Aerodynamic Design and Optimization of the Mars Airborne Remote Exploration Scout Aircraft

John C. Vassberg *
Phantom Works
The Boeing Company
Huntington Beach, CA 92647-2099, USA

Greg S. Page †
UAV Systems
ITT Advanced Engineering and Sciences
Alexandria, VA 22303, USA

Richard J. Foch ‡
Vehicle Research Section
Naval Research Lab
Washington, DC 20357-5000, USA

Antony Jameson §
Department of Aeronautics and Astronautics
Stanford University
Stanford, CA 94305-3030, USA

6 January, 2004

Abstract

The study presented in this paper demonstrates the utility of aerodynamic shape optimization technologies under development by the last author. The subject of this investigation is a flying wing configuration being developed by the second and third authors for a flight mission in the rarified Martian atmosphere. The design point for this aircraft is in the transonic regime and is at a fairly low unit Reynolds number. These flight conditions ultimately dictate that the aerodynamic optimization be based on the Reynolds-Averaged Navier-Stokes equations.

The optimization method of the present work is based on the use of the adjoint of the flow equations to compute the gradient of the cost function. This gradient is used to efficiently navigate the design space to find a local minimum. The optimizations presented herein utilize in excess of 5,000 automatically-constructed shape functions, are fully converged to a local optimum within less CPU time than that required to analyze a typical drag polar, and are conducted on very affordable, commodity off-the-shelf computer equipment. The utility of the software is demonstrated through its ease-of-use relative to an analysis run, its fast turn-around time, and through the significant improvements to aircraft performance achieved.

*AIAA Associate Fellow, Boeing Technical Fellow
†AIAA Member, Senior Staff Engineer
‡AIAA Member, Section Head
§AIAA Fellow, T. V. Jones Prof. of Engineering

Copyright ©2004 by the authors.
Published by the AIAA with permission.

Nomenclature

A	Hessian Matrix / Operator
AR	Wing Aspect Ratio $= \frac{b^2}{S_{ref}}$
b	Wing Span
$\mathcal{B}$	Shape Function Basis
CFD	Computational Fluid Dynamics
C_D	Drag Coefficient $= \frac{Drag}{q_\infty S_{ref}}$
C_L	Lift Coefficient $= \frac{Lift}{q_\infty S_{ref}}$
C_{ref}	Wing Reference Chord
$count$	Drag Coefficient Unit $= 0.0001$
$\mathcal{F}$	Surface Defining Function
$\mathcal{G}$	Gradient of Cost Function
H	Estimate of Inverse Hessian Matrix
I	Objective or Cost Function
N	Number of Design Variables
R	Flow-Equation Function
$RANS$	Reynolds-Averaged Navier-Stokes
RCS	Reaction Control System
Re	Wing Reynolds number based on C_{ref}
S_{ref}	Wing Reference Area
UAV	Unmanned Aerial Vehicle
x	Independent Spatial Variable
q	Dynamic Pressure $= \frac{1}{2}\rho V^2$
w	Flow Variable
λ	Wing Taper Ratio; Search Step Parameter
$\Lambda_{c/4}$	Wing Quarter-Chord Sweep
∞	Infinity
$\delta*$	First Variation of
$\mathcal{O}(*)$	Order of
$(*)^{-1}$	Inverse Matrix of

1 Introduction

Aircraft based remote sensing has proven to be an excellent method for large scale geologic analysis and surveying on earth. The same techniques can be applied to planetary science for planets containing an atmosphere. Aircraft-based sensors can cover a much greater surface area than rovers, at a much greater resolution than orbiting platforms, with flight-path control not possible using balloon systems. Ground-based aircraft operations usually require prepared launch sites, which are not normally available for extra-terrestrial exploration missions. Atmospheric aircraft deployment techniques, as developed by the Naval Research Labratory (NRL) during many years of unmanned aerial vehicle (UAV) research and development, can be used to eliminate the need for surface launch sites.

The Mars Airborne Remote Exploration Scout (MARES) is developed specifically for the constraints of an extra-terrestrial exploration aircraft using an atmospheric entry deployment. The overall aircraft shape maximizes the airframe size that can be packaged in a traditional conical entry shell and requires a minimum number of deployable surfaces, enhancing the system simplicity and reliability. The high angle-of-attack characteristics of the delta wing coupled with a reaction control system (RCS) allow a non-parachute assisted deployment with a minimum loss of altitude and a minimum over speed, giving greater flexibility in the selection of sites to be explored. To further minimize technical risk, MARES is powered by a reliable lightweight rocket engine. Optimizing the airframe for rocket propulsion leads to a higher cruise speed and a more structurally robust aircraft than very low-speed propeller-driven gossamer-type aircraft. The high durability enhances survival of the sensor and communications systems after landing, allowing sensor data to be stored during flight and transmitted after landing, minimizing the impact of data communication rates on sensor selection and resolution. The higher flight speeds, coupled with the low density of the Martian atmosphere, also lead to the requirement for wing design optimization in the unique environment of transonic low Reynolds number flight.

The next section addresses the basis of the present aerodynamic shape optimization method. Subsequent sections provide some history and special requirements for extra-terrestrial aircraft missions. Details of the MARES design and development efforts are included in Sections 5-6. Section 6 also highlights attributes of the present aerodynamic optimization software which are contributing factors to its utility. The paper completes with conclusions and a post-script discussion on automated design tools.

2 Aerodynamic Optimization

Traditionally the process of selecting design variations has been carried out by trial and error, relying on the intuition and experience of the designer. The degree of success with this classical approach depends directly on the level of expertise of the aerodynamic designer, which can take over a decade of apprenticeship to develop. It is not at all likely that a similar process of repeated trials in an interactive design and analysis procedure can lead to a truly optimum design. In order to take full advantage of the possibility of examining a large design space, the numerical simulations need to be combined with automatic search and optimization procedures. This can lead to automatic design methods which will fully realize the potential improvements in aerodynamic efficiency.

An approach which has become increasingly popular is to carry out a search over a large number of variations via a genetic algorithm. This may allow the discovery of (sometimes unexpected) optimum design choices in very complex multi-objective problems, but it becomes extremely expensive when each evaluation of the cost function requires intensive computation, as is the case in aerodynamic problems.

In order to find optimum aerodynamic shapes with reasonable computational costs, it pays to embed the flow physics within the optimization process. In fact, one may regard an aerodynamic shape as a device to control the flow in order to produce a specified lift with minimum drag. As a result, one can draw on concepts which have been developed in the mathematical theory of control of systems governed by partial differential equations. In particular, an acceptable aerodynamic design must have characteristics that do not abruptly vary with small changes in shape and flow conditions. Consequently, gradient-based procedures are appropriate for aerodynamic shape optimization. Two main issues affect the efficiency of gradient-based procedures; the first is the actual calculation of the gradient, and the second is the construction of an efficient search procedure which utilizes the gradient.

2.1 Gradient Calculation

For the class of aerodynamic optimization problems under consideration, the design space is essentially infinitely dimensional. Suppose that the performance of a system design can be measured by a cost function I which depends on a function $\mathcal{F}(x)$ that describes the shape, where under a variation of the design, $\delta\mathcal{F}(x)$, the variation of the cost is δI. Now

suppose that δI can be expressed to first order as

$$\delta I = \int \mathcal{G}(x)\delta\mathcal{F}(x)dx$$

where $\mathcal{G}(x)$ is the gradient. Then by setting

$$\delta\mathcal{F}(x) = -\lambda\mathcal{G}(x)$$

one obtains an improvement

$$\delta I = -\lambda \int \mathcal{G}^2(x)dx$$

unless $\mathcal{G}(x) = 0$. Thus the vanishing of the gradient is a necessary condition for a local minimum. Here, λ is a positive value that scales the step size of the search trajectory through the design space.

Computing the gradient of a cost function for a complex system can be a numerically intensive task, especially if the number of design parameters is large and if the cost function is an expensive evaluation. The simplest approach to optimization is to define the geometry through a set of design parameters, which may, for example, be the weights α_i applied to a set of shape functions $\mathcal{B}_i(x)$ so that the shape is represented as

$$\mathcal{F}(x) = \sum \alpha_i \mathcal{B}_i(x).$$

Then a cost function I is selected which might be the drag coefficient or the lift-to-drag ratio; I is regarded as a function of the parameters α_i. The sensitivities $\frac{\partial I}{\partial \alpha_i}$ may now be estimated by making a small variation $\delta\alpha_i$ in each design parameter in turn and recalculating the flow to obtain the change in I. Then

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

The main disadvantage of this finite-difference approach is that the number of flow calculations needed to estimate the gradient is proportional to the number of design variables [1]. Similarly, if one resorts to direct code differentiation (ADIFOR [2, 3]), or complex-variable perturbations [4], the cost of determining the gradient is also directly proportional to the number of variables used to define the design. Even small problems of aerodynamic shape optimization based on these approaches can require compute resources that are measured in CPU-Years, which can only be completed in reasonable elapsed time through utilization of massively-parallel computers costing millions of dollars.

A more cost effective technique is to compute the gradient through the solution of an adjoint problem, such as that developed in [5, 6, 7]. The essential idea may be summarized as follows. For flow about an arbitrary body, the aerodynamic properties that define the cost function are functions of the flowfield variables (w) and the physical shape of the body, which may be represented by the function $\mathcal{F}$. Then

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

and a change in $\mathcal{F}$ results in a change of the cost function

$$\delta I = \frac{\partial I^T}{\partial w}\delta w + \frac{\partial I^T}{\partial \mathcal{F}}\delta\mathcal{F}.$$

Using a technique drawn from 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 reevaluation of the flowfield. In order to achieve this, δw must be eliminated from the above equation. Suppose that the governing equation R, which expresses the dependence of w and $\mathcal{F}$ within the flowfield domain D, can be written as

$$R(w, \mathcal{F}) = 0. \tag{1}$$

Then δw is determined from the equation

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

Next, introducing a Lagrange multiplier ψ, we have

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

With some rearrangement

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

Choosing ψ to satisfy the adjoint equation

$$\left[\frac{\partial R}{\partial w}\right]^T \psi = \frac{\partial I^T}{\partial w} \tag{2}$$

the term multiplying δw can be eliminated in the variation of the cost function, and we find that

$$\delta I = \mathcal{G}\delta\mathcal{F},$$

where

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

The advantage is that the variation in cost function is independent of δw, with the result that the gradient of I with respect to any number of design variables can be determined without the need for additional flow-field evaluations.

In the case that equation (1) is a partial differential equation, the adjoint equation (2) is also a

partial differential equation and appropriate boundary conditions must be determined. As it turns out, the appropriate boundary conditions depend on the choice of the cost function, and may easily be derived for cost functions that involve surface-pressure integrations. Cost functions involving field integrals lead to the appearance of a source term in the adjoint equation.

