

Application of Adjoint Optimization Method to Multi-Element Rotorcraft Airfoils

Mark S. Chaffin, Sr. Engineer
Lockheed Martin Engineering & Sciences
NASA Langley Research Center
Hampton, Virginia

Abstract

An adjoint optimization method coupled with an unstructured Navier-Stokes code was applied to the case of a multi-element rotorcraft airfoil. The combined optimization tool was used to reduce the drag of the airfoil at high Mach numbers and low angles of attack without significantly reducing the maximum lift at low Mach numbers and high angles of attack.

Introduction

The application of computational fluid dynamics to airfoil design and optimization is not a recent development. A simple search of technical papers will reveal dozens of papers describing the application of computational techniques to the design of airfoils and other aerodynamic shapes.

Aerodynamic design generally refers to the modification of an airfoil or aerodynamic shape to produce a specific result, such as a specific lift coefficient. The resulting shape does not necessarily represent the optimum shape, as it might be possible to achieve a higher lift coefficient, or the same lift coefficient with lower drag. By iterating on design parameters, the optimum shape can be achieved. On the other hand, optimization refers to finding the optimum shape that minimizes a set of cost functions while maintaining a set of constraints. For example, the cost function might consist of the drag coefficient and the constraints might consist of limits placed on the airfoil deformation. The optimization method would then try to find the minimum drag coefficient for a shape within the deformation limits.

Design methods typically use computational analysis in an iterative manner to arrive at the desired design. The Direct Iterative Surface Curvature (DISC) method described by Campbell (Ref. 1) is a good example of this method. Smith and Campbell (Ref. 2) applied the DISC method to reduce the drag for a transonic aircraft. Mineck *et al.* (Ref. 3) developed a dual point design using the DISC method to reduce drag on tran-

sonic airfoils at two design points. More recently, Naramore (Ref. 4) applied the DISC method to the case of multi-element rotorcraft airfoils.

As mentioned above, optimization methods attempt to minimize one or more cost functions while maintaining a set of constraints. These methods can be expensive to use and it is possible that only a local minimum will be found. Optimization methods rely on the accurate calculation of sensitivity derivatives. These derivatives reflect the change in the cost function with respect to the design variables. The derivatives are then used to reduce the cost function. Thus, it is critical to choose the proper method for calculating the sensitivity derivatives.

However, optimization methods also have some advantages. While design methods require the knowledge of target pressures (or special rules in the case of the DISC method), optimization methods are not dependent on them. In the case of a multi-element airfoil, it is a simple matter to allow one element to be changed while the optimization function is derived from the characteristics of a separate element.

Adjoint methods allow the cost function to be coupled with the flow equations. This method is well suited to aerodynamic problems where the number of design variables is typically large. Using adjoint methods, the derivatives can be calculated for a computational cost approximately equal to the cost of a single aerodynamic solution. It is not the purpose of this paper to explain the details of adjoint methods, but rather to show the application of an adjoint method to multi-element airfoils. More information regarding adjoint methods can be found in references 5 and 6.

The motivation for this work is the desire to improve the performance of rotors through the use of high lift techniques. The application of high-lift airfoils to rotorcraft could offer increased payloads and maneuverability. There are numerous approaches to improving the performance of rotorcraft airfoils. Fixed-wing aircraft typically employ slats, flaps, and in some cases variable camber to tailor the wing for high-lift conditions. The use of a slat for fixed-wing aircraft is common practice, however its application to rotorcraft blades is much more complex. The use of actuators to deflect slats and flaps on a rotor would add a considerable amount of complexity to the rotor system.

Presented at the American Helicopter Society Vertical Lift Aircraft Design Conference, San Francisco, CA. Copyright © 2000 by the American Helicopter Society, International. All rights reserved.

