AERODYNAMIC SHAPE OPTIMIZATION OF A RENO RACE PLANE

John C. Vassberg and Antony Jameson

International Journal of Vehicle Design
Volume 28, No. 4, pp 318-338, 2002

OptiCon 2000

October 26-27, 2000

Newport Beach, CA

Aerodynamic Shape Optimization of a Reno Race Plane

John C. Vassberg *
Hydro-Aero Consulting Group
Long Beach, CA 90803, USA

Antony Jameson †
Intelligent Aerodynamics, Int'l.
Menlo Park, CA 94025, USA

26 – 27 October, 2000

Abstract

This paper chronicles the applications of aerodynamic shape optimization software during the design of a concept race plane for the annual Reno air races. The development of this aircraft began from the ground up, being an all-new design. Every major element of the airplane was 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. 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 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 under development by the authors allowed various aircraft subsystems to be routinely modified without adversely impacting development costs or schedule as the new wing designs literally occurred over night. The complete evolution of the aircraft's general layout was accomplished in a very compressed time frame. More importantly, this evolution was required in order for all of the design goals to be met.

A critical point regarding the acceptance of the authors' software within this design environment was its minimal requirements on computational resources and set-up time. Without these efficiencies, our involvement would not have played a pivotal role in the design effort, and most likely, would not have been solicited in the first place.

*Associate Technical Fellow, The Boeing Company
†T. V. Jones Prof. of Engineering, Stanford University

Nomenclature

A	Hessian Matrix / Operator
AR	Wing Aspect Ratio $= \frac{b^2}{S_{ref}} = 8.3$
AMD	Advanced Micro Devices
b	Wing Span $= 25ft$
CFD	Computational Fluid Dynamics
CD	Drag Coefficient $= \frac{Drag}{q_\infty S_{ref}}$
CL	Lift Coefficient $= \frac{Lift}{q_\infty S_{ref}}$
C_{ref}	Wing Reference Chord $= 3ft$
$count$	Drag Coefficient Unit $= 0.0001$
$\mathcal{F}$	Surface Defining Function
$FLO22$	Full-Potential Wing CFD Analysis Code
FR	Fuselage Fineness Ratio $= 5.7$
G	Acceleration Unit = Gravity $\simeq 32.2\frac{ft}{sec^2}$
$\mathcal{G}$	Gradient of Cost Function
H	Estimate of Inverse Hessian Matrix
HP	Horse Power
I	Objective or Cost Function
K	Signifying $1,000$
$KEAS$	Knots Equivalent Air Speed
M	Signifying $1,000,000$
MD	Multi-Disciplinary
MPH	Mile Per Hour
MPI	Message-Passing Interface
N	Number of Design Variables
Re	Wing Reynolds number based on C_{ref}
Re_θ	Attachment Line Reynolds number
ROI	Return On Investment
S_{ref}	Wing Reference Area $= 75ft^2$
SYN	Synthesis Code w/ Analysis & Design Modes
TAS	True Air Speed
x	Independent Spatial Variable
y	Design Variable
q	Dynamic Pressure $= \frac{1}{2}\rho V^2$
$Yehudi$	Inboard Wing Trailing-Edge Extension
λ	Wing Taper Ratio $= 0.45$
$\Lambda_{c/4}$	Wing Quarter-Chord Sweep $= 28°$
π	3.141592654...
∞	Infinity
$\delta*$	First Variation of
$\mathcal{O}(*)$	Order of
$(*)^{-1}$	Inverse Matrix of

1 Introduction

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 Furys, with power plants that produce well in excess of 1,000 HP. Miss Ashley II and Rare Bear of Figure 1 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 2. 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.

2 Optimization Overview

The performance of a complex engineering system such as an aircraft generally depends on how well a group of designers can trade the conflicting set of requirements imposed by a multitude of disciplines. Put in terms of mathematics, a system design can be defined by a large set of parameters, and the task at hand is to choose an optimal combination of these. It is almost impossible to find the best choice simply by trial and error, and this motivates the use of numerical optimization techniques.

2.1 Gradient

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 $y(x)$ that describes the design, where under a variation of the design, $\delta y(x)$, the variation of the cost is δI. Now suppose that δI can be expressed to first order as

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

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

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

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.

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. This is the case with aerodynamic simulations based on Euler or Navier-Stokes equations. For example, if one resorts to finite differences [1], direct code differentiation (ADIFOR [2, 3]), or complex-variable perturbations [4], the cost of determining the gradient is directly proportional to the number of variables used to define the design. Even small problems of aerodynamic shape optimization based on this approach 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. Nevertheless, these approaches are currently standard practice throughout the aerodynamic shape optimization community. A more cost effective technique is to compute the gradient through the solution of an adjoint problem, such as that developed by the authors [5, 6, 7].

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 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 equation above. Suppose that the governing equation R, which expresses the dependence of w and $\mathcal{F}$ within the flowfield domain D, can be written as

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

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

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

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.

