AIAA 2002–0144

Advanced Algorithms for Design and Optimization of Quiet Supersonic Platforms

Juan J. Alonso, Ilan M. Kroo
and Antony Jameson
Stanford University, Stanford, CA 94305

40th AIAA Aerospace Sciences Meeting and Exhibit
January 14–17, 2002/Reno, NV

AIAA 2002–0144

Advanced Algorithms for Design and Optimization of Quiet Supersonic Platforms

Juan J. Alonso*, Ilan M. Kroo†
and Antony Jameson‡
Stanford University, Stanford, CA 94305

This paper describes the work carried out within the Stanford University group as part of the DARPA-funded Quiet Supersonic Platform (QSP) project. The objective of our work was to develop advanced numerical methods to facilitate the analysis and design of low sonic boom aircraft. The focus of the boom reduction activities was placed on two main ideas: the shaping of the configuration and a multidisciplinary design approach to minimize its total empty weight. Accurate and efficient tools are needed to achieve these tasks. Our approach was meant to enhance the existing state-of-the-art which was developed in the 1970s and during the early stages of the High Speed Research (HSR) program. For that purpose, we designed tools that would ultimately become fully nonlinear (in contrast with current design efforts based on linear theories of the 1970s). These tools are also able to account for important tradeoffs between boom reduction and aerodynamic perormance, as well as other disciplines. Because of the difficulty of the problem, the tools were designed to support a combination of gradient and non-gradient automatic optimization techniques. As a result of our efforts, rapid turnaround boom analysis and design methods were developed. These methods permit the investigation of radical configuration changes and their effect on the ground boom signature and the L/D ratio of the aircraft. We have also created an environment for the automatic nonlinear analysis of QSP configurations based on the multiblock flow solver FLO107-MB. In addition, we have extended the adjoint-based design method to treat remote sensitivities to near-field pressure signatures, allowing for a very large number of parameters to be used in modifications of the aircraft geometry. All tools were carefully validated against existing experimental data, other boom prediction programs, and with systematic mesh refinement studies.

Introduction

THE Quiet Supersonic Platform (QSP) program was created by DARPA in the fall of 2000 in order to reassess the available technologies necessary to develop small supersonic aircraft with sufficiently low sonic boom that they may be allowed to fly supersonically over land. The program requirements stated that such an aircraft would have to be in the 100,000 lbs class, fly at a cruise Mach number of 2.4 with a range of 6,000 nautical miles, and produce an initial overpressure of less than 0.3 psf. These design requirements were issued as guidelines and were subsequently revised so that the goal would be more realistic.

Within the DARPA QSP program, a number of groups were funded to investigate topics in both the contributing technologies (RA-00-47) and in airframe integration (RA-00-48). An important part of this effort, the design of acceptable propulsion systems for this aircraft, was also included, although it had existed previously.

The Stanford University Group decided to apply its expertise in Multidisciplinary Design Optimization (MDO), advanced nonlinear analysis and design methods, and adjoint-based design methodologies to the QSP problem. The objective of our work became to develop the necessary methods and tools to facilitate the high-fidelity, multidisciplinary design of low sonic boom aircraft that can fly supersonically over land with negligible environmental impact. Although the fundamental challenge was to devise a configuration that satisfied strict sonic boom requirements, care was taken so that other mission requirements were met.

Two fundamental ways of reducing the magnitude of the sonic boom of a supersonic aircraft were considered in this research. First, a new adjoint-based, inverse design methodology for the tailoring of the footprint of the overpressures was derived and implemented to treat surface modification of complex aircraft configurations. Second, multidisciplinary optimization concepts were applied to configurations of interest to reduce the weight and boom intensity while maintaining vehicle performance constraints. This MDO environment can take advantage of accurate CFD-based

*Assistant Professor, Department of Aeronautics and Astronautics , AIAA Member
†Professor, Department of Aeronautics and Astronautics, AIAA Fellow
‡Thomas V. Jones Professor of Engineering, Department of Aeronautics and Astronautics, AIAA Fellow

aerodynamic computations and efficient gradient estimation to achieve design improvements that would not be possible with conventional methods.

Analytic or semi-analytic formulations of the sonic boom minimization problem can provide invaluable insight into the development of new concepts, theoretical limits for sonic boom improvements, and promising candidate configurations. However, alone they may not be sufficient to achieve ground signatures with initial shock strengths lower than 0.3 lbs/ft^2. Detailed numerical simulations are necessary to account for the non-linear effects present in this problem. Fortunately, due to the nature of the physics involved in this problem, the solution of the Euler equations has been shown to be sufficient at a preliminary design level. Therefore, an Euler model was used in the majority of the design calculations in this work. Since the nonlinear methods developed in this work were Navier-Stokes based, it would be easy to use RANS solutions of candidate configurations to verify that the differences with the inviscid model are negligible.

For typical cruise altitudes required for aircraft efficiency, the distance from the source of the acoustic disturbance to the ground is typically greater than 50,000 ft. A reasonably accurate propagation of the pressure signature can only be obtained with small computational mesh spacings that would render the analysis of the problem (let alone the design calculations) intractable for even the largest parallel computers. An approach that has been used successfully in the past is the use of near- to far-field extrapolation of pressure signatures based on principles of geometrical acoustics and non-linear wave propagation. These approaches, such as the F-function method of Whitham[1] or the equivalent wavefront parameter method of Thomas,[2] are based on the solution of simple ordinary differential equations for the propagation of a pressure signature from the near-field of the aircraft to the ground (far-field). They are able to take into account variable atmospheric properties and their cost of computation is practically negligible.

While these relatively simple wave propagation methods were used during the first phase of the work, more advanced propagation methods can easily be incorporated into our framework. This possibility was to be investigated during the second phase of the work.

The calculation of a detailed near-field CFD solution, has been shown to require extremely fine meshes[3] in order to obtain accurate far-field results. The number of mesh points required for these calculations is typically of the order of $5-20$ times larger than that required for a typical aerodynamic performance calculation. For a complete configuration supersonic aircraft, detailed studies have shown that around $5-10$ million mesh points are necessary. Depending on the number of design variables that are used to describe the shape of the configuration, the cost of optimization based on these large meshes can be prohibitive unless a combination of advanced adjoint-based algorithms and efficient parallel computing implementations is used. The objective of our work was to ensure that the the turnaround of these calculations fitted within a realistic design environment.

While a principal focus of our work was on sonic boom minimization, this can actually be accomplished only in the context of vehicle design. Vehicle shape modification that leads to reduced boom overpressures for a specific design often results in degraded aerodynamic performance[4] and so boom reduction must be considered as one aspect of a multidisciplinary design problem. Our group has had considerable experience in aerodynamic shape optimization and multidisciplinary aircraft design,[5–15] and the combination of these two research areas represents an opportunity for further advances.

Historically, a large body of work in sonic boom minimization already existed. In particular, most of the sonic boom propagation algorithms that are used today were developed during the 1960s and 70s. During the late 80s and early 90s, the HSR program revived the interest in low sonic boom aircraft, to the extent that quietness was emphasized during the first phase. Once it became apparent that the weight of the aircraft was such that reasonably low noise levels would be impossible to achieve, the low boom component was dropped from the HSR requirements. The work done during this period, however has been extremely helpful to jump start our efforts.

