

Pros and Cons of Airfoil Optimization ¹

Mark Drela ²

1 Introduction

Optimization has long been considered as a means to solve the aerodynamic design problem in a formal and general manner. Early work by Hicks, Murman, and Vanderplaats [1, 2] investigated this possibility for transonic airfoil flows, with the results being encouraging but showing some rather unexpected results and difficulties. These early efforts were characterized by relatively few design parameters being used, primarily due to the computational costs of black-box gradient calculation via finite-differencing and the limited available computer resources.

Recent advances in parameter gradient calculation methods, such as the adjoint method [3] and the Newton-based direct method [4], and the relentless increases in computer speed and memory capacity, have largely removed the limitation on the number of design parameters which can be employed. At first glance, this should allow “truly optimum” designs to be computed and thus lay the issue of the effectiveness of optimization to rest. The author’s recent experience indicates that this is not the case, since unforeseen difficulties and surprising results arise as the number of parameters increases.

The purpose of this paper is to investigate the behavior of constrained optimization solutions with relatively large numbers of free design parameters present. The examples will be restricted to two-dimensional viscous airfoil optimization. The difficulties which appear even in such a seemingly simple aerodynamic design problem are quite illustrative of the advantages as well as shortcomings of formal optimization as an aerodynamic design method.

2 Method Summary

The present paper focuses on the effectiveness of optimization itself, rather than on analysis and optimization algorithms. Hence, this section will be restricted to only a brief summary of the analysis and optimization methods used for the application examples.

¹Chapter in *Frontiers of Computational Fluid Dynamics, 1998*
D.A. Caughey, M.M. Hafez, Eds., World Scientific, ISBN 981-02-3707-3

²Department of Aeronautics and Astronautics, M.I.T., Cambridge, MA 02139
AIAA Associate Fellow

2.1 Analysis method

The present design/optimization method employs the viscous/inviscid MSES code as the underlying analysis solver [5]. The overall equation system

$$R(U; \alpha, M, G_k) = 0 \quad (1)$$

consisting of the interior steady Euler equations, the boundary layer equations, and the necessary coupling and boundary conditions, are solved for the flow solution U as a fully-coupled system by a direct application of the Newton method.

$$\delta U = -[\partial R/\partial U]^{-1} R \quad , \quad U \leftarrow U + \delta U \quad (2)$$

After the Newton cycle is converged, the factored Jacobian is re-used to compute flowfield sensitivities to the design parameters G_k and flow parameters α , M .

$$\begin{aligned} \partial U/\partial G_k &= -[\partial R/\partial U]^{-1} \{\partial R/\partial G_k\} \\ \partial U/\partial \alpha &= -[\partial R/\partial U]^{-1} \{\partial R/\partial \alpha\} \\ \partial U/\partial M &= -[\partial R/\partial U]^{-1} \{\partial R/\partial M\} \end{aligned} \quad (3)$$

This re-use involves only back-substitutions, and allows computation of flowfield sensitivities to a very large number of design parameters at negligible cost. Hundreds of design parameters can be handled in interactive calculations on a modest workstation.

2.2 Geometry parameterization

A suitable parameterization for the airfoil shape, defined in terms of the fractional arc length s/s_{side} on each side of the airfoil, is a summation of sinusoidal basis functions g_k which perturb the airfoil by a distance Δn normal to its current surface. The design parameters G_k are the mode amplitudes.

$$\Delta n(s) = \sum_{k=1}^K G_k g_k(s) \quad , \quad g_k(s) = \frac{1}{k} \sin(k\pi s/s_{\text{side}}) \quad (4)$$

The $1/k$ scaling factor makes all the basis functions have the same maximum slope, which theoretically doesn't affect the optimum solution, but it does appear to greatly improve the behavior of an optimization descent sequence. Other geometry bases can of course be defined, and in fact many are better suited for specific problems. The useful features of the sine basis are a guaranteed mutual orthogonality, and a uniform and predictable increase of geometric resolution with added modes, making it quite suitable for the parameter-count investigation in this paper.

2.3 Optimization method

The voluminous solution sensitivity output generated by the MSES Newton solver is applied to the optimization problem in the interactive LINDOP program [4]. This allows the designer to use the sensitivity information to interactively try out various objective functions, constraints, and design parameter sets, and to generate linearized predictions resulting from explicit parameter changes, imposed pressure distributions, or objective-function descents. The linearized predictions are displayed for all the operating points being considered, thus giving visual warning when the optimization is headed for trouble.