Here, the cost of obtaining the gradient is comparable to the cost of a single function evaluation, regardless of the dimension of the design space. The downside of this approach is that it can take man-months to develop an adjoint code for a given cost function. However, there is on-going research at Rice University to develop ADJIFOR [8] which automatically generates a discrete adjoint code from existing analysis software. While this holds promise, the resulting adjoint solvers currently take at least an order-of-magnitude more CPU time than their corresponding function evaluations. Further, the resulting adjoint codes are not 100% automatically generated, requiring the intervention of a knowledgeable programmer to finish the code construction. Nevertheless, a general guideline is that if an optimization problem has more that 25 design variables, ADJIFOR-generated gradients are more cost effective that those computed with ADIFOR. For all practical purposes, this is the case for every aerodynamic shape optimization of any substance. In the present work, the adjoint equations have been derived analytically, then approximated in discrete form.

2.2 Search

The remaining cost issue is related to finding a location in the design space where the gradient vanishes, or rather, finding 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 this search depends on the number of steps it takes to find a local minimum as well as the cost of each step. For the sake of discussion, we assume that a gradient is known.

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, $\mathcal{G}(y) = 0$, and is solved by the standard Newton iteration for nonlinear equations.

Suppose the Hessian is denoted by

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

then the result of a step δy may be linearized as

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

This is set to zero for a Newton step; therefore

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

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. In general, the cost of a quasi-Newton search scales with the dimension of the design space, thus leading to extremely large computational requirements for most engineering problems of interest.

2.3 Computational Costs

To address the issues of the search costs, the authors investigated several techniques in Reference [9] which yielded a dimension-independent convergence rate. These techniques include preconditioning a steepest-descent search direction with an implicit smoothing operator, a local step scaling set by stability, and multigrid; then 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 steps, even for problems with thousands of design variables. 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 $N > 5000$ for most of the aerodynamic shape optimization we encounter.

With the flexibility of not being constrained by the number of design variables that one can use, this has an added benefit that the optimization software can be written in such a manner that the end user is not burdened with the task of defining a set of shape functions. This is an important advance, as these user-specified shape functions are usually of the bump-function class and 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 the user to iterate with trial and error. This 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. It is ironic for systems, that boast automation, to depend so heavily on a man-in-the-loop at the finest-grain level of problem set-up. It is the authors' position that optimization software should free the engineer from mundane tasks and allow him to focus on the more global requirements of the system development. With that stated, we acknowledge that our software has not reached this level of usability; yet there exists a path to get there. Methods that require the user to input large amounts of detailed data by hand have very little chance of ever realizing this goal. For this reason, we have abandoned further development of optimization software that require user-specified design spaces, as well as trajectories based on quasi-Newton search engines. Our current focus is related to developing hierarchical rules that allow the end user to constrain the design space with minimal and intuitive input.

3 Design Overview

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 3 and a computer graphics rendering of this configuration in flight is given in Figure 4. 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 Reference [10].

4 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.

4.1 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

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

the off-design capabilities for buffet

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

the drag divergence Mach number

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

the clean wing maximum lift coefficient

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

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 $40 - 60\ lbs./ft^2$. 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 wingtip extension for a growth airplane. Conversely, a production break was included at 87% semi-span to allow a $4ft^2$ 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.

4.2 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 it's 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.

4.3 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 5. SYN88 is a wing/body Euler method which also incorporates an adjoint-based optimization procedure for aerodynamic shape design. In Figure 5, 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 onset Mach number was used. Referring to the wing-planform plot of Figure 5, 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 6. 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 7-8.

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 8 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.

4.4 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 10 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.

4.5 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 11 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.

4.6 Phase VI: Test Evaluation

Phase VI is on-going. The construction of this aircraft is to begin later this year, and a flight test will follow. Prior to first flight, wind-tunnel data will be collected on the final wing/body configuration. Fabrication of the wind-tunnel model has begun.

5 Conclusions

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 aerodynamicly 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.

All 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 this 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 design.

Acknowledgment

The aerodynamic design effort was partially funded by Renaissance Research. The first author would like to thank The Boeing Company for allowing his participation in this effort after hours.

References

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

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

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

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

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

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

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

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

[9] A. Jameson and J. C. Vassberg. Studies of alternative numerical optimization methods applied to the brachistochrone problem. In *Proceedings of OptiCON'99*, Newport Beach, CA, October 1999.

[10] 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.

