

von Karman Institute for Fluid Dynamics

Lecture Series 2016-03

**INTRODUCTION TO OPTIMIZATION AND
MULTIDISCIPLINARY DESIGN**

May 23-27, 2016

*INDUSTRIAL APPLICATIONS OF
AERODYNAMIC SHAPE OPTIMIZATION*

J. C. Vassberg¹ & A. Jameson²
¹Boeing Commercial Airplanes, USA
²Stanford University, USA

Industrial Applications of Aerodynamic Shape Optimization

John C. Vassberg *
Advanced Concepts Design Center
Boeing Commercial Airplanes
Long Beach, CA 90846, USA

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

Von Karman Institute
Brussels, Belgium
24 May, 2016

Nomenclature

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

1 Introduction

This is the second of three lectures prepared by the authors for the von Karman Institute that deal with the subject of aerodynamic shape optimization. In our first lecture we introduced some theoretical background on optimization techniques commonly used in the industry, applied these approaches to a couple of very simple model problems, compared the results of these schemes, and discussed their merits and deficiencies as they relate to the class of aerodynamic shape optimization problems the authors deal with on a regular basis. In this lecture, we illustrate how the gradient of a complex system of nonlinear partial differential equations can be obtained for about the same computational cost as that of the cost function, and we provide a set of sample applications.

In an airplane design environment, there is no need for an optimization based purely on the aerodynamics of the aircraft. The driving force behind (almost) every design change is related to how the modification

*Boeing Technical Fellow

†T. V. Jones Prof. of Engineering

improves the vehicle, not how it enhances any one of the many disciplines that comprise the design. Although we focus this lecture on the aerodynamics of an airplane, we also include the means by which other disciplines are linked into and affect the aerodynamic shape optimization subtask. Another characteristic of the problems we typically (but not always) work on is that the baseline configuration is itself within 1-2% of what may be possible, given the set of constraints that we are asked to satisfy. This is certainly true for commercial transport jet aircraft whose designs have been constantly evolving for the past half century or more. Most of the sample applications provided herein do not fall into this category. Quite often the problem can be very constrained; this is the case when the shape change is required to be a retrofitable modification that can be applied to aircraft already in service. Occasionally, we can begin with a clean slate, such as in the design of an all-new airplane. And the problems cover the full spectrum of studies in between these two extremes.

Let's note a couple of items about this setting. First, in order to realize a true improvement to the baseline configuration, a high-fidelity and very accurate computational fluid dynamics (CFD) method must be employed to provide the aerodynamic metrics of lift, drag, pitching moment, spanload, etc. Even with this, measures should be taken to estimate the possible error band of the final analyses; this discussion is beyond the scope of these lectures. Figures 1-2 illustrate the class of aircraft and the level of detail the first author addresses every day. These Navier-Stokes CFD solutions are conducted on full-up cruise configurations, complete with wing, fuselage, engine groups, empennage, flap-support fairings, and winglets. The engine groups include a pylon, nacelle, core-cowl, shelf, and bifurcation flows. Although not obvious in these images, various fillets are also included. Finally, the CFD calculations are performed at prescribed lifting conditions by altering angle-of-attack, and are trimmed to specified center-of-gravity locations by adjusting the horizontal tail incidence. This level of detail is needed to achieve an accuracy on the absolute performance of the aircraft that is within 1% of flight test data. However, this is what is required to improve the performance of the aircraft by 1-2% without a numerical optimization yielding a false positive. The second item to consider is related to the definition of the design space. A common practice is to use a set of basis functions which either describe the absolute shape of the geometry, or define a perturbation relative to the baseline configuration. In order to realize an improvement to the baseline shape, the design space should not be artificially constrained by the choice of the set of basis functions. This can be accomplished with either a small set of very-well-chosen basis functions, or with a large set of reasonably-chosen basis functions. The former approach places the burden on the user to establish an adequate design space, the later approach places the burden on the optimization software to economically accommodate problems with large degrees of freedom. Over the past two decades, the authors have focused on solving the problem of aerodynamic shape optimization utilizing a design space of very large dimension. Our principal motivation for addressing the problem of large number of design variables is two fold. The first is to provide a situation where the design space never needs to be artificially constrained. The second is to allow us the flexibility to automatically set up the design space within the optimization software at the highest dimensionality supported by the discrete numerical simulation. In doing so, the aerodynamic shape optimization software based on these concepts allow the user to run optimizations immediately after set up of the analysis inputs are complete. This speeds time-to-first-optimization and minimizes the human errors associated with defining a design space. Furthermore, aerodynamic shape optimizations based on either the Euler or Navier-Stokes equations can be run on relatively inexpensive computer equipment.

The next section provides an overview of aerodynamic optimization. We develop an efficient evaluation for the gradient; this is based on solving an adjoint equation. A brief review of the search methods we utilize are then included. Following this discussion, we present a few selected case studies. These sample applications are all on design activities that we have been involved with; they include a Mars aircraft, a Reno Racer, and an aero-structural optimization of a generic B747 wing/body configuration. In addition to these, we present some recent work based on the publicly-available, open-source SU2 software, which introduces an alternative technique for lift-constrained optimization. These examples are based on the NACA0012 airfoil and the NASA Common Research Model (CRM) Wing/Body configuration.

2 Aerodynamic Design Trades

The objective of aerodynamic design is to produce a structurally feasible shape with sufficient carrying capacity, which achieves good aerodynamic performance.

For example, consider the generic task of delivering a payload between distant city pairs. The Breguet

Range equation, which aptly applies to long-range missions of jet aircraft, is:

$$Range = \frac{ML}{D} \frac{a}{SFC} \ln \left(\frac{W_0 + W_f}{W_0} \right). \quad (1)$$

Here, M is the cruise Mach number, L & D are the aerodynamic forces of lift and drag, respectively, a is the acoustic speed, SFC is the specific fuel consumption of the engines, W_0 is the aircraft landing weight, and W_f is the weight of fuel burned during the flight. The Breguet Range equation illustrates the importance of drag prediction as a function of lift and Mach number in the context of aerodynamic design; it also provides a glimpse into the interplay between the various disciplines.

Referring to Eqn (1), one might assume that the aerodynamic efficiency of an aircraft is represented by $\frac{ML}{D}$, the propulsion efficiency is embedded in SFC , and that the structural efficiency directly impacts W_0 . Interestingly, historical trends of in-service transport aircraft indicate that very little improvement in the $\frac{ML}{D}$ metric has been accomplished in the past 50 years, until recently with the Boeing 787. Yet it would be somewhat naive to state that *no* aerodynamic advances have been made during this period. In actuality, advances in aerodynamics have better served aircraft designs when traded for improvements in other disciplines. For example, the ability to increase the thickness-to-chord ratio of a wing while maintaining $\frac{ML}{D}$ not only reduces the structural weight of the wing, it also provides additional fuel volume. In terms of Eqn (1), an aerodynamic improvement of this nature would manifest itself as a decrease in W_0 and an increase in W_f with the net result being an increase in range. Reducing the aircraft's empty weight has the added benefit of reducing the cost of the vehicle. Obviously, this aerodynamic improvement would not be apparent in the trend charts of $\frac{ML}{D}$. Ironically, the recent increase in $\frac{ML}{D}$ for the B787 is primarily due to its composite wing structure, where the aero-structural trade optimizes out at a higher wing aspect ratio than that of a traditional aluminum wing.

Assume that an airline would like to provide a service between two cities with an aircraft that, when fully loaded with payload and fuel, is 1% short on range. Since the aircraft is fuel-volume limited, the only recourse is to reduce the payload weight. In relative terms, a typical ratio of weights might have $W_f = \frac{2}{3}W_0$ and $W_{payload} = \frac{1}{6}W_0$. In this scenario, Eqn (1) shows that the operator would have to reduce the payload (read revenue) by 7.6% to recover the 1% shortfall on range. Since most airlines operate on very small margins, this service most likely will no longer be a profit-generating venture. This example illustrates that in the current business of flight, a 1% delta in aircraft performance is a significant change. While improving an aircraft's performance by 1% may not be a trivial task given the usual constraints, losing 5% is easily done if attention is not paid to details such as juncture flows, external doublers, gaps, etc.

Now consider a more typical case where the aircraft does not suffer from a shortfall on range. In round numbers, the Direct Operating Costs (DOC) of a transport aircraft can be itemized as: 50% for the cost of ownership, 20% for fuel burn, 20% for crew salaries and maintenance, and 10% for miscellaneous other items. From an airline's perspective, if the DOC of its fleet of aircraft could be reduced by 5% with a new design (while providing the same set of services to its customers), the airline would most likely retire its entire fleet and replace it with the new aircraft [1].

So how can aerodynamics be leveraged to improve the economics associated with a flight-based mission? By enabling the development of a simplified high-lift-system design that for a given $\frac{L}{D}$ and $C_{L\ max}$ reduces manufacturing and maintenance costs as well as part count. By increasing the cruise Mach number without reducing $\frac{ML}{D}$ reduces the time-dependent costs such as crew and maintenance. Also, by increasing $\frac{ML}{D}$ without penalizing the other disciplines reduces fuel burn. These are just a few examples that illustrate how aerodynamic advances would impact DOC. A common requirement for achieving these goals is the accurate prediction of drag as a function of lift, Mach number, altitude and CG position, and with power effects included in the simulations.

To push aerodynamic technologies forward, it is becoming more important that accurate drag prediction become a consistent product of the CFD community. Once this prerequisite is accomplished, the full benefits of automated aerodynamic shape optimization may begin to be realized.

With the various on-going design programs, these are exciting times for the aircraft industry. A prime example is the Blended-Wing-Body (BWB) which has established a renaissance in the design of a family of all-new aircraft [2]. This revolutionary concept is enabling aerodynamic advances in all of the above areas, and then some. It presents challenges, yet offers significant opportunities, and as a result, a 5% reduction in DOC is within grasp. Suffice it to say that aerodynamics is not a sunset technology, but rather, it is as important today as it was a century ago; only the stakes have changed.

3 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.

3.1 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 [3]. Similarly, if one resorts to

direct code differentiation (ADIFOR [4, 5]), or complex-variable perturbations [6], 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 [7, 8, 9]. 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

$$R(w, \mathcal{F}) = 0. \tag{2}$$

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

$$\left[\frac{\partial R}{\partial w} \right]^T \psi = \frac{\partial I^T}{\partial w} \tag{3}$$

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 (2) is a partial differential equation, the adjoint equation (3) 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. ***** UPDATE ***** However, there is on-going research at Rice University to develop ADJIFOR [10] 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.

3.2 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 particularly 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 implicitly 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 [11], and acts as a preconditioner which allows the use of much larger steps. To apply an implicit smoothing in the ξ_1 direction, for example, the smoothed gradient $\bar{\mathcal{G}}$ may be calculated from a discrete approximation to

$$\bar{\mathcal{G}} - \frac{\partial}{\partial \xi_1} \epsilon \frac{\partial}{\partial \xi_1} \bar{\mathcal{G}} = \mathcal{G} \quad (4)$$

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

$$\delta I = - \iint \mathcal{G} \delta \mathcal{F} d\xi_1 d\xi_3$$

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

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 4, $\mathcal{G} = 0$. (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.

3.3 Computational Costs

In order to address the issues of the search costs, a variety of techniques were investigated in Reference [12] 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	$\mathcal{O}(K)$
(Note: K is independent of N)	

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 remaining sections of this lecture are devoted to sample applications associated with five case studies: a Mars aircraft, a Reno Racer, an aero-structural optimization of a generic B747 wing/body configuration, a NACA0012 airfoil, and the NASA Common Research Model.

4 MARES Development

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 Laboratory (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 light-weight 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 initial configuration layout and baseline geometry were created by the NASA/NRL/ITT development team. Further refinement of this geometry was conducted by the authors and is documented in Reference [13]. Aerodynamic design and shape optimization of the MARES wing-surface geometry was conducted using a number of computational fluid dynamics (CFD) methods, including FLO22 [14, 15], CFL3D [16], SYN88 and SYN107P; this lecture will concentrate on the simulations associated with the synthesis codes SYN88 and SYN107P, which are capable of both analysis and design calculations.

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.

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.

S_{ref}	36.38 ft^2	AR	4.9
b	13.38 ft	λ	0.3
C_{ref}	3.28 ft	$\Lambda_{c/A}$	5.5°
X_{ref}	3.28 ft	Λ_{LE}	10.0°
Y_{ref}	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 MacBookPro notebook in a few 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: $\rho = 2.356 * 10^{-5} \text{ slugs}/ft^3$, and $\nu = 2.2517 * 10^{-7} \text{ slugs}/ft/sec$. 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 C_p 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 noticeably 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.

