



AIAA 2003-3957

**Optimization of High-Lift
Configurations Using a
Newton–Krylov Algorithm**

Marian Nemec and David W. Zingg,
*University of Toronto Institute for Aerospace Studies
4925 Dufferin Street, Toronto, ON, M3H 5T6, Canada*

**16th AIAA Computational Fluid Dynamics
Conference**

June 23–26, 2003/Orlando, Florida

For permission to copy or to republish, contact the copyright owner named on the first page.
For AIAA-held copyright, write to AIAA Permissions Department,
1801 Alexander Bell Drive, Suite 500, Reston, VA, 20191-4344.

Optimization of High-Lift Configurations Using a Newton–Krylov Algorithm

Marian Nemeč* and David W. Zingg,†

*University of Toronto Institute for Aerospace Studies
4925 Dufferin Street, Toronto, ON, M3H 5T6, Canada*

A gradient-based Newton–Krylov algorithm for aerodynamic shape optimization is applied to lift maximization of a multi-element landing configuration. The governing flow equations are the two-dimensional compressible Navier–Stokes equations in conjunction with a one-equation transport turbulence model. The objective function gradient is computed via the discrete-adjoint method. The design examples reveal a number of difficulties for numerical optimization methods when applied to high-lift design problems, which include flow solver convergence problems at stall and post-stall conditions and poor off-design performance. Strategies are presented for addressing these difficulties, and an example is provided to demonstrate the approach.

Introduction

THE principal challenge in the design of high-lift systems is the complex nature of the flow resulting from the use of configurations with multiple elements, such as flaps and slats, at high angles of attack. The dominant flow features include regions of separated flow, confluent boundary layers and wakes, and regions of supercritical flow.¹ Such flow features have a strong influence on the aerodynamic performance of the configuration due to their nonlinear effects, which must be carefully controlled in order to realize optimal performance. In addition, the design problem inherently involves multiple operating conditions, namely take-off, cruise, and landing, and is further complicated by strict geometry constraints.

Although the design of high-lift systems is a complex task, an optimized configuration can significantly improve the aerodynamic performance of an aircraft, as well as provide weight savings and reductions in mechanical complexity.² This has motivated the development of effective design strategies based on numerical optimization methods, where the selection of an optimal configuration is accomplished by a systematic and potentially fast evaluation of candidate configurations.^{3–10} In particular, a promising approach is offered by gradient-based methods,^{4,6,9,10} since significant performance improvements can be obtained in relatively few evaluations of the objective function and gradient. The validity of the optimization ultimately depends on the accuracy of the flow solver. Current algorithms for the solution of the Reynolds-averaged Navier–Stokes equations provide reasonable

predictions of aerodynamic performance for complex airfoil geometries up to stall conditions.¹

In Refs. 11–13, we presented an accurate and efficient Newton–Krylov algorithm for the optimization of single- and multi-element configurations, where the objective function gradient is computed via the discrete-adjoint approach. The flow is governed by the two-dimensional compressible thin-layer Navier–Stokes equations in conjunction with the Spalart–Allmaras turbulence model.¹⁴ The objectives of this paper are to demonstrate the performance of the Newton–Krylov algorithm for lift maximization of multi-element configurations and to identify and address difficulties that can be encountered in a practical aerodynamic design context.

Problem Formulation

The aerodynamic shape optimization problem consists of determining values of design variables X , such that the objective function \mathcal{J} is minimized

$$\min_X \mathcal{J}(X, Q) \quad (1)$$

subject to constraint equations C_j :

$$C_j(X, Q) \leq 0 \quad j = 1, \dots, N_c \quad (2)$$

where the vector Q denotes the conservative flowfield variables and N_c denotes the number of constraint equations. The flowfield variables are forced to satisfy the governing flowfield equations, \mathcal{F} , within a feasible region of the design space Ω :

$$\mathcal{F}(X, Q) = 0 \quad \forall X \in \Omega \quad (3)$$

which implicitly defines $Q = f(X)$.

For the lift-maximization problem under consideration here, the objective function is given by

$$\mathcal{J} = \left(1 - \frac{C_L}{C_L^*}\right)^2 \quad (4)$$

*NRC Research Associate, presently at NASA Ames Research Center, nemeč@nas.nasa.gov

†Professor, Senior Member AIAA, <http://goldfinger.utias.utoronto.ca/~dwz/>

Copyright © 2003 by M. Nemeč and D. W. Zingg. Published by the American Institute of Aeronautics and Astronautics, Inc. with permission.

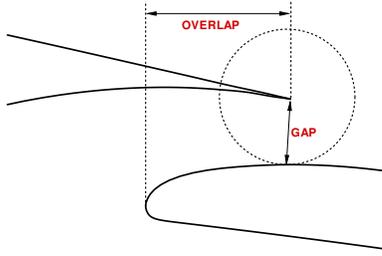


Fig. 1 Definition of gap and overlap distances

where C_L^* represents the target lift coefficient. Note that the lift-maximization problem is formulated by setting C_L^* to an unattainable value. In some problems, the robustness and convergence of the optimization can be improved by considering a sequence of lift-enhancement sub-problems with increasing values of the target lift coefficient. The weighted-sum method is used for multi-point optimization problems:

$$\mathcal{J}_m = \sum_{i=1}^{N_m} w_i \mathcal{J}_i \quad (5)$$

where N_m denotes the number of design points (typically Mach numbers or angles of attack), and w_i represents a user-assigned weight for each design point.