Linear Analysis Methods

In parallel with the development of high-fidelity nonlinear methods for aerodynamic analysis and design, a suite of linear methods was identified, developed, and integrated into a design framework. These linear methods, including classical axisymmetric approaches to boom analysis along with fully 3-D surface panel methods, are important to the present study for several reasons: they serve as fast surrogates for the emerging nonlinear methods, allowing testing of design approaches and a better understanding of the design space; they are dissipation-free, providing important data for studies of required survey locations without the complicating issues of grid resolution and shock smoothing; they provide a means for direct comparison with previous work; and they are useful for concept evaluation. During the course of last year, three aspects of the linear method development were pursued.

A classical boom analysis method (equivalent area plus propagation code) was combined with a drag estimation code and graphical interface to form an interactive QSP design tool. This rapid design tool was used to evaluate and recommend possible conceptual vehicle approaches for other QSP contractors including DTI, Northrop-Grumman, and Boeing. It has also been provided to Lockheed, Gulfstream, and Raytheon, for use as an exploratory tool. This analysis code was also integrated into an optimization framework as described later and proved useful in evaluating strategies for numerical optimization. Figure 1 shows the interactive version of this tool applied to the design of a generic low boom QSP vehicle that was later used for testing and validation of other codes.

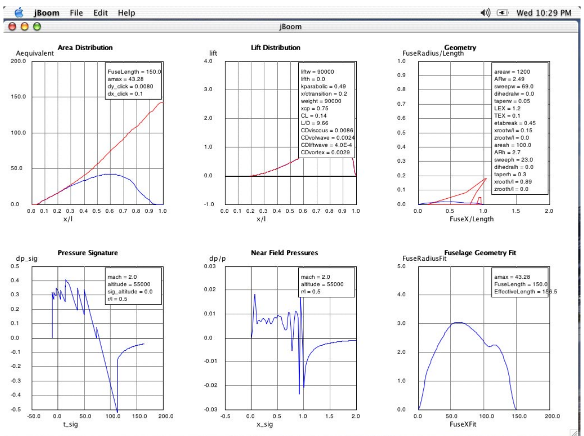

Fig. 1 jBoom interactive design code includes aero cruise performance, boom propagation, center of pressure tracking, and many useful parameters with nearly instant feedback to a designer.

The second area of linear method development included the integration and application of the high order panel code, A502. After substantial testing of several approaches to boom estimation that included the use of programs ranging from the Harris wave drag program, to constant pressure methods (Woodward-Carmichael) such as WingBody, it was determined that true 3D effects required the use of a full surface panel method with off-body surveys used to provide inputs to the propagation code. The A502 program was integrated into the Caffe design framework, providing rapid communication between the conceptual design tool, this more accurate aerodynamic analysis tool, and a suite of optimizers. This integration involved the development of several pieces of software for automated paneling of designs with multiple nonplanar lifting surfaces and extraction of results for the propagation code and for visualization of results. Figure 2 shows the web-based implementation of A502.

The final area pursued as part of the linear analysis work was undertaken as an initial check of the linear codes and to provide a comparison with the nonlinear results. These validation cases included comparisons of several nonlinear and linear codes on configurations ranging from simple bodies of revolution to more complete lifting and nonlifting designs representative of the diverse set of aircraft that might be of interest in the QSP program. Near field pressures and propagated signatures using two different propagation codes were computed and compared with previous results, includ-

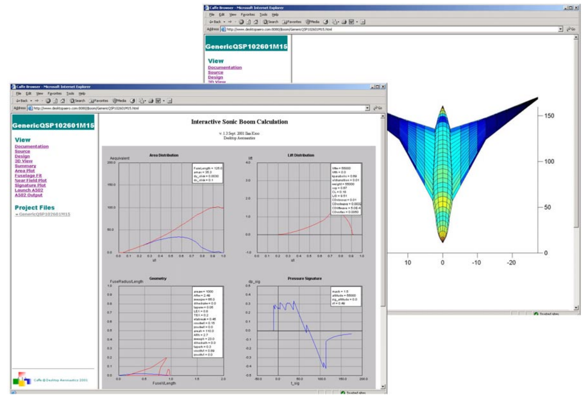

Fig. 2 A502 integration in Caffe

ing some experimental data.

Figure 3 shows a comparison between the axisymmetric code, A502, and Euler calculations for a simple parabolic body of revolution. The character and magnitude of the near field pressure distributions agreed quite well although some differences exist between the codes. This is due not only to differences in the flow physics, but also the fidelity of the geometric modeling.

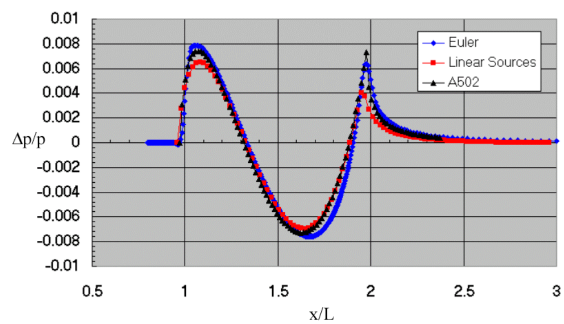

Fig. 3 Comparison of near field pressures using various aerodynamic analysis codes.

Similar comparisons with more complex configurations whose area distribution was developed using jBoom were included in the TEM-4 presentation, showing excellent agreement. With nonaxisymmetric lifting cases, the agreement is less compelling, Results can be quite sensitive to section camber and twist and to the lift carry-over across the fuselage. For low boom configurations most of the high frequency, low amplitude oscillations in the near field pressures disappear in the Euler results and are sometimes quite sensitive to paneling detail in the surface panel code.

In Figure 4 a more complete configuration was analyzed using A502, jBoom, and the Euler code CART3D. This example shows that while certain critical aspects of the aerodynamics may be captured by each of these methods, differences do exist that are important to low boom design.

Several experimental results were available to compare with these codes as well. One of these cases is

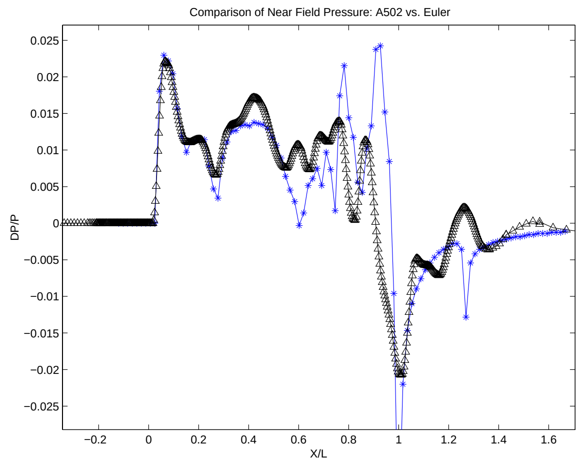

Fig. 4 Generic low boom QSP configuration. Near field signatures from linear and nonlinear codes

shown in Figure 5. The experimental data from NASA TND-7160 is compared with Euler, A502, and equivalent body computations. Agreement here is reasonable although some differences among the codes appear and the Euler result requires very fine meshes to properly capture the peak amplitudes.

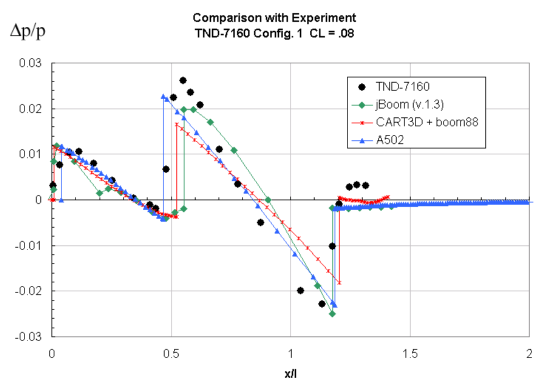

Fig. 5 Comparison of linear and Euler codes with experimental data.