The cost of solving the adjoint equation is comparable to that of solving the flow equation. Hence, the cost of obtaining the gradient is comparable to the cost of two function evaluations, regardless of the dimension of the design space. The downside of this approach is that it can take man-months to develop an adjoint code for a given cost function. However, there is on-going research at Rice University to develop ADJIFOR [8] which automatically generates a discrete adjoint code from existing analysis software. So far, however, this has not realized the same level of efficiency. In the present work, the adjoint equations have been derived analytically and then approximated in discrete form.

2.2 Search Procedure

The remaining cost issue is related to finding a location in the design space where the gradient vanishes, and hence where there is a local optimum. Normally, this search starts from a baseline design and the design space is traversed by a search method. The final state of the search may be subject to constraints imposed on the design space, yet there is no requirement that the trajectory adhere to these except at its end point. The efficiency of the search depends on the number of steps it takes to find a local minimum as well as the cost of each step.

In order to accelerate the search, one may resort to using the Newton method. Here, the search direction is based on the equation represented by the vanishing of the gradient, $\mathcal{G}(\mathcal{F}) = 0$, and is solved by the standard Newton iteration for nonlinear equations.

Suppose the Hessian is denoted by

$$A = \frac{\partial \mathcal{G}}{\partial \mathcal{F}}$$

then the result of a step $\delta\mathcal{F}$ may be linearized as

$$\mathcal{G}(\mathcal{F} + \delta\mathcal{F}) = \mathcal{G}(\mathcal{F}) + A\,\delta\mathcal{F}$$

This is set to zero for a Newton step; therefore

$$\delta\mathcal{F} = -A^{-1}\mathcal{G}.$$

The Newton method is generally very effective if the Hessian can be evaluated accurately and cheaply. Unfortunately, this is not the case with aerodynamic shape optimization.

Quasi-Newton methods estimate A or A^{-1} from the changes of $\mathcal{G}$ recorded during successive steps. For a discrete problem with N design variables, it requires N steps to obtain a complete estimate of the Hessian, and these methods have the property that they can find the minimum of a quadratic form in exactly N steps. Thus in general, the cost of a quasi-Newton search scales with the dimension of the design space. For the class of optimizations under consideration, this is an undesirable property.

Rank-2 quasi-Newton schemes have an additional cost associated with the line search of each step, where multiple function evaluations (5-10) are required to locate the minimum along a fixed search direction. This additional cost is further amplified by a requirement that each of these function evaluations be converged to tighter tolerance than that typically needed for engineering-level accuracy. As a consequence, it has been our experience that quasi-Newton searches are not particularily suited for the class of optimizations that are addressed here. However, alternative techniques have been developed that can navigate the design space quite effectively; the foundation of our search method is described next.

Efficient aerodynamic shapes are predominately smooth. This suggests a natural alternative approach to the search method. In order to make sure that each new shape in the optimization sequence remains smooth, one may smooth the gradient and replace $\mathcal{G}$ by its smoothed value $\bar{\mathcal{G}}$ in the descent process. This is equivalent to reformulating the gradient in a Sobelov space [9], and acts as a preconditioner which allows the use of much larger steps. To apply smoothing in the ξ_1 direction, for example, the smoothed gradient $\bar{\mathcal{G}}$ may be calculated from a discrete approximation to

$$\bar{\mathcal{G}} - \frac{\partial}{\partial \xi_1}\epsilon\frac{\partial}{\partial \xi_1}\bar{\mathcal{G}} = \mathcal{G} \tag{3}$$

where ϵ is the smoothing parameter. Then, if one sets $\delta\mathcal{F} = -\lambda\bar{\mathcal{G}}$, assuming the modification is applied on the surface $\xi_2 =$ constant, the first order change in the cost function is

$$\begin{aligned}
\delta I &= -\iint \mathcal{G}\delta\mathcal{F}\, d\xi_1 d\xi_3 \\
&= -\lambda \iint \left(\bar{\mathcal{G}} - \frac{\partial}{\partial \xi_1}\epsilon\frac{\partial \bar{\mathcal{G}}}{\partial \xi_1}\right)\bar{\mathcal{G}}\, d\xi_1 d\xi_3 \\
&= -\lambda \iint \left(\bar{\mathcal{G}}^2 + \epsilon\left(\frac{\partial \bar{\mathcal{G}}}{\partial \xi_1}\right)^2\right) d\xi_1 d\xi_3 \\
&< 0.
\end{aligned}$$

Thus, an improvement is assured for arbitrary choices of the smoothing parameter ϵ if λ is sufficiently small and positive, unless the process has

already reached a stationary point at which $\bar{\mathcal{G}} = 0$, and therefore, according to Equation 3, $\mathcal{G} = 0$. (Note that the λ acting on $\bar{\mathcal{G}}$ can be significantly larger than that acting on $\mathcal{G}$, with respect to stability limits.)

It turns out that this approach is extremely tolerant to the use of approximate values of the gradient, so that neither the flow solution nor the adjoint solution need be fully converged before making a shape change. This results in very large savings in the computational cost of the complete optimization process, which is discussed next.

2.3 Computational Costs

In order to address the issues of the search costs, a variety of techniques were investigated in Reference [10] using a trajectory optimization problem (the brachistochrone) as a representative model. The study verified that the search cost (i.e., number of steps) of a simple steepest descent method applied to this problem scales as N^2, where N is the number of design variables, while the cost of quasi-Newton methods scaled linearly with N as expected. On the other hand, with an appropriate amount of smoothing, the smoothed descent method converged in a fixed number of steps, independent of N. Considering that the evaluation of the gradient by a finite difference method requires $N + 1$ flow calculations, while the cost of its evaluation by the adjoint method is roughly that of two flow calculations, one arrives at the estimates of total computational cost given in Tables 1-2.

Table 1: Cost of Search Algorithm.

Steepest Descent	$\mathcal{O}(N^2)$ steps
Quasi-Newton	$\mathcal{O}(N)$ steps
Smoothed Gradient	$\mathcal{O}(K)$ steps
(Note: K is independent of N)	

Table 2: Total Computational Cost of Design.

Finite Difference Gradients + Steepest Descent	$\mathcal{O}(N^3)$
Finite Difference Gradients + Quasi-Newton Search	$\mathcal{O}(N^2)$
Adjoint Gradients + Quasi-Newton Search	$\mathcal{O}(N)$
Adjoint Gradients + Smoothed Gradient Search (Note: K is independent of N)	$\mathcal{O}(K)$

Other methods have been investigated which further improve the dimension-independent convergence rate, including multigrid and postconditioning with a Krylov subspace acceleration. Implementing these in the current aerodynamic shape optimization software consistently converges the design to a local optimum within 30-60 steps, even for problems with thousands of design variables. Moreover, because they do not require either the flow or adjoint solutions to be fully converged, complete optimizations are routinely completed with a computational cost equivalent to 2-10 converged flow solutions. As a consequence, our standard practice is to allow every discrete surface point within the CFD grid to be its own design function, aligned with the grid line emanating from the surface. The amplitude of this design variable corresponds to the signed distance from the original baseline surface node. This typically leads to design space dimensions of $N > 4000$ for three-dimensional aerodynamic shape optimizations.

Not being constrained by the number of design variables enables the optimization software to be written so that the end user is not burdened with the task of defining a set of shape functions. This added benefit is a significant improvement over the standard approach. A common practice throughout industry is to require a user to specify shape functions of the bump-function class which must be tailored for each specific application. The main reason for using bump functions is to reduce the number of design variables needed and yet obtain reasonable results. Specifying an appropriate set of bump functions for a given problem is somewhat of a black art in itself, and getting an effective set frequently requires experimentation by the user. The process is also prone to input error that may not be discovered until an optimization run fails to produce reasonable results. This is a poor environment for both user and design, as the "optimized" design will only be as good as the user is at choosing his shape-function set and accurately inputting it. Hence, the need for an expert user is mandated by this approach. Optimization software should free the engineer from mundane tasks and allow him to focus on the more global requirements of the system development. This has been achieved by the present aerodynamic shape optimization software.

The next section gives a brief history of aircraft concepts designed for missions through the Martian atmosphere.

3 History of Mars Flyers

The first engineering study of flying robotic airplanes for planetary observation was started within months of the arrival of the Viking spacecraft on Mars [11]. Some significant technology development studies were pursued, but a full aircraft mission was never developed. After a hiatus of nearly 20 years, the Mars exploration aircraft concept was revived and proposed as a Discovery Mission in 1996. Although several aircraft-based approaches were proposed for the Discovery Missions, none was selected. Again in 1999, NASA investigated the technical readiness of extra-terrestrial flight for a *Mars Flyer* mission announced by Administrator Dan Goldin to commemorate the centennial of the Wright brothers' 1903 flight, whose goal was to fly an airplane on Mars on December 17, 2003. Unfortunately that mission was cancelled before the year ended. However, as a result of the renewed interest in Mars aircraft, a design team assembled from NASA, NRL and ITT/AES personnel finally demonstrated the successful high-altitude deployment and flight of a Mars aircraft design, NASA 731, in August 2001. In June 2001, when the first Mars Scout Mission Concept Workshop was held, Mars airplane and glider missions were the largest single category of mission represented. The continuing 25-year history of aircraft designs for flight on Mars shows the underlying enthusiasm for the exploration of Mars by air, both in the public and within the technical community. Adding extra-terrestrial aerial exploration to NASA's Space Science repertoire could introduce capabilities that are not possible or practical to achieve by any other means. This includes access to the entire planet's surface, precision landing with robust hazard avoidance, and the ability to controllably traverse thousands of kilometers in close proximity to terrain features.

4 Extra-Terrestrial Missions

There are unique design requirements for any extra-terrestrial aircraft. Most of the proposed Mars airplane missions begin with deployment of the airplane from an atmospheric-entry aeroshell at an altitude typically between 10 and 20 km above the surface, following atmospheric entry, hypersonic deceleration, and descent on a parachute. A key technical challenge is the design of an airframe which can be folded for stowage into the atmospheric entry aeroshell, and is able to be autonomously unfolded in mid-air. It must then enter stable horizontal flight without exceeding the aerodynamic limits of the aircraft or losing excessive altitude.

A free-fall vertical drop deployment is often used, converting potential energy into kinetic energy to establish the required flight velocity airflow over the airframe. Within a few seconds of release, the dynamic pressure is adequate to support aerodynamic deployment of stored wing and tail panels to their normal flight configuration. A few seconds of free fall provide sufficient airspeed to allow the control surfaces to begin pitching the nose of the aircraft up into a pull-out maneuver. However, this initial dive-and-pull-out maneuver can be one of the most demanding aspects of the entire mission with its requirement of high lift coefficients in an environment of rapidly increasing flight velocity. Even small perturbations in the pull-up profile, such as can be caused by uncertainty in the knowledge of the local seasonal atmospheric density, can cause significant increases in flight speed and the possibility of exceeding the aircraft structural limits.