The design variables are the horizontal and vertical locations of slats and flaps, as well as the angle of attack of the configuration. The vertical and horizontal design variables are used to control the gap and overlap distances. We define the gap as the minimum distance between the trailing edge of the upstream element and the surface of the downstream element. The overlap is defined as the difference between the x -locations of the upstream element's trailing edge and the downstream element's leading edge, see Fig. 1. Although not considered in this work, additional design variables may include the slat and flap deflection angles, and the shape of the configuration can be controlled by the use of a B-spline parameterization¹⁵ of each airfoil surface, see Refs. 11–13 for examples.

The constraint equations, Eq. 2, are used to enforce upper- and lower-bound limits for the gap and overlap distances. These constraints are necessary in order to ensure a reasonable computational grid and are typically inactive at the optimal solution.

The governing flow equations are the compressible two-dimensional thin-layer Navier–Stokes equations in generalized coordinates:

$$\frac{\partial \hat{E}(X, \hat{Q})}{\partial \xi} + \frac{\partial \hat{F}(X, \hat{Q})}{\partial \eta} = Re^{-1} \frac{\partial \hat{S}(X, \hat{Q})}{\partial \eta} \quad (6)$$

where $\hat{Q} = J^{-1}Q = J^{-1}[\rho, \rho u, \rho v, e]^T$ is the vector of conservative dependent state variables, ξ and η are the streamwise and normal generalized coordinates, respectively, and J is the Jacobian of the coordinate

transformation from Cartesian coordinates. Vectors \hat{E} and \hat{F} represent the inviscid flux vectors, the viscous flux vector is given by \hat{S} , and Re denotes the Reynolds number. Sutherland's law is used to determine the laminar viscosity. The equations are in non-dimensional form. For further details, see Ref. 16. The turbulent viscosity is modeled with the Spalart–Allmaras turbulence model.¹⁴ All cases considered in this study are assumed to be fully turbulent, and therefore, the laminar-turbulent trip terms are not used.

Numerical Method

The aerodynamic shape optimization problem defined by Eqs. 1–3 is cast as an unconstrained problem. This is accomplished by lifting the side constraints, Eq. 2, into the objective function \mathcal{J} using a penalty method. Furthermore, the constraint imposed by the flowfield equations, Eq. 3, is satisfied at every point within the feasible design space, and consequently these equations do not explicitly appear in the formulation of the optimization problem.

The unconstrained problem is solved using the BFGS quasi-Newton method in conjunction with a backtracking line search.^{11,17} At each step of the line search, the objective function value and gradient are required in order to construct a local cubic interpolant. Note that the optimization problem is based on the discrete form of the flowfield equations. Using the discrete approach, we expect the gradient to vanish at the optimum solution. In the following sections, we present the formulation for the penalized objective function, as well as the algorithms used for the flowfield evaluation (objective function value), the gradient evaluation, and the grid-perturbation strategy.

Objective with Constraints

A penalty method is used to lift the constraints into the objective function. We re-define the objective given by Eq. 4 as

$$\mathcal{J} = \mathcal{J}_d + \mathcal{J}_p \quad (7)$$

where \mathcal{J}_d denotes the design objective of lift maximization, and \mathcal{J}_p denotes constraint penalty terms which are cast using a quadratic formulation. For example, the formulation for the upper-gap constraint (G_u) is

$$\mathcal{J}_p = \begin{cases} [1 - G_u/G_u^*]^2 & \text{if } G_u > G_u^* \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where G_u^* denotes the upper allowable limit.

Flowfield Evaluation

The spatial discretization of the flowfield equations, Eq. 6, is the same as that used in ARC2D¹⁶ and TOR-NADO¹⁸ for multi-block H-topology grids. The discretization consists of second-order centered-difference operators with second- and fourth-difference scalar artificial dissipation. The Spalart–Allmaras turbulence

model is discretized as described in Refs. 14 and 19. Overall, the spatial discretization leads to a nonlinear system of equations:

$$R(X, \widehat{\mathcal{Q}}) = 0 \quad (9)$$

where $\widehat{\mathcal{Q}}$ denotes the discrete vector of conservative dependent flow variables including the turbulence model variable $\tilde{\nu}$. Hence, at each node (j, k) within the computational domain $\widehat{\mathcal{Q}}_{j,k} = (J^{-1}\mathcal{Q})_{j,k} = J^{-1}[\rho, \rho u, \rho v, e, \tilde{\nu}]_{j,k}^T$. The turbulence model equation is scaled by J^{-1} . On multi-block grids, the block interfaces are overlapped in the streamwise direction and averaged in the normal direction. Two columns of halo points are used at the streamwise block interfaces. Although R is written as a function of the design variables, we emphasize that during a flowfield solution the design variables, and consequently the computational grid, are constants. It is important to note that TORNADO has been carefully validated for the analysis of high-lift, turbulent aerodynamic flows.^{19–21} In particular, one of the cases studied by Godin *et al.*¹⁹ showed that this flow solver is able to predict trends in the changes of the lift coefficient values due to variations in flap position.

