On Aerodynamic Optimization Under a Range of Operating Conditions

David W. Zingg,* and Samy Elias†
Institute for Aerospace Studies, University of Toronto
4925 Dufferin St., Toronto, Ontario M3H 5T6, Canada

In aerodynamic design, good performance is generally required under a range of operating conditions, including off-design conditions. This can be achieved through multipoint optimization. The desired performance objective and operating conditions must be specified, and the resulting optimization problem must be solved in such a manner that the desired performance is achieved. Issues involved in formulating multipoint optimization problems are discussed. A technique is proposed for automatically choosing sampling points within the operating range and their weights in order to obtain the desired performance over the range of operating conditions. Examples are given involving lift-constrained drag minimization over a range of Mach numbers. Trade-offs and their implications for the formulation of multipoint problems are presented and discussed.

I. Introduction

Several algorithms have been developed that can efficiently perform aerodynamic shape optimization.[1–6] The designer specifies an objective, operating conditions, constraints, and a set of parameters that define the range of possible geometries. The optimization algorithm finds the values of the geometric parameters that minimize the objective function while satisfying the constraints. Numerical optimization has four primary advantages over the traditional approach to aerodynamic design, i.e. cut and try driven by the designer's expertise:

1. potentially faster;
2. more likely to achieve a truly optimal design;
3. forces the designer to specify the design problem, i.e. the objective, operating conditions, constraints, and geometric parameters, carefully and completely;
4. provides insight into the nature of the design space and the trade-offs between various competing objectives and operating points.

Virtually all aerodynamic components must perform efficiently over a range of operating conditions. Optimization at a single operating point invariably leads to poor off-design performance. Therefore, the optimization problem must be posed such that a range of operating conditions and off-design performance requirements are included in either the objective function or the constraints. The outcome of the optimization is greatly dependent on the details of how this is done, and consequently the designer must give careful consideration to precisely what is required, consistent with item 3 in the list above. For example, higher expectations of off-design performance typically compromise on-design performance. Fortunately, numerical optimization can aid in the problem specification, consistent with item 4 above.

Based on the above discussion, it is clear that numerical optimization will in general not proceed directly from problem specification to the optimal design. Rather, the problem specification will evolve iteratively

*Professor, Senior Canada Research Chair in Computational Aerodynamics, Senior AIAA Member, dwz@oddjob.utias.utoronto.ca.
†Graduate Student

based on feedback provided by the optimization results. In this paper, we investigate various means of posing two-dimensional airfoil optimization problems under a range of operating conditions. The objective is twofold: 1) to identify a successful strategy for multipoint optimization, i.e. an automated procedure for designing an airfoil to produce specified optimal multipoint performance, and 2) to provide some insight into the trade-offs involved in different priorities among various objectives in order to aid designers in applying numerical optimization algorithms.

II. Multipoint Aerodynamic Shape Optimization

An airfoil can be expected to perform well under a range of Mach numbers and lift coefficients. In addition, there are typically off-design requirements such as high lift at low Mach number and minimizing shock strength at dive conditions (low angle of attack, high Mach number). Multipoint optimization involves sampling the range of operating conditions at a sufficient number of operating points that optimal performance is achieved throughout the desired range. As discussed by Li at al.,[7] there are several different ways to optimize over a range of operating conditions.[a] Three possible approaches are as follows:

- find the design with the best worst-case performance, i.e. the minimax strategy;
- find the design that produces constant performance over the desired range of conditions;
- optimize some weighted integral of the performance over the desired range.

Combinations of these strategies are also possible. For example, the designer may wish to optimize performance over a range of conditions while constraining the performance under off-design conditions.

Given that the overall objective is presumably some combination of maximizing return on investment and minimizing risk, the task of the designer is to determine which of these three approaches is most appropriate for a given design. However, this may not be apparent a priori, and several iterations may be needed to determine the best approach in a given context.

Once the designer has specified how performance is to be optimized, the numerical optimization problem must be posed such that this is achieved. If we consider aerodynamic optimization at fixed lift over a range of Mach numbers, how are the sampling points, i.e. the individual Mach numbers within the range, selected? Drela[4] shows that the number of sampling points is related to the number of degrees of freedom in the design space. Similarly, how is each operating point weighted in the composite objective function? The sampling points and their weights can be determined by trial and error, but an automated procedure is preferred and is further discussed below.