This free-fall and pullout flight release profile was employed for the high-altitude demonstration flight of NASA 731 on August 9, 2001. Here, an unpowered Mars aircraft design with a wingspan of 2.75 m was dropped nose-down from a helium balloon at an altitude of 31.4 km over Tillamook, OR. The aircraft autopilot was programmed to follow a pre-computed and optimized pull-out trajectory, followed by transition to a steady-state descending glide. Figure 1 is taken from video recorded by the wing tip mounted camera as NASA 731 flew in steady-state flight at approximately 25.5 km altitude. This flight successfully demonstrated that steady-state flight in Mars-like atmospheric conditions is possible.

5 MARES General Design

Using the flight experience gained from the NASA 731 test, members of the design team conceived a new tailless aircraft configuration and deployment approach to minimize the altitude and airspeed excursions during the deployment. This design also improves payload survivability during landing on the planet surface.

By redesigning the aircraft configuration to incorporate the high angle-of-attack stability characteristic of a delta wing configuration and adding a reaction control system (RCS), the MARES aircraft configuration is able to provide a deployment profile that greatly minimizes over speed and altitude loss, yielding a much lower risk mission.

Figure 2 shows the altitude and velocity data from the recorder on board NASA 731 and compares that data with the computed values for the deployment trajectory of the baseline MARES configuration. The blue(dark) traces are the altitude loss

from the initial drop point, and the magenta(gray) traces are the corresponding indicated airspeed. The NASA 731 data is depicted with symbols, while the MARES profiles are shown as solid lines. An indicated airspeed of 95 knots is a true airspeed of nearly Mach 1 at the release altitude. The recorded data clearly shows the vertical overshoot of NASA 731 and its close approach to sonic speeds; this is characteristic of a vertical free-fall and pull-out deployment. The MARES aircraft is capable of a smooth asymptotic approach to its cruise altitude and Mach number as shown by solid-line profiles in Figure 2.

The MARES aircraft configuration is shown stowed in the entry aeroshell in Figure 3 and in its flight configuration in Figures 4-5.

Deployment of the MARES begins with the folded aircraft being released from the aeroshell, where it falls in a flat attitude enabled by the natural stability provided by the delta planform and the location of the center of mass, which is augmented by the zero axial velocity control capability provided by the RCS. The main propulsion system solid rocket motor is ignited to begin transition of the aircraft into forward flight without requiring a nose-down body attitude. As the forward flight speed reaches approximately 30 m/s, the drag forces acting on the vertical stabilizer surfaces are sufficient to begin the aerodynamic deployment of the outer wing panels. As the forward speed approaches the cruise flight speed (150 m/s), the RCS is no longer needed and flight control is transferred to conventional aerodynamic control surfaces.

At the completion of the data collection flight, the MARES aircraft will be pitched up into a nose-high (deep-stall) attitude. The delta planform coupled with the RCS will maintain the aircraft in a stable attitude and descent rate until reaching a predetermined altitude above the ground, when an additional z-axis thruster will fire to minimize vertical velocity at the moment of impact. The MARES will contact the ground ventral fins first, in a nose-high attitude. The fins are designed to crush and provide additional energy absorption during impact to provide a relatively soft landing.

6 MARES Development

The initial configuration layout and baseline geometry were created by the NASA/NRL/ITT development team. Further refinement of this geometry is being conducted by the first and last authors. Aerodynamic design and shape optimization of the MARES wing-surface geometry is being conducted using a number of computational fluid dynamics (CFD) methods, including FLO22 [12, 13], CFL3D [14], SYN88 and SYN107P; this report will concentrate on the simulations associated with the synthesis codes SYN88 and SYN107P, which are capable of both analysis and design calculations.

Because of the constraints associated with packaging the aircraft in the conical entry shell, the configuration has evolved into a flying-wing design, with folding wings. When deployed, the outboard wing planform has a quarter-chord sweep of only 5.5 degrees. Figure 6 provides the planform of MARES, and the following list specifies reference quantities and sweep angles.

Table 3: MARES Planform Quantities.

$Sref$	$36.38\ ft^2$	AR	4.9
b	$13.38\ ft$	λ	0.3
$Cref$	$3.28\ ft$	$\Lambda_{c/4}$	5.5°
$Xref$	$3.28\ ft$	Λ_{LE}	10.0°
$Yref$	$1.51\ ft$	$\Lambda_{LE.\Delta}$	50.0°

In order to scope the problem and work out any issues with the baseline geometry definition, SYN88 was initially applied. SYN88 is an inviscid aerodynamic shape optimization code based on the Euler equations and corresponding adjoint formulae. The results of this Euler-based optimization are provided in Figures 7-8. While the outcome of this exercise is somewhat academic given the nature of the task at hand, nonetheless, it is worth noting that SYN88 is capable of solving this class of problem on a notebook computer in about 30 minutes of CPU time. Utilizing this code to sort out any possible issues with the problem set-up can greatly accelerate the time to final RANS-based design.

The final detailed aerodynamic design of the MARES wing is accomplished through application of SYN107P by minimizing total drag at a fixed lift coefficient, while maintaining the spanload and airfoil thickness distributions of the baseline configuration. The thickness constraint is imposed as a greater than or equal to condition. The design flight condition for the MARES mission is: $M = 0.65$, $C_L = 0.62$, and $Re = 170K$, where Reynolds number is based on reference chord. The flow is assumed to be fully turbulent over the wing surface. The rarified Martian atmosphere design condition also has that the density and viscosity are: $\rho = 2.356 * 10^{-5}\ slugs/ft^3$, and $\nu = 2.2517 * 10^{-7}\ slugs/ft/sec$. Our initial analysis of the baseline configuration indicated that the original selection of airfoils (with leading-edge droops and aft-camber reflex modifications to trim the aircraft) yielded an undesirable, strong, unswept shock system.

SYN107P is an aerodynamic synthesis method based on the Reynolds-Averaged Navier-Stokes (RANS) equations and corresponding continuous-adjoint formulae. SYN107P runs in parallel on distributed-memory computer systems by using the Message-Passing Interface (MPI). The work performed herein utilizes a Linux cluster of 8 nodes, where each node is comprised of dual AMD Athlon 1.7GHz processors and 2GB RAM. The nodes are interconnected with a switched 100BaseT network. Running on 8 processors (half the cluster), a RANS analysis requires less than 30 minutes of elapsed time; optimizations complete in less than 5 hours. These wall-clock times correspond to grid dimensions of (257 x 64 x 49) and with a design space defined with 5,313 basis functions.

Figure 9 provides the history of total drag during the Navier-Stokes optimization process. The baseline wing-only configuration has a cruise drag of 592 counts, while the optimized design converges to a cruise drag of 480 counts. Figure 10 gives the corresponding history of the lift-to-drag ratio, which starts at about 10.4 and finishes at 12.8; an improvement of 23%. Figure 11 over-plots the chordwise pressure distributions of the baseline geometry (solid curves) and optimized shape (dashed curves). It is evident from this comparison that the wing leading-edge pressure peaks are being reduced and moved forward, both of which act to soften the adverse pressure gradient of the upper-surface recovery path. The effects of the upper-surface changes are also illustrated in the side-by-side comparison of isobars of Figure 12. In addition, a lower-surface leading-edge peak just outboard of the baseline wing midspan region has been eliminated by the optimizer, which reshaped the Cp distribution to be monotonically favorable to about 35% chord.

Figure 13 gives a comparison of the drag loops for the baseline and optimized wings. It is often useful to view pressure distributions in this manner to better understand how changes in geometry and pressures interplay to effect an improvement in drag. For example, the thrust lobes (negative areas of the drag loops) of the optimized design are noticably larger than those of the baseline configuration.

Although the drag minimization of the present study is conducted as a single-point optimization, Figure 14 illustrates that the optimized design clearly out-performs the baseline at all lifting conditions; the solid symbol data is at the cruise design point. This result is especially interesting as both planform and spanload are common to the initial and final designs, hence their associated induced drags must be comparable. This leaves shock drag, profile drag, and possibly skin-friction drag as the components contributing to the performance improvements. Upon closer inspection of the results, the improvements are made in both the shock and profile drag components. Figure 15 provides the lift curves for the baseline and redesigned configurations. This figure indicates that the boundary-layer health of the baseline has been greatly improved with the optimized design, hence the related improvement in profile drag.

Several baseline and optimized geometry airfoil sections are decomposed into thickness and camber distributions and are compared in Figures 16-18. Although it is counter-intuitive to increase airfoil thickness when attempting to reduce the drag of a wing, that is precisely what SYN107P did at the wing root. In addition, the camber levels across the span have been reduced, and in particular, the leading-edge droop of the mid-to-outboard sections have been significantly reduced.

Regarding the utility of SYN107P as an aerodynamic shape optimization design tool, this exercise highlights several key attributes that benchmark the state-of-the-art. A few of these are as follows.

- Problem set-up time is essentially only that which is required to prepare an input deck for an analysis run. Since the set of basis functions is automatically created within the code, the user expends absolutely no effort on this task, and an optimization run to minimize drag at a fixed lift is as simple as changing the number of design cycles in the input deck from a 0 (for analysis-only) to a non-zero number (say 50) for an optimization to be performed. If the optimization is an inverse design, then the user will have to specify the target pressure distribution.

- The automatically set-up design space is essentially infinitely dimensional. To be precise, the continuum design space is represented by the highest possible discrete space that can be supported by the surface-defining grid. The key attribute here is that the design space of the optimization is not artificially nor arbitrarily constrained. This can occur when a user is required to specify a very low dimension design space with a much reduced set of basis functions.

- The cost of a fully converged optimization is less than or equivalent to that of about 10 converged analyses. To emphasize this point, consider the computer time required to generate the drag polars of Figure 14; this represents about twice the computer time as that needed for the drag minimization run. Further, both analysis and optimization capabilities run adequately fast on very affordable computer equipment, allowing multiple optimizations to be performed within a single day.

- Thickness constraints relative to the initial geometry are automatically set-up by the code and can be globally controlled by one input parameter.

The four properties listed above have been developed to the stage that they can be implemented into a general setting, such as that for an unstructured-mesh RANS aerodynamic shape optimization method. In addition to these, there are a few other niceties related to SYN107P. These are made possible because this code is tailored for very specific aircraft geometries, namely wing-out-of-a-wall, or wing-body configurations. These features include the following.

- Automatic grid generation is available for both Euler and RANS simulations.
- A variety of spanload constraints are available, which are either based on the initial distribution or a blending of elliptic and linear spanloads.
- Calculations can be performed at specified lifting conditions.

All of the properties listed above have enabled the authors to perform the present detailed design study for the MARES wing under minimal funding levels.