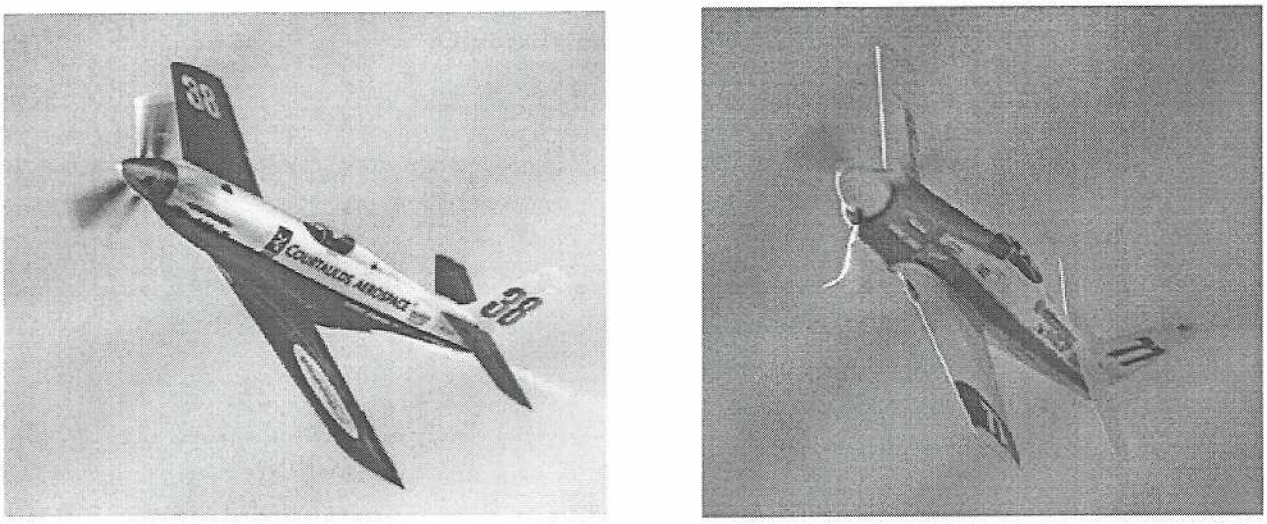

Figure 1: Miss Ashley II and Rare Bear en Route.

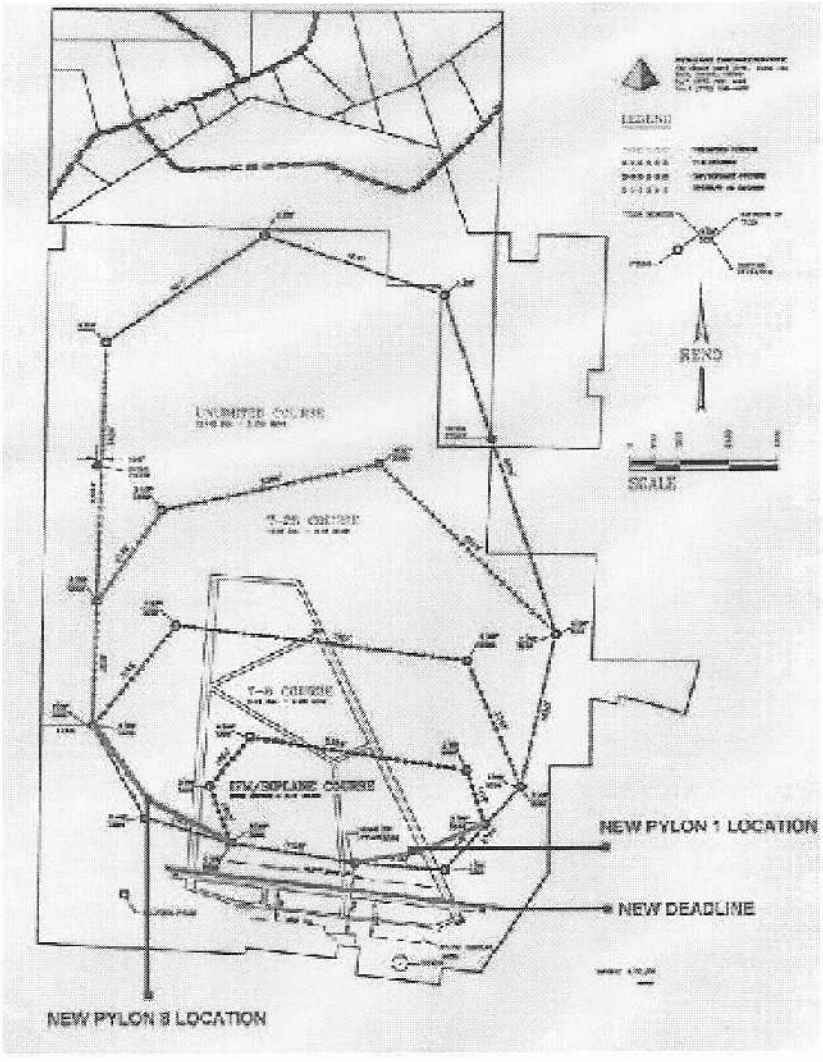

Figure 2: Reno Race Course Layout.

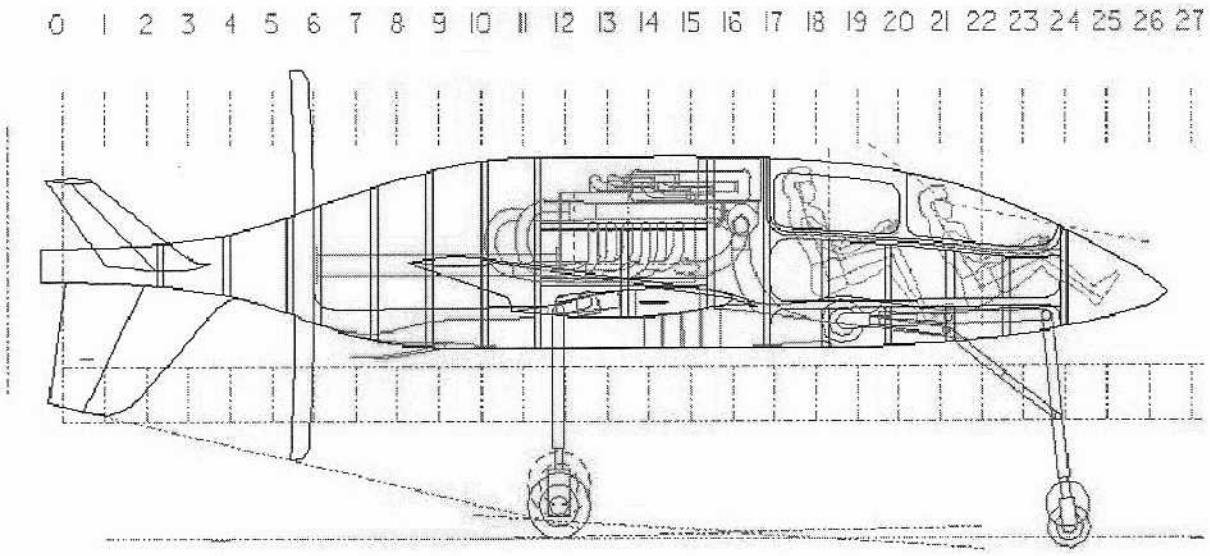

Figure 3: Side View of Body-Prop Design.