III. Algorithm Description

Multipoint aerodynamic optimization is studied using the Newton-Krylov approach of Nemec and Zingg.[5] The geometry is parameterized through B-splines. The B-spline control points and the angle of attack are the design variables. The compressible Navier-Stokes equations are solved with a Newton-Krylov method in which the linear system arising at each Newton iteration is solved using the generalized minimal residual method (GMRES) preconditioned with an incomplete lower-upper (ILU) factorization with limited fill. The Spalart-Allmaras turbulence model is used to compute the eddy viscosity. The gradient is calculated using the discrete-adjoint method; solution of the adjoint equation is accomplished through the same preconditioned Krylov method. Geometric constraints are added to the objective function as penalty terms. A new set of design variables is computed using a quasi-Newton optimizer in which an estimate of the inverse Hessian based on the BFGS (Broyden-Fanno-Goldfarb-Shannon) rank-two update formula is used to compute a search direction.[8] If the initial step does not produce sufficient progress toward the minimum, the step size is determined using a line search, which terminates when the strong Wolfe conditions are satisfied.[8] Each time a new shape is calculated, the initial grid is perturbed using a simple algebraic technique. For a complete description of the algorithm, see Nemec.[9]

[a] Li et al. use the term "robust optimization" to describe optimization under a range of operating conditions.

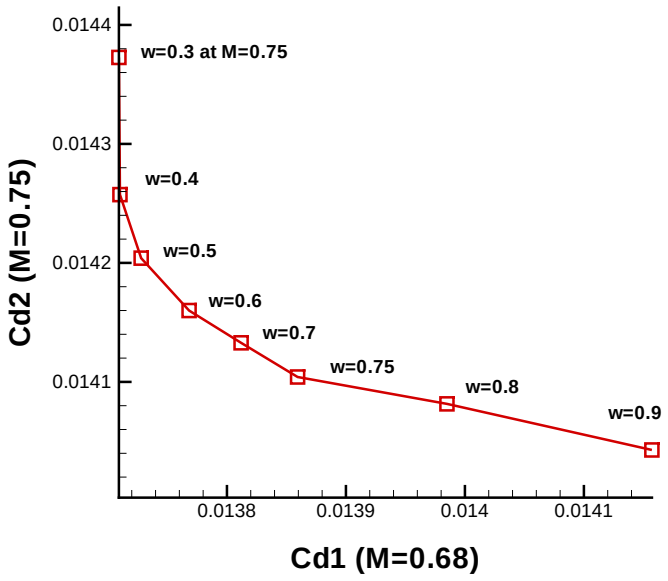

Figure 1. The Pareto front for the two-point lift-constrained drag minimization.

IV. A Simple Two-Point Example

In order to illustrate the basic ideas, we begin with a simple two-point lift-constrained drag minimization problem. The objective function is given by:

$$J = \begin{cases} \omega_{\mathrm{L}} \left(1 - \frac{C_{\mathrm{L}}}{C_{\mathrm{L}}^*}\right)^2 + \omega_{\mathrm{D}} \left(1 - \frac{C_{\mathrm{D}}}{C_{\mathrm{D}}^*}\right)^2 & \text{if } C_{\mathrm{D}} > C_{\mathrm{D}}^* \\ \omega_{\mathrm{L}} \left(1 - \frac{C_{\mathrm{L}}}{C_{\mathrm{L}}^*}\right)^2 & \text{otherwise} \end{cases} \tag{1}$$

where C_{L}^* is a target lift coefficient, C_{D}^* is a target drag coefficient, and ω_{L}, ω_{D} are weights. Choosing a target lift coefficient that is attainable and a target drag coefficient that is unattainable with suitable weights leads to lift-constrained drag minimization with the lift constraint treated as a penalty term.

The target lift coefficient is 0.715; the Reynolds number is 9 million. Fully turbulent flow is assumed, i.e. transition is assumed to occur at the leading edge. The initial geometry is the RAE 2822 airfoil. Fifteen B-spline control points are used to parameterize the geometry. Three control points are frozen at the leading edge and two at the trailing edge. Hence there are eleven design variables, including the angle of attack. The following thickness constraints are imposed: $t/c \geq 0.0253$ at $x/c = 0.01$, $t/c \geq 0.121$ at $x/c = 0.25$, $t/c \geq 0.002$ at $x/c = 0.99$. The two operating points are $M = 0.68$ and $M = 0.75$, where M is the freestream Mach number. Of course, minimizing the drag at two operating points with eleven design variables will not produce any sort of optimum performance over a range of Mach numbers. However, this simple problem shows some interesting trade-offs that apply to problems with more sampling points. Furthermore, multipoint optimization can be considered as a subset of multi-objective optimization, and therefore we can make use of a Pareto front to gain an understanding of the trade-offs between the two operating conditions.