7 Conclusions

The utility of SYN107P has been demonstrated on the Mars Airborne Remote Exploration Scout. This investigation has redesigned the baseline wing shape to improve its aerodynamic performance. The redesign has yielded a 112-count reduction in total drag at its design point, and increased the lift-to-drag ratio 23% from 10.4 to 12.8. The cruise flight condition of the MARES is $M = 0.65$, $C_L = 0.62$, and $Re = 170K$. Drag polars for the baseline and redesigned wing show that the performance improvement extends over all lifting conditions and is well behaved at off-design points. All SYN107P calculations (analyses and optimizations) are performed on very affordable computer equipment with fast turn-around times, due in part to parallel processing. (In fact, a serial version of SYN107P will currently reproduce the present Navier-Stokes optimization on a 2GHz notebook computer within 27 hours.) Furthermore, essentially no additional set-up effort is required to run optimizations once an analysis input deck is defined. The effectiveness of this aerodynamic shape optimization has been made possible because the developers have addressed the adjoint-based gradient calculation and the design-space traversal processes from the perspective of what is required in the infinite-dimensional continuum, then taken the simple step towards discrete space.

8 Post Script

As can be seen from this study, aerodynamic shape optimization can significantly streamline the design process. The present method has now been successfully applied in a variety of projects, including the McDonnell-Douglas MDXX [15], the NASA HSCT studies, the Boeing Blended-Wing-Body project, the Beech Premier [16], and a Reno Racer [17].

There are many different examples on how aircraft design teams have utilized the rapidly provided information of aerodynamic shape optimization to make improvements to their aircraft configurations. The diversity of these examples illustrate the artistic and creative nature of the thought processes by the design teams. It is through these unpredictable paths in design direction that dramatic improvements of the multi-disciplinary systems are accomplished. Further, because unforeseen directions may be required to accomplished the design goals, it is highly unlikely that the designers will be replaced by a comprehensive, multi-disciplinary optimization (MDO) method. On the other hand, there are very well established dependencies between certain disciplines that can and should be coupled for MDO.

Although the search method is only guaranteed to find a local minimum, it turns out in practice that our aerodynamic optimizations are yielding results that are in the neighborhood of the known lower bounds for aerodynamic drag, as determined by optimum span loading, flat-plate skin-friction, and minimum wave drag. Furthermore, in practice, the designer's goal is not to determine the absolute best design, but rather, is tasked to make the most improvement to a design in a fixed amount of time specified by program schedule.

Aerodynamic shape optimization will not replace the judgement and insight of the aircraft designers. Rather, it should properly be viewed as an enabling tool that allows the designers to focus their efforts on the creative aspects of aircraft design, by relieving them of the need to spend large amounts of time exploring small variations. By intelligent choice of the cost function to measure the aerodynamic performance and perhaps also the deviation from a desired pressure architecture, one can essentially eliminate the need to carry out detailed section design. Instead, the designers can concentrate their attention on large scale parameters such as wing span, area and sweep, knowing that the optimization process will improve the performance for any given choice of these parameters.

Acknowledgments

This work has benefited from the generous support of the Naval Research Lab. Computing resources were provided by Intelligent Aerodynamics, International. The first author thanks The Boeing Company for allowing his participation in this project.

References

[1] R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15:407–412, 1978.

[2] C. Bischof, A. Carle, G. Corliss, A. Griewank, and P. Hovland. Generating derivative codes from Fortran programs. *Internal report MCS-P263-0991*, Computer Science Division, Argonne National Lab. and Center of Research on Parallel Computation, Rice Univ., 1991.

[3] L. L. Green, P. A. Newman, and K. J. Haigler. Sensitivity derivatives for advanced CFD algorithm and viscous modeling parameters via automatic differentiation. *AIAA paper 93-3321*, 11th AIAA Computational Fluid Dynamics Conference, Orlando, Florida, 1993.

[4] W. K. Anderson, J. C. Newman, D. L. Whitfield, and E. J. Nielsen. Sensitivity analysis for the Navier-Stokes equations on unstructured meshes using complex variables. *AIAA paper 99-3294*, Norfolk, VA, June 1999.

[5] A. Jameson, L. Martinelli, J. J. Alonso, J. C. Vassberg, and J. Reuther. Simulation based aerodynamic design. *IEEE Aerospace Conference*, Big Sky, MO, March 2000.

[6] A. Jameson. Optimum aerodynamic design using control theory. *Computational Fluid Dynamics Review*, pages 495–528, 1995.

[7] A. Jameson. Optimum aerodynamic design using CFD and control theory. *AIAA paper 95-1729*, AIAA 12th Computational Fluid Dynamics Conference, San Diego, CA, June 1995.

[8] A. Carle, M. Fagan, and L. L. Green. Preliminary results from the application of automated adjoint code generation to CFL3D. *AIAA paper 98-4807*, 1998.

[9] A. Jameson, L. Martinelli, and J. C. Vassberg. Using computational fluid dynamics for aerodynamics - A critical assessment. *ICAS Paper 2002-1.10.1*, 23rd International Congress of Aeronautical Sciences, Toronto, Canada, September 2002.

[10] A. Jameson and J. C. Vassberg. Studies of alternative numerical optimization methods applied to the brachistochrone problem. In *CFD Journal, Vol. 9, No. 3*, pages 281–296, Kyoto, Japan, October 2000. Japan Society of CFD.

[11] A concept study of a remotely piloted vehicle for Mars exploration. *JPL Technical Report*, Jet Propulsion Laboratory, August 1978.

[12] A. Jameson. Transonic potential flow calculations using conservative form. In *Proceedings of AIAA 2nd Computational Fluid Dynamics Conference*, pages 148–161, June 1975.

[13] P. A. Henne and R. M. Hicks. Wing analysis using a transonic potential flow computational method. *NASA-TM 78464*, July 1978.

[14] S. L. Krist, R. T. Biedron, and C. L. Rumsey. CFL3D user's manual, version 5.0. *NASA Report* , NASA Langley Research Center, Hampton, VA, November 1996.

[15] A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamic design using the Navier-Stokes equations. *Theoret. Comput. Fluid Dynamics*, 10:213–237, 1998.

[16] A. Jameson, L. Martinelli, J. J. Alonso, J. C. Vassberg, and J. Reuther. Simulation based aerodynamic design. *Computational Fluid Dynamics of the 21st Century*, 78:135–178, July 2000. A Symposium in honor of Professor Satofuka's 60th anniversary.

[17] J. C. Vassberg and A. Jameson. Aerodynamic shape optimization of a Reno race plane. *Int'l Journal of Vehicle Design*, 28(4):318–338, 2002. Special Issue on: Design Sensitivity and Optimization.

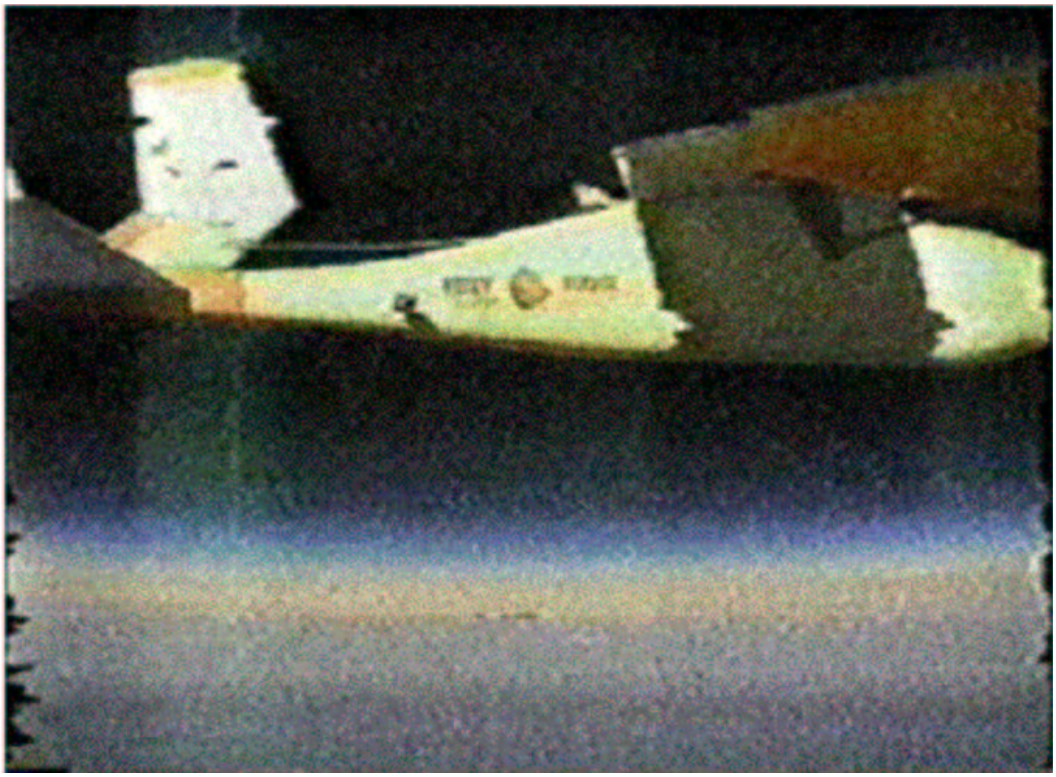

Figure 1: NASA 731 in Flight.

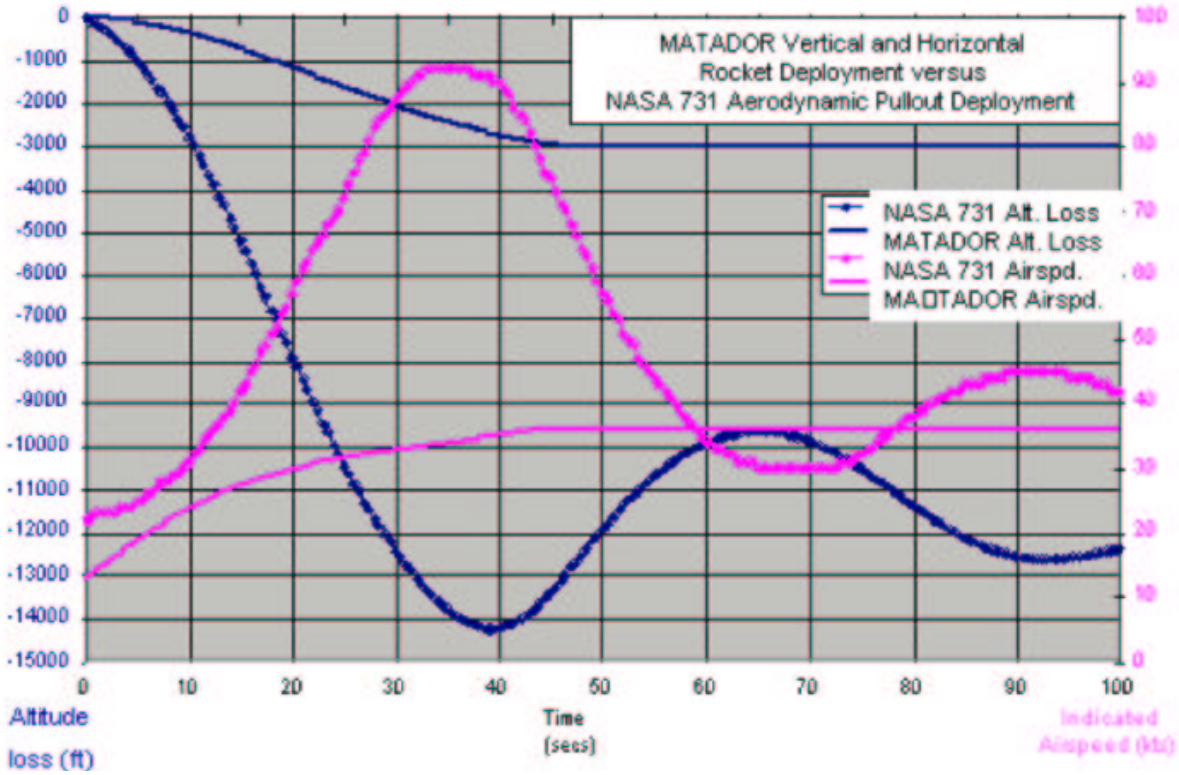

Figure 2: Comparison of NASA 731 and MARES Flight Profiles during Deployment.