On the other hand, a fixed slat or flap would offer a more simple solution. Noonan *et al.* (Ref. 7) present two-dimensional experimental data for two rotorcraft airfoils with slats. Also, Noonan *et al.* (Ref. 8) presented experimental data for a model helicopter rotor using high lift devices. A large challenge to such a concept comes from the requirement for the design to perform well at the off-design conditions. A slatted airfoil will perform well at high angles of attack and low to moderate Mach numbers (retreating side), but the performance will suffer greatly as the angle of attack is reduced and the Mach number increases (advancing side).

This paper will describe the process and results from using an adjoint optimization method coupled with an unstructured grid Navier-Stokes code to improve the performance of a multi-element rotorcraft airfoil. The combined optimization tool was used to reduce the drag of the airfoil at high Mach numbers and low angles of attack while still retaining the improvement in lift at low Mach numbers and high angles of attack.

Computational Tools

The computational tools used for this work consisted of an unstructured grid Navier-Stokes code, an optimization routine, and various scripts which drove the process. These tools were compiled by Dr. Kyle Anderson of the NASA Langley Research Center. Descriptions of the various components follow.

FUN2D

The unstructured grid Navier-Stokes code, FUN2D (Ref. 9) is node based and uses an implicit upwind flow solver. The inviscid fluxes are obtained using Roe's approximate Riemann solver and the viscous terms are evaluated with a central-difference formulation. The Spalart-Allmaras (Ref. 10) turbulence model was used for all the cases reported in this paper.

Optimization Routine

The optimization tool used for this work uses a discrete adjoint formulation to obtain sensitivity derivatives on unstructured grids (Ref. 5). The optimization algorithm, KSOPT, (Ref. 11) along with a mesh movement algorithm is coupled with the flow solver. It is not the goal of this paper to explain in detail these methods. Rather, it is the purpose of this paper to show the degree to which such optimization tools have evolved and how they can be used for rotorcraft applications.

Tool Usage

The first step to using the optimization software is to decide the manner in which the airfoil surface

can be modified. A computer program is used to fit a spline on the airfoil surface. The user controls the number of control points, the order of the spline, and the number of final points on the airfoil by means of a graphical user interface. Figure 1 shows a sample geometry and the control points on the splined geometry. A trade-off is usually required in this step between the number of control points used to define the spline and the number of control points (design variables) that will be allowed to move. The more control points used to define the spline, the more accurately the airfoil will be defined. However, the more control points (design variables) that are allowed to move, the greater the expense will be during the optimization runs. The next step in the process is to specify which control points will be allowed to move and how far. The vertical lines passing through the control points on the lower surface represent the limits of vertical movement allowed for these control points during optimization.

Once the movement limits have been set on the control points, the full grid can be generated. The grids were obtained using the AFLR grid generator (Ref. 12) obtained from Mississippi State University. Scripts are used to simplify the gridding. Figures 2 and 3 show the full grid and the grid around the slat, respectively. After the grid has been generated, FUN2D is run until the solution is well converged. At this time, various preprocessors are run prior to the start of the optimization. There are three cost functions currently available with the optimization tool. They are lift, drag, and pressure. The user is allowed to specify a target lift, drag, and/or pressure distribution. The cost function becomes the difference between the target and actual values. The user can also specify the weighting of each cost function, so that they can have equal weighting, or one can take a higher priority to another.

In all the cases shown in this paper, the target drag coefficient was set to 0.0, and the weighting function was set to 1.0. In other words, the optimization routines were asked to modify the airfoil such that the drag coefficient became zero. The weighting factors for the lift and pressure coefficients were set to 0.0, i.e., there was no target lift or pressure distribution. In some cases, it was necessary to place restrictions on the local curvature of the airfoil surface. The current implementation of the optimization tool operates on the total lift and drag. Thus, the effects of changes to the slat on the main element are taken into account in the optimization process.