5 Reno Racer

Once a year, Reno, Nevada plays host to an air show like none other; this of course is the Reno air races. Spectators from all over the world converge on this remote site to witness man and machine compete with one another in a series of races which culminate with the unlimited class events. The only hard rule required of an unlimited class race plane is that it be a propeller design and powered by a piston engine. Most entries are modified warbirds of WW-II vintage such as North American P-51 Mustangs, or Hawker Sea Furies, with power plants that produce well in excess of 1,000 HP. Miss Ashley II and Rare Bear of Figure 19 are representative of the unlimited class Reno racer. These aircraft are now over 50 years old and have very little service life left. To be competitive in the unlimited class, these aircraft see less than one hour between engine rebuilds, and these overhauls can cost upwards of \$250K. Furthermore, the historic value of flying WW-II fighters has increased so much that they are becoming unsuited for racing use. If unlimited class racing is to continue through the next decade, new race plane designs are required.

A goal of a new unlimited class design would be to significantly push the performance envelope of propeller-aircraft technologies. The mission of this special-purpose vehicle is to minimize the lap time around the 8.3 mile unlimited race course at Reno, depicted as the largest oval in Figure 20. Previous winners of this race have achieved average speeds around the oval course in excess of 450 MPH. The design requirements of the current development effort called for an average race speed of about 550 MPH. Because of this average velocity and the geometry of the race course, the aircraft pulls 4Gs about 60% of the time, 1G about 20% of the time, and transitions between these loads the remaining 20% of the time. Furthermore, because of the nature of this racing environment, 7G maneuvers are typically encountered to avoid mid-air collisions. While there are many other factors to consider during the design of this class an aircraft, the aforementioned loading conditions set the stage for the aerodynamic shape optimization of this vehicle. More precisely, the mission emphasizes that multi-point optimizations be performed over a weighted range of lifting conditions and on-set Mach numbers.

The design objectives of the Reno racer are given with respect to a standard day at the race location which is at 5000' MSL and ISA +20° C. The top speed in straight and level flight is to exceed 600 MPH TAS. The average lap speed around the 1999 unlimited race course is to exceed 550 MPH TAS. The aircraft is to be capable of sustaining a 9G maneuver load, subject to a 5G gust load; yielding a 14G limit load with a 1.5X safety factor. Roll rate should exceed 200° per sec at 350 KEAS. Stall speed should be less than 90 KEAS. Landing distance should not exceed 1500', dead stick. Note that some of these requirements are more stringent than the performance of some state-of-the-art jet fighters.

The design requirements state that the aircraft must be piston powered and propeller driven. Engine power-to-weight ratio for reliability at continuous output should be about 2.5 HP/lb. for a turbocharged piston engine with gear reduction and other accessories. The stability & control is to be provided by a manual, unboosted system with positive static & dynamic margins that exceed current unlimited-class race planes. There should be minimal change in stability between power on and off. For crew provisions, the design allows dual pilots in a tandem seating arrangement, with seats inclined 30° for G tolerance, and include MIL-SPEC oxygen and G-suit connections. Low altitude ejection for both pilots is also required.

The development of this aircraft began from the ground up, as an all-new design. Every major element of the airplane had to be engineered. This included the airplane's general layout, a unique propulsion system, the aerodynamic designs of the wing, fuselage and empennage, as well as the efficient integration of these and other subsystems. Although a unique propulsion system eventually became our baseline design, several systems were considered. These included a conventional tractor propeller with a front-mounted engine, and two mid-engine designs – one with a pusher prop aft of the tail and the other a body-prop design. While a tractor design is much more conventional, the design requirements favored a mid-engine concept. Avoiding propeller strike for the pusher design during rotation was a major issue. While there were many other factors that played into our decision, the body-prop design became our baseline configuration. With the propeller mounted aft of the wing, this concept also provided the possibility of promoting laminar flow on the forward fuselage and wing surfaces. A side view of the body-prop's general layout is provided in Figure 21 and a computer graphics rendering of this configuration in flight is given in Figure 22. The highest risk item of

this design is definitely related to engineering the structure to accommodate the load path between the tail and center wing box.

One can see from the general layout that the vertical tail (rudder) is rigged downward instead of in a normal upward position. This was done for two reasons; the first to provide a skid at the rudder tip to prevent propeller strike, and the second to keep it in clean air during a high, positive G maneuver.

The complete aircraft design effort was conducted by a very small team; the authors were tasked with the responsibility to design the outer-mold-line surfaces of the wing and fuselage components.

During this multi-disciplinary design effort, the general layout of the body-prop concept race plane evolved as the design team better understood how to maximize the performance of the integrated system. Normally, global changes such as those encountered are very disruptive during the design of a high-performance, transonic wing. However, utilization of the aerodynamic shape optimization software developed by the authors allowed various aircraft subsystems to be routinely modified without adversely impacting development costs or schedule; new wing designs occurred over night. Our ability to perform new optimizations over night, on affordable computers, was a key factor which allowed this form of simulation-based aerodynamic design work to be embraced by the rest of the design team. There were two other members of the team that are world-class wing designers, and they would have quickly relieved us of this duty if we were not providing quality designs in a timely manner. The complete evolution of the aircraft's general layout was accomplished in a very compressed time frame: our aerodynamic shape optimizations played a pivotal role in this achievement. More importantly, this evolution was required to meet all of the design goals imposed on the team by our sponsor. For more detailed information on this aircraft design, see References [17]-[18].

6 Reno Racer Wing Design

Any successful system design effort must accommodate a changing set of requirements as the designers of the various subsystems learn more about how their individual efforts impact and are affected by the actions of the other designers. This was certainly the case with this wing design as we integrated it with the fuselage, propulsion system, stability & control, manufacturing and overall packaging. Most of the assembled team has worked closely together for more than a decade. Our style has been to allow the individuals of the team to gravitate towards the work items that they feel most comfortable with, however, each member loosely participates in all concurrent activities in progress. On this project, participation was usually in the form of daily discussions regarding the overall design of the aircraft. By disseminating everybody's findings on a very frequent basis, the group as a whole began to understand how best to maximize the system's performance. Hence, this was very much a multi-disciplinary (MD) effort with the team members exploring the MD design space for increasingly better aircraft designs. These informal design reviews also provided regular sanity checks such that a poor design direction was never ventured too far. A characteristic of this dynamic design environment was that the design constraints at the subsystem levels were constantly and rapidly changing. New constraints can be in unexpected directions, and trying to program these dynamic changes in an MDO code, in a timely manner, can be quite daunting. This is where the man-in-the-loop belongs. The interface between humans can easily adapt to the ever-changing design requirements such that pertinent information continues to be shared across the appropriate disciplines.

This group regularly works on tight schedules and under small budgets. As a result, the most cost-effective tools are used at every stage of a design effort. Initially, when the design is not very well understood, design charts and rules-of-thumb dominate the effort. As the design begins to evolve, and these methods no longer add value to the direction of the group, linear methods are drawn into the tool set. Then, as the ROI of linear methods begins to reduce, they are replaced by non-linear tools – starting with the simple and finishing with the most sophisticated. Using the right tool at the right time helps manage costs and schedule, and allows the final designs to be competitive at the highest level. This approach is in contrast to those efforts that start using the most sophisticated numerical tools from the onset of the design activity.

6.1 Phase I: Conceptual Layout

The design of the wing geometry occurred in several phases; the duration of each of the first five phases lasted from 1 day to 1 week long. In most cases, there was a lapse between phases, as time was required for the team to digest the evolution of the aircraft design and formulate new ideas to investigate.

The basic requirements defined in phase I were based on conceptual methods and design charts. These requirements included the general layout of the wing (planform & thickness distribution), the design cruise condition

$$M = 0.77, \quad CL_{Total} = 0.32, \quad Ren = 14.5M,$$

the off-design capabilities for buffet

$$CL_{Buffet} = 0.64 \quad \text{at} \quad M = 0.72,$$

the drag divergence Mach number

$$M_{dd} = 0.80 \quad \text{at} \quad CL = 0.1,$$

the clean wing maximum lift coefficient

$$CL_{maxCW} \geq 1.6 \quad \text{at} \quad M = 0.2,$$

and a pad on the divergence Mach number to allow room for growth in out years.

The conceptual methods set the wing area, $S_{ref} = 75ft^2$, to provide a wing loading range of 40 – 60 lbs./ft². Despite the high dynamic pressures of the racing environment and the opportunity for very high wing loading, the stall speed requirement sized the wing area. The maneuver loads, sustained turn rate, and gust loads required that the wing have no buffet at $CL_{Total} = 0.64, M = 0.72$. A trade study of wing thickness, sweep and taper ratio was made using NACA SC(2) airfoils as a baseline. From this, a section thickness of 13.5% at the wing root and 12% at the tip was chosen, combined with a quarter-chord sweep of 28° to meet the $M_{dd} = 0.8$ requirement. An aspect ratio of $AR = 8.3$ and a taper ratio of $\lambda = 0.45$ were chosen to allow a wing-tip extension for a growth airplane. Conversely, a production break was included at 87% semi-span to allow a 4ft² reduction in wing area if ever needed. A planform Yehudi (inboard chord extension) was incorporated into the wing trailing edge to accommodate the main landing gear. Inclusion of this Yehudi also helped reduce the wing downwash angle of the flow entering the propeller. The wing is a two-spar design with spars at 15% and 65% chord, and is augmented with a secondary spar behind the main gear wells that parallel the Yehudi trailing edge. This secondary spar provides structural support at the main gear pivots. A one-piece wing box construction will be used to reduce weight and complexity.

6.2 Phase II: Rough Detailed Design

The baseline wing of phase II was defined using airfoil sections derived from NACA 64 sections, scaled to conform to the planform and thickness distribution established in phase I. Some cursory 2D aerodynamic optimizations were performed on these sections to better tailor their characteristics for the initial design conditions; the 2D conditions and geometry transformations used for this effort were based on simple-sweep theory. SYN103 was run in Euler, drag-minimization mode for this 2D design effort. The remaining unspecified geometric quantity for the wing was its twist distribution. To set this, FLO22 was used to provide the span load of the wing. This code, which solves the three-dimensional transonic potential flow equation, has been extensively used since its inception in 1976. (For reference, FLO22 runs in about 5 seconds on an AMD Athlon 850 MHz PC.) While FLO22 is a wing-only CFD code, pseudo-fuselage effects were included in the present work. The first pseudo-body influence is its acceleration of the on-set Mach number at a critical station on the wing; typically this is around 50%-60% semi-span. Running the isolated fuselage geometry in a surface-panel method and interrogating the flowfield velocity at the critical wing station determines this acceleration. The second pseudo-body effect is how the presence of the fuselage at an angle of attack warps the flowfield's local angle of attack as a function of span location. The third pseudo-body influence is related to the carry-over lift of the wing's circulation onto the fuselage. This ratio is defined as CL_{Total}/CL_{Wing} , is 1.22 for this configuration and was determined by running a surface-panel method on the wing/body combination. These pseudo-body effects are included in FLO22's wing-only solution by running the exposed wing in the code at the wing's CL, at a higher Mach number and re-referencing the results back to the original Mach, and adding a delta-twist distribution to the wing to simulate the flowfield warping. Using this procedure, a twist distribution was specified that yielded a near-elliptic span loading. This initial design was done very rapidly, covering only a two-day period, and provided a point to start the 3D design effort.

The initial FLO22 analyses indicated that the wing design requirements could be satisfied; the initial wing had a Mach capability of 0.775 at $CL_{Total} = 0.3$. However, there was serious concern with the body effects of the fuselage's low fineness ratio. The team was relatively sure that the baseline wing would have problems near the root region because of the atypical contouring of the fuselage geometry.