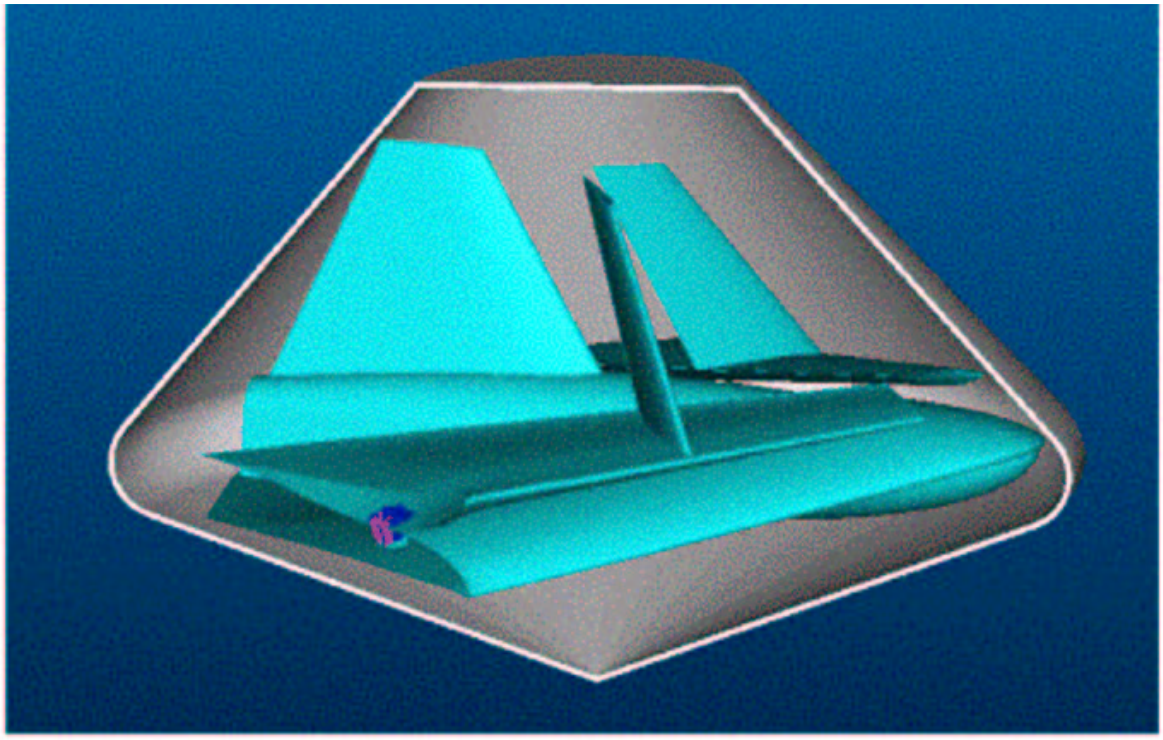

Figure 3: MARES Packaging in the Aerodynamic-Shell Capsule.

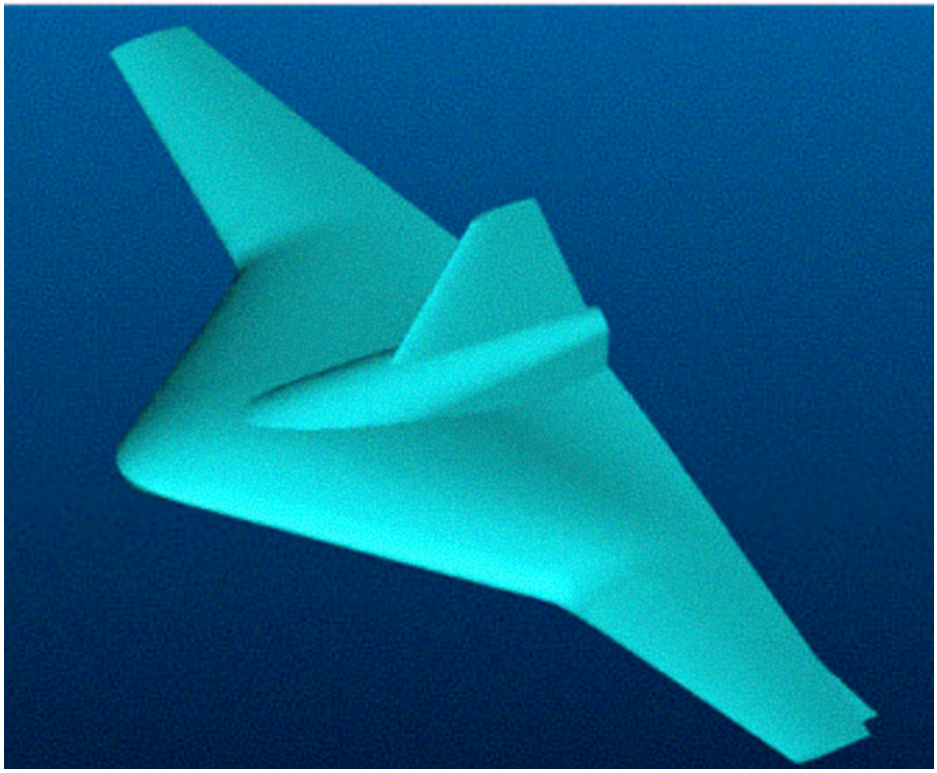

Figure 4: MARES Configuration in Flight, Top-View Rendering.

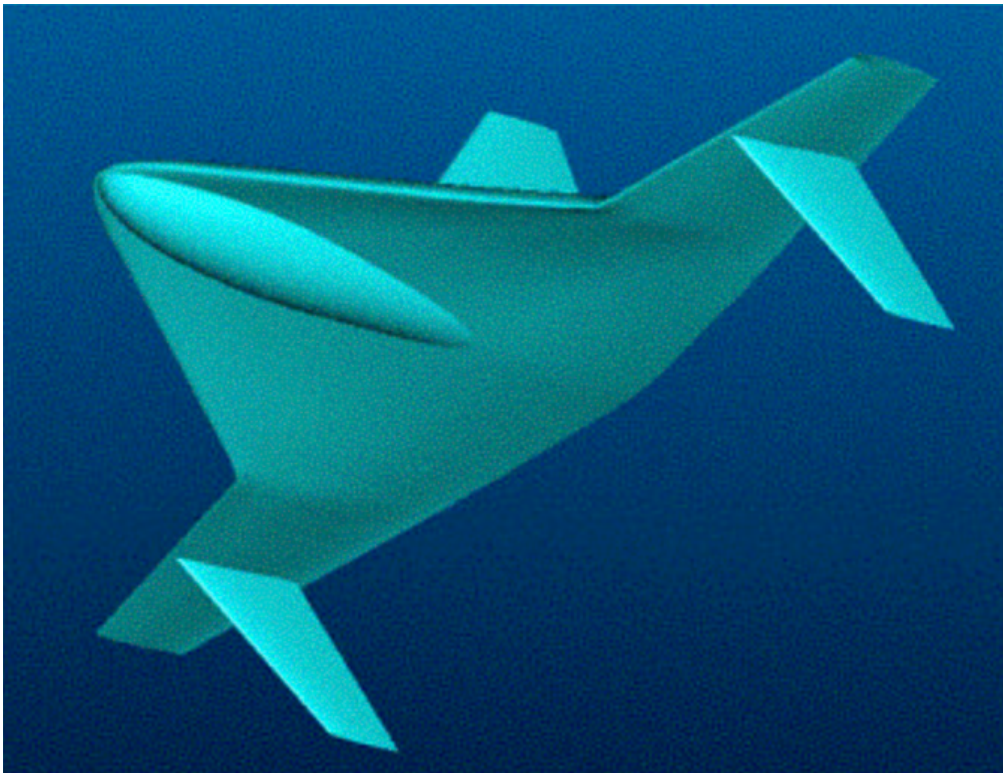

Figure 5: MARES Configuration in Flight, Bottom-View Rendering.

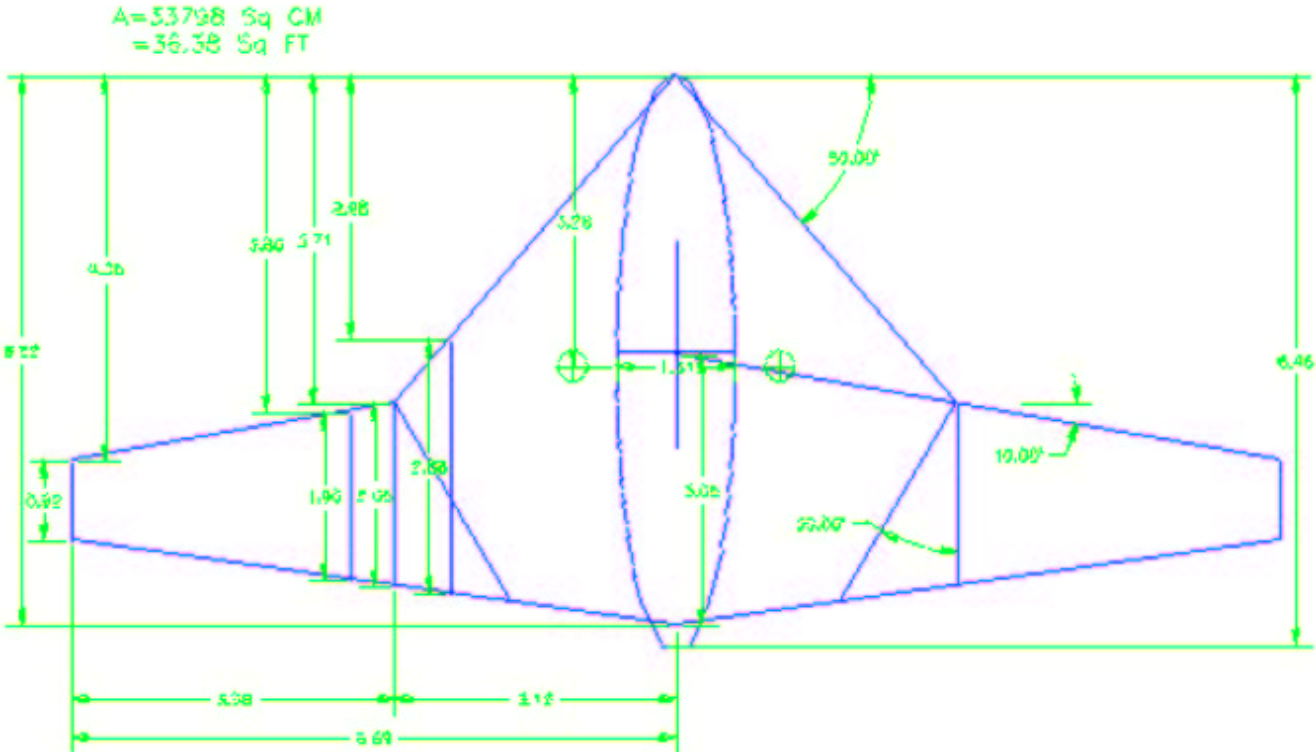

Figure 6: MARES General Planform Layout.