Scripts are used to control the operation of FUN2D, KSOPT, and the associated mesh deformation routines. The user specifies the number of optimization cycles to perform. The flow field is saved at the end of each cycle to allow the user to monitor the solution as it progresses. In general, the optimizer might be run sev-

eral times to arrive at a final solution. Runs were made on various workstations. Typical run times for a 15 cycle run on a 195 Mhz workstation were on the order of 24-36 hours depending on the size of the grid and the number of design variables. Often, the optimization routines were restarted for an additional 15 cycles. The optimization process was stopped when it became evident that no further reduction in drag was occurring. The grids used for the optimization runs typically contained approximately 23,000 nodes, or 45,000 cells. Usually, 10-15 design variables were used during each design.

RC(6)-08 Airfoil

The airfoil chosen for modification in this study was configuration 106 of the RC(6)-08 airfoil (Ref 7). This airfoil consists of a main element and a leading edge slat and was derived from the RC(6)-08 single element airfoil by the addition of a slot. The two airfoils are shown in Figure 4. Before the airfoil can be optimized, it is important to understand some of its characteristics.

Because the leading edge slat is fixed, a key driver to the airfoil design is the reduction of drag experienced by the airfoil while at low angles of attack and moderate to high Mach numbers. These conditions would be experienced on the advancing side of the rotor, and would be most severe at the tip region. At the same time, it is important to maintain the high lift benefits at high angles of attack and low Mach numbers experienced on the retreating side of the rotor. To demonstrate the drag characteristics, Figure 5 shows a drag polar for the baseline configuration 106 at a Mach number of 0.40. The figure shows results from the unstructured Navier-Stokes code, FUN2D, as well as experimental data. Of particular note is the behavior near zero lift. As the lift decreases with decreasing angle of attack, the drag begins to increase. As will be shown later, this drag increase is primarily due to the boundary layer separating on the lower surface of the slat. However, if the angle of attack is further decreased, the drag level drops off. Examination of the computational flow field showed that this drop off in drag happens when the recirculation region caused by the separation on the slat extends back to the main element. When this occurs, the entire gap behind the slat becomes one recirculation region, effectively cutting off flow through the slat.

Design Procedure

There were some constraints placed on the modification of the airfoil. The greatest constraint was that the main element could not be modified. The reason for this was primarily to make future wind tunnel testing more practical. Designing a new slat while keeping the main element constant would allow various slat designs

to be tested in the wind tunnel more readily. The other constraints were more benign; namely, it was desirable that the thickness and camber of the slat did not change by a large amount.

Examination of the flow field about the slat for a typical low angle of attack condition illustrates the main problem with the baseline configuration. Figure 6 shows Mach number contours for such a condition. In this figure, the darker colors represent regions of low Mach number. The large dark region below the slat shows the large region of separated flow that leads to the drag rise at low angles of attack. One solution to the problem would be to rotate the slat nose up about its trailing edge. This would allow the boundary layer on the slat to remain attached longer. This was the first step taken in the design procedure. Rotating the slat about its trailing edge generated a series of slats. In all, six new slats were generated, one for each degree of rotation. Each new configuration was evaluated at a Mach number of 0.60, an angle of attack of 8.0 degrees, and a Reynolds number of 7×10^6 which was near the baseline C_{lmax} . It was found that the lift coefficient reached a maximum for this condition when the slat was rotated 6 degrees leading edge up relative to the baseline configuration.

The next step was to use the optimization procedure on each of the new configurations. Remembering that the drag increase occurs near zero lift and at a relatively high Mach number, the optimization procedure was carried out at an angle of attack of 0.75 degrees, a Mach number of 0.65, and a Reynolds number of 7.5×10^6 . While this does not represent the worst case, certainly a portion of the rotor will experience lower angles of attack and higher Mach numbers, it was felt to be a good compromise.

This work served as the first application of the optimization method to such a complex case and more than a few unanticipated problems were found. For example, it was found that if the method was pushed too hard, for instance 0 degrees angle of attack, the design goals would be met by generating a slat that would not be physically practical. As pointed out earlier, once the separation completely bridges the slot a reduction in drag is seen. Because the method is only asked to reduce the drag, at times the optimization routine would actually cause the flow to separate to such an extent that the slot was effectively closed. It usually accomplished this by growing what could best be referred to as a "tooth" on the lower surface of the slat. In cases like this, limits were placed on the surface curvature.