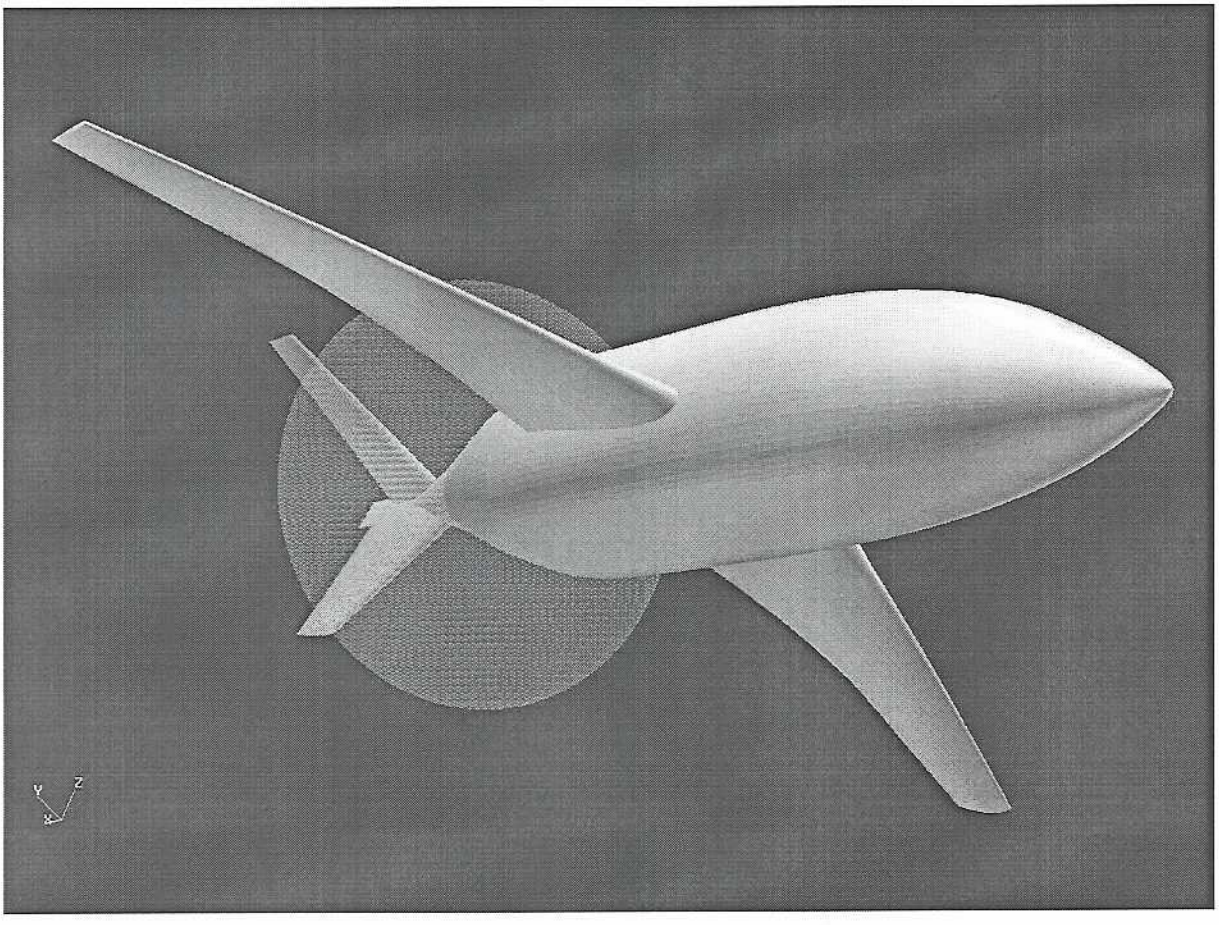

Figure 4: Rendering of Body-Prop Design in Flight.