One MSES/LINDOP cycle constitutes one descent step in design space. The well-known BFGS method [6] is used to generate the sequence of descent directions. The descent is continued until the objective function refuses to decrease further to within 0.00005 in C_D . The number of descents required for this is typically several times the number of free design parameters.

3 Low Reynolds Number Airfoil Application

The DAE-11 airfoil was designed by the author in 1987 specifically for the *Daedalus* human powered aircraft, using traditional inverse-method and direct geometry manipulation techniques. It is a second-generation airfoil, redesigned from its predecessor airfoils whose performance was partially verified in flight tests [7]. Hence, it makes an interesting candidate for possible further improvement via numerical optimization.

The key requirement for a human powered aircraft airfoil is to achieve minimum drag at the design flight lift coefficient. It is constrained by a structural thickness requirements at the spar locations, and by a number of other minor geometric requirements which influence the weight of the wing's secondary structure. Here, the focus will be on the primary requirements.

3.1 One-point optimization

The following optimization problem embodies the low drag requirement.

$$\text{minimize} \quad \mathcal{F}(G_k, \alpha) \equiv C_D \quad (5)$$

Using 40 G_k DOFs with no geometric constraints present (only C_L is held fixed) a physically unrealizable airfoil results soon after some number of optimization descent steps are taken, as shown in Figure 1. The airfoil becomes very thin, with the first trouble sign being the appearance of a re-entrant trailing edge. This type of result from an unconstrained optimization is of course entirely expected. Suitable thickness constraints must be imposed based on structural

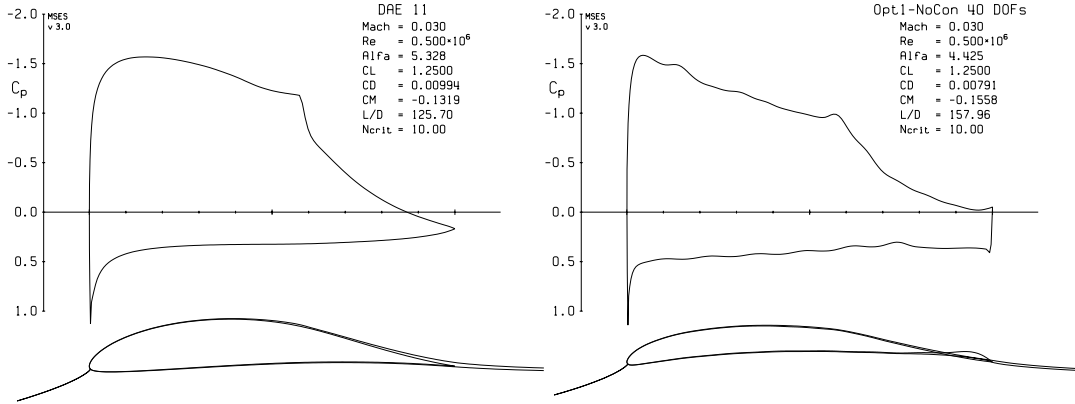


Figure 1: Baseline DAE-11 airfoil, and partially-optimized unconstrained airfoil.

considerations, and in this case the trailing edge and leading edge angles must be explicitly constrained as well. The following constraints have been found suitable for this problem after a number of optimization attempts.

$$\begin{aligned}
 C_L &= 1.25 & \theta_{TE} &= 6.25^\circ & (t/c)_{0.33c} &= 0.128 \\
 C_M &= -0.133 & \theta_{LE} &= 180^\circ & (t/c)_{0.90c} &= 0.014
 \end{aligned} \tag{6}$$

The C_M constraint has also been found necessary because the optimizer tends to strongly drive it more negative, which then has a large detrimental impact on the wing structural weight in a human-powered aircraft. The specified angles, thicknesses, and C_M are the same as those of the starting DAE-11 airfoil. The leading edge angle $\theta = 180^\circ$ constraint simply matches the surface slopes between the top and bottom surface at the airfoil nose, and is necessary to prevent the appearance of a sharp chisel-type leading edge. The C_L constraint effectively eliminates α as a degree of freedom, while the C_M constraint and each of the four geometric constraints removes one geometric degree of freedom. The total effective number of free design parameter DOFs is therefore $K - 5$. Using 100 geometric DOFs is typically required to generate a practically arbitrary airfoil, although a smaller number may be adequate if the starting airfoil is reasonably close to the final design, as in this case.

All constraints are imposed explicitly in LINDOP by augmenting the objective function \mathcal{F} using Lagrange multipliers. The net effect is to project all changes in the design space onto the admissible constrained subspace.