The next decision to make was how to modify the slat surface itself. The shape of the lower surface near the leading edge is a critical factor to controlling

the characteristic of the boundary layer for low angle of attack conditions. However, this region is also critical to the upper surface boundary layer for high angle of attack cases due to the shift in the stagnation point. This is another reason why the first step to optimizing the slat was to start with the low angle of attack case. Trying to optimize the slat at high angles of attack often resulted in changes to the leading edge region that negatively affected the lower surface boundary layer at low angles of attack. On the other hand, the effect on the upper surface conditions at high angles of attack was generally small when the leading edge was optimized for lower angles of attack. Another consideration in deciding where to modify the slat was the expense of the computation. The expense of the method is directly related to the number of control points that are allowed to move during the optimization. The more control points (design variables) that are allowed to be moved, the larger the design space becomes, but at the expense of longer design cycles. Also, the introduction of more design variables to the method increases the chances a nonphysical slat will result and typically requires the use of surface curvature constraints.

For each of the seven (baseline plus six rotated) slats, the optimization tool was allowed to adjust the lower surface in order to reduce drag. In some cases, limits were placed on the surface curvature of the slat to yield a smooth surface. Approximately 30 optimization cycles were run for each slat. As mentioned above, the method typically required 24-36 hours for a 15 cycle run, thus the total cost ranged from 48-72 hours per slat. Both the drag coefficient and the degree of flow separation on the slat were used to determine when to end the optimization procedure.

Final Analysis

Once the slats were optimized, FUN2D was run to determine the maximum lift at a Mach number of 0.40 as well as the drag at zero lift for a Mach number of 0.65. Figure 7 shows this comparison. Of the seven optimized slats, the 2 degree rotated slat and the 6 degree rotated slat offered advantages over the others. The 2 degree rotated slat had the highest lift at low Mach numbers and reduced the drag at zero lift and high Mach numbers. On the other hand, the 6 degree rotated slat had the lowest drag values at zero lift and high Mach numbers, but the maximum lift at low Mach numbers was sacrificed to achieve this. While it is apparent that further rotating the slat might yield a further reduction in drag, it would also result in a further decrease in the maximum lift.

The better two of the new slat designs along with the baseline are shown in Figure 8. The two slat designs were then analyzed over a range of Mach num-

bers that they would be likely to experience on a rotor. The slat designs were analyzed at Mach numbers of 0.40, 0.60, and 0.80. These Mach numbers represent the potential extremes and the average conditions that the airfoil would see on a rotor.

Mach 0.40

The Mach 0.40 case represents a fairly benign test of the slat aerodynamics. The Reynolds number for these cases was 5×10^6 . At this low Mach number, compressibility is not a factor except at higher angles of attack. As the Figure 9 illustrates, the 2 degree rotated slat yielded the highest c_l at this Mach number. The baseline (configuration 106) performed almost as well. The figure also shows good correlation between the experiment and FUN2D. The 6 degree rotated slat suffers from a reduced c_{lmax} for this Mach number.

Figure 10 shows the drag coefficient for the three airfoils. While FUN2D under predicts the drag at low values of c_l and over predicts it at higher values, the overall trend is captured. This figure shows the effect of the modifications on the drag. The result of the optimization is to reduce the drag near zero lift as well as to flatten out the drag polar in general. The 6 degree rotated slat performs better than the baseline and the 2 degree rotated slat until c_l values of approximately 1.2.

Figure 11 shows the pitching moment coefficient for the three airfoils. The pitching moment coefficient was calculated about the effective quarter chord of the slat/main element system. FUN2D does a good job of predicting the pitching moment for the baseline. Here again, the largest change is seen in the 6 degree rotated slat. For this configuration, the magnitude of the pitching moment is reduced significantly over a wide range of lift.

