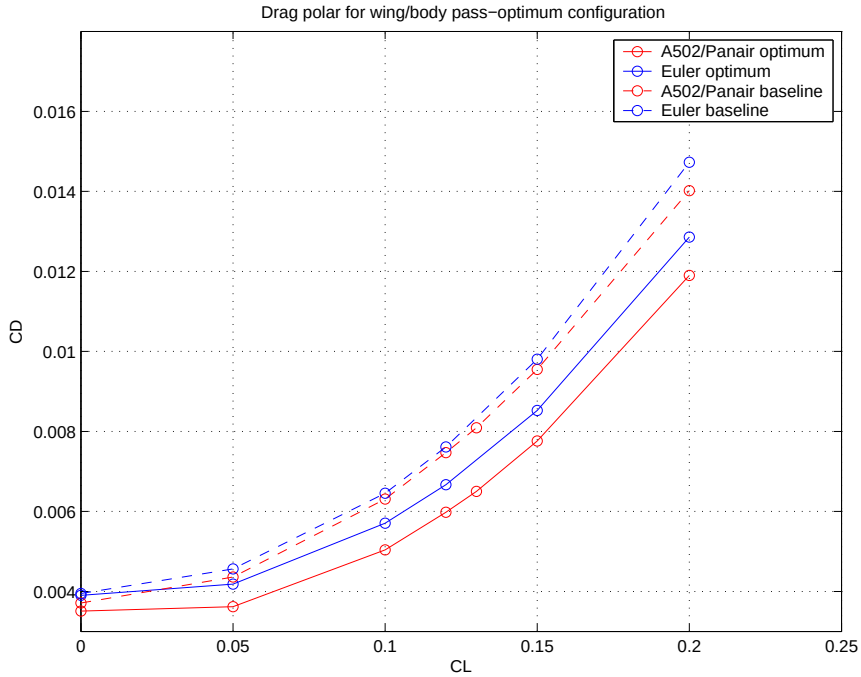

Figure 16. Drag polar of baseline configuration when wing section changes with cambers and twists

C. PASS Optimizations Using Response Surface Fits

Changes were made as necessary to the PASS code to incorporate the C_D fits. For more details of how this was accomplished, the reader is referred to.[20]

1. PASS Optimized Configuration for Minimum MTOW

In this section we present the results of the optimization that used the version of PASS that had been enhanced with the response surface fits created with our multi-fidelity approach. The focus was on minimizing the MTOW of the aircraft while meeting all of the mission requirements.

Views of the fuselage layout, a top view of optimized configuration and the resulting mission profile with key values are shown in Figure 17. Figure 18 shows the side-by-side comparison of the baseline and optimized configurations. The increase of the wing inboard sweep and decrease of the wing AR are noticeable in the wing planform changes. From the front view, we can see that the flat wing of the baseline configuration has received both camber and twist. Minor changes in several fuselage sections have also occurred.

There are several points of interest that must be noted. The response surface fit predicts a lower C_D and a higher C_L (improving L/D significantly) when compared with the direct Euler analysis. As a result of the quality of the response surfaces our optimizations believe that they have achieved a higher performance aircraft than they really have. In fact, PASS predicts an inviscid cruise C_D of 0.0074 at a design C_L of 0.1397, which leads to a high inviscid L/D of 18.774. Euler validation of this configuration results in an inviscid C_D that is larger by 13% with a value of 0.0087.

After careful investigation, the reason for these differences becomes obvious: the Euler analyses used to create the C_D response surfaces are clustered around design C_L's of 0.12 to 0.13. In the meanwhile, the optimizer has determined that a larger C_L would be beneficial to the design (in the neighborhood of 0.14). For this reason, the amount of high-fidelity data points around the final design region is small and the