6.3 Phase III: Aero Optimization

In phase III, the first step was to assess the issues existing with the baseline wing geometry, designated Shark1, as it integrates with the fuselage. This analysis was performed using SYN88, and is illustrated in Figure 23. SYN88 is a wing/body Euler method which also incorporates an adjoint-based optimization procedure for aerodynamic shape design. In Figure 23, pressure distributions at several stations on the wing are provided. Adhering to standard aerodynamic practices, the pressure coefficient of the subplots are presented with the negative axis upward. The area trapped by the upper- and lower-surface pressure-distribution curves is equivalent to the local sectional lift coefficient. Each subplot is linked graphically to its corresponding location on the wing depicted in the center of the figure. Also included on the wing-planform plot are the upper-surface isobars of the first solution which is depicted by the solid curves in the perimeter subplots. A shock is evident with a concentration of contour lines in the isobar image and corresponds to a sharp discontinuity in the pressure-distribution subplots. The quantities in the legend of this figure correspond to the wing forces. The drag listed is only the inviscid drag (induced+shock). Recall that the design lift was $CL_{Total} = 0.32$ and that the carry-over lift ratio was 1.22 for this configuration. Hence, the wing lift is $CL_{Wing} = CL_{Total}/1.22 = 0.27$. The other item to note is that this analysis was performed at a Mach number of $M = 0.78$ rather than $M = 0.77$. The reason for this increase in freestream Mach number was due to the acceleration of the flowfield near the wing root from the propulsion system. Methods based on actuator-disc and blade-element theories determined this acceleration to be $\Delta M \simeq 0.01$. Since the wing-root region was of utmost concern at this stage in the design, the full level of propeller effects on the on-set Mach number was used. Referring to the wing-planform plot of Figure 23, notice the strong shock that unsweeps as it nears the side-of-body. The main purpose of a swept-back wing is to reduce the normal Mach number of the flow into a shock, however, if the shock unsweeps, this benefit is lost. As suspected, the contouring of the fuselage cross-sections had an adverse effect on the wing aerodynamics, unfortunately it was worse than expected. The inviscid drag (induced+shock) of the wing was $CD_{WingINV} = 180$ counts for the baseline configuration.

Phase III continued by running SYN88 in drag-minimization mode, constraining the wing modifications to be thicker everywhere than the baseline geometry to maintain structural depth; the fuselage geometry was frozen. Initially, these were single-point optimizations at the 4G design condition, just to scope the potential benefit. (For reference, a SYN88 wing/body analysis takes about 15 minutes on a Sony VAIO notebook computer with a 500 MHz Pentium II chip; a single-point optimization takes about 75 minutes.) Eventually, all optimizations were migrated to triple-point designs that considered a range of lifting conditions at the design Mach number. This range corresponded to variation and persistence of G loads being pulled during a lap of the race course. The design Mach number corresponded to an average speed around the track. Within 30 design cycles, SYN88 dropped the wing's inviscid drag from 180 counts to 104 counts. The results of this optimization are illustrated in Figure 24. Although fairly large improvements were realized, we felt we could do better if the fuselage contour near the wing trailing edge was allowed to be modified. Several concurrent changes to the aircraft's general layout were being considered. The team was forming new ideas as the complete system integration was beginning to be better understood. The changes that were directly related to the wing design were the fuselage reshaping and a trailing-edge planform blending that would allow more room for stowing the landing-gear structure. The planform modifications were made to the current wing and three additional fuselages were defined that stretched it by 1, 2 and 4 feet aft of the wing-root mid-chord and consistent with the engine packaging requirements. In fact, the trailing-edge modification was also done in a manner to help alleviate the shock-unsweep problem, as well as accommodate the landing gear. This planform change proved to be beneficial as another triple-point drag minimization was performed, which dropped the wing's inviscid drag from 148 counts to 98 counts at the design point. For clarification, this wing redesign was done with the original fuselage to establish a new base for the parametric study stretching the fuselage. Repeating similar triple-point optimizations on the 1-, 2- and 4-foot fuselage extensions provided sufficient data to show that the shock unsweep problem could be completely eliminated with a 2-foot fuselage stretch. This optimization reduced the wing's inviscid drag from 92 counts to 74 counts within 30 design cycles; the resulting wing geometry was designated Shark52. The pressure distributions and drag loops for Shark52 at $M = 0.78$ and $CL_{Wing} = 0.27$ are shown in Figures 25-26.

It should be emphasized that within the course of one week, the wing geometry had evolved from one that produced 180 counts of inviscid drag to Shark52, which only had 74 counts at the design point. During this week, the wing planform changed and the fuselage length stretched. This is an extremely large improvement accomplished in a very compressed time! Furthermore, the database of CFD solutions $\mathcal{O}(100)$ had grown

large enough that very informed modifications to the configuration could be made. This included the wing-planform change to better stow the landing gear, as well as the fuselage reshaping to eliminate the shock unsweep issue. Figure 26 shows a side-by-side comparison of the pressure isobars of the Shark52 wing and the initial wing that clearly illustrates the reduction of the shock strength across the entire wing. This improvement was a result of all of these important changes to the configuration.

The final aspect of phase III was to rebalance the aircraft. This required a 6-inch fuselage stretch forward of the wing to compensate for the 2-foot stretch of the after-body. Once the fuselage geometry was frozen, the wing pressures were polished by running SYN88 in inverse-design mode. A drag build-up of this design showed that, at 550 MPH, the maximum L/D of the complete aircraft was estimated to be 14.78. It occurs at $CL_{Total} = 0.49$, which corresponds to a 7G maneuver.

6.4 Phase IV: Laminar Flow

The team began to kick around the idea of a laminar-flow design. Quick calculations on the attachment-line Reynolds number, Re_θ , indicated that the Tollmien-Schlichting waves would decay rather than amplify. The possibility of having runs of laminar flow was achievable. The dilemma, however, was could laminar flow be achieved in the field? The primary mission of this plane occurs just above ground level where bug strikes are sure to occur, thus contaminating the wing's leading edge. We decided to investigate whether or not the wing's pressure distributions could be tailored to have favorable gradients for up to 40% chord without adversely affecting the aerodynamic performance of the fully-turbulent wing design. If it could, then the resulting design would be adopted, yet without taking credit for laminar-flow drag reductions.

Phase IV concentrated on promoting laminar flow on the wing without degrading the performance of the wing if it was fully turbulent as compared with the fully-turbulent design of Shark52. This objective was not limited to the design point, but rather was expanded to include a Mach number range $M \geq 0.74$ and a lift range of $CL_{Wing} \leq 0.27$. The first task was to compute the viscous flow about the Shark52 configuration at various flow conditions. This was accomplished using SYN107P, a wing/body Navier-Stokes method for analysis and design. The design Reynolds number was $Re = 14.5M$, based on the reference chord. (For reference, SYN107P runs in parallel under MPI; on a 16 processor AMD Athlon 650 MHz cluster, an analysis takes about 30 minutes of wall-clock time, while an optimization of 30 design cycles takes less than 3 hours.) Starting with the computed pressure distributions of Shark52, a series of inverse designs were performed, also with SYN107P. It was easier to redesign the wing at the higher Mach number and accommodate the requirements at $M = 0.74$, rather than the other way around. This study was completed with the wing geometry designated SharkNS7. Figure 28 illustrates the pressure distributions for SharkNS7 at $M = 0.78$ and a lift range of

$$0.18 \leq CL_{Wing} \leq 0.34.$$

At the design point $CD_{Wing} = 128$ counts which is composed of $CD_{WingForm} = 77$ counts and $CD_{WingSkinFriction} = 51$ counts.

Note that favorable gradients exist on both upper and lower surfaces for about the first 30%-40% chord, depending on span location and lifting condition. On the upper surface the shock will trigger transition provided any attachment line contamination from the fuselage boundary layer is removed by a notch-bump, and $Re_\theta < 200$. Re_θ varied from approximately 125 just outboard of the fuselage to around 80 at the wing tip. The amount of laminar run on the upper surface increases as Mach increases due to the shock moving aft on the airfoil as well as the pressure gradients becoming more favorable. At race conditions the wing should have an appreciable extent of laminar flow, provided the surface of the wing is smooth and free of particulate contamination. The estimated benefit of the laminar flow runs is between 10 and 20 counts of drag reduction, depending on Mach number. This level of drag reduction increases the aircraft's performance by an additional 5%, which is a significant improvement.

6.5 Phase V: Final Touches

The first task of phase V was to establish an appropriate leading-edge radius distribution, tailored for low-speed characteristics, without really changing the wing pressure distributions at the cruise design conditions. This modification was accomplished with local, explicit geometry perturbations. An additional modification to the wing thickness distribution was also done. After these changes were incorporated into SharkNS7, the final wing was analyzed to verify that these geometry changes did not adversely effect the pressure distributions. When overlaid on the same plot, the curves nearly appeared as one.

Finally, clean wing CL_{max} , CL_{maxCW} , was computed to ensure the wing satisfied the required clean wing stall speed. The design requirement was to provide $CL_{maxCW} > 1.6$ at $M = 0.2$. This was determined by finding the flow condition where the wing's Cp_{min} distribution reached an empirically-determined critical value. The final wing provided a $CL_{maxCW} = 1.64$, just meeting the requirement.

The final wing geometric characteristics are shown in Figure 29 which illustrates the half-thickness and camber distributions of the root and outboard airfoil sections. Note that the root airfoil is 13% thick @ 31.5% chord and has -1.2% camber. The outboard airfoil is 11.5% thick @ 38.6% chord and has +1% camber. While these thicknesses are about 0.5% thinner of that specified in the conceptual design stage, it was the team that reduced this thickness, not the optimization exercises.

This concludes our discussion on the Reno Racer. Next, we will review an aero-structural optimization of a generic 747 wing/body configuration.

7 Generic 747 Wing/Body

In this section, a simplified structural weight model is directly coupled with the high-fidelity aerodynamic shape optimization procedure. Further, the planform of the aircraft wing is allowed to be redesigned. The original work on this case study is documented in detail in Reference [19].

In an airplane design, there are many criteria to be satisfied: multiple design points, fuel distributions, stability-and-control effects, aero-elastics, etc. Some of the general criteria that must be met in the design of any efficient transonic wing include:

1. Good drag characteristics (parasite, induced, compressibility) over a range of lift coefficients, i.e. $C_{L_{design}} \pm 0.1$ at M_{cruise} .
2. No excess penalties for installation of nacelle-pylons, fairing, etc.
3. Sufficient buffet boundary margin at cruise lift coefficients (1.3g margin).
4. No pitch-up tendencies near stall or buffet.
5. Maintain control surface effectiveness.
6. No unsatisfactory off-design performance.
7. Sufficient fuel volume for design range.
8. Structurally efficient (to minimize weight).
9. Sufficient space to house main landing gear.
10. Compatible with the high-lift system.
11. Consistent with airplane design for relaxed static stability.
12. Manufacturable at a reasonable cost.

7.1 Cost Function

In this optimization, the two relevant disciplines are aerodynamics and structural weight. Therefore we minimize a combination of aircraft drag and wing weight. This optimization not only makes the design more realistic, but we can also relax some of the constraints on thickness. Hence, we mainly target the minimization of

$$I = \alpha_1 C_D + \alpha_2 C_W. \quad (5)$$

Here α_1 and α_2 are properly chosen weighting constants, and $C_W = \frac{weight}{q_\infty S_{ref}}$ is a non-dimensional weight coefficient. This choice of cost function emphasizes the trade-off between aerodynamics and structures. For design of a long range transport aircraft, Eqn (5) actually centers on the idea of improving the range of the aircraft.

Consider the well known Breguet range equation which provides a good first estimate of the range of the airplane

$$R = \frac{V}{C} \frac{L}{D} \ln \left(\frac{W_e + W_f}{W_e} \right) \quad (6)$$

where V is the speed, C is the specific fuel consumption of the engine, L/D is the lift to drag ratio, W_e is the landing weight, and W_f is the weight of the fuel burnt. During the last few decades, the means to improve the efficiency of an airplane has centered around reducing the fuel consumption C of the engine, increasing VL/D , and reducing the airplane weight. The last two methods together imply that the constants α_1 and α_2 in Eqn (5) can be estimated from the range Eqn (6).

7.2 Structural Weight Model

Wing weight is directly given by the wing structure, which is sized by aerodynamic load, allowable deflection, and failure criteria such as buckling. Different methods to estimate wing weight have been proposed, ranging from empirical expression to complex Finite Element Analysis. At the detailed design level, once the wing structure has been laid out, structure engineers use finite element analysis to size the interior structure to satisfy various criteria. Then the material weight of each component is added up to calculate the structural weight.

When the detailed structure layout is not known, which is usually the case for conceptual and preliminary design levels, it is very difficult to predict the structural weight correctly. We select an empirically-influenced simple structure model that can be expressed analytically. Here, the wing structure is modeled by a box beam whose primary structural material is the upper and lower box skin. The skin thickness (t_s) varies along the span and resists the local bending moment caused by the wing lift. Then, the structural wing weight W_{wing} can be calculated on the basis of the skin material.