Figure 12 shows the lift to drag ratio for the three airfoils. Correlation between the experiment and FUN2D is not as good as was seen in previous figures. This is a result of the over prediction of drag and the under prediction of lift by FUN2D. However, this figure shows that the 6 degree rotated slat holds a slight advantage over the 2 degree rotated slat.

Figure 13 shows Mach contours for the three airfoils at 0.75 degrees angle of attack. In these figures, the darker shades represent low Mach numbers, while the lighter colors indicate the highest value of Mach numbers. Figure 13(a) shows the baseline airfoil Mach contours. This figure illustrates the main cause of increased drag at lower angles of attack, namely separation on the lower surface of the slat. Figure 13(b) shows the 2 degree rotated slat Mach contours. Notice that after the slat is rotated and optimized to reduce drag, the amount of separation is seen to decrease substantially.

Figure 13(c) shows the 6 degree rotated slat Mach contours. Reduction of the flow separation was easier to achieve with this slat due to the larger amount of slat rotation.

Mach 0.60

At the higher Mach number of 0.60, the flow field becomes more interesting. Shocks become a significant factor and a prime driver of separation. The Reynolds number for these cases was 7×10^6 .

Figure 14 shows the lift coefficient of the three airfoils. Again, good correlation between FUN2D and experiment is seen for the baseline airfoil. After the slat is rotated and optimized, the maximum c_l is seen to increase. It is interesting to note that the maximum c_l is approximately the same for both the 2 degree and the 6 degree rotated slats. Another effect is the reduction in the non-linearity of the lift curve slope at low angles of attack. The zero lift angle of attack becomes less positive as the slat is rotated upwards and optimized.

Figure 15 shows the drag coefficient for the airfoils. Reasonable correlation is seen between FUN2D and the experiment. This figure also shows the large improvement in the drag that is achieved with the optimized slats at low to moderate values of c_l . In particular, the 6 degree rotated slat performs better than the baseline and the 2 degree rotated slat up to a c_l of approximately 0.75. The drag rise near zero lift is evident in all three airfoils, however it is delayed the longest in the 6 degree rotated slat.

The pitching moment coefficient shown in Figure 16 again shows that the largest change is seen in the 6 degree rotated slat. All the airfoils experience a large variation in pitching moment as lift changes, but over the middle range of lift values, the pitching moment magnitude is reduced for the 6 degree rotated slat.

Figure 17 shows the lift to drag ratio. Again, the maximum L/D is seen to occur for the 6 degree rotated slat, although it does drop off more sharply than the other two airfoils due to the higher drag at higher lift values.

Figure 18 shows the Mach contours for the three airfoils at a Mach number of 0.60 and an angle of attack of 0.75 degrees. Notice that the degree of separation for the baseline slat has increased significantly from the previous Mach number. Increased separation is also seen on the other two slats, but because of the optimization, the amount of separation has been reduced. As previously mentioned, the rotated slats were optimized at a Mach number of 0.65 and an angle of attack of 0.75 degrees. It should also be pointed out that the redesign of the slat has eliminated the shock and separation on the main element that were present in the baseline con-

figuration.

Mach 0.80

The highest Mach number at which the airfoils were analyzed for was 0.80. The Reynolds Number for this Mach number is 9×10^6 . At this Mach number, shocks dominate the flow field. These shocks lead to large regions of separation. The optimization of the slat at this Mach number is very challenging. In general, it was felt that it would not be wise to design a rotor such that the slat would operate at this Mach number. Generally, the portions of the blades exposed to high Mach numbers are not required to generate high values of lift and the penalty paid in the price of drag (torque) would probably not be worth the gain in lift.

Figure 19 shows the lift coefficient at the Mach number of 0.80. The correlation between the experiment and FUN2D for the baseline is not as good as what was seen in the previous cases. FUN2D under predicted the lift for the baseline at higher angles of attack. The effect of the slat rotation is also more dramatic for this case than the previous cases. Both of the rotated slats yield a higher c_{lmax} as well as a steeper lift curve slope.