Eq. 9 is solved in a fully-coupled manner, where convergence to steady state is achieved using the preconditioned GMRES algorithm in conjunction with an inexact-Newton strategy.^{11,22} The main components include matrix-free GMRES(40) and a block-fill incomplete LU (BFILU) preconditioner. The matrix-vector products required at each GMRES iteration are formed with first-order finite-differences. Right preconditioning is used, and the preconditioner is based on an approximate-flow-Jacobian matrix. The level of fill for most cases is 2 [BFILU(2)], but difficult multi-element cases may require BFILU(4). The approximate-factorization algorithm of ARC2D in diagonal form^{16,18} in conjunction with a subiteration scheme¹⁴ for the turbulence model equation is used to reduce the initial residual by three orders of magnitude in order to avoid Newton startup problems.

The approximate flow-Jacobian used for the preconditioner is identical to the flow-Jacobian matrix, $\partial R/\partial \widehat{\mathcal{Q}}$, except for the treatment of the artificial-dissipation coefficients.¹¹ Hence, the preconditioner contains the contributions from all components of the residual vector, namely inviscid and viscous fluxes, boundary conditions, block interfaces, and the turbulence model. The artificial-dissipation coefficients, which include the spectral radius and the pressure switch, are assumed to be constant with respect to the flowfield variables. Furthermore, the preconditioning matrix is formed with only second-difference dissipation, but the second-difference coefficient is combined with the fourth-difference coefficient as follows,

$$d_l^{(2)} = d_r^{(2)} + \phi d_r^{(4)} \quad (10)$$

where the subscript r denotes the contribution from the right-hand side, and the subscript l denotes the resulting left-hand side value used in forming the preconditioner. This modification does not affect the steady-state solution or the convergence of the Newton iterations. Fast convergence of GMRES is obtained with the value of ϕ set to 6.0, which has been determined through numerical experiments.

Eq. 10 improves the diagonal dominance of the preconditioning matrix and reduces the work and storage requirements of the incomplete factorization. This approach is similar to the ‘diagonal shift’ strategy suggested by Chow and Saad.²³ The present preconditioning matrix is a compromise between a preconditioner based on a first-order upwind discretization of the flowfield equations and a preconditioner based on the actual second-order discretization. This novel ‘intermediate’ preconditioner is significantly more effective than either of these more commonly used approaches.

Gradient Evaluation

Using the discrete-adjoint method, the expression for the gradient, \mathcal{G} , of the objective function, $\mathcal{J}[X, \mathcal{Q}(X)]$, is given by

$$\mathcal{G} = \frac{d\mathcal{J}}{dX} = \frac{\partial \mathcal{J}}{\partial X} - \psi^T \frac{\partial R}{\partial X} \quad (11)$$

where we reduce the vector of design variables, X , to a scalar in order to clearly distinguish between partial and total derivatives. For problems with multiple design variables, it may be helpful to note that \mathcal{G} and $\partial \mathcal{J}/\partial X$ are $[1 \times N_D]$ row vectors, ψ is a $[N_F \times 1]$ column vector, and $\partial R/\partial X$ is a $[N_F \times N_D]$ matrix, where N_D represents the number of design variables and N_F represents the number of flowfield variables. We assume that the implicit function $\mathcal{Q}(X)$ defined by Eq. 9 is sufficiently smooth even in the presence of flow discontinuities such as shock waves.^{24–26}

The vector ψ represents adjoint variables, which are given by the adjoint equation:

$$\frac{\partial R}{\partial \mathcal{Q}}^T \psi = \frac{\partial \mathcal{J}}{\partial \mathcal{Q}}^T \quad (12)$$

This is a large, sparse, linear system of equations that is independent of the design variables. The GMRES strategy from the flow solver is adopted to solve the adjoint equation. Fast solutions are obtained with GMRES(85) and for the preconditioning matrix we use BFILU(6) and $\phi = 3.0$. Multi-element airfoil cases with complex flowfields may require $\phi = 6.0$, which improves the robustness of the adjoint solver. Due to the transpose on the left-hand-side of Eq. 12, the matrix-free approach used in the flow solver is not possible for the adjoint equation. The flow-Jacobian matrix is stored explicitly, where we include the contribution from the spectral radius, but we treat the pressure

switch associated with the artificial dissipation scheme as a constant.

The remaining terms in Eqs. 11 and 12, namely the objective function sensitivities $\partial\mathcal{J}/\partial X$ and $\partial\mathcal{J}/\partial Q$, as well as the residual sensitivity $\partial R/\partial X$, are evaluated using centered differences. The use of centered differences for the evaluation of the partial derivative terms is not computationally expensive. For example, the centered-difference formula for the residual sensitivities is given by

$$\frac{\partial R}{\partial X_i} = \frac{R(X + he_i, Q) - R(X - he_i, Q)}{2h} \quad (13)$$

where

$$h = \max(\epsilon \cdot |X_i|, 1 \times 10^{-8}) \quad (14)$$

and $i = 1, \dots, N_D$. The i th unit vector is denoted by e_i , and a typical value of ϵ is 1×10^{-5} . It is important to realize that Eq. 13 involves two evaluations of only the residual vector per design variable and *not* two flow-field solutions. Note that the evaluation of residual sensitivities includes the evaluation of grid sensitivities, since the design variables do not explicitly appear in the residual equations except for the angle of attack design variable.