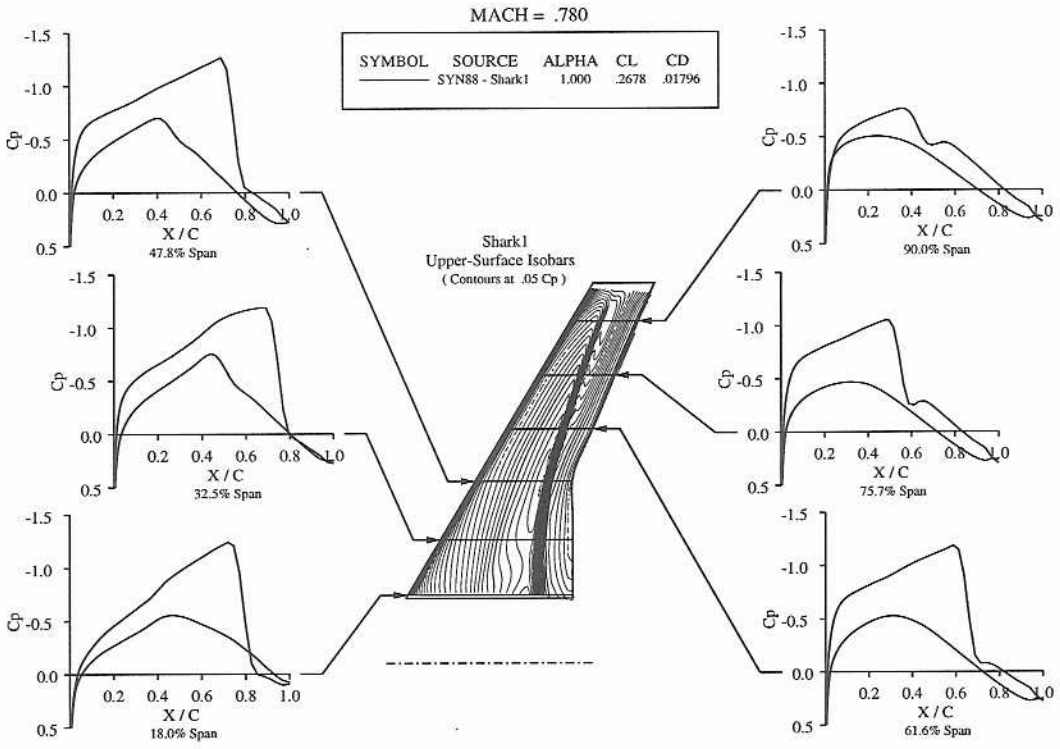

Figure 5: Pressure Distributions of Shark1 Baseline Wing.

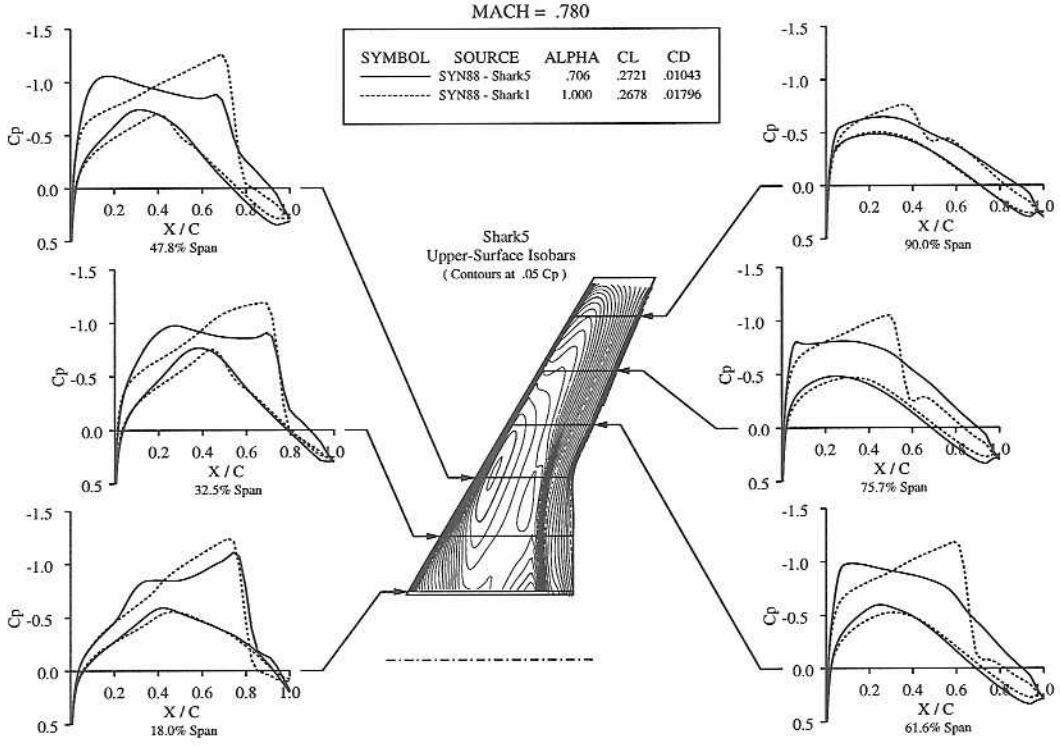

Figure 6: Comparison of Shark5 Wing on Baseline Fuselage with Baseline Configuration.