As could be expected, the drag coefficient shown in Figure 20 is higher than for the previous cases. The drag rise for the baseline case is over predicted with FUN2D, although the drag values at low lift compare well. With that in mind, the improvements with the rotated slats appear to be encouraging. The drag rise for the 6 degree rotated slat is delayed until a c_l of approximately 0.25.

Unlike the previous cases, Figure 21 shows that all three of the airfoils produce essentially the same pitching moment coefficient. Fair correlation between experiment and FUN2D is shown for the baseline airfoil at low values of c_l .

Figure 22 again shows that the 6 degree rotated slat yields a higher lift to drag ratio than the other two airfoils. Again, it should be noted that FUN2D under predicts L/D for the baseline.

Finally, Figure 23 shows the Mach number contours for the three airfoils at a Mach number of 0.80 and an angle of attack of 0.75 degrees. Notice that the separation on the baseline slat extends almost to the main element. There is also a large degree of separation on the 2 degree rotated slat. In this case, a shock is present on the main element which induces the flow field to separate. On the 6 degree rotated slat, the separation is significant; however, the flow remains attached on the main element, thus helping to reduce the overall drag.

Conclusions

An optimization package for two-dimensional airfoils has been used to improve the performance of a multi-element rotorcraft airfoil. The method uses a discrete adjoint formulation to calculate the sensitivity derivatives along with optimization and mesh movement algorithms. The unstructured Navier-Stokes solver, FUN2D, is used in concert with the optimization tools. The results show the ability of the package to reduce the drag on the airfoil by reducing the separation behind the slat at high Mach numbers and low angles of attack.

Acknowledgments

This work was supported, in part, by the National Aeronautics and Space Administration under Contract No. NAS1-96014. The author would also like to thank Dr. Kyle Anderson of the NASA Langley Research Center for providing and assisting with the use of the optimization routines, and to Kevin Noonan of the U.S. Army for his many helpful discussions.

References

1. Campbell, R. L., "An Approach to Constrained Aerodynamic Design with Application to Airfoils," NASA TP 3260, 1992.
2. Smith, L. A., Campbell, R. L., "Applications of a Direct/Iterative Design Method to Complex Transonic Configurations," NASA TP 3234, 1992.
3. Mineck, R. E., Campbell, R. L., Allison, D. O., "Application of Two Procedures for Dual-Point Design of Transonic Airfoils," NASA TP 3466, September, 1994.
4. Narramore, J. C., McCroskey, W.J., Noonan, K. W., "Design and Evaluation of Multi-Element Airfoils for Rotorcraft," presented at the American Helicopter Society 55th Annual Forum, Montreal, Quebec, Canada, May 25-27, 1999.
5. Nielsen, E. J., Anderson, W. K., "Aerodynamic Design Optimization on Unstructured Meshes Using the Navier-Stokes Equations," AIAA Paper 98-4809.
6. Anderson, W. K., Venkatakrishnan, V., "Aerodynamic Design Optimization on Unstructured Grids with a Continuous Adjoint Formulation," AIAA Paper 97-0643.
7. Noonan, K. W., Allison, D. O., Stanaway, S., "Investigation of a Slotted Rotorcraft Airfoil at Mach Numbers from 0.20 to 0.88 at Full-Scale Reynolds Numbers," Presented at the American Helicopter Society Aeromechanics Specialists Conference, San Francisco, CA, January 1994.
8. Noonan, K. W., Yeager, W. T., Singleton, J. D., Wilbur, M. L., Mirick, P. H., "Evaluation of Model Helicopter Main Rotor Blade with Slotted Airfoils at the Tip," presented at the American Helicopter Society 55th Annual Forum, Montreal, Quebec, Canada, May 25-27, 1999.
9. Anderson, W. K., Bonhaus, D. L., "An Implicit Upwind Algorithm for Computing Turbulent Flows on Unstructured Grids," Computers and Fluids, Vol. 23, No. 1, pp 1-21, 1994.
10. Spalart, P. R., Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," AIAA Paper 92-0439, January 1991.
11. Wrenn, G. A., "An Indirect Method for Numerical Optimization Using the Kreisselmeier-Steinhauser Function," NASA CR 4220, March 1999.
12. Marcum, D. L., Weatherill, N. P., "Unstructured Grid Generation Using Iterative Point Insertion and Local Reconnection," AIAA Journal, Vol. 33, No. 9, pp 1619-1625, September 1995.