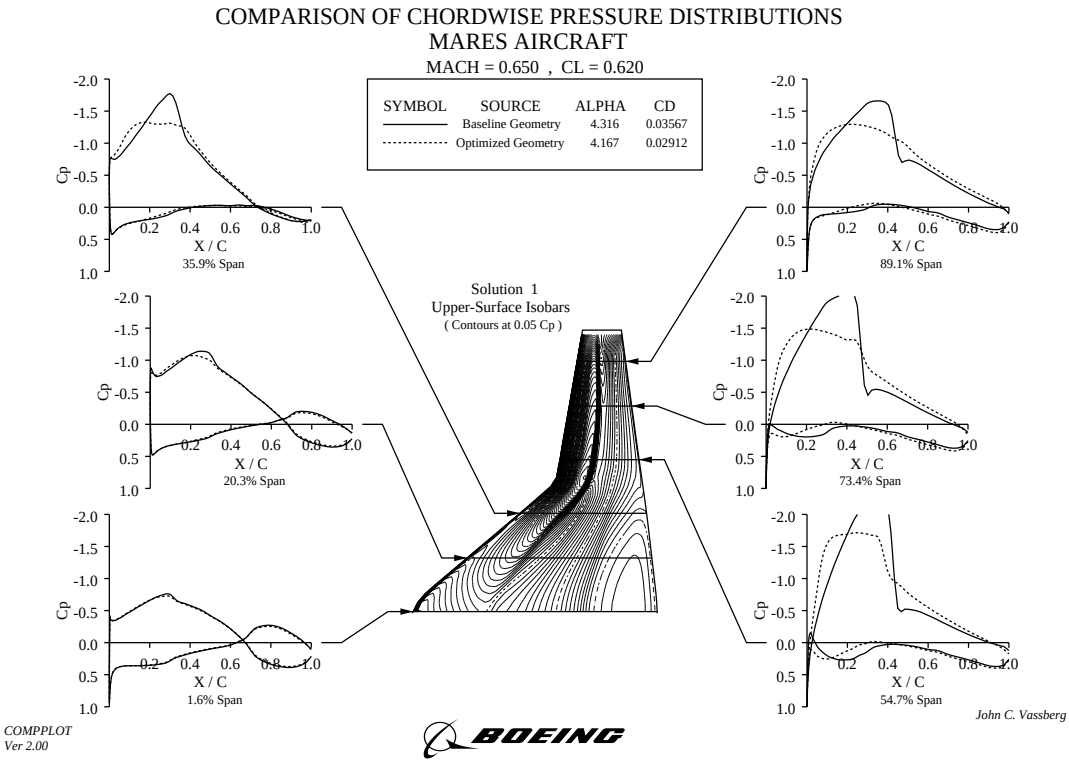

Figure 7: Comparison of Baseline and Euler Optimized Wing Pressure Distributions.

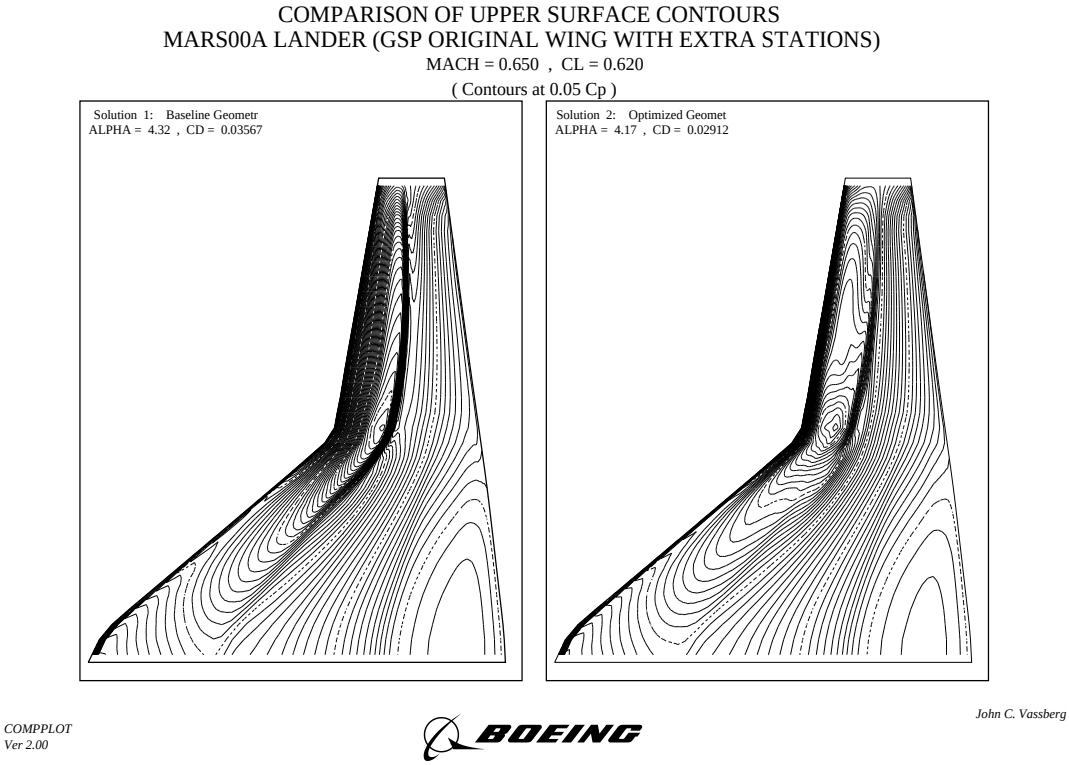

Figure 8: Comparison of Baseline and Euler Optimized Wing Pressure Contours.

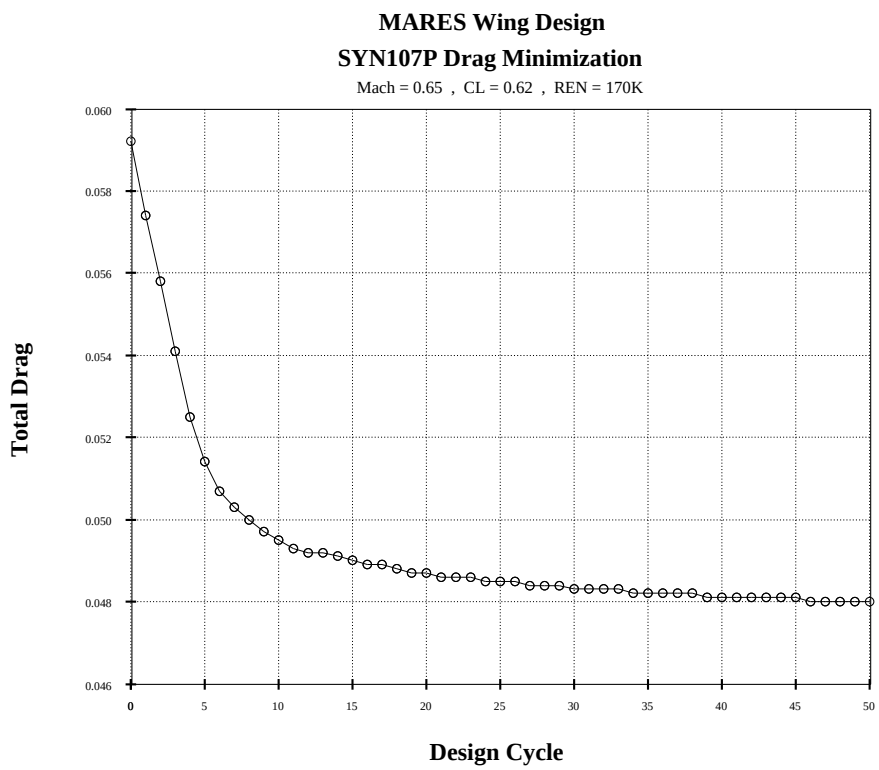

Figure 9: History of Drag Minimization during Navier-Stokes Optimization.

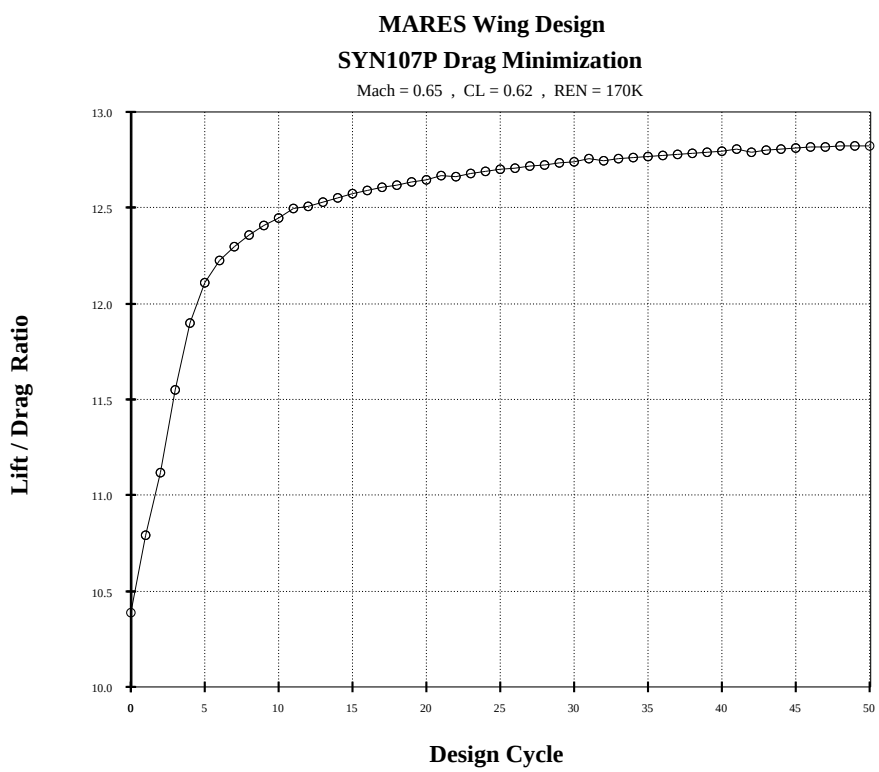

Figure 10: History of Lift-to-Drag Ratio during Navier-Stokes Optimization.