Nonlinear Integrated Boom Analysis Environment

A major outcome of our work for the QSP program has been the development of a fast integrated tool, QSP107, for sonic boom prediction based on fully nonlinear CFD analyses. This tool couples the multiblock Euler and Navier Stokes flow solver FLO107-MB[16] to an H-mesh generator adapted from the HFLO4 code of Jameson and Baker,[17] and to the PC Boom software for far field propagation developed by Wyle Associates. The H-type mesh, which is generated automatically from the geometric definition, can handle arbitrary wing-fuselage-canard configurations. The grid is further adjusted to have higher resolution in the areas where shock waves and expansions are present, and its grid lines slanted at the Mach angle to maximize the resolution of the pressure signature at distances of the order of one body length. The user may specify the location of an arbitrary cylindrical surface where the near field signature is provided as an input to a modified version of PC Boom which propagates a full three-dimensional signature to the ground along all rays that reach the ground. This allows for the calculation of arbitrary cost functions (not only ground-track initial overpressure) that involve weighted integration of the complete sonic boom footprint. The flow solver combines advanced multigrid procedures and a preconditioned explicit multistage time stepping algorithm which allows full parallelization. Consequently the integrated tool provides fully nonlinear simulations with very rapid turn around time. Using a mesh with 3 million mesh points we can obtain a complete flow solution and ground signature prediction in 15 minutes, using 16 processors on a Beowulf cluster made up of 1.2Ghz AMD Athlon processors.

Additionally, a geometry generation module for QSP107 has been developed. This module allows for the generation of a complete sonic boom footprint using a number of parameters to describe the geometry. This is the only input required to QSP107; all other procedures are fully automated. Figure 7 shows a flowchart with the various modules that form part of QSP107. Figure 8 below presents a view of the typical outcome of QSP107 including a bottom view of the Mach number distribution and the symmetry plane wave pattern.

During the development of our integrated fast nonlinear boom prediction tool, QSP107, we carried out extensive validation studies in order to obtain a better understanding of the sensitivity of the result to parameters such as the number of mesh points, flow solution convergence, and location of the cutting plane for the near field signature. These studies were reported on during TEM-1 and TEM-2. We became aware that these sensitivities vary substantially with different configurations, and it seemed extremely desirable to validate QSP107 against an alternative nonlinear method. We identified the MIM3D software developed by Mike Siclari, formerly of Northrop-Grumman, as the most advanced such method.

Accordingly we asked Mike Siclari to join our program as a consultant, and obtained MIM3D from NASA. MIM3D (Multigrid Implicit Marching 3D) assumes a fully supersonic flow field, and obtains a solution by marching in the streamwise direction with an implicit scheme based on trapezoidal integration. This requires the simultaneous solution of the flow field at all the points in each new cross-plane as the solution is advanced downstream. A multigrid method developed in collaboration with Jameson is used to calculate the cross-plane solutions, typically requiring about 100 iterations for each marching step. MIM3D has an integrated mesh generator using a conical topology designed to maximize the resolution away from the body, and it also provides coupling to the NF Boom propagation code. The space marching procedure yields solutions with quite fast turnaround, of the order of 1-2 hours on a workstation, but limits MIM3D to fully supersonic flow fields, whereas QSP107 can treat mixed supersonic and subsonic flows. Moreover, the cost of the time marching scheme in QSP107 is not much greater than the space marching procedure in MIM3D, because the multigrid procedure in QSP107 produces adequately converged solutions in 200-300 time steps, while MIM3D requires of the order of 100 iterations in each cross-plane. When run on a parallel computer QSP107 can in fact provide substantially faster turnaround.

However, MIM3D has provisions in its mesh generation procedure for including engines, while QSP107 could not automatically treat engines unless it were provided with an externally generated mesh. Thus we regard the tools as complementary, and for that purpose we are using MIM3D to compare the results of the two codes on a suite of representative test cases. At the same time we also plan to pursue further work on the use of unstructured meshes for the analysis of complete configurations, and in order to address the problem of engine integration.

Optimization Problem Formulation

The first step in addressing QSP design as an MDO task involves the formulation of the problem. This includes identification of suitable design variables, constraints, and objective(s). In an initial attempt at understanding some of the issues involved in the problem formulation for this program, the rapid turnaround linear analysis methods were integrated with several optimization methods and various design problems were investigated. Variables associated with fuselage tailoring, wing location, and planform shape were selected as design parameters with bounds on their values and other constraints imposed on the problem. In order to simplify computation of near-field pressures and volume wave drag, the equivalent area distribution was represented with a Fourier series, the coefficients of which were taken as design variables. This produced smooth, rapid function evaluations, but since these design variables were related in a complex way to the payload constraints, and since the sensitivity of other constraints was also related in very nonlinear ways to these coefficients, the optimization problem was very difficult to solve. Alternatively, when specific stations on the fuselage were selected as control points in an Akima spline representation of the fuselage radius distribution, the problem became much more tractable. This is simply an example of the sensitivity of optimization to problem formulation and the need for experience in this problem domain to achieve an effective formulation.

Similar difficulties were encountered in the selection of objective and constraints. Initially a composite objective function consisting of a weighted combination of maximum overpressure and L/D was selected along with constraints of cabin dimensions, c.g. location, and other parameters. The optimizer achieved a multi-shock solution with 4-5 shocks, each with a maximum overpressure of 0.5 psf. This indicated that the scheme was functioning correctly, but this was not consistent with the QSP goals. When maximum overpressure was replaced by initial boom overpressure in excess of 0.3 psf, the optimizer produced a design with the desired 0.3 psf initial overpressure and a very reasonable L/D. Closer examination showed that the initial shock was followed only a few milliseconds later by a much stronger shock. The objective was changed to include a measure sensitive to overpressure in the first 10 milliseconds of the signature. This too was exploited by the optimizer, and after some experimentation an objective that included maximum overpressure, initial shock strength, aft shock strength, and finally some measure of noise (dbA) produced more acceptable optimal solutions.

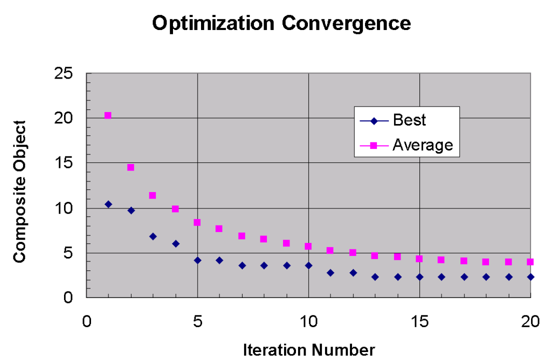

Fig. 6 Iteration history for evolutionary algorithm. Composite objective includes metrics related to boom and aero performance.

Figure 6 illustrates progress made by the optimizer using this composite objective in 8 design variables, including fuselage, wing planform, and wing location.

Optimization Methodology

The minimization of sonic boom with simple constraints is a problem that is usually handled by a simplified parameterization of the f-function and inverse design. Because other design considerations are important to this problem and are not easily incorporated into an inverse problem, an obvious alternative is to employ nonlinear optimization and direct analysis. This has been undertaken in the past, but has generally led to unspectacular reductions in signature. Indeed, in many of these references, it is difficult to see significant improvement in the optimized designs. One of the reasons for frequent failure of search methods in boom minimization is the highly nonlinear character of the relation between geometric design variables and signature-based objective. Experience with the interactive design code provides a clear indication of difficulties with conventional optimization. The sudden appearance of multiple shocks, the coalescence of shocks, and the change in sensitivity of shock peaks depending on the strength of other peaks, leads to one of the most difficult aerodynamic optimization problems that the authors have encountered. The design space is generally filled with many local optima and, depending on the details of the aero analysis and propagation code, may exhibit discontinuities in objective or constraint gradients.