Figure 1. Splined representation of the airfoil elements prior to the generation of the unstructured grid. Vertical lines on lower surface represent the movement limits for the optimization routine.




Figure 2. Unstructured volume grid.






Figure 3. Grid details at slat.



RC(6)-08



RC(6)-08 Configuration 106

Figure 4. RC(6)-08 airfoils.




Figure 5. A typical drag polar for the RC(6)-08, Configuration 106 airfoil, $M = 0.40$.




Figure 6. Mach number contours for the RC(6)-08 Configuration 106 at $M = 0.6$, $\alpha = 0.75^\circ$, $Re = 7 \times 10^6$.

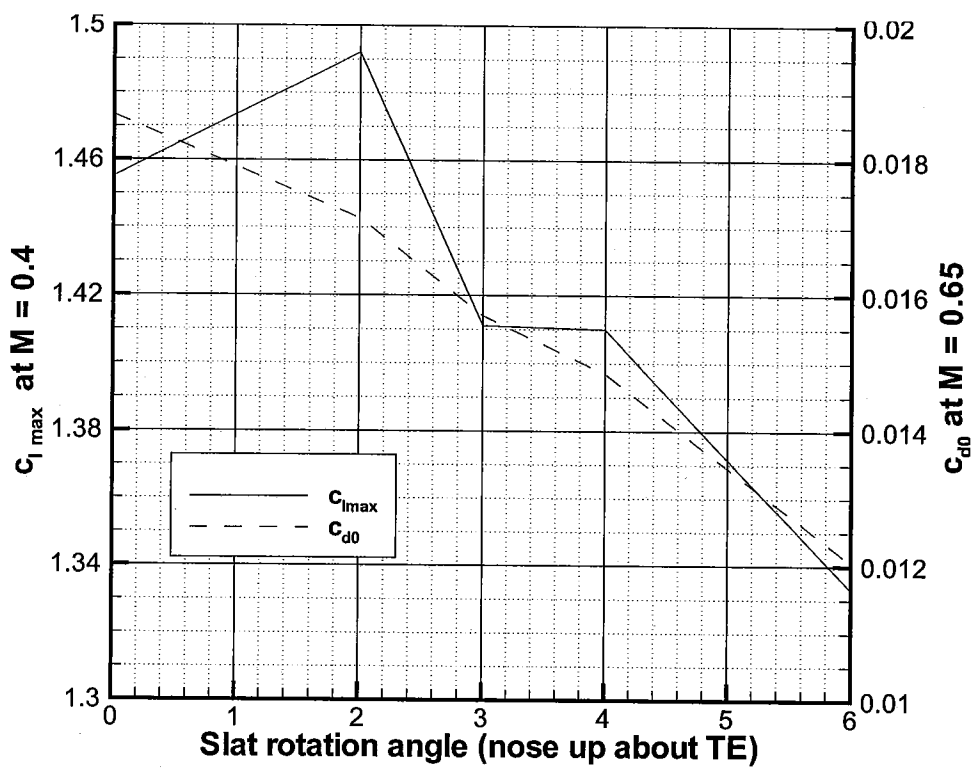


Figure 7. Performance of optimized slats.




Figure 8. Baseline and optimized slats.