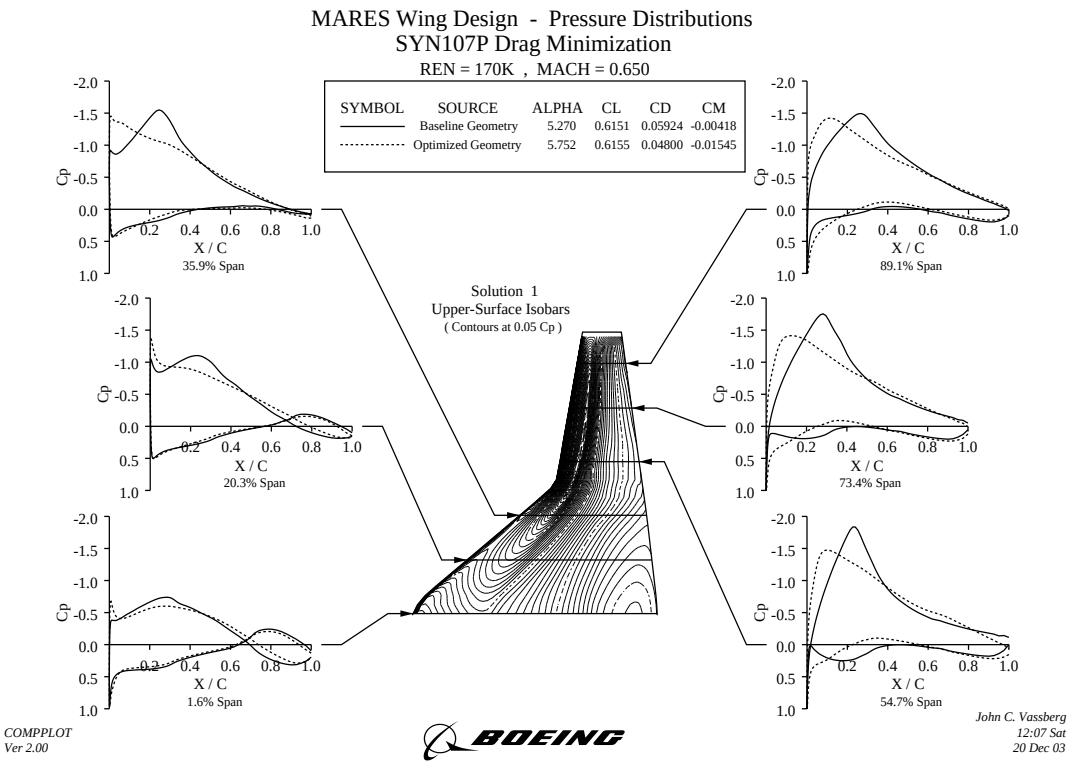

Figure 11: Comparison of Baseline and Navier-Stokes Optimized Wing Pressure Distributions.

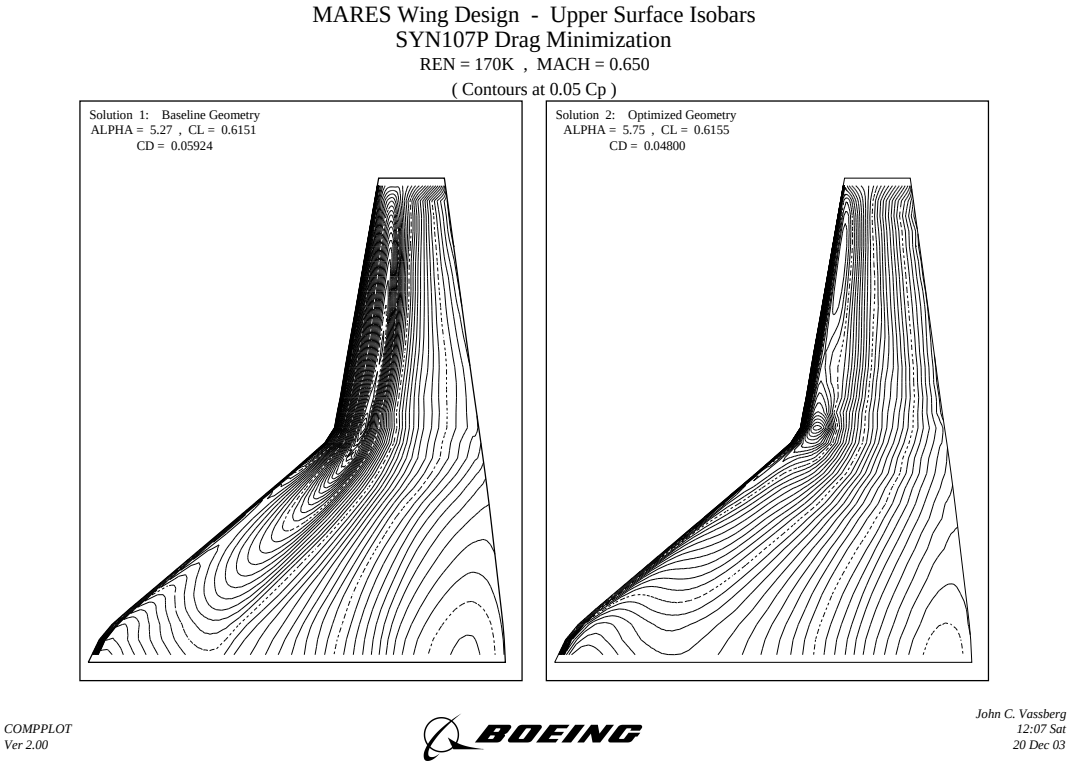

Figure 12: Comparison of Baseline and Navier-Stokes Optimized Wing Upper-Surface Isobars.

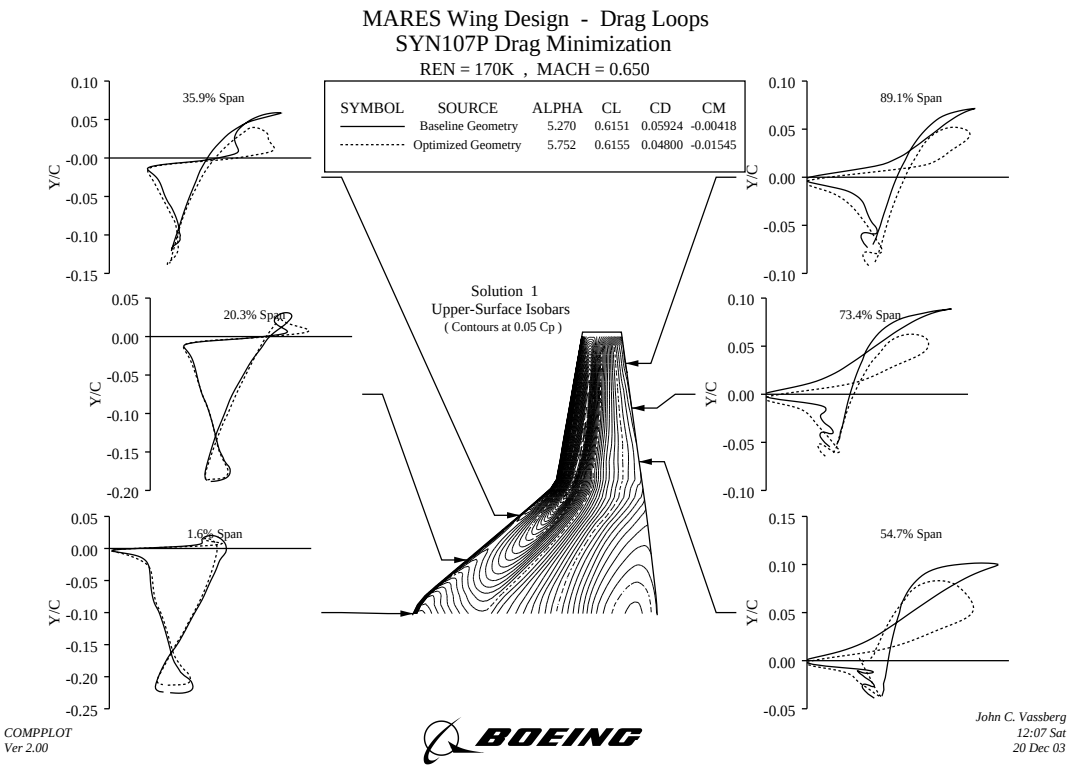

Figure 13: Comparison of Baseline and Navier-Stokes Optimized Wing Drag Loops.

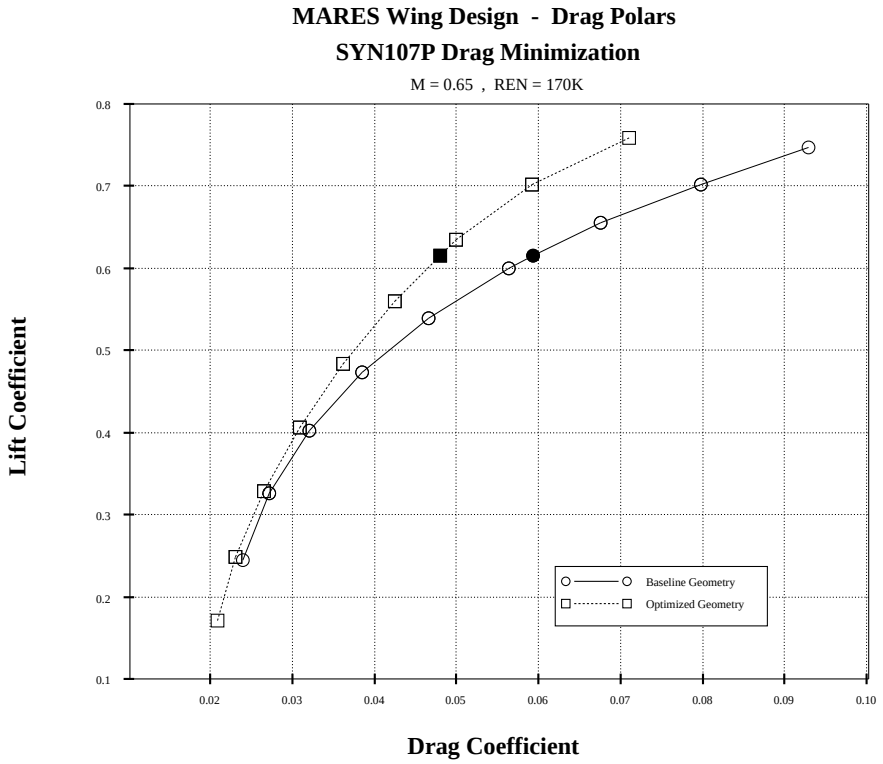

Figure 14: Comparison of Baseline and Navier-Stokes Optimized Wing Drag Polars.