To better understand the character of the design space, and search methods that are effective for this problem, several optimization problems were investigated in the first phase of the program. As described in the previous section several problem formulations were investigated, but several different search methods were also employed. The conclusion of this study was that for the complete design problem (including geometric design variables, multidisciplinary constraints, and boom/performance objectives), the presence of multiple local minima, created by the physics of the problem itself or introduced by noise in the computations, prevented successful application of conventional gradient-based search methods. Optimization codes with which the authors are quite familiar, including workhorse SQP methods such as SNOPT, achieved little progress on this problem, for the reasons mentioned above.

Alternative optimization methods that do not require gradient information (e.g. nonlinear simplex) showed somewhat better performance, but consistently converged on local minima. This led to exploration of less efficient but more robust schemes such as genetic algorithms and other stochastic approaches. The GA was most effective, although very inefficient, consistently finding a better solution than could be achieved by the other optimizers, but at the cost of many thousands of function evaluations even for 8-10 variable problems. The results shown in figure 6 are typical of those achieved with this method. The figure shows the best and average fitness in a population of 50 individuals, with convergence usually achieved after about 30-50 generations. Since these nondeterministic algorithms must generally be run several times, such design problems are feasible only with very rapid analysis codes or with highly parallel implementations.

This suggests that one approach to these problems is to combine robust but inefficient search methods with powerful parallel computational facilities. This may prove useful, but remains limiting as the number of design variables is increased. A more subtle approach is to combine rapid linear methods and robust search strategies with higher fidelity nonlinear analysis and gradient based design.

Coupled Adjoint Design Method

This section presents an overview of the coupled adjoint design method for sonic boom reduction. The main objective of this derivation is to formally present the method that allows for the calculation of sensitivities of cost functions based on the ground signature to an arbitrarily large number of design parameters with a single flow and adjoint solution. The discussion in this section uses the nomenclature presented in Figure 9. The Figure clearly shows the aircraft flying at a cruise altitude h, the ground plane, and an arbitrarily located near-field plane which creates the interface between the CFD solution and the far-field propagation method. Note that the external boundary of the CFD mesh may extend beneath the near field plane for numerical convergence issues.

Assume that the external shape of the configuration is given by $\mathcal{F}(\vec{b})$, where $\vec{b}$ is the vector of design variables such that $\vec{b} \in R^N$. These variables can include anything from global surface arrangement, position, and shape details. The three-dimensional space around the configuration is discretized using a mesh of appropriate density. This mesh will be used for both the flow and adjoint solutions that are detailed below.

At a pre-specified distance below the aircraft, and still within the CFD mesh, the location of a near-field plane can be seen. This plane is the effective interface between the CFD solution and the wave propagation program. At the near-field plane, the flow solution w_0 can be extracted. Given these initial conditions, w_0, the propagation altitude, and the altitude-dependent atmospheric properties $\rho(z), p(z), T(z)$, the propagation method produces a flow solution at the ground plane we are interested in, w_G, which can be used to determine any of a variety of measures of sonic boom impact such as overpressure, rise time, impulse, perceived noise level, etc.[18]

Regardless of the type of cost function that we decide to optimize, we will be looking to minimize a scalar function $I = I(w_G)$ using any of a variety of gradient-based optimization algorithms. For this purpose, it is necessary to obtain sensitivities of the cost function, I, to the design variables in the vector $\vec{b}$, $\frac{\partial I}{\partial \vec{b}}$.

There are several ways in which the gradient vector can be obtained, but their dependence on the number of design variables in the problem, N, can be quite different. Our objective is to find $\frac{\partial I}{\partial \vec{b}}$ in a manner which is independent of N so that design iterations can be performed at a minimum computational cost.

In order to emphasize the power of the coupled adjoint procedure, we analyze the calculation of all sensitivities using finite differencing first.

Finite Difference Method

Given that an analysis capability of the sort we have described above is readily available, the most straightforward way of computing design sensitivities is via the traditional finite difference method.

The simplest approach to optimization is to define the geometry, $\mathcal{F}$, through a set of design parameters, which may, for example, be the weights b_i applied to a set of shape functions $\mathcal{S}_i(x)$ so that the shape is represented as

$$f(x) = \sum b_i \mathcal{S}_i(x).$$

Then, a cost function I is selected which might, for example, be the drag coefficient, the lift to drag ratio, or the impulse of the sonic boom signature, and I is regarded as a function of the parameters b_i. The sensitivities $\frac{\partial I}{\partial b_i}$ may now be estimated by making a small variation δb_i in each design parameter in turn and recalculating the flow to obtain the change in I. Then

$$\frac{\partial I}{\partial b_i} \approx \frac{I(b_i + \delta b_i) - I(b_i)}{\delta b_i}.$$

The gradient vector $\frac{\partial I}{\partial \vec{b}}$ may now be used to determine a direction of improvement. The simplest procedure is to make a step in the negative gradient direction by setting

$$\vec{b}^{n+1} = \vec{b}^{n} - \lambda \frac{\partial I}{\partial \vec{b}},$$

so that to first order

$$I + \delta I = I + \frac{\partial I^T}{\partial \vec{b}} \delta \vec{b} = I - \lambda \frac{\partial I^T}{\partial \vec{b}} \frac{\partial I}{\partial \vec{b}}.$$

If we assume that the computational cost of a single CFD solution is represented by t_{CFD}, while the cost of the near- to far-field propagation procedure is given by t_{WAVE}, the total cost of computation of a single evaluation of the gradient vector, C_{FD}, is given by

$$C_{FD} = (N+1)(t_{CFD} + t_{WAVE})$$

Clearly, this cost is directly proportional to the total number of design variables in the problem. Although the procedure is quite simple to implement, the sonic boom minimization problem requires hundreds of design variables to represent the shape of the configuration and therefore this approach is not realistic. However, the finite difference method is often used in our work to validate the results of more sophisticated sensitivity analysis methods, and in situations in which the total cost of the analysis is virtually negligible.

Remote Sensitivity Method

As a first step towards decreasing the cost of optimization, we can use existing adjoint technology on each of the modules that make up the sonic boom analysis procedure. Note that the cost function of interest is given by

$$I = I(w_G),$$

and therefore, the total sensitivity vector can be constructed as follows:

$$\frac{\partial I}{\partial \vec{b}} = \frac{\partial I}{\partial w_G} \frac{\partial w_G}{\partial w_0} \frac{\partial w_0}{\partial \vec{b}}. \quad (1)$$

The first term in this expansion can be easily obtained analytically. The second and third terms in Equation 1 are slightly more complicated since they represent the Jacobian matrices of the flow solution and signature propagation subproblems.

The calculation of the last term in Equation 1 using an adjoint approach is both more involved and computationally expensive than the signature propagation procedure. Firstly, the governing equations that relate variations in $\vec{b}$ to variations in w_0 are partial differential equations, not ODEs. Secondly, the computational cost of solving the direct CFD problem is orders of magnitude larger than for the wave propagation method, requiring high-performance parallel computers to fit within the preliminary design cycle.

During the last few years[7,9,19,20] our group has develop adjoint methodologies for the computation of aerodynamic shape sensitivities with very promising results. The methods developed have been successfully used in studies of new aircraft, some of which are flying today. Before proceeding, a short description of the adjoint method, its requirements and advantages are given below.