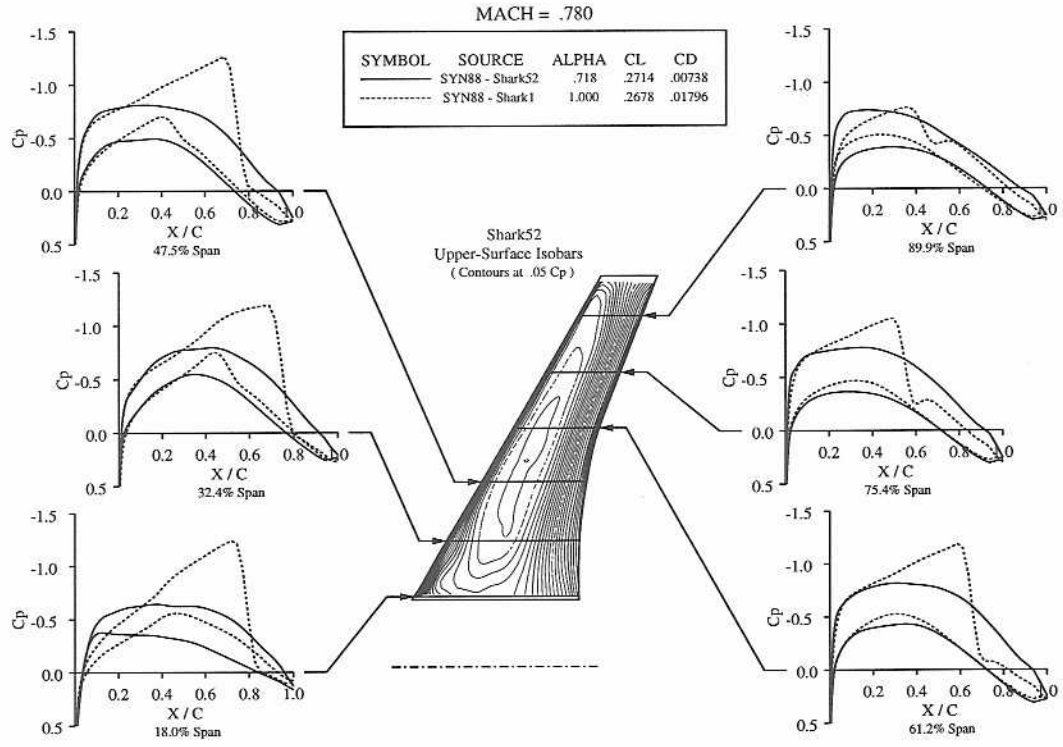

Figure 7: Comparison of Shark52 Wing on Stretched Fuselage and Baseline Configuration.

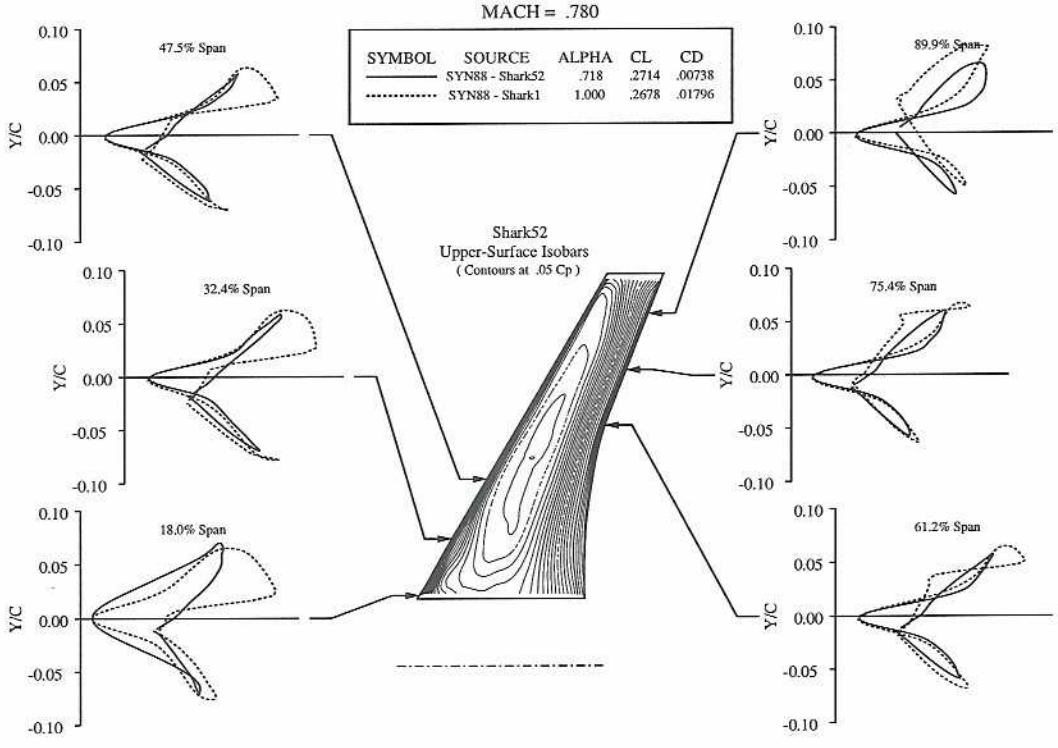

Figure 8: Comparison of Shark52 and Shark1 Wing Drag Loops.