Figure 2 shows the optimized airfoil C_D and geometry which results from using 5, 10, 20, and 40 net degrees of freedom. Both the airfoil shape and the C_D appear to asymptote, indicating that the 40-DOF result is in fact the true solution to the optimization problem as posed, at least in the local minimum sense. The C_D is reduced from 0.0994 in the original DAE-11 airfoil to 0.0836 in the optimized airfoil – a rather large 13% reduction. However, calculation of




Figure 2: One-point optimized airfoil geometry and objective function versus number of design DOFs.

the entire drag polar of each airfoil tells a very different story. As the number of DOFs is increased, the drag reduction is attained over an ever-narrower C_L range. Figure 3 shows that the polar curve takes on a cusped form, so that the benefit range shrinks to nearly zero. Hence, the “optimized” airfoils actually get considerably worse in a practical sense with increasing number of DOFs.




Figure 3: Polars for original DAE-11 airfoil and 1-point optimized airfoils.

In retrospect, this dramatic behavior is not surprising. Airfoil design — and in fact most aerodynamic design — is *fundamentally* driven by tradeoffs, which almost always include off-design performance. If such tradeoffs are not considered, as in this 1-point optimization example, very poor results are almost certain to occur by chance if anything — the number of poor airfoils vastly exceeds the number of good airfoils! In this case, the optimizer raises a bump on the surface to “fill” the transitional separation bubble, effectively reducing the mixing and associated drag penalty which occurs when a bubble undergoes

transition and reattachment [7]. However, the bubble location changes with C_L , so this “optimized” airfoil shape is effective for only the sampled C_L point seen by the optimizer. For lower C_L values, the shear layer does not undergo transition at the bump, but instead separates off the bump in a laminar state, forming an even bigger and more lossy separation bubble than would occur without the bump. For higher C_L values, transition runs forward ahead of the bump much faster than it would otherwise, very quickly precipitating a rapid rise in skin friction drag, which then causes a rapid thickening of the downstream boundary layer and quickly precipitates stall.

The real deficiency here is not with the optimization technique, which gives the demonstrably correct answer, but rather with the formulation of the optimization problem itself. The simple 1-point drag minimization, even with a number of real geometric constraints determined by considerable trial and error, still does not embody the real design requirements of the airfoil. Unfortunately, this shortcoming is not at all obvious at the outset.

3.2 Two-point optimization

To remedy the obvious deficiency in the 1-point design, a 2-point optimization is formulated by replacing the objective function (5) with

$$\mathcal{F}(G_k, \alpha) \equiv \frac{1}{3} C_D|_{C_L=1.0} + \frac{2}{3} C_D|_{C_L=1.5} \quad (7)$$

so that the two ends of the expected operating C_L range are now sampled. The 1:2 weighting between the two operating points has been determined to be necessary so that the upper part of the drag polar is not compromised excessively by the less important lower part. The same geometric constraints (6) are used, but of course two separate C_L constraints are used for the two operating points. The objective function now decreases less quickly with number of DOFs than the 1-point case, and 60 DOFs are needed to asymptote reasonably well to the true optimum solution. The evolution of the drag polars with increasing DOFs is shown in Figure 4. The striking feature immediately apparent is that the peaky local-optimization polar shape persists, but this now occurs at the two sampled C_L values. The airfoil geometry now has two bumps, each one at the bubble location at the sampled operating points. The bumps are not as pronounced as in the 1-point example, but they still show a significant drag penalty away from the sampled C_L values as the scalloped drag polar in Figure 4 shows. Again, the “optimized” airfoil is inferior in a practical sense to the starting DAE-11 airfoil, although its deficiencies are less severe than in the 1-point optimized airfoil.




Figure 4: Polars for original DAE-11 airfoil and 2-point optimized airfoils.

3.3 Six-point optimization

Carrying the multi-point optimization concept further, a 6-point objective function is defined as

$$\begin{aligned} \mathcal{F}(G_k, \alpha) \equiv & \frac{5}{45} C_D|_{C_L=0.8} + \frac{6}{45} C_D|_{C_L=1.0} + \frac{7}{45} C_D|_{C_L=1.2} \\ & + \frac{8}{45} C_D|_{C_L=1.4} + \frac{9}{45} C_D|_{C_L=1.5} + \frac{10}{45} C_D|_{C_L=1.6} \end{aligned} \quad (8)$$

so that the sampled C_L points $0.8 \dots 1.6$ now span slightly beyond the expected $1.0 \dots 1.5$ operating range of the airfoil to give some margin at the ends. The evolution of the design with added modes is now even more gradual, and 90 modes are required to nearly asymptote to the optimum design as shown in Figure 5.




Figure 5: Six-point optimized airfoil geometry and objective function versus number of design DOFs.