In order to reduce the computational costs, there are advantages in formulating both the inverse problem and more general aerodynamic problems within the framework of the mathematical theory for the control of systems governed by partial differential equations.[9,21,22] A wing, for example, is a device to produce lift by controlling the flow, and its design can be regarded as a problem in the optimal control of the flow equations by variation of the shape of the boundary. If the boundary shape is regarded as arbitrary within some requirements of smoothness one must use the concept of the Frechet derivative of the cost with respect to a function. Using techniques of control theory, the gradient can be determined indirectly by solving an adjoint equation which has coefficients defined by the solution of the flow equations. The cost of solving the adjoint equation is comparable to that of solving the flow equations. Thus the gradient can be determined with roughly the computational cost of two flow solutions, independently of the number of design variables, N.

For the flow about an airfoil or wing, the aerodynamic properties which define the cost function are functions of the flow-field variables (w) and the physical location of the boundary, which may be represented by the function $\mathcal{F}$, say. Then

$$I = I(w, \mathcal{F}),$$

and a change in $\mathcal{F}$ results in a change

$$\delta I = \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta \mathcal{F}, \qquad (2)$$

in the cost function. Using control theory, the governing equations of the flowfield are introduced as a constraint in such a way that the final expression for the gradient does not require re-evaluation of the flowfield. In order to achieve this δw must be eliminated from (2). Suppose that the governing equation R which expresses the dependence of w and $\mathcal{F}$ within the flowfield domain D can be written as

$$R\left(w, \mathcal{F}\right) = 0. \qquad (3)$$

Then δw is determined from the equation

$$\delta R = \left[\frac{\partial R}{\partial w}\right]\delta w + \left[\frac{\partial R}{\partial \mathcal{F}}\right]\delta \mathcal{F} = 0. \qquad (4)$$

Next, introducing a Lagrange Multiplier ψ, we have

$$\begin{aligned} \delta I &= \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta \mathcal{F} - \psi^T\left(\left[\frac{\partial R}{\partial w}\right]\delta w + \left[\frac{\partial R}{\partial \mathcal{F}}\right]\delta \mathcal{F}\right) \\ &= \left\{\frac{\partial I^T}{\partial w} - \psi^T\left[\frac{\partial R}{\partial w}\right]\right\}\delta w + \left\{\frac{\partial I^T}{\partial \mathcal{F}} - \psi^T\left[\frac{\partial R}{\partial \mathcal{F}}\right]\right\}\delta \mathcal{F}. \end{aligned}$$

Choosing ψ to satisfy the adjoint equation

$$\left[\frac{\partial R}{\partial w}\right]^T \psi = \frac{\partial I}{\partial w} \qquad (5)$$

the first term is eliminated, and we find that

$$\delta I = \mathcal{G}\delta \mathcal{F}, \qquad (6)$$

where

$$\mathcal{G} = \frac{\partial I^T}{\partial \mathcal{F}} - \psi^T\left[\frac{\partial R}{\partial \mathcal{F}}\right].$$

The advantage is that (6) is independent of δw, with the result that the gradient of I with respect to an arbitrary number of design variables can be determined without the need for additional flow-field evaluations. In the case that (3) is a partial differential equation, the adjoint equation (5) is also a partial differential equation and appropriate boundary conditions must be determined.

After making a step in the negative gradient direction, the gradient can be recalculated and the process repeated to follow a path of steepest descent until a minimum is reached. In order to avoid violating constraints, such as a minimum acceptable wing thickness, the gradient may be projected onto the allowable subspace within which the constraints are satisfied. In this way one can devise procedures which must necessarily converge at least to a local minimum, and which can be accelerated by the use of more sophisticated descent methods such as conjugate gradient or quasi-Newton algorithms.

As was mentioned above, the cost of obtaining the gradient of I with respect to an arbitrary number of design variables, N, is independent of the number of design variables. However, notice that this is only true for a single scalar function, I.

Since the last term in Eq. 1 refers to the sensitivity of cost functions in the near-field plane to changes on the surface of the aircraft, we typically refer to this implementation of the adjoint as the *Remote Sensitivity Method*. This method has been implemented and tested and is the subject of an accompanying paper.[23] In particular, this method allows for efficient modifications of the surface geometry to recover a near field pressure distribution which is known to translate into desirable shapes of the farfield signature.

If this remote sensitivity method is to be used for inverse design work, it is important to be able to specify realistic near field target pressure distributions that result in ground signatures with the necessary properties. For this purpose, we have developed a finite-difference based optimization loop around the wave propagation algorithm that allows us to recover the near field signature that results in the closest agreement with a specified target at the ground.

A realistic example of this procedure involved a generic QSP configuration obtained from our linear design methods that produced a ramp shape ground boom with the first shock strength at less than 0.3 psf and the second shock strength at less than 0.5 psf. The flight conditions were $M_\infty = 2$, flight altitude $= 55,000 ft$, angle of attack $= 3.5°$, and $R/L = 0.5$. As shown in Figure 10b, due to nonlinear effects that are accounted for by the Euler equations, the initial ground boom signature in *blue* is no longer a ramp shape. In this design example, instead of achieving a complete optimization, the current boom has been improved to satisfy the criterion used for the linear design. Figure 10b shows a reduction in the shock strength and the ramp shape (approximated with multiple shocks) was recovered in 500 wave propagation design iterations. As illustrated in Figure 10a, this case has produced a resulting near field pressure signature which could be achievable by modifications of the surface shape.

Coupled Adjoint Sensitivity Method

The ideal situation for the design of low sonic boom aircraft is one where the computation of the complete gradient vector, $\frac{\partial I}{\partial b}$ has a cost which is independent of N, the number of design variables. The main focus of our work has been to develop such an algorithm. The fundamentals of this approach are outlined in this section.

We proceed by developing separate adjoints for both the CFD and wave propagation modules and we conclude by showing how these two adjoint problems can be coupled and solved without need for an inner iteration.

Consider the near- to far-field wave propagation problem first. Let's consider a coordinate along the direction of information propagation, τ, such that $\tau = \tau_0$ corresponds to the location of the near-field plane, while $\tau = \tau_g$ coincides with the ground plane. Along this coordinate, the governing equations of the wave propagation can be conceptually expressed as

$$\frac{dw}{d\tau} = f(w), \text{ with initial conditions: } \quad w(\tau = \tau_0) = w_0 \tag{7}$$

where w are the flow variables that participate in the propagation problem, and $f(w)$ can be a complicated function, similar in nature to a Whitham F-function, that depends on the particular method of propagation chosen.