Grid-Movement Strategy

As the position of a flap or a slat changes during the optimization process, the location of the grid nodes is adjusted from the baseline configuration to conform to the new configuration. In Ref. 11, we use an algebraic grid-perturbation strategy that preserves the distance to the outer boundary and relocates the grid nodes in the normal direction proportional to the distance from the airfoil boundary. When the optimization problem involves the horizontal and vertical translation of a slat or a flap, the use of this strategy can result in significantly skewed grid cells near the boundary.

In order to improve the quality of the modified multi-block grids, we present a new grid-perturbation strategy given by

$$y_k^{\text{new}} = y_k^{\text{old}} + \frac{\Delta y}{2} [1 + \cos(\pi S_k)] \quad (15)$$

where Δy represents the airfoil shape change. S_k is the normalized arclength distance given by

$$S_k = \frac{1}{L_g} \sum_{i=2}^k L_i \quad k = 2, \dots, k_{\text{max}} - 1 \quad (16)$$

where $S_1 = 0$, L_i is the length of a segment between nodes k and $k - 1$, and L_g is the grid-line length from the body to the outer boundary.

An example is presented in Fig. 2, which shows the grid near the trailing edge of a main element and the leading edge of a flap for a two-element configuration. The flap is re-positioned such that the gap and

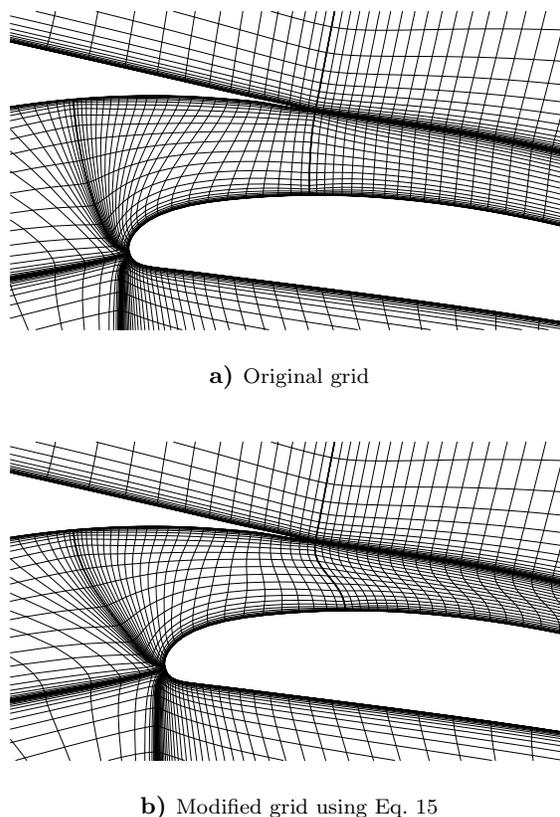


Fig. 2 Grid distortion resulting from a flap displacement

overlap are reduced by roughly $0.4\%c$ and $1.0\%c$, respectively. Note that only blocks in the immediate neighbourhood of the flap are modified, and that the orthogonality of grid lines near the flap is maintained. For additional details see Ref. 27, where a numerical study is performed that compares the influence of grid-perturbation strategies on the values of the resulting aerodynamic coefficients.

Results and Discussion

The optimization problem under consideration consists of increasing the maximum lift coefficient of a two-element configuration based on the NACA 23016 airfoil. The flap deflection angle is 30° , which is representative of a landing setting and remains constant throughout the optimization. The initial gap and overlap distances are $2.2\%c$ and $2.7\%c$, respectively. These settings were selected without the aid of formal optimization techniques. The freestream conditions are $M_\infty = 0.15$ and $Re = 5.39 \times 10^6$. The flow is assumed to be fully turbulent. The main flow features include a substantial region of separation on the upper surface of the flap and a region of recirculating flow in the cove of the main element. The presence of these features is essentially independent of the angle of attack. A multi-block H-topology structured grid is used with roughly 50,000 nodes and an off-wall spacing of

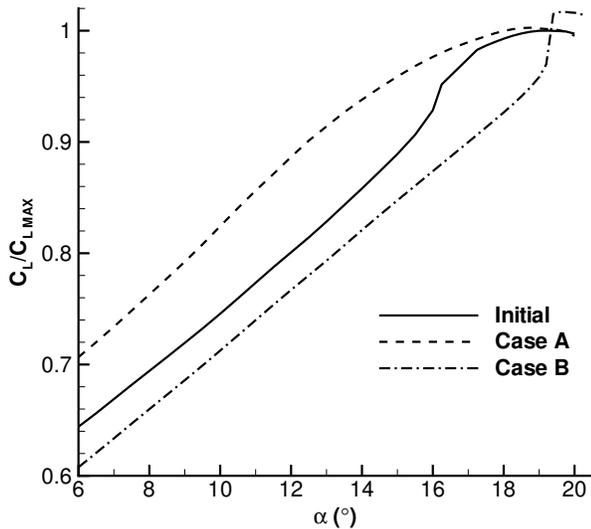


Fig. 3 C_L as a function of the angle of attack (normalized by the maximum C_L of the initial configuration)