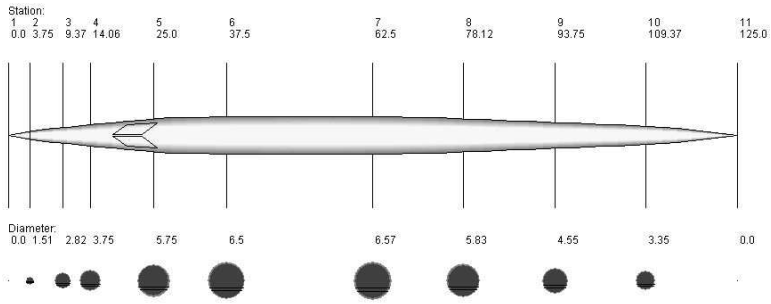

(a) Fuselage layout and stations

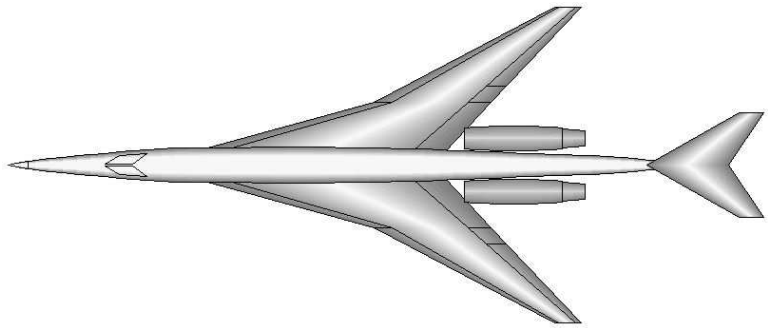

(b) Top view of configuration

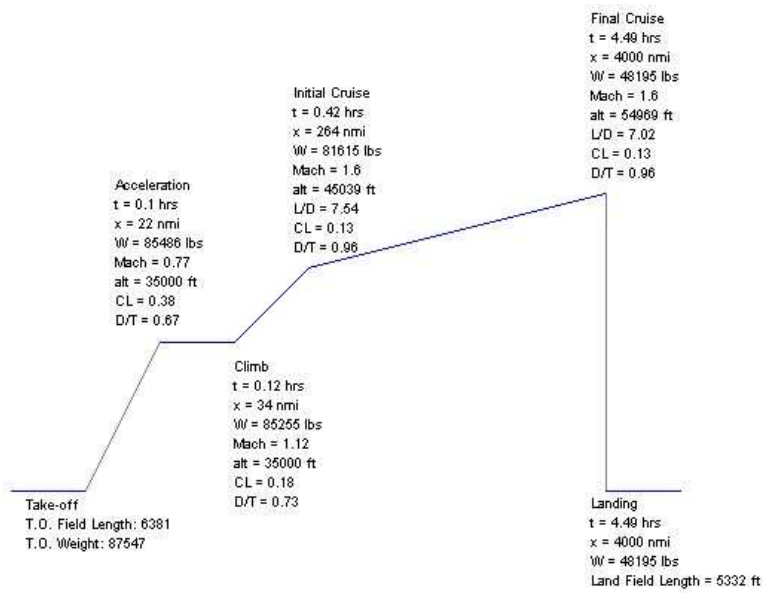

(c) Mission profile

Figure 17. Summary of optimized configuration

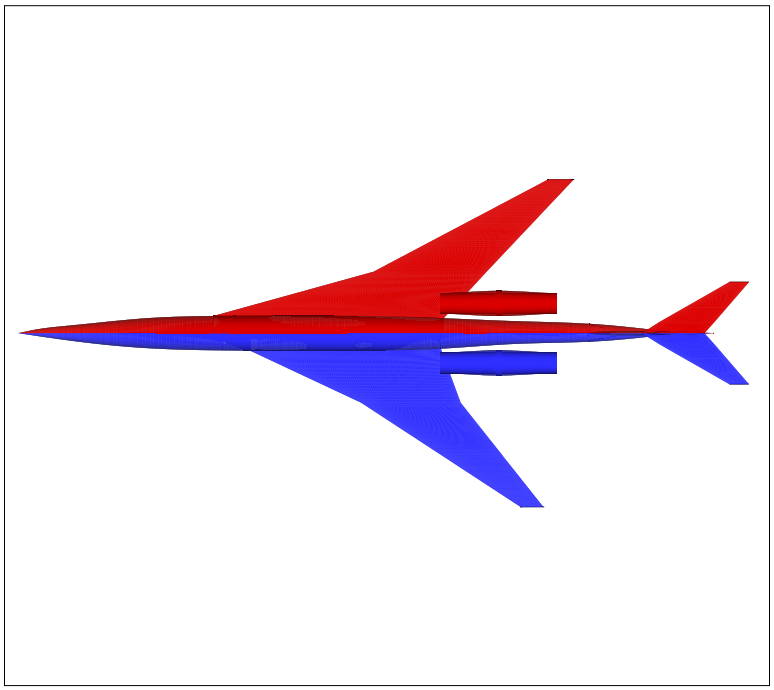

(a) Top view of baseline and optimized configurations

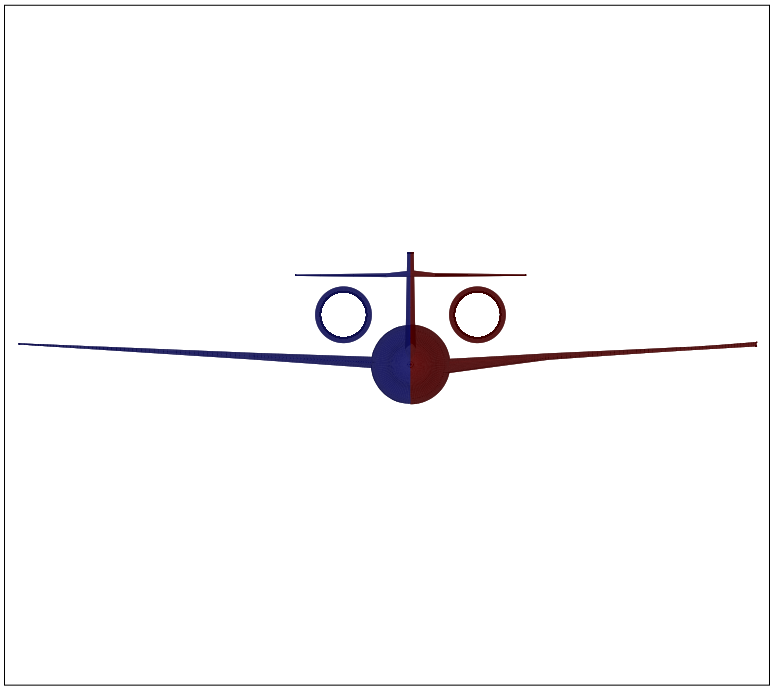

(b) Front view of baseline and optimized configurations

Figure 18. Comparison of shapes of baseline and optimized configuration (blue: baseline, red: optimized).

quality of the fit is reduced. This seems to indicate that additional Euler evaluations in the neighborhood of $C_L = 0.14$ would be needed to improve the agreement between the predicted and achieved designs.

After realizing that the richness of our fits was poor around the 0.14 value of lift coefficient, other optimizations were run that were forced to fly at lower levels of C_L (around 0.12). Surprisingly, the design parameters for those designs were very similar to the previous one already presented with $C_L = 0.1397$. The corresponding Euler validation shows similar differences from the response surface results, and for that reason we have omitted those results here.

In comparison with previous work[20] it appears clear that the number of high-fidelity function evaluations required for a 23-dimensional design space is more in the neighborhood of 500-1000 than in the 200 range that we have used here.

2. *High-Fidelity Validation of Optimization Results*

The aerodynamic performance of the configurations predicted by PASS combined with the response surface fits should be validated with our high-fidelity tool, AirplanePlus. As mentioned above, although the results of the Euler validation are slightly different from what the fit predicted, the optimized configuration shows a good improvement in aerodynamic performance while satisfying all the mission requirements. A comparison of surface pressure distributions (for both the lower and upper surfaces and in side view) is shown below in Figures 19, 20, and 21 The reader should notice a lightly larger fuselage radius around the nose area and increased wing inboard sweep which has reduced the shock strength significantly.