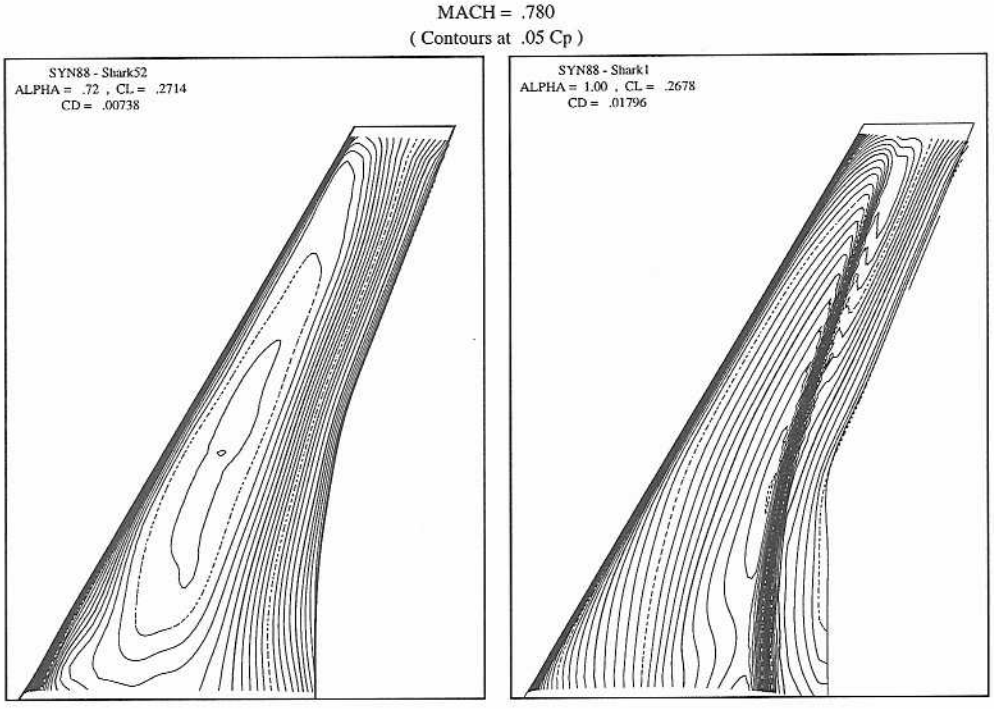

Figure 9: Comparison of Shark52 and Shark1 Wing Pressure Contours.

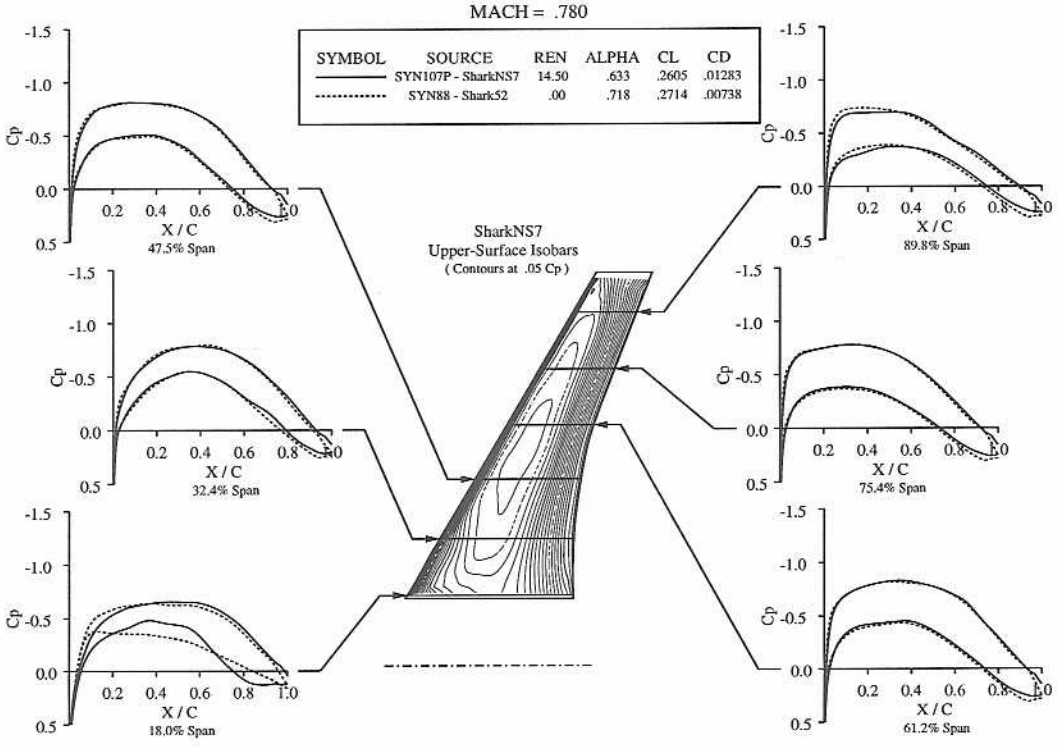

Figure 10: Result of Navier-Stokes Inverse Design.