$2 \times 10^{-6}c$. The present grid can be expected to produce low numerical errors in lift, but numerical errors in drag could be significant.²⁸

The primary difficulty of high-lift design is the need to obtain sufficiently converged flow solutions at high angles of attack. As the configuration evolves during the optimization, large changes in the geometry can cause flow separation and stall, which can result in inadequate convergence and even divergence of the flow solver. In order to avoid flow-solver convergence problems, we first perform the lift maximization at a moderate and fixed angle of attack, and then examine the impact of this optimization on the value of the maximum lift coefficient:

Case A Lift maximization at a fixed angle of attack with two design variables, namely the horizontal and vertical translation of the flap. The angle of attack (α) is set to 12° , and the target lift coefficient, C_L^* in Eq. 4, is set to a value that represents a roughly 11% increase from the lift coefficient of the initial configuration at $\alpha = 12^\circ$.

Case A converges to optimal gap and overlap values of $1.7\%c$ and $1.5\%c$, respectively. The lift coefficient for a range of angles of attack is shown in Figure 3. A significant lift increase is obtained at the design condition of $\alpha = 12^\circ$. The increase in lift is due to flow reattachment on the upper surface of the flap. The improvement in the maximum lift coefficient, however, is negligible.

In order to increase the maximum lift coefficient, we consider the following optimization problem:

Case B Lift maximization with the angle of attack as a design variable, in addition to the two translation design variables. The target lift coefficient,

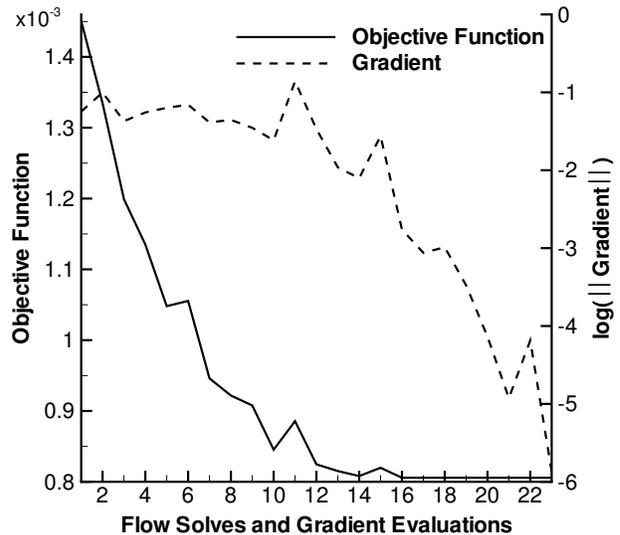


Fig. 4 Optimization convergence history for Case B

C_L^* in Eq. 4, is set to a roughly 5% increase of the maximum lift coefficient of the initial configuration.

Note that the optimization is performed very close to stall conditions, and therefore flow-solver convergence difficulties are likely to occur. These difficulties are avoided by the use of the following strategy. Since the flow solver will not converge from freestream conditions to conditions near maximum lift of the initial configuration, the flow is first solved at a moderate angle of attack and then re-converged with gradually increasing angles of attack until maximum lift is approached. This solution is then used as the starting guess for the optimization. Furthermore, for design-variable updates that result in poor convergence of the flow solver, the update stepsize is reduced by a factor of two, and the previous converged solution is used to “warm-start” the flow solver. Using this strategy, the flow equations are converged at least eight orders of magnitude for all examples presented in this work.

Figure 4 shows the optimization convergence history for case B. The gradient is reduced by five orders of magnitude, indicating that at least a local optimum has been attained. The objective function converges to a non-zero value, which indicates that the target maximum lift coefficient could not be achieved. The optimization convergence history for case A is similar.

The lift coefficient for a range of angles of attack is shown in Figure 3. Case B converges to a relatively large gap of $2.85\%c$, and a small overlap of $0.7\%c$. As shown in Figure 3, this configuration results in a 1.7% increase in the maximum lift coefficient at a slightly higher angle of attack when compared with the initial configuration. The lift curve shows an abrupt increase in slope just prior to reaching maximum lift, which is

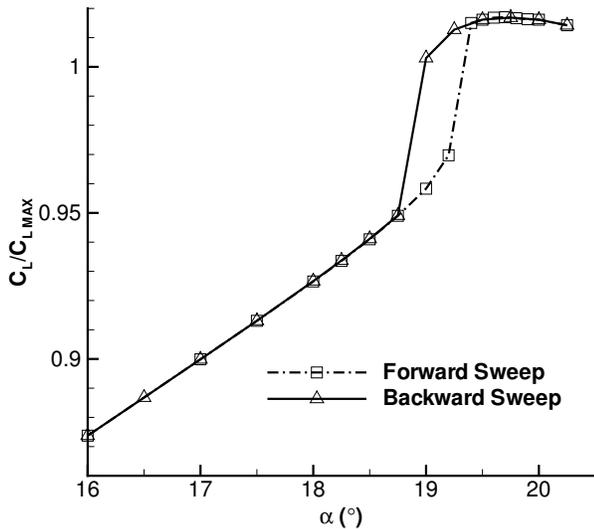


Fig. 5 Hysteresis near maximum lift for case B

due to reattachment of the flow on the upper surface of the flap. This highly nonlinear behaviour, which is also present to some degree in the initial configuration, could lead to undesirable aircraft handling qualities. Furthermore, Figure 5 reveals a hysteresis effect in the region of the abrupt lift change. One flow solution is obtained as the angle of attack is increased (labelled as forward sweep) and a different flow solution is obtained as the angle of attack is decreased (labelled as backward sweep). Note that the flow solver is warm-started from the previous solution during an angle-of-attack sweep.