We have also created a multiblock mesh with nearly 6 million nodes for the optimized configuration that we intend to use in future work for complete configuration adjoint designs. Given that this mesh was at hand, it provided us with a unique opportunity to cross-validate all of the Euler results that had been, up to then, produced with our unstructured flow solver, AirplanePlus. A view of the coarsened meshes on the surface of the configuration and the symmetry plane can be seen in Figure 22. The results of a drag polar validation using both codes (the mesh for AirplanePlus had nearly 1.6 million nodes) can be seen in Figure 23. The results cannot be more satisfying since they provide nearly identical solutions throughout the range of C_Ls. This is important because the airfoils across the span of the configuration have rounded leading edges, but, because of the low thickness-to-chord ratios, it is quite hard to put enough grid resolution around the leading edge using nearly isotropic unstructured meshes. With the multiblock approach, anisotropic cells are easily created around the leading edge and can resolve the effects of leading edge curvature rather nicely. This means that the unstructured Euler solutions are just as capable of doing so. As an aside, we had thought earlier on that some of the discrepancies between the Euler solvers and A502 were due to the inability of the Euler solver to capture (with a coarse leading edge mesh) the leading edge suction. This drag polar comparison seems to indicate that this is not the case.

D. Adjoint Optimization

The final section of this paper involves the result of wing-body adjoint-based optimizations using SYN87-SB where the only design variables in the problem correspond to a detailed parameterization of the twist and camber distributions on the wing. Everything else remains unchanged from the values of the optimized configuration. For this purpose a drag minimization calculation at a $C_L = 0.14$ was carried out. The twist and camber distributions of the wing were parameterized at 7 defining stations with 18 design variables (leading and trailing edge droop, twist and 15 camber Hicks-Henne bumps) for a total of 126 design variables. The optimization was allowed to run for 50 design iterations, at which point it was stopped (although full convergence had not been reached, the results were making very small drag improvements, of the order of a tenth of a count of drag). At the design C_L the computed value of the drag for the baseline wing-body configuration (using a block-structured mesh with $257 \times 64 \times 49$ nodes) was 72.02 counts. In 50 design iterations and through twist and camber changes, the drag of the wing-body configuration decreased by almost 5.5 counts.