Consider the box structure of a swept wing whose quarter-chord sweep is Λ and its cross-section A-A is as shown in Figure 30. The skin thickness t_s , structure box chord c_s , and overall thickness t vary along the span, such that the local stress is equal to the maximum allowable stress everywhere. The maximum normal stress due to bending at section z^* is:

$$\sigma = \frac{M(z^*)}{t t_s c_s}.$$

The corresponding structural box-beam weight is:

$$\begin{aligned} W_{\text{wing}_{\text{box}}} &= \rho_{\text{mat}} g \int_{\text{structural span}} 2t_s c_s dl \\ &= 2 \frac{\rho_{\text{mat}} g}{\sigma \cos(\Lambda)} \int_{-\frac{b}{2}}^{\frac{b}{2}} \frac{M(z^*)}{t(z^*)} dz^* \\ &= 4 \frac{\rho_{\text{mat}} g}{\sigma \cos(\Lambda)} \int_0^{\frac{b}{2}} \frac{M(z^*)}{t(z^*)} dz^*, \end{aligned}$$

and

$$\begin{aligned} C_{W_b} &= \frac{W_{\text{wing}_{\text{box}}}}{q_{\infty} S_{\text{ref}}} \\ &= \frac{\beta}{\cos(\Lambda)} \int_0^{\frac{b}{2}} \frac{M(z^*)}{t(z^*)} dz^*, \end{aligned} \quad (7)$$

where

$$\beta = \frac{4\rho_{\text{mat}} g}{\sigma q_{\infty} S_{\text{ref}}},$$

ρ_{mat} is the material density, and g is the acceleration due to gravity.

The bending moment can be calculated by integrating pressure toward the wing tip. Ignoring the end effects due to the rotated axis of the box beam,

$$M(z^*) = - \int_{z^*}^{\frac{b}{2}} \frac{p(x, z)(z - z^*)}{\cos(\Lambda)} dA$$

$$= - \int_{z^*}^{\frac{b}{2}} \oint_{\text{wing}} \frac{p(x, z)(z - z^*)}{\cos(\Lambda)} dx dz.$$

Thus

$$C_{W_b} = \frac{-\beta}{\cos(\Lambda)^2} \int_0^{\frac{b}{2}} \int_{z^*}^{\frac{b}{2}} \oint_{\text{wing}} \frac{p(x, z)(z - z^*)}{t(z^*)} dx dz dz^*. \quad (8)$$

However, C_{W_b} must be expressed as $\int_{\mathcal{B}} d\mathcal{B}_\xi$ in the computational domain, or $\int \int dx dz$ in a physical domain to match the adjoint boundary term.

To switch the order of integral of Eqn (8), introduce a Heaviside function

$$H(z - z^*) = \begin{cases} 0, & z < z^* \\ 1, & z > z^* \end{cases}$$

Then Eqn (8) can be rewritten as

$$C_{W_b} = \frac{-\beta}{\cos(\Lambda)^2} \int_0^{\frac{b}{2}} \oint_{\text{wing}} p(x, z) K(z) dx dz, \quad (9)$$

where

$$\begin{aligned} K(z) &= \int_0^{\frac{b}{2}} \frac{H(z - z^*)(z - z^*)}{t(z^*)} dz^* \\ &= \int_0^z \frac{z - z^*}{t(z^*)} dz^* \end{aligned}$$

Finally, to account for the weight of other wing material such as ribs, spars, webs, stiffeners, leading and trailing edges, slats, flaps, main gear doors, primer and sealant, we multiply C_{W_b} by a correction factor $K_{\text{corr,b}}$.

Moreover, statistical correlation over the range of aircraft type indicates a relationship between W_{wing}/S and $W_{\text{wing,box}}/S$ as a linear function, shown in Figure 31. Therefore we add another term to account for area-dependent wing weight

$$C_{W_s} = \frac{1}{S_{\text{ref}}} \oint_{\mathcal{B}} |S_{22}| d\xi_1 d\xi_3, \quad |S_{22}| = \sqrt{S_{2i} S_{2i}}, \quad (10)$$

along with a correction factor $K_{\text{corr,s}}$. Thus

$$C_W = K_{\text{corr,b}} C_{W_b} + K_{\text{corr,s}} C_{W_s}. \quad (11)$$

7.3 Wing Weight Estimation

From our discussion in section 7.2, the expression of wing weight contains terms that correspond to bending-load-carrying material and two correction factors; $K_{\text{corr,b}}$ and $K_{\text{corr,s}}$. To estimate these correction factors, we can rewrite the expression (11) as

$$\frac{W_{\text{wing}}}{S} = \alpha_1 \frac{W_b}{S} + \alpha_2 \quad (12)$$

where W_{wing} is the total weight of the wing, W_b is the weight of the wing box, and S is a gross wing area. When there is information from more than two aircraft, we can use the “least-square” curve-fitting strategy to calculate α_1 and α_2 .

To compute α_1 and α_2 , we solve a system of two equations and two unknowns. These equations have the form of Eqn (12) but with different values of $\frac{W_{\text{wing}}}{S}$ and $\frac{W_b}{S}$. This gives

$$\alpha_1 = 1.30 \quad \text{and} \quad \alpha_2 = 6.03 \text{ (lb/ft}^2\text{)}.$$

However, the total wing weight estimation will be inexact because of the limitation of aircraft number. Finally, $K_{\text{corr,b}}$ and $K_{\text{corr,s}}$ can be computed by

$$K_{\text{corr,b}} = \alpha_1 \quad \text{and} \quad K_{\text{corr,s}} = \frac{\alpha_2}{q_\infty}$$

7.4 Planform Design Variables

From the trade study, the parameters that lead up to a basic design which satisfies the general design criteria include:

- wing shape
- span
- aspect ratio
- airfoil types
- area
- sweep
- taper ratio
- airfoil thickness

Because some of these parameters do not define the wing geometry uniquely, we employ another set of design parameters that still represent these parameters but can be extracted from surface mesh points. Here we model the wing of interest using six planform variables: root chord (c_1), mid-span chord (c_2), tip chord (c_3), span (b), leading-edge sweep(Λ), and wing thickness ratio (t), as shown in Figure 32. This choice of design parameters will lead to an optimum wing shape that will not require an extensive structural analysis and can be easily manufactured.

7.5 Maximizing Range

The choice of α_1 and α_2 in Eqn (5) greatly affects the optimum shape. We can interpret α_1 and α_2 as how much emphases we give to drag and wing weight. If $\frac{\alpha_1}{\alpha_2}$ is high, we focus more on minimizing the drag than the weight and we tend to get an optimum shape that has low C_D but high C_W .

An intuitive choice of α_1 and α_2 can be made by considering the problem of maximizing range of an aircraft. Consider the Breguet range Eqn (6)

$$R = \frac{V L}{C D} \ln \left(\frac{W_e + W_f}{W_e} \right)$$

where W_e is the gross weight of the airplane without fuel and W_f is weight of fuel burnt.

If we take

$$\begin{aligned} W_1 &= W_e + W_f = \text{fixed} \\ W_2 &= W_e \end{aligned}$$

then the variation of the weight can be expressed as

$$\delta W_2 = \delta W_e.$$

With fixed $\frac{V}{C}$, W_1 , and L , the variation of R can be stated as

$$\begin{aligned} \delta R &= \frac{V}{C} \left(\delta \left(\frac{L}{D} \right) \ln \frac{W_1}{W_2} + \frac{L}{D} \delta \left(\ln \frac{W_1}{W_2} \right) \right) \\ &= \frac{V}{C} \left(-\frac{\delta D}{D} \frac{L}{D} \ln \frac{W_1}{W_2} - \frac{L}{D} \frac{\delta W_2}{W_2} \right) \\ &= -\frac{V L}{C D} \ln \frac{W_1}{W_2} \left(\frac{\delta D}{D} + \frac{1}{\ln \frac{W_1}{W_2}} \frac{\delta W_2}{W_2} \right) \end{aligned}$$

and

$$\begin{aligned} \frac{\delta R}{R} &= - \left(\frac{\delta C_D}{C_D} + \frac{1}{\ln \frac{W_1}{W_2}} \frac{\delta W_2}{W_2} \right) \\ &= - \left(\frac{\delta C_D}{C_D} + \frac{1}{\ln \frac{C_{W_1}}{C_{W_2}}} \frac{\delta C_{W_2}}{C_{W_2}} \right). \end{aligned}$$

If we minimize the cost function defined as

$$I = C_D + \alpha C_W,$$

where α is the weighting multiplication, then choosing

$$\alpha = \frac{C_D}{C_{W_2} \ln \frac{C_{W_1}}{C_{W_2}}}, \quad (13)$$

corresponds to maximizing the range of the aircraft.

7.6 Optimization Results

In these calculations the flow is modeled by the Reynolds Averaged Navier-Stokes equation, with a Baldwin Lomax turbulence model. This turbulence model was considered sufficient because the design point at cruise has predominately-attached flow. The nominal cruise condition of the generic 747 wing/fuselage combination has a Mach number of 0.85 and a lift coefficient of $C_L = 0.45$. On the corresponding (256x64x48) grid, the wing is represented by 4,224 surface mesh points and six planform variables. All coefficients are calculated with a fixed reference area based on the baseline configuration. Thus an increase in skin friction due to an increase in wetted area will appear as an increase in the skin-friction drag coefficient.

As a reference point, we first modified only the wing sections to eliminate the shock drag, while keeping the planform fixed. Figure 33 shows the redesigned calculation. In 30 design cycles, the drag was reduced from 137 counts to 127 counts (7.3% reduction), while the weight remained roughly constant.

Next, we implement both section and planform optimization in a viscous redesign, using an inviscid redesign as a starting point. Figure 34 shows the effect of allowing changes in sweepback, span, root chord, mid-span chord, and tip chord. The parameter α_2/α_1 is chosen to maximize the range of the aircraft. In 30 design cycles, the drag was reduced to 117 counts (14.5% reduction from the baseline), while the dimensionless structure weight was decreased from 546 counts to 516 counts (6.1% reduction), which corresponds to a reduction of 4,800 lbs. The planform changes are shown in Figure 35. This viscous redesigned wing has less drag and structural weight than the fixed-planform viscous redesigned wing.

When we compare this planform with the redesign by inviscid optimization, as seen in Figure 36, we can see that the effect of viscosity is to shrink the area of the inviscidly redesigned planform. This trend is to be expected because skin friction drag varies roughly linearly with the area. By reducing the area, we can reduce the skin friction drag.

The results from the viscous planform optimizations yield large drag reduction without structural weight penalty in a meaningful way. They show the following basic trends:

- Increase wing span to reduce vortex drag,
- Reduce sweep but increase section-thickness to reduce structural weight,
- Use section optimization to minimize shock drag.

Although the suggested strategy tends to increase the wing area, which increases the skin friction drag, the pressure drag drops at a faster rate, dominating the trade-off. Overall, the combined results yields improvements in both drag and weight.

7.7 Pareto Front

In order to present the designer with a wider range of choices, the problem of optimizing both drag and weight can be treated as a multi-objective optimization problem. In this sense one may also view the problem as a “game”, where one player tries to minimize C_D and the other tries to minimize C_W . In order to compare the performance of various trial designs, designated by the symbol X in Figure 37, they may be ranked for both drag and weight. A design is un-dominated if it is impossible either to reduce the drag for the same weight or to reduce the weight for the same drag. Any dominated point should be eliminated, leaving a set of un-dominated points which form the Pareto front. In Figure 37, for example, the point Q is dominated by the point P (same drag, less weight) and also the point R (same weight, less drag). So point Q will be eliminated. The Pareto front can be fit through the points P, R and other dominating points, which may be generated by using an array of different values of α_1 and α_2 in the cost function to compute different optimum shapes. With the aid of the Pareto front the designer will have freedom to pick the most useful design.

The problem of optimizing both drag and weight can be treated as a multi-objective function optimization. A different choice of α_1 and α_2 will result in a different optimum shape. The optimum shapes should not dominate each other, and therefore lie on the Pareto front. The Pareto front can be very useful to the designer because it represents a set which is optimal in the sense that no improvement can be achieved in one objective component that does not lead to degradation in at least one of the remaining components.

Figure 38 shows the effect of the weighting parameters (α_1, α_2) on the optimal design. As before, the design variables are sweepback, span, chords, section thickness, and mesh points on the wing surface. In Figure 38, each point corresponds to an optimal shape for one specific choice of (α_1, α_2) . By varying α_1 and α_2 , we capture a Pareto front that bounds all the solutions. All points on this front are acceptable solutions, and choosing the final design along this front depends on the nature of the problem and several other factors. The optimum shape that corresponds to the maximum Breguet range is also marked in the figure. In this test case, the Mach number is the current normal cruising Mach number of 0.85. We allowed section changes together with variations of sweep angle, span length, chords, and section thickness. Figure 33 shows the baseline wing. Figure 34 shows the redesigned wing. The parameter $\frac{\alpha_2}{\alpha_1}$ was chosen according to formula (13) such that the cost function corresponds to maximizing the range of the aircraft. Here, in 30 design iterations the drag was reduced from 137 counts to 117 counts and the structural weight was reduced from 546 counts (88,202 lbs) to 516 counts (83,356 lbs). The large reduction in drag is the result of the increase in span from 212.4 ft to 231.7 ft, which reduces the induced drag. The redesigned geometry also has a lower sweep angle and a thicker wing section in the inboard part of the wing, which both reduce the structural weight. Moreover the section modification prevents the formation of shock. The baseline and redesigned planforms are shown in Figure 36, together with the planform which resulted from inviscid optimization. Overall, the redesign with variation planform gives improvements in both aerodynamic performance and structural weight, compared to the previous optimization with a fixed planform.