The results presented in Figure 3 for case B demonstrate the typical difficulties encountered in optimization problems that consider only a single operating condition. Studies by Drela,²⁹ Li *et al.*,³⁰ and Huyse *et al.*³¹ reveal similar difficulties in the design of airfoils for transonic flow conditions. For the present problem, we seek a compromise between the improvement in maximum lift, i.e. case B, and the well-behaved lift curve that was obtained for case A. Consequently, the following two-point optimization problem is considered:

Case C Weighted sum of cases A and B.

The objective function is given by Eq. 5, where the weights w_1 and w_2 are set to 0.125 and 0.875, respectively. Weight w_1 is associated with the objective based on case A, while w_2 is associated with the objective based on case B. The selection of these weights required a few trial optimizations. It is important to note that only the flap translation design variables are used for the first design condition, while the angle of attack and the flap translation design variables are used for the second design condition.

For case C, the optimization converges to gap and overlap distances of 1.87% c and 1.29% c , respectively,

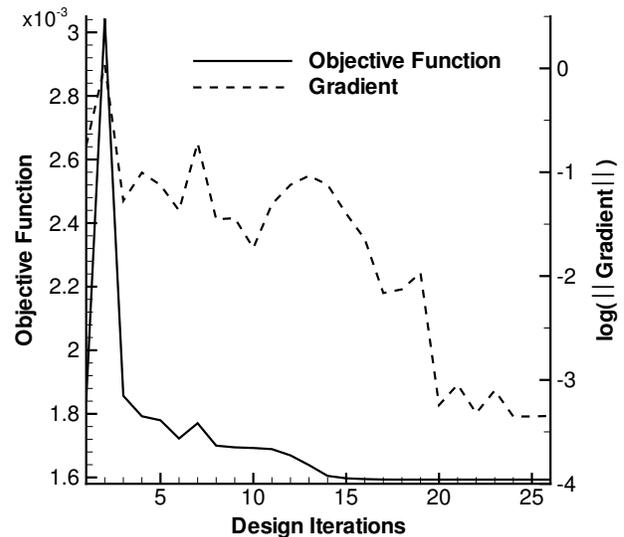


Fig. 6 Optimization convergence history for case C

in 26 design iterations. The gradient is reduced by 3.5 orders of magnitude, as shown in Fig. 6. Since this is a two-point optimization problem, each design iteration corresponds to two flow and gradient evaluations. Figure 7 shows the lift curve for the optimal configuration. The maximum lift coefficient is increased by 0.8% over the initial configuration. Furthermore, the undesirable behaviour of the lift coefficient shown in Figure 3 for case B and the initial configuration is eliminated, resulting in a smooth lift curve and a significantly higher lift coefficient throughout the angle-of-attack range. Pressure distributions for the initial and case C configurations at low, moderate, and high angles of attack are shown in Fig. 9. The region of separated flow on the flap is reduced, which results in a higher lift and lower drag for the optimal configuration. The pitching moment for the optimal configuration is moderately increased due to the aft movement of the flap, as shown in Fig. 8; however, the variation in the moment coefficient is reduced due to the smoother lift curve.

Conclusions

The results demonstrate that gradient-based methods can provide an effective strategy for high-lift design problems. Key reasons for the effectiveness of the present Newton–Krylov algorithm are the use of a warm-start procedure for the flow solver, where the design-variable updates are progressively reduced to ensure convergence near stall conditions, and a careful formulation of the objective function that includes multiple operating conditions. Future work should concentrate on the assessment of the physical models, including the treatment of laminar-turbulent transition, in order to establish the validity of the optimal designs.

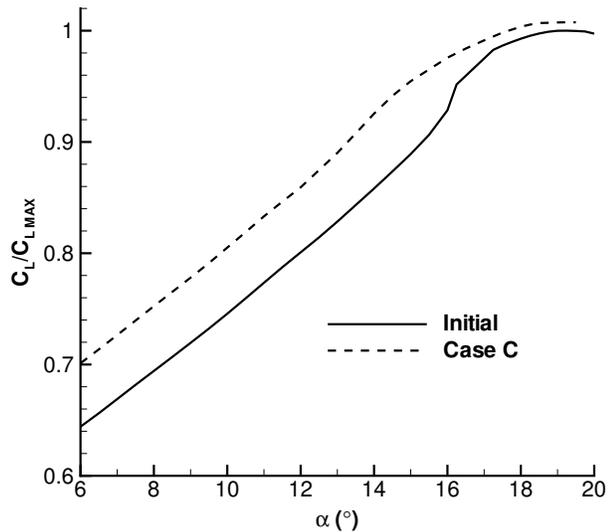


Fig. 7 C_L as a function of the angle of attack for case C (normalized by the maximum C_L of the initial configuration)