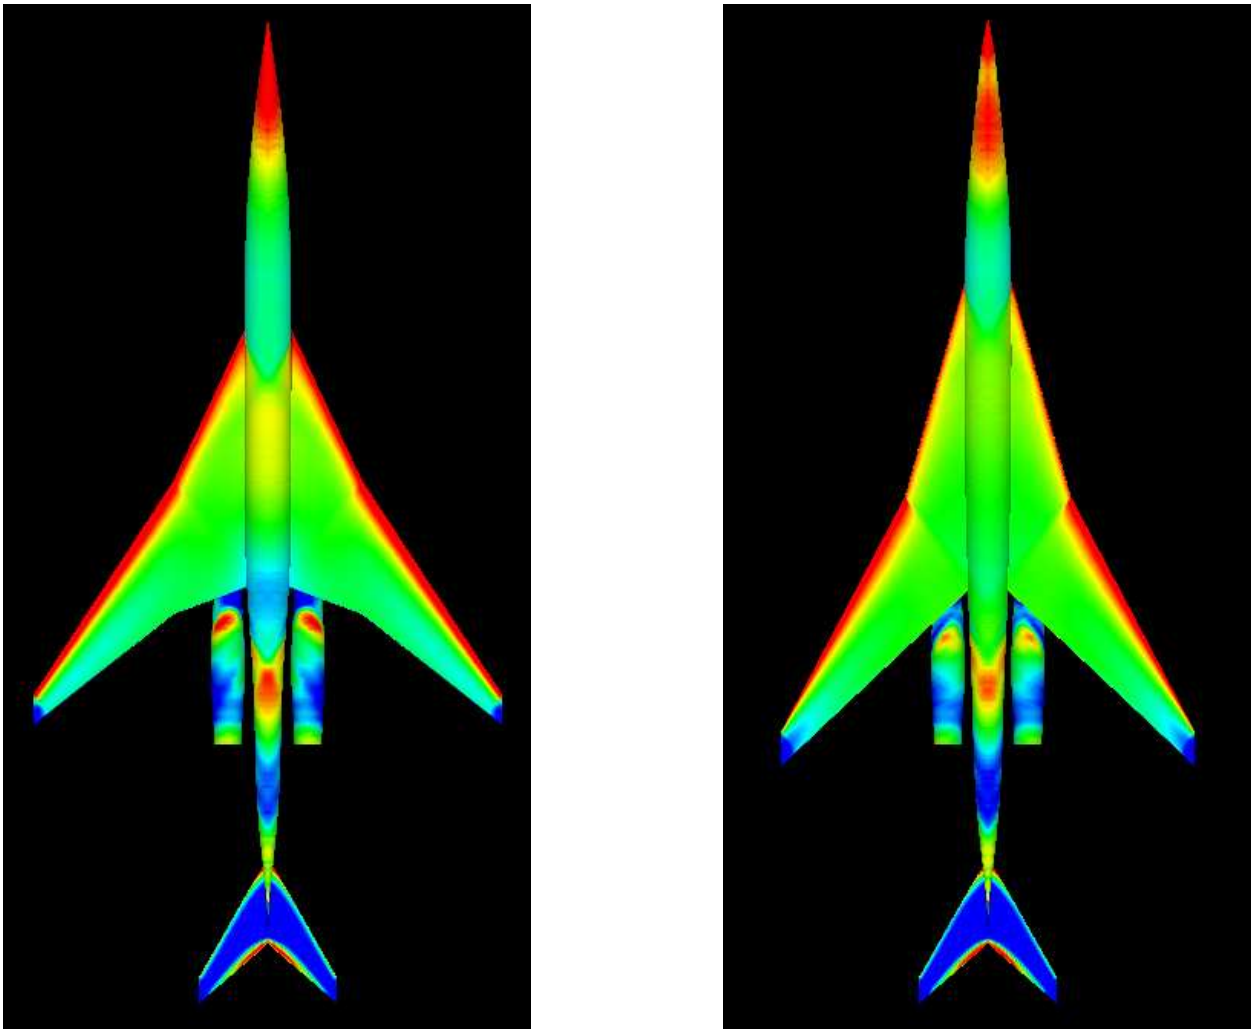

(a) baseline configuration

(b) optimized configuration

Figure 19. Pressure distribution plots - lower surface.