The polar of this 6-point optimized airfoil is compared with the starting DAE-11 polar in Figure 6. The peaky behavior around each sampled point is still present, but to a much lesser degree than before. The geometry of this airfoil is rather striking, however. Figure 7 shows the geometries for the 1-point, 2-point, and 6-point optimized airfoils, showing one, two, and six bumps at the separation bubble locations at the sampled operating point. The inviscid C_p distribution of the 6-point airfoil in Figure 8 shows the severity of its six bumps. Note also that the inviscid C_p is somewhat “noisy” over most of the airfoil due to slight surface curvature irregularities. This is typical in problems with many geometric DOFs, since such small-scale irregularities have virtually no aerodynamic penalty and hence are invisible to the optimizer.




Figure 6: Polars for original DAE-11 airfoil and 6-point optimized airfoils.




Figure 7: Geometries of optimized airfoils, showing surface bumps.

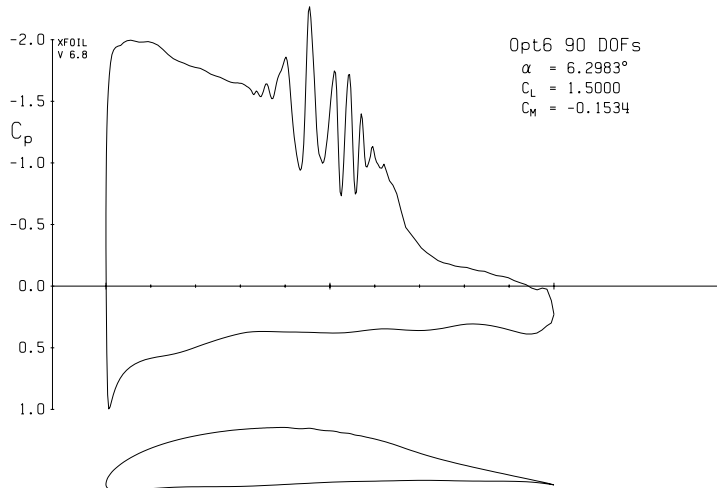


Figure 8: Inviscid C_p distribution for 6-point optimized airfoil.

3.4 Observations

In the 1-point, 2-point, and 6-point optimization examples above, the optimizer manipulated the geometry at the smallest physical scale which has a significant impact on the objective function — in this case the transitional separation bubble. If presented with sufficient geometric DOFs, the extent to which the optimizer performs such manipulation is startling. The 6-point airfoil with the six distinct bumps is surely not what was expected, but the result makes perfect sense in retrospect. The viscous C_p distribution shown in Figure 9 for one of the sampled points at $C_L = 1.5$ shows that the bumps in the laminar region are mostly submerged under the free shear layer. Each bump protrudes just enough from the surface to barely reach the shear layer, and is positioned along the surface so that it “catches” the shear layer as it undergoes transition and allows it to reattach without significant drag-producing mixing of the stagnant fluid under the shear layer. The computed velocity profiles at the bumps show this clearly in Figure 10. At C_L values midway between the sampled points, there is no bump to catch the shear layer, which must instead reattach by mixing out the laminar bubble in the valley, which carries a drag penalty. This is responsible for the “scalping” in the drag polar in Figure 6. Surely, this rather silly multi-point optimum solution could not have been foreseen by even the most experienced aerodynamicist, and illustrates the almost devious cleverness exhibited by an optimizer which is armed with a large number of design DOFs.

Given the likelihood that the localized bumps can be eliminated with denser operating point sampling in the C_L range, optimization appears to have produced a slight improvement over the baseline DAE-11 airfoil over the lower part of the polar, as Figure 6 indicates. Away from the bumps, the airfoils shown in Figure 7 are clearly converging to a unique shape. However, the geometric de-

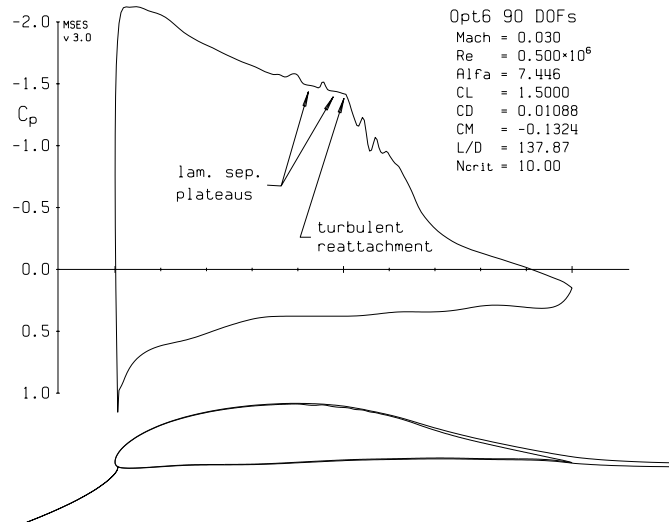


Figure 9: Viscous C_p for 6-point optimized airfoil at sampled point $C_L = 1.5$




Figure 10: Velocity profiles for 6-point optimized airfoil at $C_L = 1.5$, showing second bump “catching” reattaching shear layer.