8 NACA0012-ADO Airfoil

Our next example optimization is based on a modified NACA0012 airfoil section. Here, we have adopted the sharp trailing-edge closure of the Aerodynamic Design Optimization Discussion Group's (ADO-DG) Case 1 airfoil [20]. We introduce an alternative method of performing lift-constrained optimizations based on the capabilities of the open-source SU2 software [21, 22]. SU2 includes a variety of popular continuous adjoints, such as that for lift, drag, and pitching moment. We also provide comparisons of gradients based on the continuous adjoint, the discrete adjoint, and that through finite differencing. The discrete adjoint is automatically generated using the CodiPack environment under development by Gauger and Albring [23]. The implementation of CodiPack into SU2 is clean and efficient, and has proven to be very useful when introducing new cost functions into our daily industrial optimization studies. It has always been our experience that the cost of solving a continuous adjoint is comparable to solving the flow equations. The discrete adjoint provided by SU2 based on CodiPack also exhibits this level of efficiency, however, memory requirements are significantly larger than that for the continuous adjoint. The optimizer utilized for all optimizations of this test case is the Sequential Least-Squares Programming (SLSQP) of the SciPy package.

8.1 Aerodynamic Characteristics

In order to develop a lift-constrained drag minimization, whether it is for an airfoil or an aircraft, we will first review some basic aerodynamic characteristics for these class of lifting geometries. Figure 39 illustrates a typical relationship between lift and angle-of-attack, α . Here, the lift coefficient, C_L , is nearly a linear function of α . The airfoil design point for this inviscid test case is $M = 0.8$ and $\alpha = 1.25^\circ$. These conditions yield a design lift coefficient of $C_L^* \approx 0.347$. This data point is highlighted in the figure with a large circle. Also at the design point, the slope of the lift curve is $C_{L\alpha}^* \approx 0.274/^\circ$. We will show how to utilize this slope to drive both the analysis and adjoint solutions for fixed- C_L optimizations. Finally, the problem statement for this airfoil optimization is given as:

$$\text{Minimize } C_D \text{ , s.t. } C_L = C_L^* . \quad (14)$$

To keep this example extremely simple, we do not include any of the usual constraints on geometric thickness, symmetry or aerodynamic pitching-moment, and yet, we have crafted this example problem such that it exhibits all behaviours necessary to teach the alternative technique for lift-constrained optimization.

Let's now review the drag polar of Figure 40, represented by the black line with solid circle symbols. By observation, note that C_D is an increasing function of C_L . As a consequence, optimizers will naturally try to reduce lift in an attempt to reduce drag. For a fixed- α optimization, this effect is embedded in the gradient of drag, in that a "hidden" part of the shape change primarily works to alter the lift, and hence to change the drag. This part of the fixed- α drag gradient is undesirable, and it may take 10s of design cycles before an optimizer can compensate for it accordingly. Hence, we seek a better approach.

8.2 Lagrangian Dual Cost Function

We introduce an alternative optimization problem, based on a Lagrangian dual cost function, defined as:

$$\mathcal{L} \equiv C_D - \lambda C_L. \quad (15)$$

Here, λ is a Lagrange multiplier which physically represents the inverse of the slope of the drag polar at the design point, i.e., $\lambda = \frac{dC_D}{dC_L}$. However, since λ is not known *a priori*, it must be determined during the course of optimization. Furthermore, λ may change with each airfoil throughout the design cycles. Figure 40 includes this dual cost function, depicted as a blue line with solid square symbols. Notice that the Lagrangian naturally minimizes at the design point. Hence, $\frac{d\mathcal{L}}{dC_L} = 0$ at C_L^* . Our problem statement for the Lagrangian dual is now given by:

$$\text{Minimize } \mathcal{L} \text{ , s.t. } C_L = C_L^*. \quad (16)$$

When we set up our design space, we include α as a design variable. Since $C_L(\alpha)$ is a well behaved differentiable function, $\frac{d\mathcal{L}}{dC_L} = \mathcal{L}_\alpha C_{L\alpha}$, which in turn implies that $\mathcal{L}_\alpha = 0$ at C_L^* .

In a similar manner as the process we introduced in our first lecture for the Spider-Fly problem with an active constraint, we first project the flow analysis into its allowable design space by adjusting α to drive C_L to C_L^* . This yields a converged flow solution at the design lifting condition. Then, we adjust λ during the adjoint solution by driving $\mathcal{L}_\alpha \rightarrow 0$ at C_L^* . This provides a converged adjoint solution for the Lagrangian dual cost function with an appropriate value of λ . However, since neither α nor λ are known *a priori*, they must be determined with a driver of some sort. We discuss our approach to developing these drivers next.

α -Driver

A Newton iteration that updates α to seek $C_L = C_L^*$ is given by:

$$\alpha^{new} = \alpha^{old} - (C_L - C_L^*)/C_{L\alpha}. \quad (17)$$

Application of this driver is illustrated by Figure 41. Here, α is initialized to be 0° , and is first updated at 4000 iterations, then every 2000 iterations after that. A fixed value of $C_{L\alpha} = 0.274/^\circ$ is used throughout. (This value is estimated by the curve of Figure 39, but any reasonable estimate will do, provided it is greater than or equal to half the actual slope.) In this figure, the blue line represents the history of α (left axis), while the red line depicts the history of C_L (right axis) as it approaches $C_L^* \approx 0.347$. Note that α converges to $\alpha^* = 1.25^\circ$ as expected.

λ -Driver

Now we will review what is required to drive λ to its appropriate value during the convergence of the adjoint solution of the Lagrangian dual cost function. Since $\mathcal{L}_\alpha = 0$ at C_L^* , we use a Newton iteration to update λ as given by:

$$\lambda^{new} = \lambda^{old} - \mathcal{L}_\alpha / \frac{d\mathcal{L}_\alpha}{d\lambda}. \quad (18)$$

Hence, we need to determine an appropriate value for $\frac{d\mathcal{L}_\alpha}{d\lambda}$. If we differentiate Eqn (15), first by λ , and then by α , we get:

$$\begin{aligned} \frac{d\mathcal{L}}{d\lambda} &= -C_L, \\ \frac{d^2\mathcal{L}}{d\lambda d\alpha} &= \frac{d\mathcal{L}_\alpha}{d\lambda} = -C_{L\alpha}, \end{aligned}$$

and our Newton iteration driving λ becomes:

$$\lambda^{new} = \lambda^{old} + \mathcal{L}_\alpha / C_{L\alpha}. \quad (19)$$

Interestingly, the drivers of both Eqns (17 & 19) utilize $(\pm) C_{L\alpha}$ as the slope for their Newton solves. Figure 42 provides the convergence histories of λ and \mathcal{L}_α for the NACA0012 airfoil. Here, λ is initialized to zero and first updated at 4000 iterations, then every 2000 iterations after that. The blue line represents λ (left axis), while the red line depicts \mathcal{L}_α (right axis). Since the results are converged at 4000 iterations for $\lambda = 0$, we note that $\frac{dC_D}{d\alpha} \approx 1.052/rad \approx 0.0184/^\circ$.

Before leaving the discussion on these drivers, we note that the process illustrated herein with updates at 4000/2000 iterations is far from the most efficient option. It was strictly utilized here to cleanly show the cause-and-effect of these step changes. A much more efficient technique would be to update α and λ every 1000 iterations, achieving fully converged results well within 5000 total iterations. In practice, converged solutions based on these drivers can be as efficient as solutions with fixed α or λ .

8.3 Gradient Comparisons

The SU2 simulations use a (1025x1025) conformal mesh of Vassberg and Jameson [24]. This structured grid is converted to an unstructured mesh, purely comprised of quadrilateral elements for the SU2 simulations. As this level of resolution is far too dense to show, instead, a coarsened version of this mesh with dimensions of (257x257) is provided in Figure 43.

The design space for this example problem utilizes α and 19 evenly-spaced Henne-Hicks bump functions for both the lower and upper surfaces, respectively, and ordered from leading-to-trailing edges. This gives a total of 39 design variables, ranging in index from 0 to 38. Here δ_0 is α , while δ_1 - δ_{19} are bump amplitudes on the lower surface, and δ_{20} - δ_{38} are bump amplitudes on the upper surface. Per SU2 convention, a positive amplitude pushes the H-H bump outward of the flowfield domain, or into the airfoil geometry. Remember this when viewing the various gradients of SU2.

Continuous adjoint gradients for drag at fixed- α and fixed- C_L have been incorporated into our version of SU2. Discrete adjoint gradient for drag at fixed- α is also available. To help verify that the adjoints give accurate gradients, a python script is utilized which cycles through the set of design variables to approximate the sensitivities with forward-finite-differences; this is only possible because this case is a 2D problem with a minimal number of design variables.

Figure 44 provides gradients of the various forms for comparison. In this figure, black lines/symbols depict the gradient of C_D at fixed- α , while the red lines/symbols provide the gradient at fixed- C_L . The solid lines represent continuous adjoint gradients, while the open symbols depict discrete or finite-differenced gradients. As we noted earlier, there is a substantial difference between the fixed- α and fixed- C_L gradients. This difference represents the "hidden" portion of the fixed- α gradient that affects lift to alter drag. Closer inspection of these gradients reveals that the fixed- C_L gradient works to add aft camber to the airfoil, while the fixed- α gradient attempts to remove it. This is consistent with the aforementioned statement that the fixed- α gradient will attempt to dump lift in order to reduce drag.

Next, we will show how the performance of an optimization to minimize drag is improved by using the fixed- C_L gradient.

8.4 Optimization Results

Figures 45-52 provide a complete set of results of the NACA0012 airfoil optimizations. Figure 45 illustrates a comparison of airfoil sections. Here, the NACA0012 baseline is depicted with a red line, the optimum airfoil based on the fixed- C_L approach is shown as a blue line, and the optimum based on fixed- α is given as a green line. Notice that the camber level of the Opt.Fixed- C_L airfoil is greater than that of the Opt.Fixed- α geometry.

Figure 46 provides a comparison of the pressure distributions, using the same color code as before. Notice that the leading-edge pressure peak of the Opt.Fixed- C_L airfoil is lower than that of the Opt.Fixed- α geometry. This implies that the angle-of-attack, α , of Opt.Fixed- C_L is less than that of Opt.Fixed- α . Indeed, the angle-of-attacks for these optimum airfoils are $\alpha = 0.678^\circ$ and 1.25° , respectively.

From these figures, one may ask why the two optimum airfoils are noticeably different from each other, both in geometric shape as well as in pressure distributions. The answer is that their respective trajectories through the design space, traveling from the baseline to their end states, take different paths. This occurs

because the gradients of drag are distinctly different, as previously shown in Figure 44. Different trajectories can lock into different local minima. Here, however, we specifically setup this test case to allow for an infinite number of *zero-drag* possibilities. The baseline NACA0012 airfoil, at its design point, has a drag coefficient of $C_D = 0.02233$, or 223.3 counts. Since there are no constraints on airfoil thickness, the optimization process keeps thinning the airfoil sections during the design cycles until no shock-drag is produced. Once the optimizations achieve a shock-free flow, they have nowhere better to go, and realize they have converged.

Referring again to Figure 46, one can verify that both optimums have no shocks. When this state is reached, for an inviscid 2D flow, all physical drag vanishes. Numerically, however, the discrete drag reaches some small value which is indicative of the error of discretization. For this problem, on this mesh, and for shock-free flows, the minimum discrete drag coefficient appears to be about 0.00007, or 0.7 counts.

Figures 47-48 provide colored isobars of field pressures about the NACA0012 and the Opt.Fixed- C_L airfoil sections, respectively. (For reference, at $M = 0.8$ the critical pressure coefficient is $C_P^* = -0.43464$.) Note the strong shock on the upper surface of the NACA0012 and the weak shock on its lower surface. Now observe the shock-free flow of the Opt.Fixed- C_L geometry. Here, the pocket of supersonic flow over the upper surface exhibits an isentropic recompression as it travels downstream. The overall flow features of the Opt.Fixed- α airfoil are generically similar that those of the Opt.Fixed- C_L geometry.