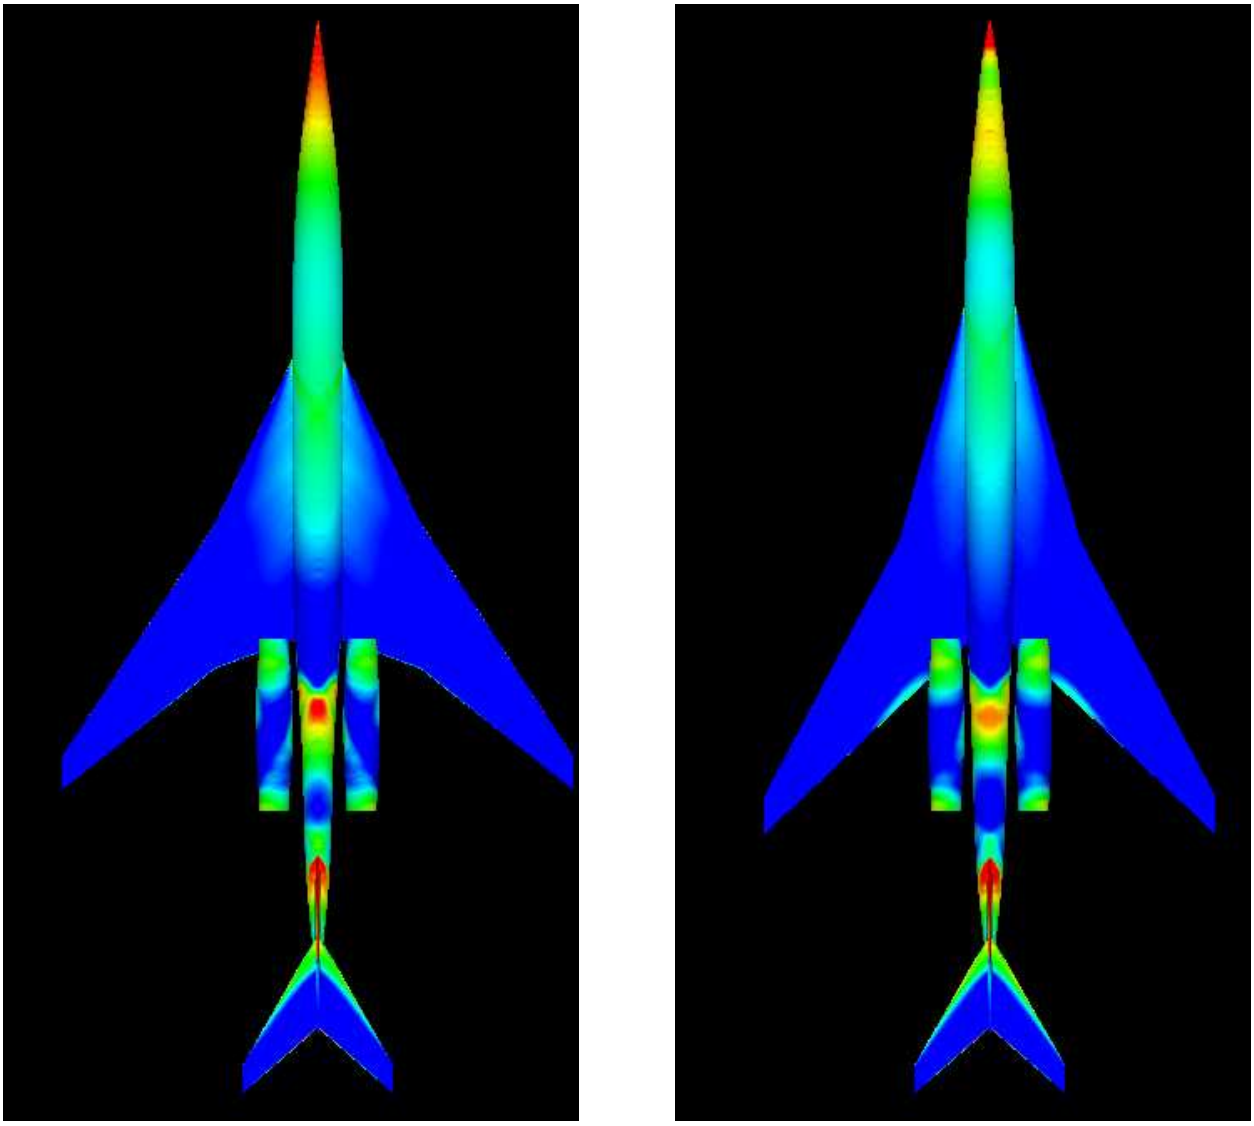

(a) baseline configuration

(b) optimized configuration

Figure 20. Pressure distribution plots - upper surface.