The first variation of Equation 7 can be derived as follows:

$$\frac{d}{dx}\delta w = \frac{df}{dw}\delta w, \text{ subj. to ic. } \quad \delta w(\tau = \tau_0) = \delta w_0. \tag{8}$$

Let's assume a cost function of the form

$$I = I(w_G) = I(w(\tau = \tau_G)), \tag{9}$$

which can easily represent all of the potential cost functions in the sonic boom minimization problem. The first variation of the cost function can be expressed as

$$\delta I = \frac{\partial I}{\partial x}\delta w_G, \tag{10}$$

which can be augmented by the variation of the governing equations of propagation 8 without change to produce

$$\delta I = \frac{\partial I}{\partial w}\delta w_G - \int_{\tau_0}^{\tau_G} \phi^T \left(\frac{d}{d\tau}\delta w - \frac{df}{dw}\delta w \right) d\tau, \tag{11}$$

and where ϕ is a costate variable or Lagrange multiplier that can take on any value along the wavefront propagation direction. Integrating Equation 11 by parts we obtain

$$\begin{aligned} \delta I &= \frac{\partial I}{\partial w}\delta w_G - \phi^T(\tau_G)\delta w_G + \phi^T(\tau_0)\delta w_0 \\ &+ \int_{\tau_0}^{\tau_G} \left(\frac{d\phi^T}{d\tau}\delta w + \phi^T \frac{df}{dw}\delta w \right) d\tau. \end{aligned} \tag{12}$$

Since ϕ is completely arbitrary, we are free to choose it so that it satisfies the adjoint equation of the wave propagation method

$$\frac{d\phi}{d\tau} + \left[\frac{df}{dw} \right]^T \phi = 0, \tag{13}$$

subject to boundary conditions

$$\phi(\tau_G) = \left[\frac{\partial I}{\partial w} \right]^T \Bigg|_{\tau = \tau_G}, \tag{14}$$

and then, the expression for the gradient in Equation 12 simplifies to

$$\delta I = \phi^T(\tau_0)\delta w_0. \tag{15}$$

We still require the values of the variation of the flow solution, δw_0, at the near-field plane with respect to the shape design variables, δb. Here we apply the adjoint method for the flow calculation in a manner similar to that described in the previous section.

For flows governed by the Euler equations of fluid motion, it proves convenient to denote the Cartesian coordinates and velocity components by x_1, x_2, x_3 and u_1, u_2, u_3, and to use the convention that summation over $i = 1$ to 3 is implied by a repeated index i. Then, the three-dimensional Euler equations may be written as

$$\frac{\partial w}{\partial t} + \frac{\partial f_i}{\partial x_i} = 0 \quad \text{in } D, \tag{16}$$

where

$$w = \left\{ \begin{array}{c} \rho \\ \rho u_1 \\ \rho u_2 \\ \rho u_3 \\ \rho E \end{array} \right\}, \quad f_i = \left\{ \begin{array}{c} \rho u_i \\ \rho u_i u_1 + p\delta_{i1} \\ \rho u_i u_2 + p\delta_{i2} \\ \rho u_i u_3 + p\delta_{i3} \\ \rho u_i H \end{array} \right\} \tag{17}$$

and δ_{ij} is the Kronecker delta function. Also,

$$p = (\gamma - 1)\,\rho \left\{ E - \frac{1}{2}\left(u_i^2\right) \right\}, \tag{18}$$

and

$$\rho H = \rho E + p \tag{19}$$

where γ is the ratio of the specific heats.

Consider a transformation to coordinates ξ_1, ξ_2, ξ_3 where

$$K_{ij} = \left[\frac{\partial x_i}{\partial \xi_j} \right], \quad J = \det(K), \quad K_{ij}^{-1} = \left[\frac{\partial \xi_i}{\partial x_j} \right],$$

and

$$Q = JK^{-1}.$$

The elements of Q are the coefficients of K, and in a finite volume discretization they are just the face areas of the computational cells projected in the x_1, x_2, and x_3 directions. Also introduce scaled contravariant velocity components as

$$U_i = Q_{ij} u_j.$$

The Euler equations can now be written as

$$\frac{\partial W}{\partial t} + \frac{\partial F_i}{\partial \xi_i} = 0 \quad \text{in } D, \tag{20}$$

where

$$W = Jw,$$

and

$$F_i = Q_{ij} f_j = \begin{bmatrix} \rho U_i \\ \rho U_i u_1 + Q_{i1} p \\ \rho U_i u_2 + Q_{i2} p \\ \rho U_i u_3 + Q_{i3} p \\ \rho U_i H \end{bmatrix}.$$

In order to link the near-field CFD solution with the wave propagation problem, let's define a cost function for the CFD solution as the following "boundary" integral

$$J = \int_{\mathcal{B}} \phi^T w_0 \, d\xi_{\mathcal{B}},$$

such that its first variation coincides with the expression for the gradient of the propagation algorithm in Equation 15. That is

$$\delta J = \int_{\mathcal{B}} \phi^T \delta w_0 \, d\xi_{\mathcal{B}}.$$

Define the Jacobian matrices

$$A_i = \frac{\partial f_i}{\partial w}, \quad C_i = Q_{ij} A_j. \tag{21}$$

The Euler equations (20) in the steady state can be written as

$$\frac{\partial}{\partial \xi_i} Q_{ij} f_j = 0, \tag{22}$$

and their first variation, an equation for δw, can then be obtained as

$$\frac{\partial}{\partial \xi_i} \left(\delta Q_{ij} f_j + Q_{ij} \frac{\partial f_j}{\partial w} \delta w \right) = 0. \tag{23}$$

The cost function can then be easily augmented by Equation 23 to produce

$$\begin{aligned} \delta J &= \int_{\mathcal{B}} \phi^T \delta w_0 \, d\xi_{\mathcal{B}} - \int_{\mathcal{D}} \psi^T \frac{\partial}{\partial \xi_i} \delta Q_{ij} \, f_j \, d\xi \\ &- \int_{\mathcal{D}} \psi^T \frac{\partial}{\partial \xi_i} Q_{ij} \frac{\partial f_j}{\partial w} \, \delta w \, d\xi, \end{aligned} \tag{24}$$

whose last term can be integrated by parts to yield

$$\begin{aligned} \delta J &= \int_{\mathcal{B}} \phi^T \delta w_0 \, d\xi_{\mathcal{B}} - \int_{\mathcal{D}} \psi^T \frac{\partial}{\partial \xi_i} \delta Q_{ij} \, f_j \, d\xi \\ &- \int_{\mathcal{B}} \psi^T n_i Q_{ij} \frac{\partial f_j}{\partial w} \, \delta w_0 \, d\xi_{\mathcal{B}} \\ &+ \int_{\mathcal{D}} \frac{\partial \psi^T}{\partial \xi_i} Q_{ij} \frac{\partial f_j}{\partial w} \, \delta w \, d\xi, \end{aligned} \tag{25}$$

Since Equation 24 should hold for an arbitrary choice of the costate vector ψ, we can choose it as the solution of the adjoint equation

$$\frac{\partial \psi}{\partial t} - C_i^T \frac{\partial \psi}{\partial \xi_i} = 0 \quad \text{in } D, \tag{26}$$

subject to "boundary" conditions

$$\psi^T n_i Q_{ij} \frac{\partial f_j}{\partial w} = \phi.$$

These choices eliminate all but the second term in Equation 25 which does not depend on the variations of either the CFD flow solution or the wave propagation program, thus rendering the computation of the sensitivities of the overall cost function I with respect to an arbitrary number of design variables completely independent of N.

Notice that the "boundary" conditions on the CFD adjoint are not true boundary conditions. The near-field plane, as mentioned earlier, is inside the extent of the CFD mesh. Therefore, the boundary conditions we have been discussing will appear as source terms to the CFD adjoint problem. The location of these source terms can be obtained with simple interpolation techniques.

In sum, since $\delta J = \delta I$, the expression for the gradient can be shown to be

$$\delta J = - \int_{\mathcal{D}} \psi^T \frac{\partial}{\partial \xi_i} \delta Q_{ij} \, f_j \, d\xi,$$

which requires the solution of the CFD adjoint equation, ψ. In order to obtain ψ, we must have solved the adjoint of the wave propagation problem subject to appropriate boundary conditions so that ϕ can be used in the "boundary" conditions of the CFD adjoint.