Figures 49-50 provide the histories of drag coefficients during the optimizations. The blue solid line corresponds to the fixed- C_L optimization, while the green chain-dot line represents the fixed- α history. This plot is provided on a semi-log scale as the level of drag is reduced by over two-and-a-half orders-of-magnitude from start to finish. Notice that the fixed- α optimization lags the fixed- C_L optimization. The fixed- C_L optimization is fully converged in 13 design cycles, whereas the fixed- α optimization requires 16 design cycles. The performance benefit is actually more substantial than this, as the fixed- α optimization requires two adjoints (C_D and C_L) to be solved per design cycle, while the fixed- C_L optimization only requires one (\mathcal{L}). Since the computational cost to evaluate a flow solution is essentially the same as that of an adjoint solution, the total cost of an optimization is directly proportional to the number of evaluations required. Figure 50 illustrates that the fixed- C_L optimization requires 32 evaluations for convergence, while the fixed- α optimization demands 52 solutions. This confirms the anticipated 2:3 ratio in optimization costs.

In order to better understand why the fixed- α optimization lags behind, let's review the C_L histories provided by Figures 52. Notice that the history of lift for the fixed- α optimization clearly indicates that, initially, it is dumping lift in an attempt to reduce drag. This trend continues for the first 4 design cycles and 13 function evaluations. At this point, the optimizer has gathered enough information to know that it has to compensate for missing the constraint on lift, and in fact, it over compensates before beginning to lock in on the target value. Even after 10 design cycles and 34 function evaluations, the lift coefficient of the fixed- α optimization is only semi-converged to the constraint value of $C_L^* \approx 0.347$. On the other hand, the fixed- C_L optimization starts out and remains at the constraint value for lift throughout.

The fixed- α optimization presented herein was done so without α being a design variable. We have conducted an analogous optimization which includes α as a design variable, and the results exhibit substantially the same trends. The only notable difference is that the C_L history is able to better lock onto the C_L^* constraint after 10 design cycles. This is enabled with the α design variable, and the final angle-of-attack for this optimum is $\alpha \approx 1.16^\circ$.

This concludes our discussion on the NACA0012 test case and the alternative approach to conducting lift-constrained optimizations. The next section continues to illustrate the differences between the fixed- α and fixed- C_L gradients, however, on a much more industrial-strength problem.

9 NASA CRM Wing/Body

In this section we showcase SU2's ability to handle large-scale engineering problems of interest. Although we do not perform yet another optimization, we continue our discussion on the differences between fixed- α and fixed- C_L gradients. The example of this section is taken from the Sixth AIAA Drag Prediction Workshop (DPW-VI) Test Cases. For more information about DPW-VI, as well as to gain access to the CRM geometry and various grids, please go to the DPW-VI website [25].

The concept of the NASA Common Research Model (CRM) was conceived by Vassberg and Wahls [26]. The CRM is representative of a contemporary transonic transport aircraft design. The high-speed cruise configuration includes fuselage, wing, nacelle, pylon, and horizontal-tail components. Figure 53 provides an image of the CRM wing/body configuration. Note the detail of the fuselage geometry at the cockpit region,

as well as with the wing-body fairing. Also notice the side-of-body treatment of the after-body. This is to ensure that the horizontal-tail remains sealed against the fuselage over its full range of incidence motion. The CRM was originally intended to be used primarily for CFD validation efforts, but it has grown far beyond that in utility. In addition to NASA's CRM wind-tunnel model, other organizations around the world have built their own CRMs, including JAXA, ONERA, and the University of Washington. It is being used as the basis for icing studies on swept wings, as well as for the NASA Juncture Flow Model (JFM), currently under development. Boeing has been recently supplied a low-speed, high-lift geometry definition to NASA; this should become available to the public domain later this year. For additional information on the NASA Common Research Model, including freely available wind-tunnel data and structural models, please refer to the CRM website [27].

Available from the DPW-VI website are complete sets of grid families for a number of grid types. Each grid family is populated with grids that range in size by an order-of-magnitude. These families have been carefully constructed to be appropriate for grid-resolution studies. In particular, we downloaded the *Tiny* unstructured mesh of merged elements for the CRM Wing/Body (WB AE2.75) configuration as provided by NASA's GeoLab group. The designation AE2.75 indicates that this grid is about the wing with an aero-elastic deflection under the loads produced at the $\alpha = 2.75^\circ$ flow condition, which closely corresponds to the CRM's design point of $M = 0.8$, $C_L = 0.5$, and $Ren = 5$ million. This unstructured grid (ironically named *Tiny*) is comprised of 20,472,098 field nodes, and 83,578,942 field elements. The field elements consist of 65,402,114 tetrahedra, 119,340 pyramids, and 18,057,488 triangular prisms. The flowfield boundary is defined with 1,242,346 triangles and 34,222 quadrilaterals, hence, 655,395 nodes reside on the boundary. This grid is sufficiently large to demonstrate our point that SU2 is capable of industrial-strength applications.

Figure 54 provides the surface pressures for the CRM WB configuration at $M = 0.8$, $C_L = 0.5$, and $Ren = 5$ million. In this figure, as well as those to follow, the upper surface of the CRM is shown at the top of the image, while its lower surface is shown at the bottom of the image. Supersonic flow is depicted with yellow-to-red coloring, while stagnation flow is shown as dark blue. Note the λ shock structure on the wing upper surface. Also notice how the flow tends towards stagnation at the fuselage nose, cockpit windshields, extremes of the wing-body fairing, and at the leading and trailing edges of the wing's lower surface.

Figures 55-56 provide a comparison of the sensitivities of drag, C_D , with geometry change for the fixed- α and fixed- C_L gradients, respectively. As shown, these gradients represent more than half-a-million design variables. Recall SU2's convention regarding gradients. Drag increases when you push the geometry inward at positive sensitivities (green-yellow-red) and pull the geometry outward at negative sensitivities (green-aqua-blue). At first glance, these gradients look very similar, and in fact, they do share many similarities. However, upon closer inspection, one will notice some differences too. Let's begin with the upper surface. Here, the fixed- α gradient is generally more negative (darker blue) than that of the fixed- C_L . The sensitivity at the shock is clearly more positive for the fixed- C_L gradient and the λ structure is more evident. Now, let's inspect the lower surface. With focus near the trailing-edge, notice that the fixed- α gradient is positive (yellow-red), while that of the fixed- C_L gradient is negative (aqua-blue). The combination of all of the aforementioned differences indicate that, in order to reduce drag, the fixed- α gradient is trying to de-camber the wing more so than does the fixed- C_L gradient. Hence, the fixed- α gradient dumps lift to reduce drag. This is consistent with what we learned in the NACA0012 example.

Figure 57 provides the sensitivity of lift, C_L , with geometry change at fixed α . Generally speaking, the upper surface is predominately negative (aqua-blue), while the lower surface is substantially positive (yellow-red). This confirms that adding camber to the wing increases lift, or conversely, de-cambering it dumps lift at fixed α .

Figure 58 provides the sensitivity of pitching moment, C_{My} , with geometry change at fixed α . We note that SU2 adheres to the convention that a positive C_{My} represents a nose-up pitching moment. Generally speaking, the upper surface is predominately positive (yellow-red), while the lower surface is substantially negative (aqua-blue). This is especially true for the wing outboard of the planform break, where the center of sectional lifts are behind the moment reference center. The reason the C_{My} and C_L sensitivities appear almost opposite of each other is that any increase in lift behind the moment reference center results in a decrease in pitching moment.

This concludes our discussions for this Lecture.

10 Conclusions

In this lecture, sample applications of aerodynamic shape optimizations are provided through five case studies in aircraft design. These include the design of a Mars aircraft, a Reno Racer, a generic 747-class wing/body configuration, an NACA0012 airfoil, and the NASA Common Research Model. The generic wing/body case included a simple structure-weight model to allow aero-structural optimizations to be performed.

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, SYN107P will currently reproduce the present Navier-Stokes optimization on a current MacBookPro notebook computer within an hour.) 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.

Aerodynamic shape optimization methods were successfully applied to the design of a new unlimited class Reno race plane. Very significant performance gains were achieved in very compressed time. Utilization of this software also allowed global changes to occur at the aircraft level without adversely affecting the efforts to aerodynamically design its high-performance, transonic wing as new designs could be performed over night. Normally, such major changes would have had a very disruptive effect on the design of the wing. Yet, the evolution of the general layout was necessary for all of the design goals to be achieved.

Coupled aerodynamic-structural optimizations were performed on a generic 747-class wing/body configuration. In this case study, a simple structure-weight model was developed and incorporated into the aerodynamic shape optimization process. The cost function in this investigation is a blending of drag and structure-weight coefficients. The coefficients of blending were determined to maximize range, as well as were allowed to vary to provide a pareto front of optimum designs.

Application of the open-source SU2 software suite has been demonstrated on the NACA0012 airfoil and the CRM wing/body configuration. Enhancements of SU2 include an automatically-generated discrete adjoint based on CodiPack, and an alternative fixed- C_L drag gradient based on a Lagrangian dual cost function for lift-constrained optimizations.

All 3D optimizations utilized thousands of design variables and were carried out on affordable computers systems. These were key to the success of this simulation-based design effort.

Exercises such as these have been very beneficial to the authors, as we get a better understanding and appreciation of the chaos and schedule pressures that exist in real-world design environments. We are also pleased that the aerodynamic shape optimization software have had dramatic and positive effects on the outcome of the final designs.

11 Post Script

As can be seen from these studies, aerodynamic shape optimization can significantly streamline the design process. In addition to the case studies presented herein, the present method has now been successfully applied in a variety of projects, including the McDonnell-Douglas MDXX [28], the NASA HSCT studies, the Boeing Blended-Wing-Body project, the Beech Premier [29], the original design of the NASA Common Research Model [26], and many more recent and on-going aircraft development efforts.

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 accomplish 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.

References

- [1] J. C. Vassberg and R. D. Gregg. Overview of aerodynamic design for transport aircraft. *Presentation*, First M.I.T. Conference on Computational Fluid and Solid Mechanics, Cambridge, MA, June 2001.
- [2] R. H. Liebeck. Design of the Blended-Wing-Body subsonic transport. *Wright Brothers Lecture, AIAA paper 2002-0002*, Reno, NV, January 2002.
- [3] R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15:407–412, 1978.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] A. Jameson. Optimum aerodynamic design using control theory. *Computational Fluid Dynamics Review*, pages 495–528, 1995.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.

- [13] J. C. Vassberg, G. S. Page, R. J. Foch, and A. Jameson. Aerodynamic design and optimization of the mars airborne remote exploration scout. *AIAA Paper 2004-0401*, 42nd AIAA Aerospace Sciences Meeting and Exhibition, Reno, NV, January 2004.
- [14] A. Jameson. Transonic potential flow calculations using conservative form. In *Proceedings of AIAA 2nd Computational Fluid Dynamics Conference*, pages 148–161, June 1975.
- [15] P. A. Henne and R. M. Hicks. Wing analysis using a transonic potential flow computational method. *NASA-TM 78464*, July 1978.
- [16] 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.
- [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.
- [18] E. Ahlstrom, R. Gregg, J. Vassberg, and A. Jameson. G-Force: The design of an unlimited class Reno racer. *AIAA paper 2000-4341*, 18th AIAA Applied Aerodynamics Conference, Denver, CO, August 2000.
- [19] K. Leoviriyakit. Wing planform optimization via an adjoint method. *Stanford University Thesis*, February 2006.
- [20] S. Nadarajah. Drag minimization of the NACA0012 in transonic inviscid flow. ADO DG Sharepoint file: *twodimensional.pdf*, July 2013.
- [21] F. Palacios, T. D. Economon, A. D. Wendorff, and J. J. Alonso. Large-scale aircraft design using SU2. *AIAA Paper 2015-1946*, 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL, January 2015.
- [22] T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso. SU2: An open-source suite for multiphysics simulation and design. *AIAA Journal*, 54(3):828–846, March-April 2016.
- [23] T. Albring, B. Zhou, and N. Gauger. An aerodynamic design framework based on algorithmic differentiation. *Bulletin*, ERCOFTAC, 2015.
- [24] J. C. Vassberg and A. Jameson. In pursuit of grid convergence for Two-dimensional Euler solutions. *AIAA Journal of Aircraft*, 47(4):1152–1166, July-August 2010.
- [25] 6th AIAA CFD Drag Prediction Workshop. <http://www.aiaa-dpw.larc.nasa.gov>, June 2016.
- [26] J. C. Vassberg, M. A. DeHaan, S. M. Rivers, and R. A. Wahls. Development of a Common Research Model for applied CFD validation studies. *AIAA Paper 2008-6919*, 26th AIAA Applied Aerodynamics Conference, Honolulu, HI, August 2008.
- [27] NASA Common Research Model. <http://www.commonresearchmodel.larc.nasa.gov>, 2016.
- [28] A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamic design using the Navier-Stokes equations. *Theoret. Comput. Fluid Dynamics*, 10:213–237, 1998.
- [29] 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.