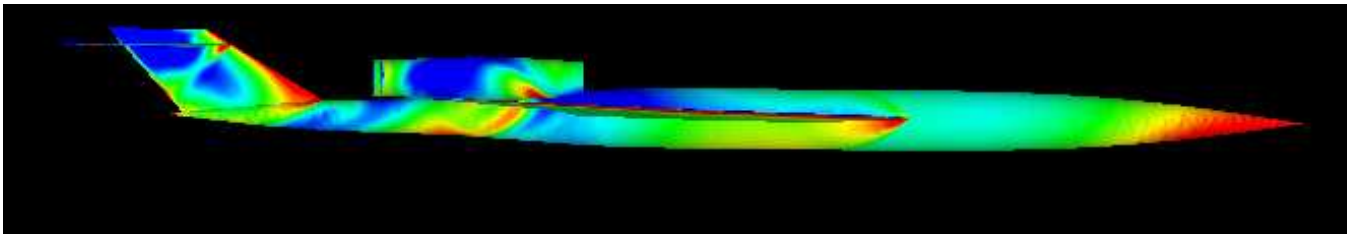

(a) baseline configuration

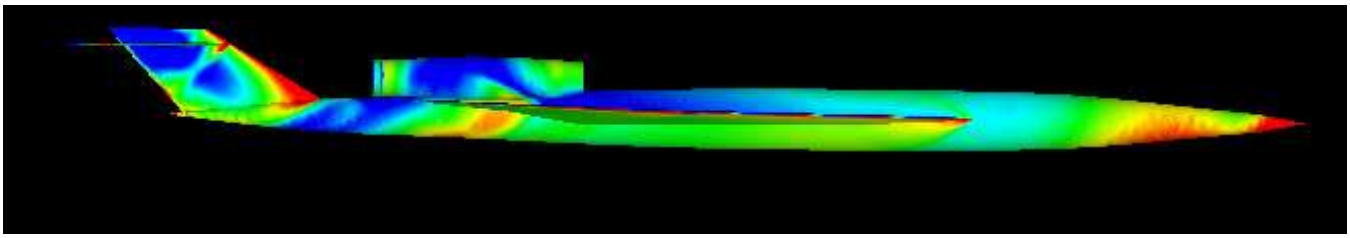

(b) optimized configuration

Figure 21. Pressure distribution plots - side view.

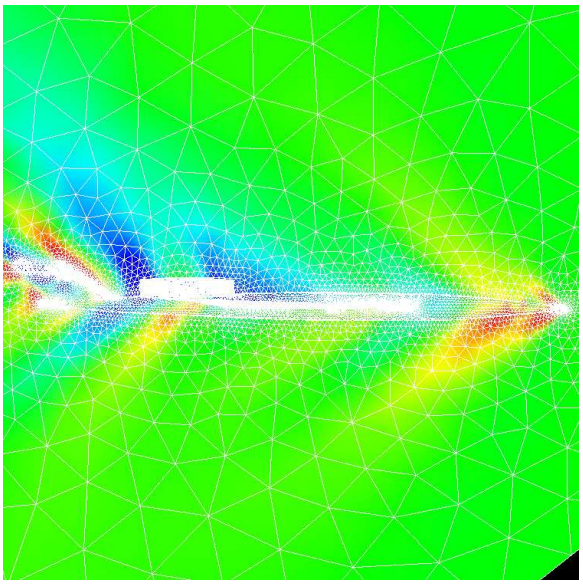

(a) Unstructured mesh

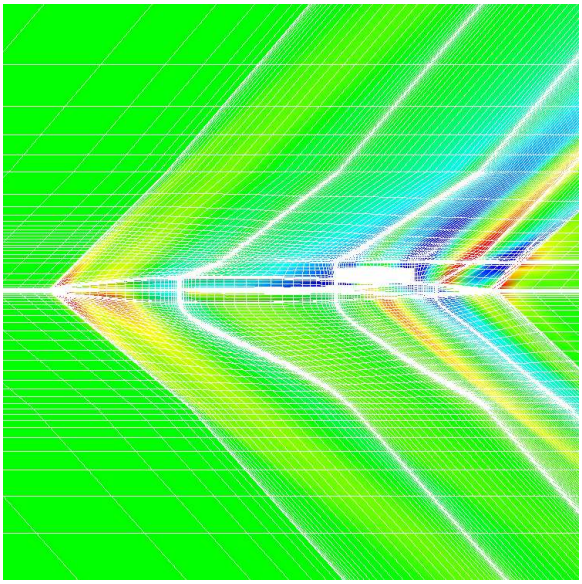

(b) Stuctured multi-block mesh

Figure 22. Mesh topologies for AirplanePlus and FLO107-MB analyses.