tails of this shape are not particularly compatible with the structural techniques employed in human-powered aircraft. The highly cambered trailing edge, the distinct concavity on the upper surface near $x/c = 0.7$, and the bottom-surface inflection points would all be quite difficult to implement. Since there was no way to predict the appearance of these particular features at the outset, a new optimization problem with the appropriate constraints would likely have to be constructed and solved. The conclusion is that in an engineering setting, using optimization for airfoil design is still an iterative cut-and-try undertaking. But compared to the traditional inverse techniques, the cutting-and-trying is not on the geometry, but rather on the precise formulation of the optimization problem.

4 Transonic Airfoil Application

In this example, the well-known RAE-2822 airfoil is optimized, starting from the baseline Case 13a of reference [8]. This case is partially into the drag rise Mach range, but short of shock-induced separation. Hence it is in an inefficient but

not unrealistic flight condition, and significant improvement might be expected upon optimization redesign. The optimization calculations are all performed at the same $C_L = 0.733$ and $Re = 2.7 \times 10^6$ as the starting RAE-2822 case, and the same $M = 0.74$ is used for the 1-point optimization case.

4.1 One-point optimization

The objective of reducing drag is embodied by the following plausible optimization problem, again using the sine-mode coefficients G_k and angle of attack α as the design parameters. The leading and trailing edge angles and a local thickness are imposed as reasonable geometric constraints.

$$\text{minimize} \quad \mathcal{F}(G_k, \alpha) \equiv C_D \quad (9)$$

$$\text{subject to} \quad \begin{aligned} C_L &= 0.733 & \theta_{TE} &= 8.62^\circ \\ (t/c)_{0.35c} &= 0.1206 & \theta_{LE} &= 180^\circ \end{aligned} \quad (10)$$

Figure 11 shows the baseline analysis solution compared with experimental data, and the 1-point optimized airfoil. Although a 40% drag reduction is obtained and the local thickness has been maintained, the airfoil is not suitable for a jet transport application since the average thickness over the typically wide spar box has decreased considerably. The local thickness constraint therefore does not represent the real structural requirements, although again this was not obvious at the outset.

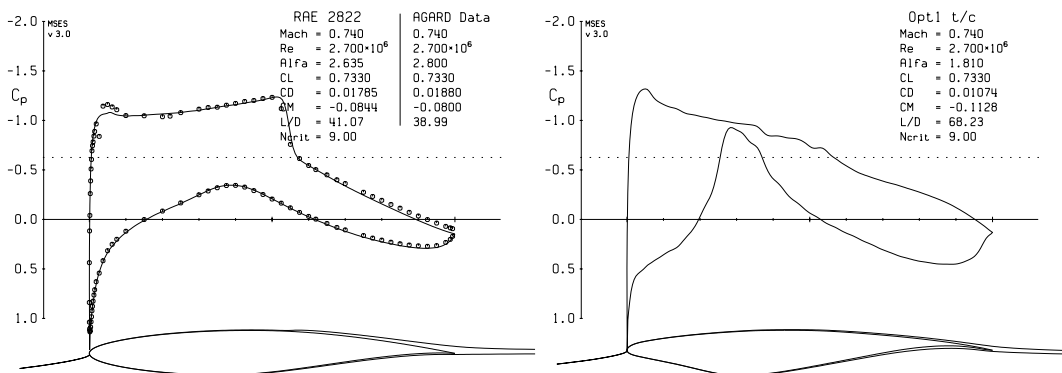


Figure 11: C_p distributions for baseline RAE-2822 airfoil, and 1-point optimized airfoil with local thickness constraint at $x/c = 0.35$.