Figure 1: CFD simulation of a commercial transport with aft-mounted engines.




Figure 2: CFD simulation of a commercial transport with wing-mounted engines.




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.

MARES Wing Design - Pressure Distributions
 SYN107P Drag Minimization
 REN = 170K , MACH = 0.650




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

MARES Wing Design - Upper Surface Isobars
 SYN107P Drag Minimization
 REN = 170K , MACH = 0.650
 (Contours at 0.05 Cp)




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




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



Figure 19: Miss Ashley II and Rare Bear en Route.

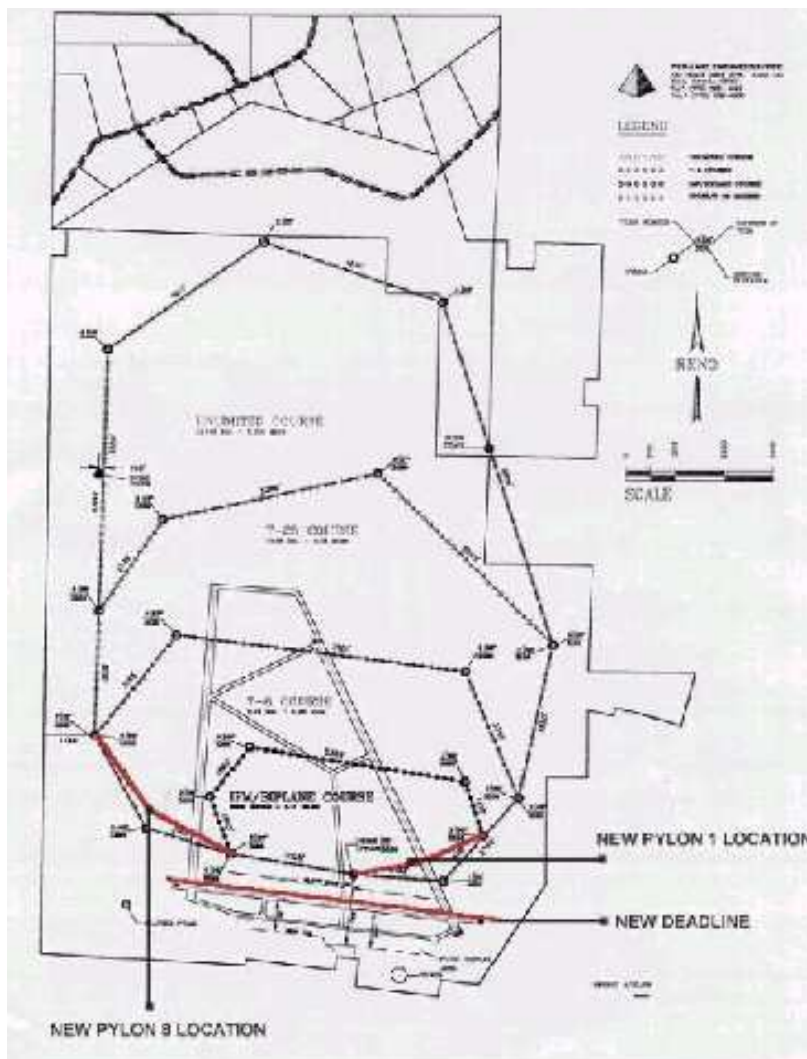


Figure 20: Reno Race Course Layout.




Figure 21: Side View of Body-Prop Design.




Figure 22: Rendering of Body-Prop Design in Flight.




Figure 23: Pressure Distributions of Shark1 Baseline Wing.




Figure 24: Comparison of Shark5 Wing on Baseline Fuselage with Baseline Configuration.




Figure 25: Comparison of Shark52 Wing on Stretched Fuselage and Baseline Configuration.




Figure 26: Comparison of Shark52 and Shark1 Wing Drag Loops.




Figure 27: Comparison of Shark52 and Shark1 Wing Pressure Contours.




Figure 28: Result of Navier-Stokes Inverse Design.

Final Shark Wing

Airfoil Geometry -- Camber & Thickness Distributions




Figure 29: Final Wing Airfoil Geometry - Thickness & Camber Plots.




Figure 30: Structural Model for a Swept Wing.

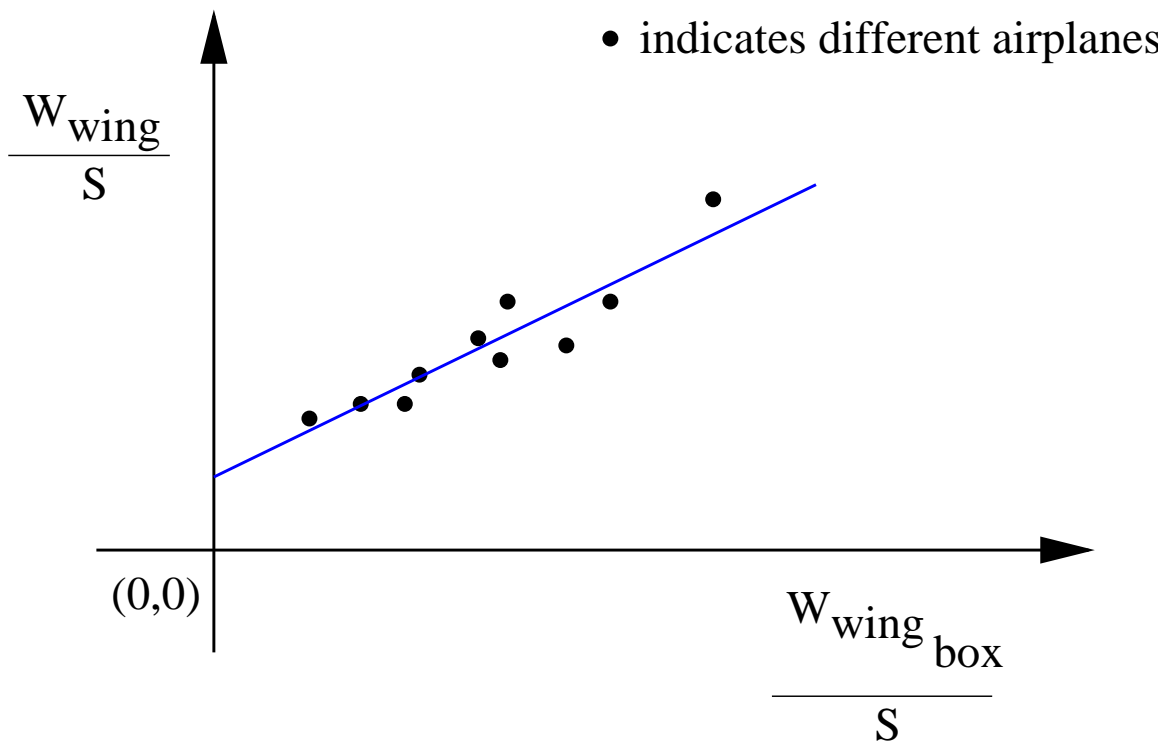


Figure 31: Statistical Correlation of Total Wing Weight and Box Weight.





Figure 32: Wing Planform Design Variables.



Baseline (Dashed) / Redesign (Solid).

Figure 33: Wing-Section Optimization of Generic 747 at Fixed Baseline-Planform.

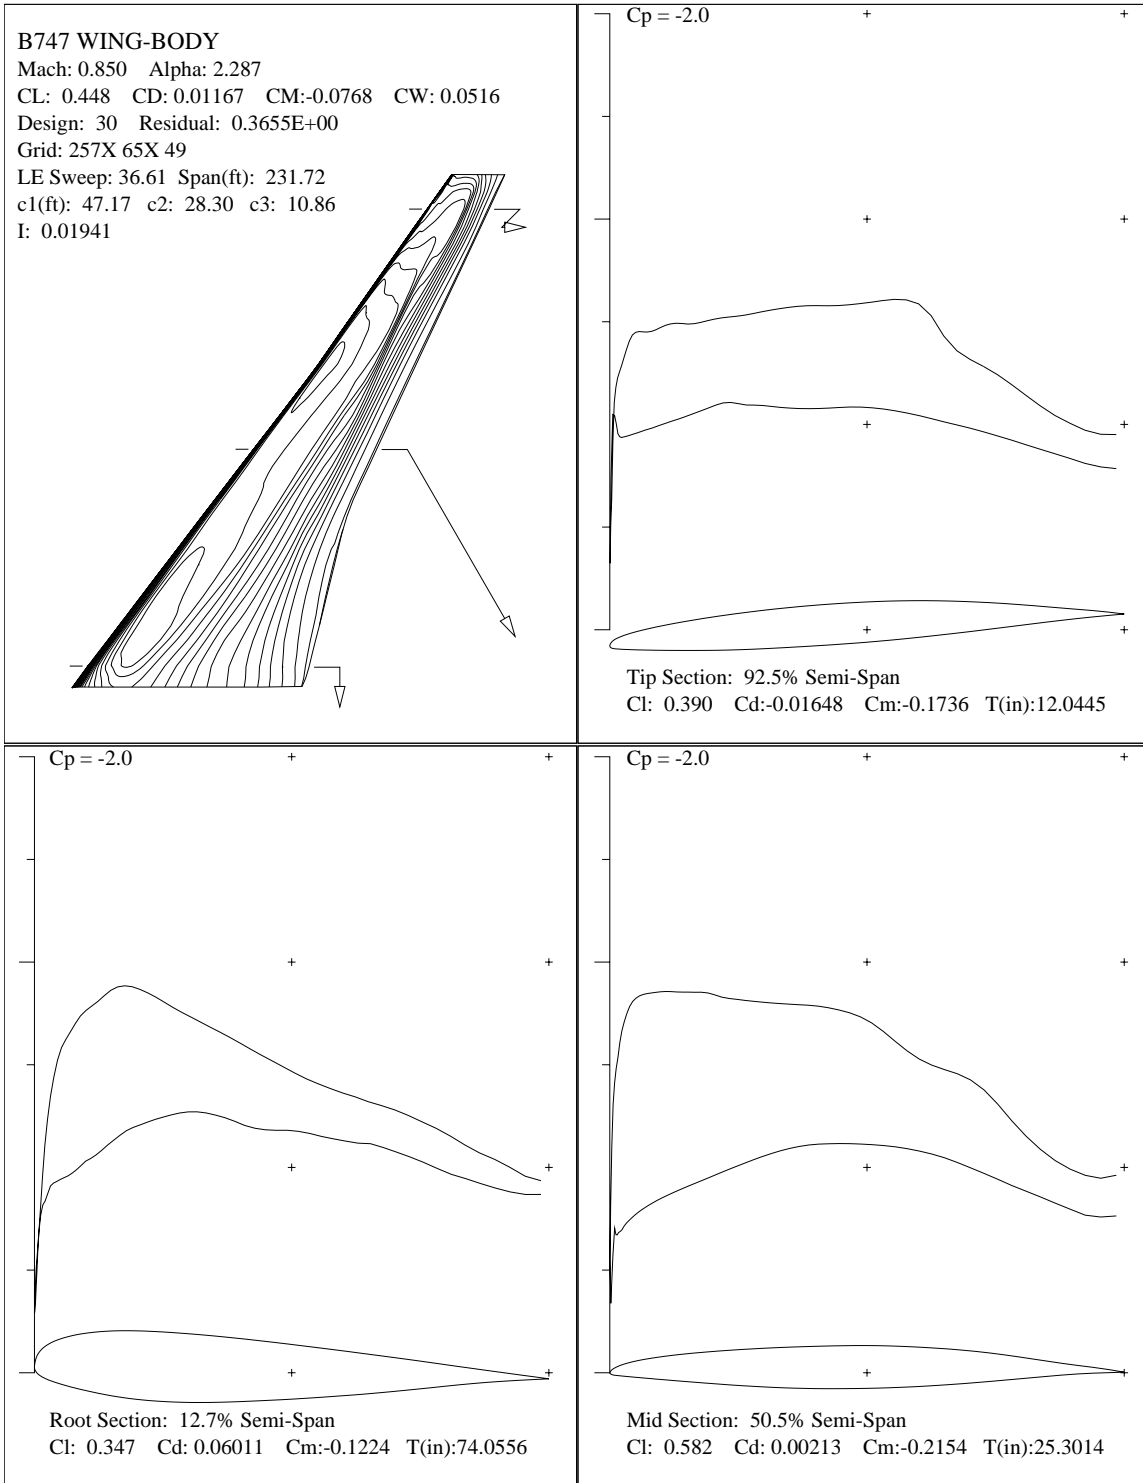


Figure 34: Complete Optimization of Generic 747 to Maximize Breguet Range.