Conclusions

After one year of effort, we have developed a variety of linear and nonlinear methods for the analysis of aircraft with low sonic boom. These include very fast turnaround linear methods and fast turnaround nonlinear methods that must complement each other in an effective design framework. In addition, we have developed multidisciplinary optimization methods based on the fast linear analysis techniques, which can be used in combination with non-gradient based optimization algorithms to determine the overall configuration of the aircraft. In the process, we have learned important lessons regarding the setup of the optimization problem, the parameterization of the design space, and the ability of various optimization algorithms to succeed in such a difficult design problem. Finally, we have also developed an efficient, accurate, nonlinear coupled adjoint method for the inexpensive calculation of the sensitivity of ground signatures to modifications in the aircraft shape. The intention is to use this method after the overall aircraft configuration has been optimized using linear methods to recover the shape that produces the results originally expected. Much work remains to be done to make this type of design environment a reality.

Acknowledgments

The authors acknowledge the support of the DARPA Tactical Technology Office under the direction of Dr. Richard Wlezien for the completion of this work.

References

[1]G. B. Whitham. The flow pattern of a supersonic projectile. *Commun. Pure Appl. Math*, V(3):301–348, 1952.

[2]C. Thomas. Extrapolation of sonic boom pressure signatures by the waveform parameter method. *NASA TN D-6832*, 1972.

[3]S. E. Cliff and S. D. Thomas. Euler/experiment correlations of sonic boom pressure signatures. *AIAA paper 91-3276*, September 1991.

[4]K. J. Plotkin. Review of sonic boom theory. *AIAA paper 89-1105*, 12th AIAA Aeroacoustics Conference, San Antonio, TX, April 1989.

[5]J. J. Reuther, A. Jameson, J. J. Alonso, M. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers: Part I. *Journal of Aircraft*, 36(1):51–60, 1999.

[6]J. J. Reuther, A. Jameson, J. J. Alonso, M. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers: Part II. *Journal of Aircraft*, 36(1):61–74, 1999.

[7]A. Jameson, L. Martinelli, and N. Pierce. Optimum aerodynamic design using the navier-stokes equations. *AIAA paper 97-0101*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.

[8]A. Jameson. *Optimum Aerodynamic Design Using Control Theory, Computational Fluid Dynamics Review 1995*. Wiley, 1995.

[9]A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing, also ICASE Report No. 88-64*, 3:233–260, 1988.

[10]I. M. Kroo. Decomposition and collaborative optimization for large scale aerospace design. In N. Alexandrov and M. Y. Hussaini, editors, *Multidisciplinary Design Optimization: State of the Art*. SIAM, 1996.

[11]R. D. Braun and I. M. Kroo. Development and application of the collaborative optimization architecture in a multidisciplinary design environment. In N. Alexandrov and M. Y. Hussaini, editors, *Multidisciplinary Design Optimization: State of the Art*, 1996.

[12]I. M. Kroo. The aerodynamic design of oblique wing aircraft. *AIAA paper 86-2624*, 1986.

[13]I. M. Kroo, J. Gallman, and S.C. Smith. Aerodynamic and structural studies of joined wing aircraft. *Journal of Aircraft*, 1991.

[14]A. Van der Velden and I. M. Kroo. The sonic boom of an oblique flying wing. *Journal of Aircraft*, 1994.

[15]I. Kroo. Aerodynamic concepts for future aircraft. *AIAA paper 99-3524*, 30th AIAA Fluid Dynamics Conference, Norfolk, VA, June 1999.

[16]J. Reuther, A. Jameson, J. J. Alonso, M. J. Rimlinger, and D. Saunders. Contrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers. *AIAA paper 97-0103*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.

[17]A. Jameson and T.J. Baker. Multigrid solution of the Euler equations for aircraft configurations. *AIAA paper 84-0093*, AIAA 22th Aerospace Sciences Meeting, Reno, Nevada, January 1984.

[18]R. Seebass and B. Argrow. Sonic boom minimization revisited. *AIAA paper 98-295*, 2nd AIAA Theoretical Fluid Mechanics Meeting, Albuquerque, NM, June 1998.

[19]A. Jameson. Automatic design of transonic airfoils to reduce the shock induced pressure drag. In *Proceedings of the 31st Israel Annual Conference on Aviation and Aeronautics, Tel Aviv*, pages 5–17, February 1990.

[20]A. Jameson, L. Martinelli, J. J. Alonso, J. C. Vassberg, and J. J. Reuther. Simulation based aerodynamic design. Technical report, IEEE Aerospace Conference, Big Sky, MN, March 2000.

[21]J. L. Lions. *Optimal Control of Systems Governed by Partial Differential Equations*. Springer-Verlag, New York, 1971. Translated by S.K. Mitter.

[22]A. Jameson. Computational aerodynamics for aircraft design. *Science*, 245:361–371, 1989.

[23]S. Nadarajah, A. Jameson, and J. J. Alonso. An adjoint method for the calculation of remote sensitivities in supersonic flow. *AIAA paper 2002-0261*, AIAA 40th Aerospace Sciences Meeting, Reno, NV, January 2002.