In lieu of the local t/c constraint in (10), the r.m.s. strain ϵ_{rms} (per unit bending moment and unit material modulus) is imposed to more realistically account for the requirements of a wide spar box.

$$\left[\frac{\oint \epsilon^2 t^2 ds}{\oint t^2 ds} \right]^{1/2} \equiv \epsilon_{\text{rms}} = 14.26 \quad (11)$$

The integrals are taken around the airfoil perimeter, and are weighted by the local skin thickness t . The local bending-related unit extensional strain ϵ is defined in terms of the local airfoil surface point y and the local bending inertia and bending centroid location y_o .

$$\epsilon = \frac{y - y_o}{\oint (y - y_o)^2 t ds} \quad , \quad y_o = \frac{\oint y t ds}{\oint t ds}$$

The skin thickness t is chosen to be nonzero only over the chordwise extent of the spar box $0.15 \leq x/c \leq 0.65$. In effect, the r.m.s. strain constraint enforces roughly an average airfoil thickness over the spar box, and the remainder of the airfoil is structurally irrelevant.

Figure 12 shows the airfoil which results from the optimization using a net 40 DOFs. In contrast to the local thickness constraint, the r.m.s. strain constraint now enforces a reasonable thickness over the extent of the spar box, but a severe bottom surface concavity is produced in front of the spar, and a lesser one in back. Although it adds aerodynamically advantageous bottom loading, the front concavity may not be feasible in a practical sense, so that the optimization problem as posed *still* may not embody the actual requirements of the design problem. Again, this shortcoming was not apparent at the outset. In any

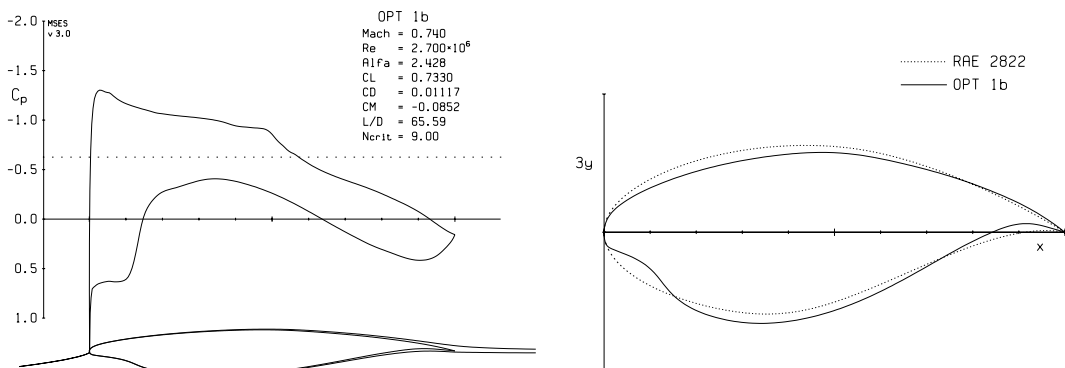


Figure 12: C_p distributions for 1-point optimized airfoil with r.m.s. strain constraint over $0.15 \leq x/c \leq 0.60$, and optimized airfoil compared with baseline RAE-2822 airfoil.

case, the optimized airfoil shows fully isentropic flow with no shock wave and no wave drag. The net C_D has been reduced from 0.0178 to 0.0112 — a 37% reduction. This is partially illusory, however. The entire Mach number sweep in Figure 13 shows that the drag reduction is realized primarily in the vicinity of the sampled operating point at $M = 0.74$, although here the degree of local optimization is not as extreme as in the low Reynolds number airfoil case. The drag is increased considerably at lower Mach numbers, a behavior which is well known in shockfree transonic airfoils. The geometric feature responsible

for this behavior is the slight bump centered roughly at $x/c = 0.50$ which produces a gradual recompression over the aft part of the supersonic zone and thus eliminates the shock. Associated with the bump is a flattened forward upper surface which produces a strong expansion and a strong shock at lower Mach numbers, thus increasing drag. At Mach numbers above the sampled $M = 0.74$, the flow expands strongly past the bump and results in a strong shock and rapid drag rise.




Figure 13: Mach sweep for 1-point optimized and baseline RAE-2822 airfoils.

The behavior of the optimizer in this example is in some ways similar to its behavior in the low Reynolds number optimization case. The optimizer exploits the flow features at the smallest geometric scale which is resolvable by the available geometry perturbation modes. In this case the physical scale manipulated is the shock/boundary-layer interaction region.

4.2 Two-point optimization

To reduce the point-optimized nature of the 1-point design, a 2-point optimization is defined by replacing the objective function (9) with

$$\mathcal{F}(G_k, \alpha) \equiv \frac{1}{3} C_D|_{M=0.68} + \frac{2}{3} C_D|_{M=0.74} \quad (12)$$

so that some reasonable range of Mach numbers is now sampled. The same C_L and r.m.s. strain, and geometric angle constraints are used. The resulting airfoil is shown in Figure 14, and the entire Mach sweep is shown in Figure 15. Some “localized” optimization is clearly evident. Two precompression bumps are now present on the airfoil — one at the shock foot location at each of the two sampled Mach numbers. No improvement has occurred above the $M = 0.74$ sampled point.