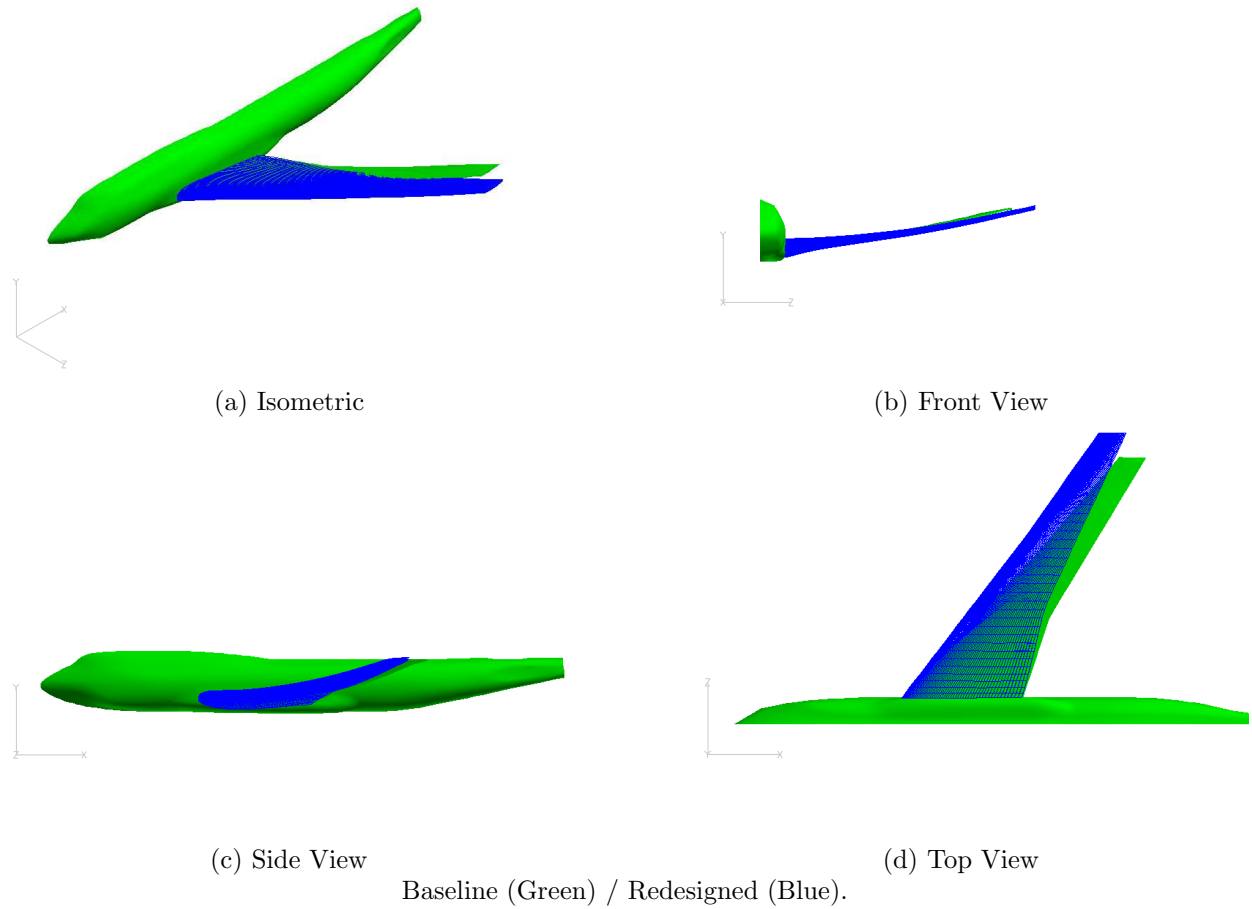


Figure 35: Geometry Changes of Complete Navier-Stokes Optimization.




Figure 36: Comparison of Euler-Redesigned (Red) and NS-Redesigned (Blue) Planforms.




Figure 37: Cooperative Game Strategy with Drag and Weight as Players.




Figure 38: Pareto Front; Ratios of $\frac{\alpha_2}{\alpha_1}$ Indicated.




Figure 39: NACA0012 Lift Curve; $M = 0.8$, $\alpha^* = 1.25^\circ$, $C_L^* \approx 0.347$.




Figure 40: NACA0012 Drag Polar; $M = 0.8$, $\alpha^* = 1.25^\circ$, $C_L^* \approx 0.347$, $C_D^* \approx 0.02233$.

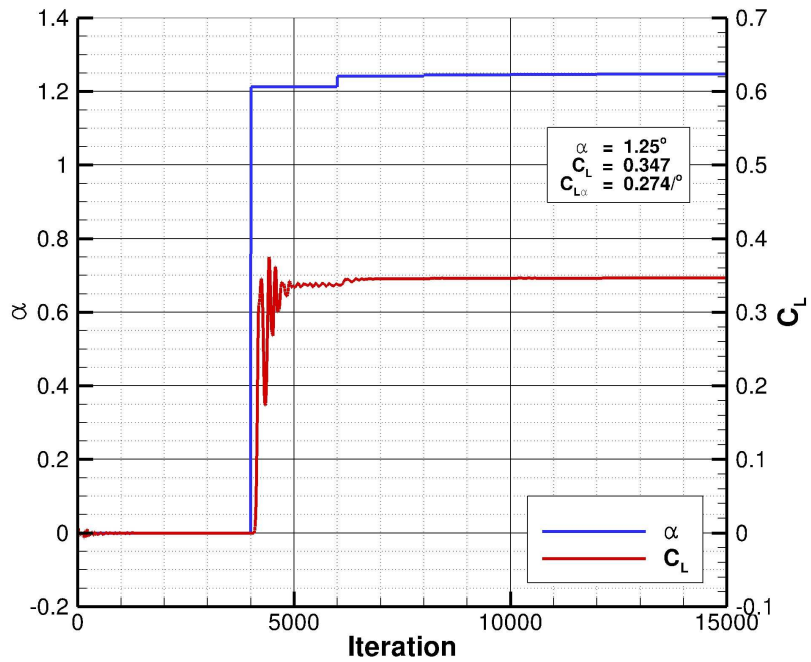


Figure 41: NACA0012 Alpha Convergence; $M = 0.8$, $C_L^* \approx 0.347$.




Figure 42: NACA0012 Lambda Convergence; $M = 0.8$, $\alpha = 1.25^\circ$, $C_L \approx 0.347$.




Figure 43: NACA0012 Airfoil (256x256) Conformal Mesh.




Figure 44: Comparison of NACA0012 Gradients; $M = 0.8$, $\alpha = 1.25^\circ$, $C_L \approx 0.347$, $C_D^* \approx 0.02233$.




Figure 45: Comparison of NACA0012 and Optimum Airfoil Sections; $M = 0.8$, $C_L \approx 0.347$.

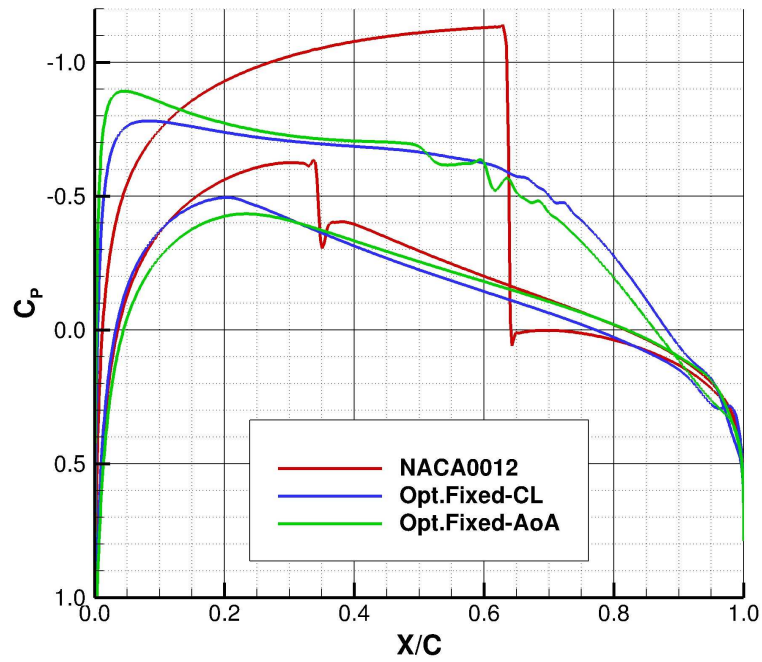


Figure 46: Comparison of NACA0012 and Optimum Airfoil Pressure Distributions; $M = 0.8$, $C_L \approx 0.347$.




Figure 47: NACA0012 Flow Field Pressures; $M = 0.8$, $C_L \approx 0.347$, $C_D^* \approx 0.02233$.




Figure 48: Opt.Fixed- C_L Flow Field Pressures; $M = 0.8$, $C_L \approx 0.347$, $C_D^* \approx 0.00007$.




Figure 49: NACA0012-to-Optimum Drag History; $M = 0.8$, $C_L^* \approx 0.347$, $C_D : 0.02233 \rightarrow 0.00007$.




Figure 50: NACA0012-to-Optimum Drag History; $M = 0.8$, $C_L^* \approx 0.347$, $C_D : 0.02233 \rightarrow 0.00007$.

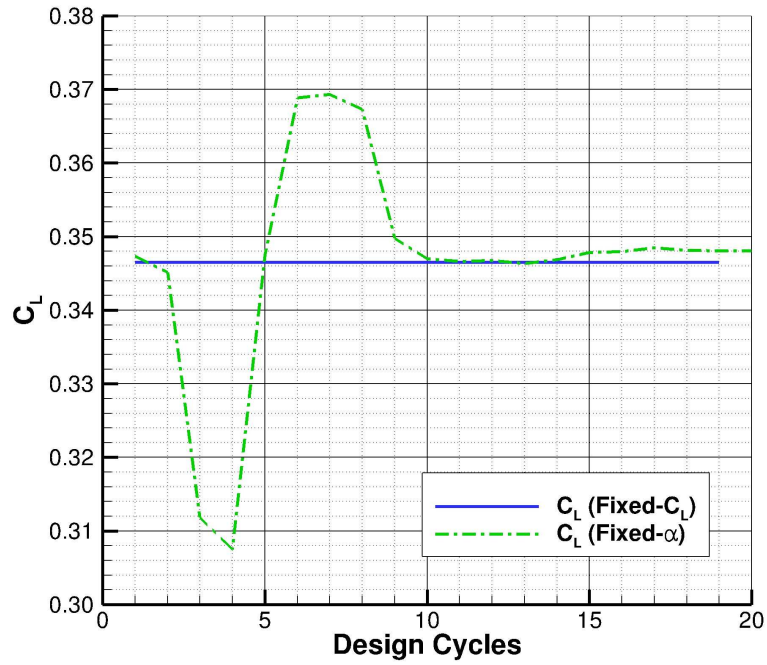


Figure 51: NACA0012-to-Optimum Lift History; $M = 0.8$, $C_L^* \approx 0.347$, $C_D : 0.02233 \rightarrow 0.00007$.




Figure 52: NACA0012-to-Optimum Lift History; $M = 0.8$, $C_L^* \approx 0.347$, $C_D : 0.02233 \rightarrow 0.00007$.




Figure 53: NASA Common Research Model (CRM) Wing/Body Configuration.




Figure 54: Common Research Model Surface Pressures; $M = 0.85$, $C_L = 0.5$, $Re_n = 5$ million.




Figure 55: CRM WB C_D Sensitivities at Fixed α ; $M = 0.85$, $C_L = 0.5$, $Ren = 5$ million.




Figure 56: CRM WB C_D Sensitivities at Fixed C_L ; $M = 0.85$, $C_L = 0.5$, $Ren = 5$ million.




Figure 57: CRM WB C_L Sensitivities; $M = 0.85$, $C_L = 0.5$, $Re_n = 5$ million.




Figure 58: CRM WB C_M Sensitivities; $M = 0.85$, $C_L = 0.5$, $Re_n = 5$ million.