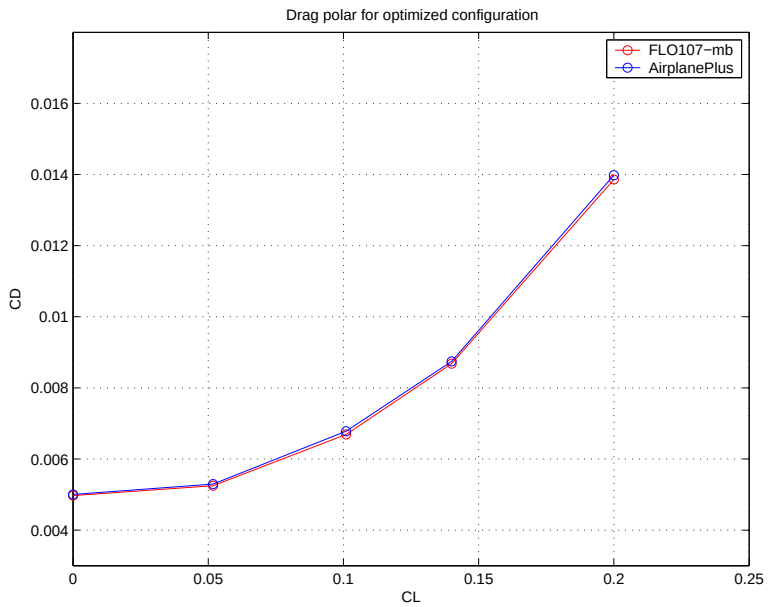

Figure 23. Drag polars using FLO107-MB and AirplanePlus for optimized configuration

This kind of optimized result could then be fed through the PASS procedure to find out the resulting improvement in performance that turns out form this additional design work.

VII. CONCLUSIONS AND FUTURE WORK

In this paper we have presented and evolution of the methodology[20] for the design of supersonic jets using a multi-fidelity approximation to the response of the vehicle 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, although only the C_D was considered in this work. 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.

The combination of a hierarchy of models for the prediction of the aerodynamic performance appears to be quite effective but one needs to pay attention to the number of function evaluations of the higher-fidelity solvers in order to ensure that the overall error in the response surface fits is low in all areas of the design space.

The two-level design procedure which complements PASS with high-fidelity, adjoint-based designs of the PASS-optimized configurations appears to produce significant performance improvements when compared to the Euler evaluation of the PASS results. We intend to pursue this effort a bit further in order to determine whether, in general, once can carry out PASS-based designs solely with low-fidelity tools while ensuring that an adjoint-based redesign can recover designs whose performance is close to that provided by these "best-case" scenario tools.

Further work will also expand this procedure to treat more systematically the boom problem. In particular an area that needs to be addressed is the creation of accurate response surface fits for various measures of the loudness of the sonic boom signature.

References

[1]National 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.

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

[3] H. 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.

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

[5]J. 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.

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

[7]M. Chan, "A Coupled-Adjoint Method for High-Fidelity Aero-Strucural Optimization," PhD thesis, Stanford University, Stanford, CA 94305, October 2002.

[8]D. 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.

[9]Thomas 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

[10]Jaroslaw 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.

[11] J. 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.

[12]J. 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.

[13]Resit 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

[14]A. Jameson "Advances in Aerodynamic Shape Optimization," *3rd International Conference on Computational Fluid Dynamics (ICCFD3)*, Toronto, Canada, July 12-16, 2004

[15]N. 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*

[16]S. 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.

[17]S. 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.

[18]S. 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, Neveda, Jan 2004.

[19]H. 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.

[20]S. I. Choi, J. J. Alonso, I. M. Kroo, and M. Wintzer "Multi-fidelity Design Optimization of Low-boom Supersonic Business Jets," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA 2004-4371 Albany, NY, September 2004.

[21]H. 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.

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

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

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

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

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

[27]J. J. Alonso, J. R. R. A. Martins, J. J. Reuther, R. Haimes, 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.

[28]J. J. Alonso, P. LeGresley, E. van der Weide, J. J. R. R. Martins "pyMDO: A Framework for High-Fidelity Multi-Disciplinary Optimization," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA 2004-4480 Albany, NY, September 2004.

[29] *Centaur System Users' Manual*, *http://www.centaursoft.com*

[30]A. 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

[31]Carmichael,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.

[32]A. Jameson and J. J. Alonso Automatic Aerodynamic Optimization on Distributed Memory Architechtures *AIAA Paper 96-0409*, 34th Aerospace Sciences Meeting, January, 1996, Reno, Nevada

[33]A. Jameson and N. A. Pierce and L. Martinelli Optimum Aerodynamic Design using the Navier Stokes Equation *Theoretical and Computational Fluid Dynamics*, vol. 10, pp. 213–237, 1998.

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

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

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