Figure 14: C_p distributions for 2-point optimized airfoil at the two sampled Mach numbers, and comparison of optimized geometry with baseline RAE-2822 airfoil.




Figure 15: Mach sweep for 2-point optimized and baseline RAE-2822 airfoils.

4.3 Four-point optimization

To improve the localized nature of the 2-point design, a 4-point optimization problem is defined as follows.

$$\begin{aligned}
 \mathcal{F}(G_k, \alpha) &\equiv \frac{1}{7} C_D|_{M=0.68} + \frac{1}{7} C_D|_{M=0.71} \\
 &\quad + \frac{2}{7} C_D|_{M=0.74} + \frac{3}{7} C_D|_{M=0.76}
 \end{aligned} \tag{13}$$

A denser sampling over the Mach number range is now used, with larger weights applied to the upper part of the range to reflect its greater importance. Note also that a higher Mach of $M = 0.76$ is now sampled in an attempt to reduce the drag above the highest $M = 0.74$ point sampled previously.

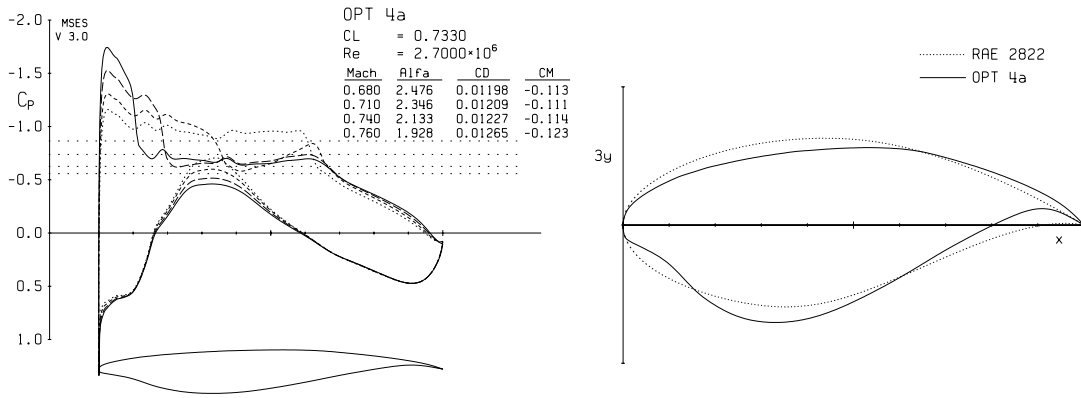


Figure 16: C_p distributions for 4-point optimized airfoil at the four sampled Mach numbers, and comparison of optimized geometry with baseline RAE-2822 airfoil.




Figure 17: Mach sweep for 4-point optimized and baseline RAE-2822 airfoils.

Aside from the front underside concavity, the rather thin trailing edge for $x/c > 0.9$, and the more-negative C_M , the results in Figure 17 show that this 4-point optimized airfoil appears to be an attractive improvement over the baseline RAE-2822. The increase in the drag-divergence Mach number is most significant. Although the airfoil contour in Figure 16 appears to be quite smooth on the upper surface, the finite number of sampled points still results in significant bumps at each of the four sampled-point shock locations. This is vividly displayed on the Mach wave plot shown in Figure 18, with an expansion fan emanating from the three forward bumps which are fully contained in the supersonic region at the fourth operating point sampled at $M = 0.76$. These shock-managing bumps employed by the optimizer are clearly much less severe than the reattachment-managing bumps in the low Reynolds number DAE-11 example. Nevertheless, they clearly are artifacts of the finite number of sampled

operating points, and would not be present in the ideal case where the Mach number is sampled in a nearly continuous manner.




Figure 18: Mach waves for 4-point optimized airfoil at $M = 0.76$, with indicated bumps at shock positions at lower Mach number conditions.