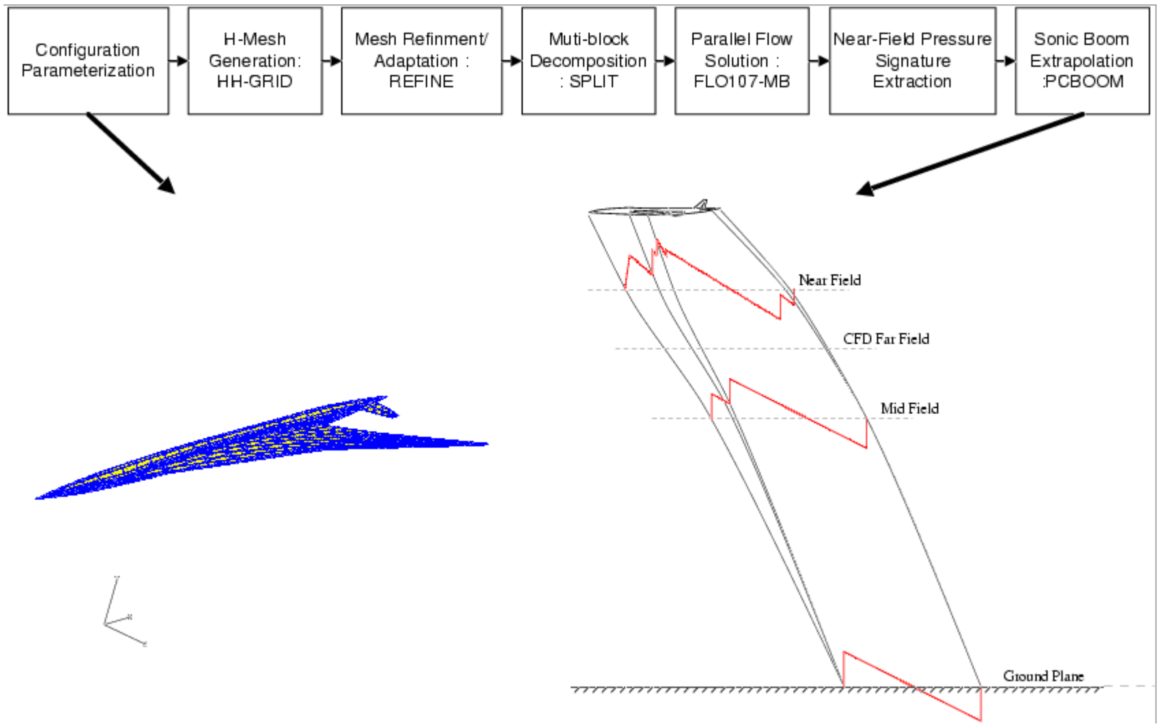

Fig. 7 A Flowchart of the Various Modules of QSP107

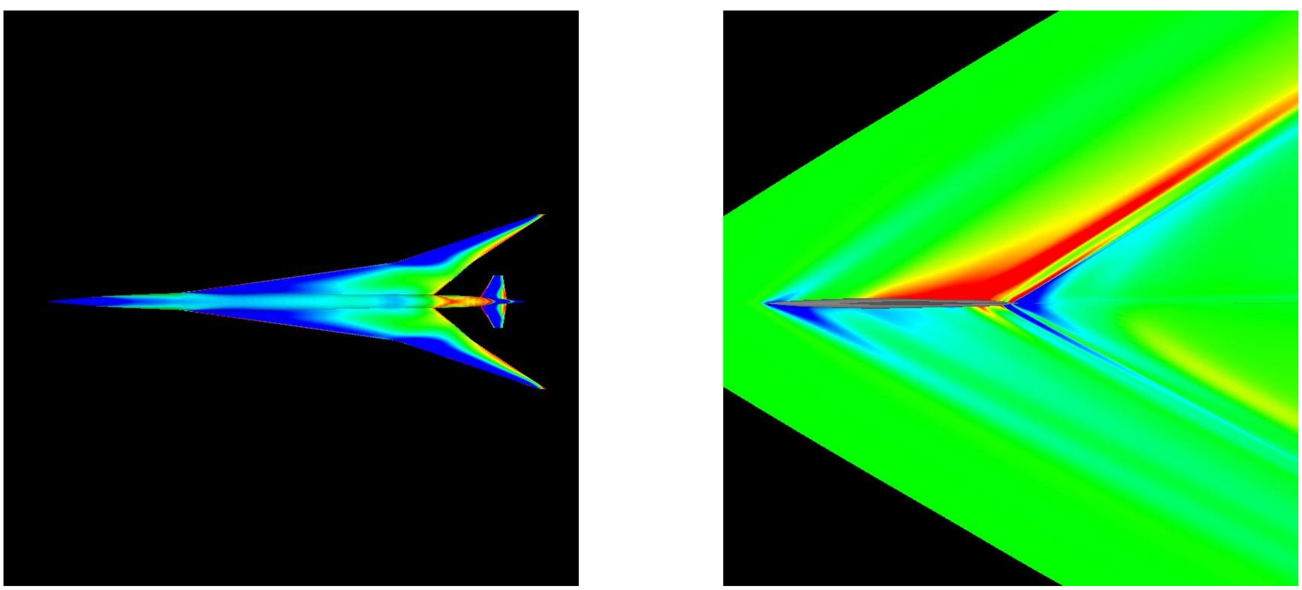

a) Bottom View of Mach Number Distribution

b) Symmetry Plane Wave Pattern

Fig. 8 Results of QSP107 for a Generic QSP Configuration

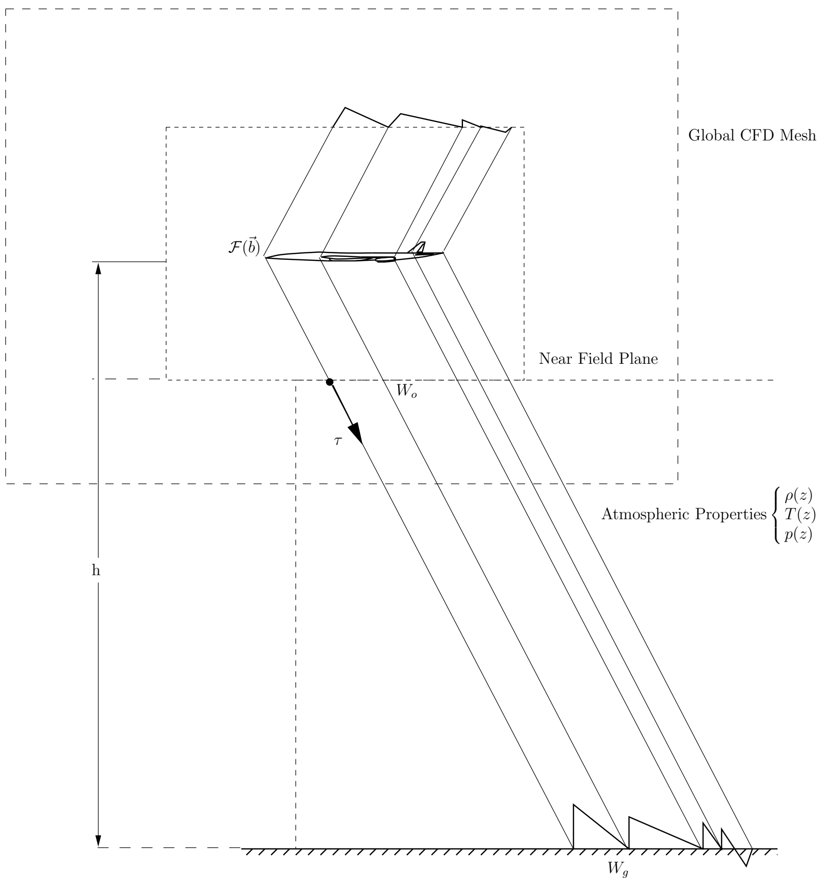

Fig. 9 Schematic of Sonic Boom Minimization Setup with Nomenclature.

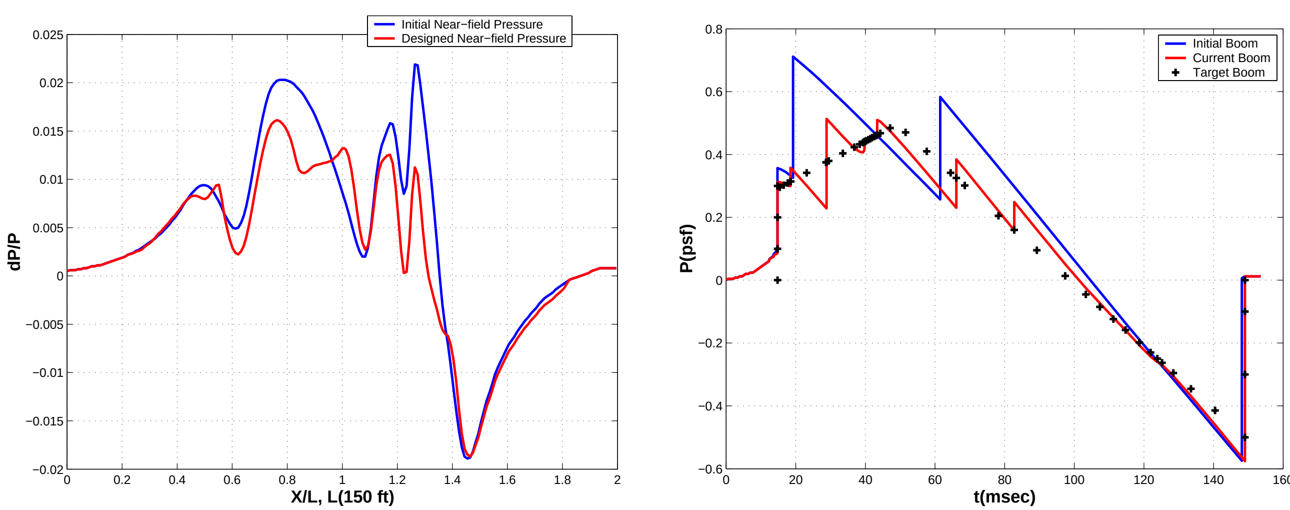

a) Near Field Pressure

b) Ground Pressure

Fig. 10 Generic QSP Configuration Nearfield to Ground Inverse Design Procedure.