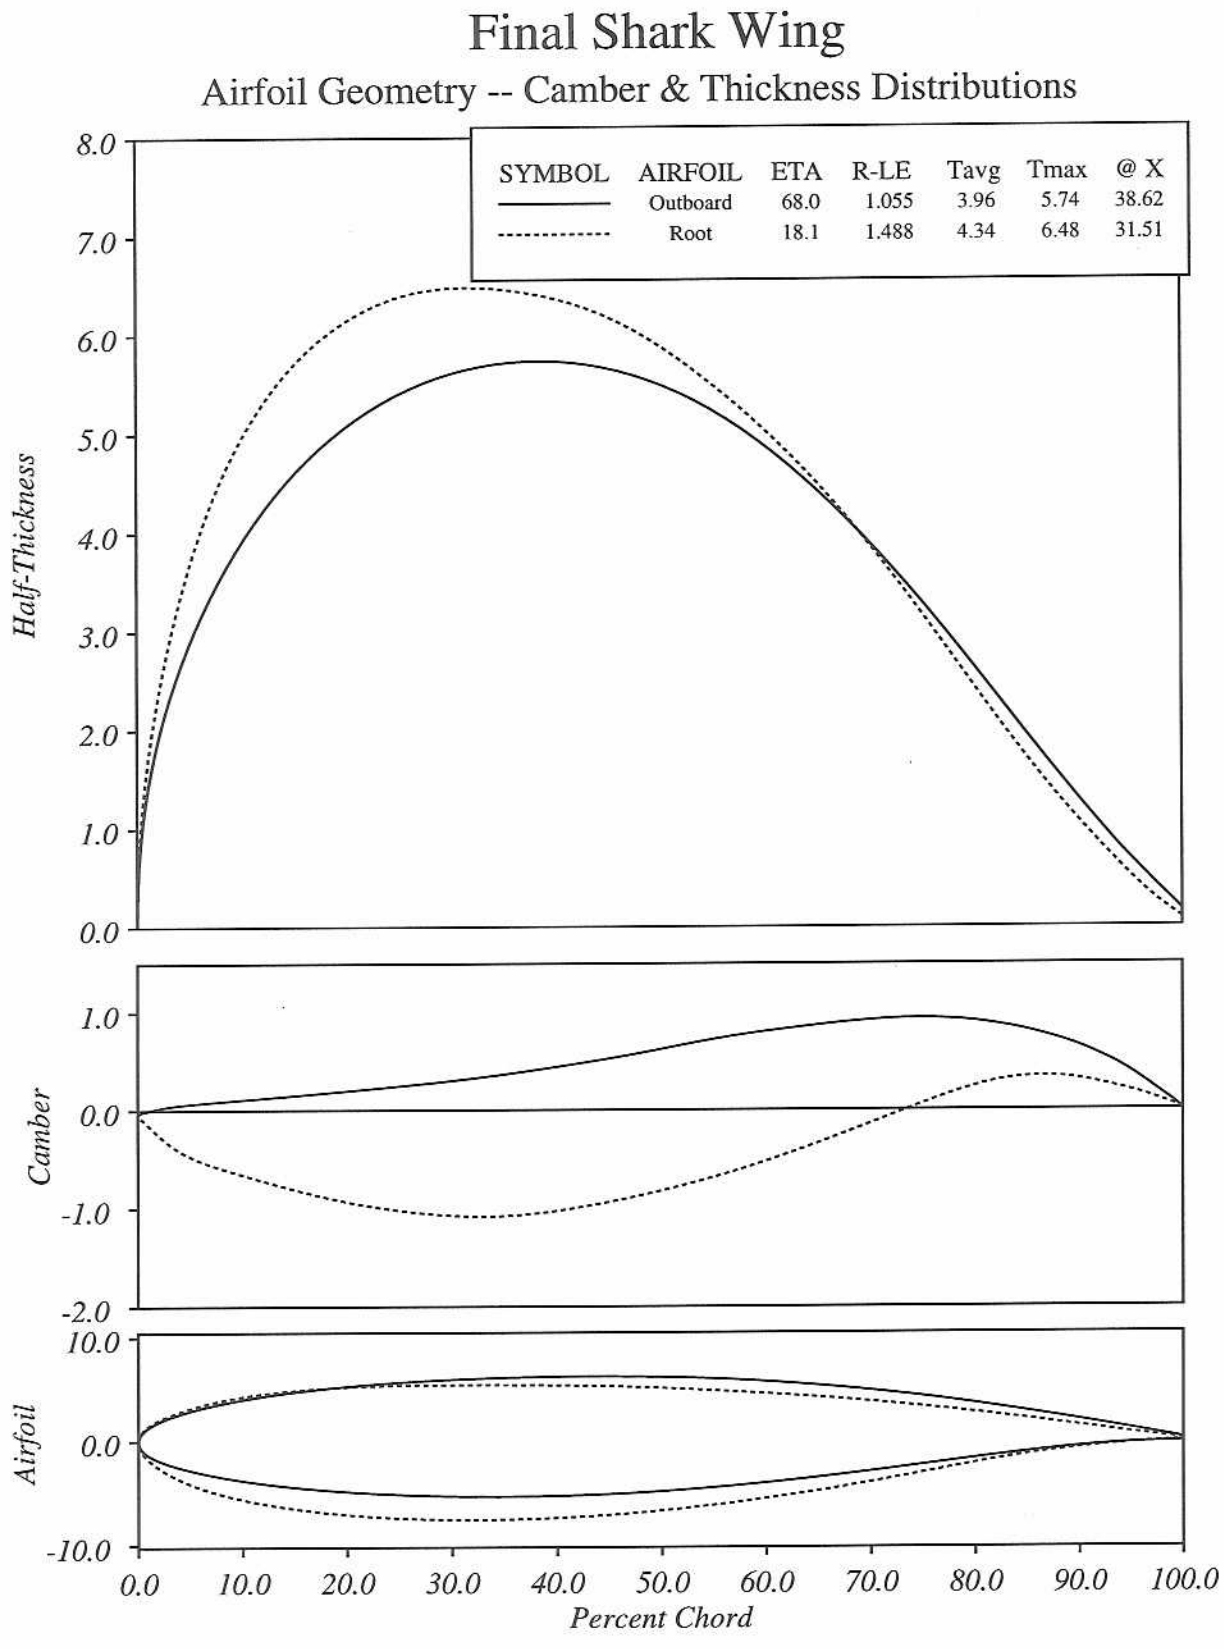

Figure 11: Final Wing Airfoil Geometry - Thickness & Camber Plots.