5 Conclusions

The examples presented clearly illustrate the numerous pitfalls which can easily appear in even simple-looking airfoil optimization problems. The appearance of small-scale geometric irregularities in both the low Reynolds number airfoil and the transonic airfoil examples could be suppressed by increasing the number of sampled operating points, so that additional bumps appear at the intermediate locations, ultimately all blending into a smooth surface. This argument supports the hypothesis that for a smooth geometry, it is necessary to have

$$\text{number of operating points sampled} = \mathcal{O}(\text{number of design parameters})$$

since an increase in the number of design parameters reduces the length scale at which the optimizer can exploit the flow, which then must be matched by a proportional increase in the number of sampled points to control this exploitation. Of course, geometric constraints such as on surface curvature could instead be imposed to control this exploitation, but this is in effect a reduction in the effective number of free geometric design parameters. Since the cost increases linearly with the number of sampled operating points, the ability to perform optimization with large numbers of design DOFs is an extremely expensive proposition. It stresses the need to reduce the effective number of design DOFs to an absolute minimum, whether this reduction is done explicitly by choosing

a particular DOF count, or implicitly by application of constraints. Unfortunately, limiting the number of DOFs has the drawback that it is not clear what the appropriate number should be for any given problem, or what the appropriate geometric constraints should be. Also, there is always the possibility that limiting the number of DOFs will produce a significantly sub-optimal design.

These conclusions have direct implications not only for 2-D airfoil optimization, but also for large-scale 3-D optimization problems with numerous design DOFs. There is no reason to hope that 3-D problems will be immune to such difficulties. Based on the present optimization examples and experience with similar other cases, the following observations are offered.

5.1 Airfoil Optimization — Cons

- The objective function and constraints which effectively embody a practical design problem are not knowable at the outset without an extensive experience base for closely related problems.
- If presented with sufficient design mode resolution, an optimizer will readily (and annoyingly) manipulate and exploit the flow at the smallest significant physical scales present. The examples exhibited such manipulation on the scale of a transitional separation bubble, and on the scale of a shock/boundary-layer interaction zone.
- Detailed geometry manipulation at the sampled-condition locations of small-scale mobile features like separation bubbles and shocks tends to produce improved performance only near the sampled operating conditions. Away from the sampled conditions, the manipulated geometry is very inappropriate and a sometimes severe degradation in performance results.
- Manipulation at the smallest scales can be discouraged by multi-point optimization. The small mobile flow features then appear “smeared” to the optimizer and are exploited less severely.
- Increasing the net number of geometric DOFs appears to require a corresponding increase in the number of operating points sampled by the objective function. Near-continuous sampling of the operating space may be required in the theoretical limit of a general airfoil design problem with a very large number of DOFs — a very expensive proposition.
- The most suitable operating points to be actually sampled in multi-point optimization are not apparent a priori. From limited experience, sampling somewhat beyond the expected operating range appears to be best.

- The point weights used in multi-point optimization are arbitrary, and their appropriate values cannot be easily estimated without prior experience.
- Optimized aerodynamic shapes are usually “noisy” and usually require a posteriori smoothing.

5.2 Airfoil Optimization — Pros

- Optimization gives a rapid indication of possible directions for improvement when intuition with traditional inverse techniques is exhausted, or when direct or inverse geometric cut-and-try is impractical.
- Optimization indicates possible design improvement paths when unusual non-aerodynamic design drivers are present, provided these drivers are precisely quantified. The r.m.s. strain constraint is one such example.
- Optimization is well-suited for design problems which inherently involve multiple operating points, since multiple operating points will likely have to be sampled anyway. This is in contrast to single-point problems which must be artificially converted to multi-point optimization problems to suppress undesirable local optimization. Such single-point problems are often better handled with traditional inverse methods.
- Optimization is well-suited for design problems where the effective number of design DOFs is limited because of externally-imposed constraints. Such problems are less prone to undesirable geometric manipulation at small geometric scales.

References

- [1] R.M. Hicks, E.M. Murman, and G.N. Vanderplaats. An assessment of airfoil design by numerical optimization. NASA TM X-3092, July 1974.
- [2] R.M. Hicks and G.N. Vanderplaats. Application of numerical optimization to the design of supercritical airfoils without drag-creep. SAE Paper 760477, 1977.
- [3] A. Jameson. Aerodynamic design via control theory. ICASE Report 88-64, Nov 1988.
- [4] M. Drela. Design and optimization method for multi-element airfoils. AIAA Paper 93-0969, Feb 1993.
- [5] M. Drela and M.B. Giles. Viscous-inviscid analysis of transonic and low Reynolds number airfoils. *AIAA Journal*, 25(10):1347–1355, Oct 1987.

- [6] P.E. Gill, W. Murray, and M.H. Wright. *Practical Optimization*. Academic Press, San Diego, 1981.
- [7] M. Drela. Low-Reynolds number airfoil design for the MIT Daedalus prototype: A case study. *Journal of Aircraft*, 25(8):724–732, Aug 1988.
- [8] P.H. Cook, M.A. McDonald, and M.C.P. Firmin. Aerofoil RAE 2822 pressure distributions and boundary layer and wake measurements. In *Experimental Data Base for Computer Program Assessment, AR-138*. AGARD, 1979.