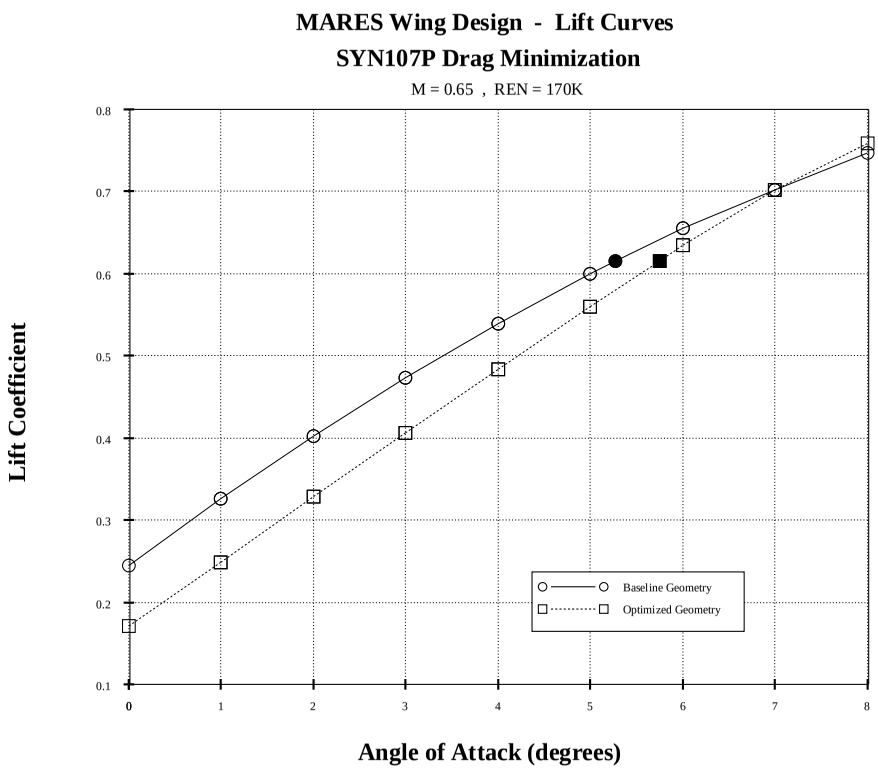

Figure 15: Comparison of Baseline and Navier-Stokes Optimized Wing Lift Curves.

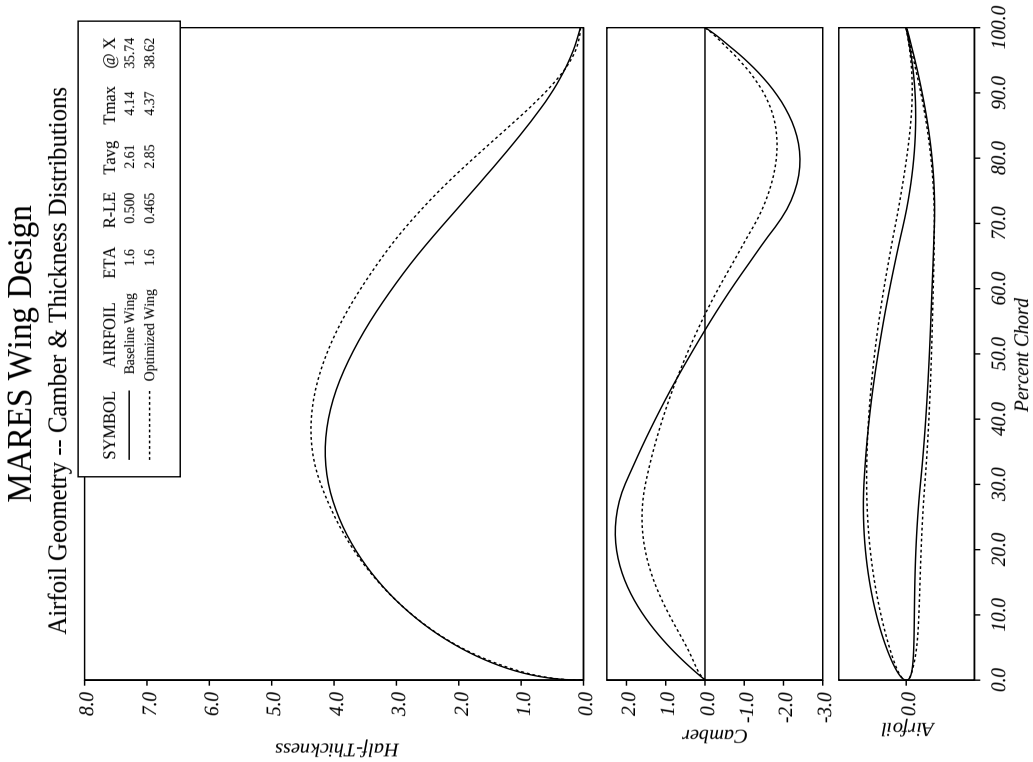

Figure 16: Comparison of Baseline and Navier-Stokes Optimized Wing Root Airfoil Sections.

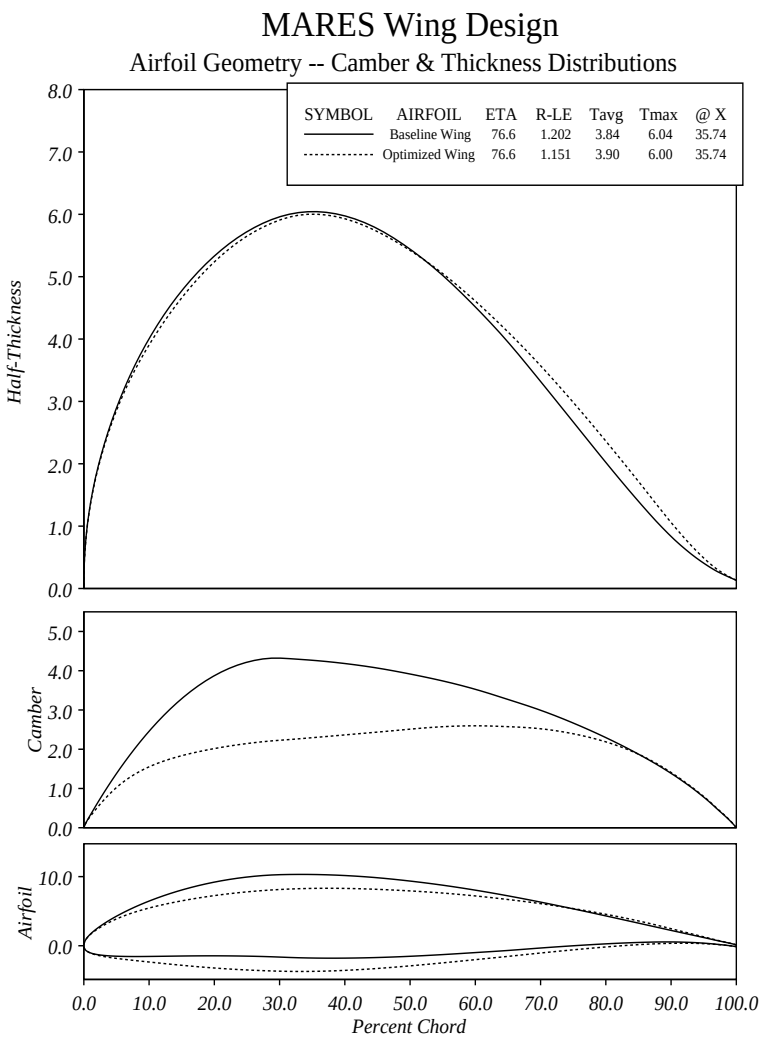

Figure 18: Comparison of Baseline and Navier-Stokes Optimized Wing Outboard Airfoil Sections.

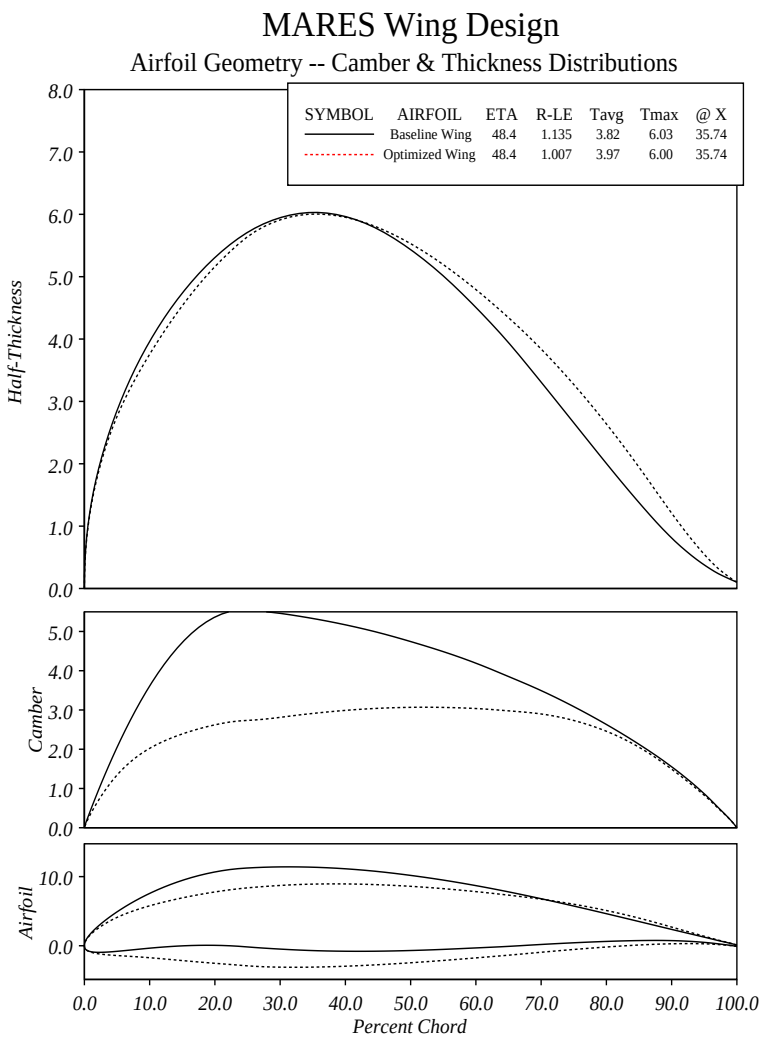

Figure 17: Comparison of Baseline and Navier-Stokes Optimized Wing Mid-Span Airfoil Sections.