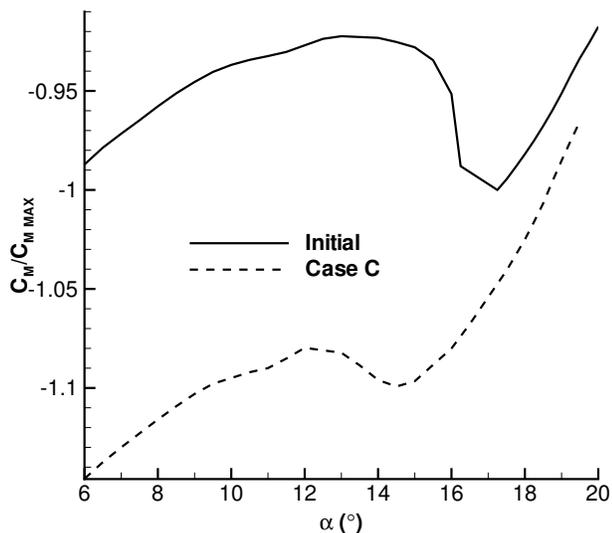


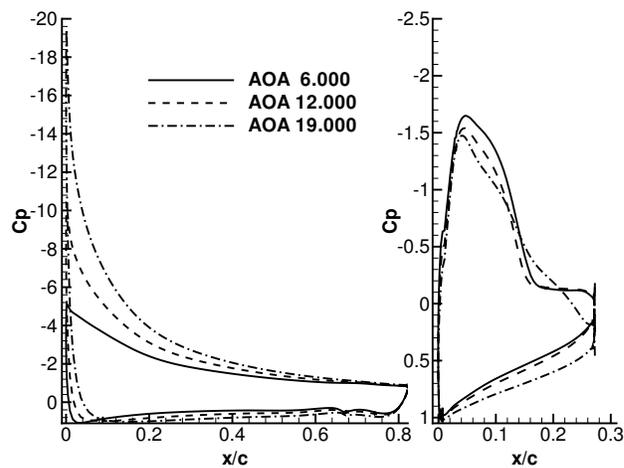
Fig. 8 C_M as a function of the angle of attack for case C (normalized by the maximum $|C_M|$ of the initial configuration)

Acknowledgments

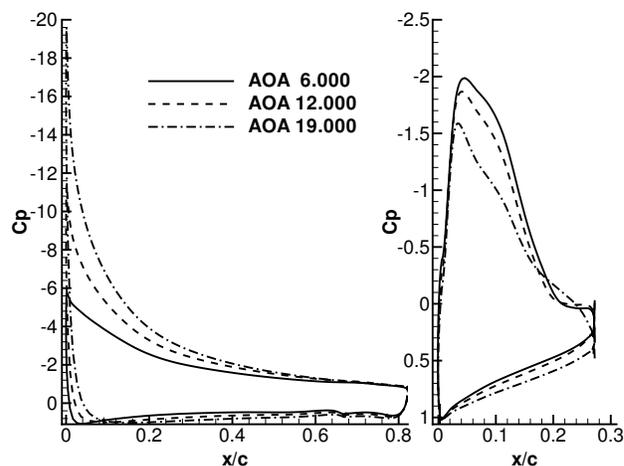
This research was supported by the Natural Sciences and Engineering Research Council of Canada, Bombardier Aerospace, Pratt & Whitney Canada, and an OGS grant from the Government of Ontario.

References

- ¹Rumsey, C. L. and Ying, S. X., "Prediction of High Lift: Review of Present CFD Capability," *Progress in Aerospace Sciences*, Vol. 38, No. 2, 2002, pp. 145–180.
- ²van Dam, C. P., "The Aerodynamic Design of Multi-Element High-Lift Systems for Transport Airplanes," *Progress in Aerospace Sciences*, Vol. 38, No. 2, 2002, pp. 101–144.
- ³Drela, M., "Design and Optimization Method for Multi-Element Airfoils," AIAA Paper 93-0969, 1993.



a) Initial configuration



b) Case C configuration

Fig. 9 Pressure distributions for main element and flap

⁴Eyi, S., Lee, K. D., Rogers, S. E., and Kwak, D., "High-Lift Design Optimization Using Navier–Stokes Equations," *Journal of Aircraft*, Vol. 33, No. 3, 1996, pp. 499–504.

⁵Greenman, R. M. and Roth, K. R., "Minimizing Computational Data Requirements for Multi-Element Airfoils Using Neural Networks," *Journal of Aircraft*, Vol. 36, No. 5, 1999, pp. 777–784.

⁶Alexandrov, N. M., Nielsen, E. J., Lewis, R. M., and Anderson, W. K., "First-Order Model Management with Variable-Fidelity Physics Applied to Multi-Element Airfoil Optimization," AIAA Paper 2000-4886, September 2000.

⁷Duvigneau, R. and Visonneau, M., "Shape Optimization of Incompressible and Turbulent Flows Using The Simplex Method," AIAA Paper 2001-2533, Anaheim, CA, Sept. 2001.

⁸Wild, J., "Validation of Numerical Optimization Method of High-Lift Multi-Element Airfoils based on Navier–Stokes–Equations," AIAA Paper 2002-2939, 2002.