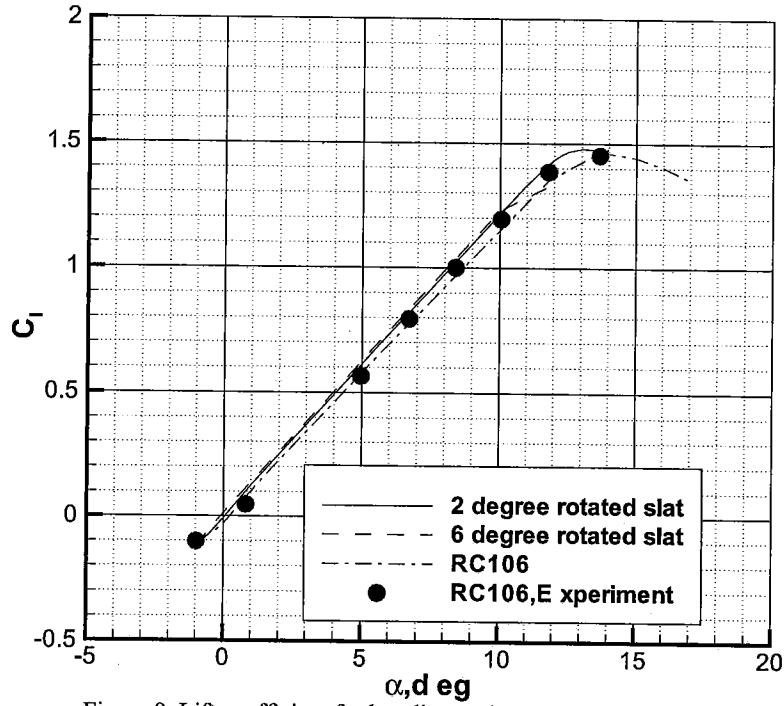


Figure 9. Lift coefficient for baseline and optimized slats at $M = 0.4$.

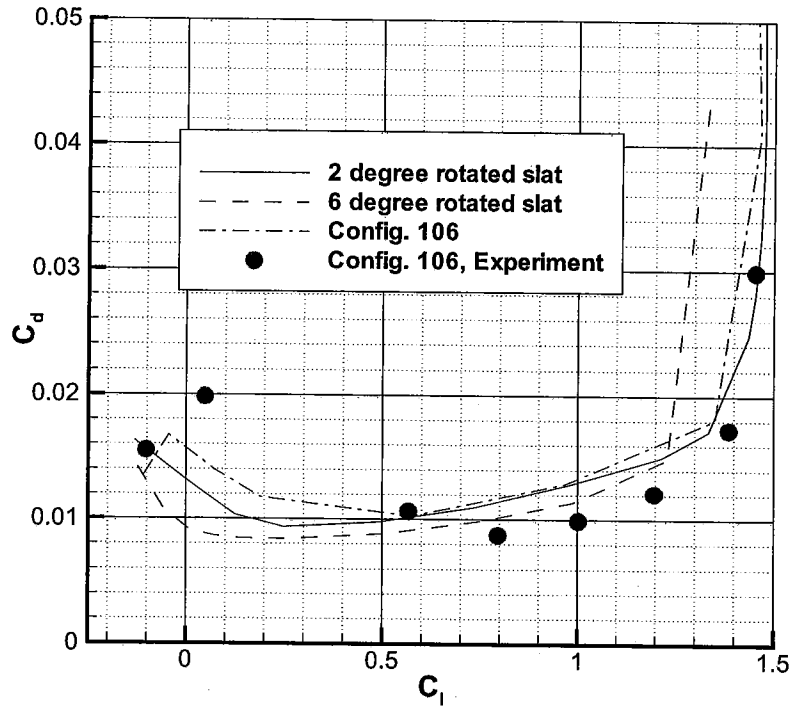


Figure 10. Drag polar for baseline and optimized slats at $M = 0.4$.




Figure 11. Pitching moment coefficient for baseline and optimized slats at $M = 0.4$.




Figure 12. L/D for baseline and optimized slats at $M = 0.4$.