⁹Nielsen, E. J. and Anderson, W. K., "Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes," *AIAA Journal*, Vol. 40, No. 6, 2002, pp. 1155–1163.

¹⁰Kim, S., Alonso, J. J., and Jameson, A., "Design Optimization of High-Lift Configurations Using a Viscous Continuous Adjoint Method," AIAA Paper 2002-0844, Jan. 2002.

- ¹¹Nemec, M. and Zingg, D. W., "Newton–Krylov Algorithm for Aerodynamic Design Using the Navier–Stokes Equations," *AIAA Journal*, Vol. 40, No. 6, 2002, pp. 1146–1154.
- ¹²Nemec, M., Zingg, D. W., and Pulliam, T. H., "Multi-Point and Multi-Objective Aerodynamic Shape Optimization," AIAA Paper 2002–5548, Atlanta, GA, Sept. 2002.
- ¹³Nemec, M. and Zingg, D. W., "From Analysis to Design of High-Lift Configurations Using a Newton–Krylov Algorithm," Paper 173, ICAS 2002, Toronto, ON, Sept. 2002.
- ¹⁴Spalart, P. R. and Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," AIAA Paper 92–0439, Jan. 1992.
- ¹⁵Farin, G. E., *Curves and Surfaces for Computer-Aided Geometric Design: A Practical Guide*, Academic Press, San Diego, 4th ed., 1997.
- ¹⁶Pulliam, T. H., "Efficient Solution Methods for the Navier–Stokes Equations," Tech. rep., Lecture Notes for the von Kármán Inst. for Fluid Dynamics Lecture Series: Numerical Techniques for Viscous Flow Computation in Turbomachinery Buildings, Brussels, Belgium, Jan. 1986.
- ¹⁷Nocedal, J. and Wright, S. J., *Numerical Optimization*, Springer–Verlag, New York, 1999.
- ¹⁸Nelson, T. E., Zingg, D. W., and Johnston, G. W., "Compressible Navier–Stokes Computations of Multielement Airfoil Flows Using Multiblock Grids," *AIAA Journal*, Vol. 32, No. 3, 1994, pp. 506–511.
- ¹⁹Godin, P., Zingg, D. W., and Nelson, T. E., "High-Lift Aerodynamic Computations with One- and Two-Equation Turbulence Models," *AIAA Journal*, Vol. 35, No. 2, 1997, pp. 237–243.
- ²⁰Nelson, T. E., Godin, P., De Rango, S., and Zingg, D. W., "Flow Computations for a Three-Element Airfoil System," *Canadian Aeronautics and Space Journal*, Vol. 45, No. 2, 1999, pp. 132–139.
- ²¹Fejtek, I., "Summary of Code Validation Results For a Multiple Element Airfoil Test Case," AIAA Paper 97–1932, June 1997.
- ²²Pueyo, A. and Zingg, D. W., "Efficient Newton–Krylov Solver for Aerodynamic Computations," *AIAA Journal*, Vol. 36, No. 11, 1998, pp. 1991–1997.
- ²³Chow, E. and Saad, Y., "Experimental Study of ILU Preconditioners for Indefinite Matrices," *Journal of Computational and Applied Mathematics*, Vol. 86, 1997, pp. 387–414.
- ²⁴Gunzburger, M. D., "Introduction into Mathematical Aspects of Flow Control and Optimization," *Inverse Design and Optimization Methods*, 1997–05, von Karman Institute For Fluid Dynamics, Brussels, Belgium, April 1997.
- ²⁵Giles, M. B. and Pierce, N. A., "An Introduction to the Adjoint Approach to Design," *Flow, Turbulence and Combustion*, Vol. 65, No. 3/4, 2000, pp. 393–415.
- ²⁶Giles, M. B. and Pierce, N. A., "Analytic Adjoint Solutions for the Quasi-One-Dimensional Euler Equations," *Journal of Fluid Mechanics*, Vol. 426, 2001, pp. 327–345.
- ²⁷Nemec, M., *Optimal Shape Design of Aerodynamic Configurations: A Newton–Krylov Approach*, Ph.D. thesis, University of Toronto, 2003, <http://oddjob.utias.utoronto.ca/marian/>.
- ²⁸De Rango, S. and Zingg, D. W., "Higher-Order Spatial Discretization for Turbulent Aerodynamic Computations," *AIAA Journal*, Vol. 39, No. 7, 2001, pp. 1296–1304.
- ²⁹Drela, M., "Pros & Cons of Airfoil Optimization," *Frontiers of Computational Fluid Dynamics 1998*, edited by D. A. Caughey and M. M. Hafez, World Scientific, Singapore, 1998, pp. 363–381.
- ³⁰Li, W., Huyse, L., and Padula, S., "Robust Airfoil Optimization to Achieve Drag Reduction Over a Range of Mach Numbers," *Structural and Multidisciplinary Optimization*, Vol. 24, No. 1, 2002, pp. 38–50.
- ³¹Huyse, L., Padula, S., Lewis, R. M., and Li, W., "Probabilistic Approach to Free-Form Airfoil Shape Optimization Under Uncertainty," *AIAA Journal*, Vol. 40, No. 9, 2002, pp. 1764–1772.