We use the following composite objective function:

$$J = wJ_{M=0.75} + (1 - w)J_{M=0.68} \tag{2}$$

where $J_{M=0.75}$ is the objective function given by Eq. 1 evaluated at $M = 0.75$, $J_{M=0.68}$ is the same objective function evaluated at $M = 0.68$, and w is a weight that controls the relative importance of the two operating points. The Pareto front that results from varying this weight is shown in Figure 1. The corresponding drag coefficient values are given in Table 1. Several important aspects of this particular problem are revealed. Although in principle each point on the front is an equally valid optimum, one can see that the extreme values

Table 1. C_D **for various values of** w**.**

w	C_D at M=0.75	C_D at M=0.68
0.1	0.015325	0.013580
0.3	0.014373	0.013709
0.4	0.014257	0.013710
0.5	0.014204	0.013728
0.6	0.014160	0.013768
0.7	0.014133	0.013812
0.75	0.014104	0.013859
0.8	0.014082	0.013985
0.85	0.014067	0.014004
0.9	0.014043	0.014157

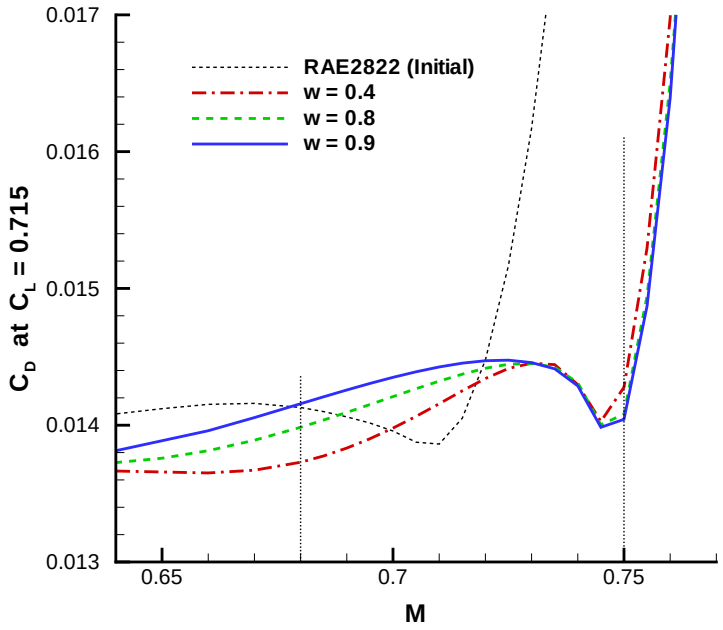

Figure 2. Variation of drag coefficient with Mach number at $C_L = 0.715$ **for various values of** w**.**

of w, i.e. $w < 0.4$ and $w > 0.8$ are not good choices. For example, increasing w from 0.3 to 0.4 produces a significant reduction in the drag coefficient at $M = 0.75$ with almost no penalty in the drag coefficient at $M = 0.68$. Furthermore, the design with equal drag at the two operating points, which is also the minimax solution,[b] is achieved with w roughly equal to 0.86. Unless there is a very good reason to require that the two drag coefficients be equal, this is not a very good choice, since the relatively low drag at $M = 0.75$ is obtained at the expense of high drag at $M = 0.68$. Between $w = 0.85$ and $w = 0.90$, the drag coefficient at $M = 0.68$ is increasing over six times faster with an increase in w than that at $M = 0.75$ is decreasing. For example, reducing w from 0.85 to 0.75 reduces the drag coefficient at $M = 0.68$ by 0.000145 while increasing the drag coefficient at $M = 0.75$ by only 0.000037. This is further illustrated in Figure 2, which shows that reducing the drag at $M = 0.75$ through an increase in w penalizes the drag over a large portion of the Mach number range. This is an example of a problem in which the designer might change the initial specification of the problem requiring equal drag at the two Mach numbers when the trade-offs are better understood.

[b]for the two specified operating points, not the whole Mach number range

V. Automated Selection of Sampling Points and Weights

Next we consider a more practical example in which lift-constrained drag minimization is performed over a range of Mach numbers from 0.68 to 0.76. The target lift and drag coefficients are $C_L^* = 0.733$ and $C_D^* = 0.01$ with weights $\omega_L = 1.0$ and $\omega_D = 0.1$. The RAE 2822 airfoil is parameterized using 25 B-spline control points. One B-spline control point is frozen at the leading edge and two at the trailing edge, so there are 23 design variables, including the angle of attack. The following thickness constraints are imposed with a weight of unity: $t/c \geq 0.0253$ at $x/c = 0.01$, $t/c \geq 0.121$ at $x/c = 0.35$, $t/c \geq 0.0137$ at $x/c = 0.924$, $t/c \geq 0.001516$ at $x/c = 0.99$. The Reynolds number is again nine million, and the flow is fully turbulent.

We assume initially that four sampling points will be sufficient and that a constant drag coefficient is desired over the Mach number range. The initial sampling points are $M = 0.68,\ 0.70667,\ 0.73333,\ 0.76$, an even distribution, with a composite objective function given by

$$J = w_1 J_{M=0.68} + w_2 J_{M=0.707} + w_3 J_{M=0.733} + w_4 J_{M=0.76} \tag{3}$$

with $\sum w_i = 1$. With all four sampling points weighted equally, the difference between the highest drag coefficient (at $M = 0.76$) and the lowest drag coefficient (at $M = 0.68$) is over ten counts.

In order to achieve a constant drag coefficient over the specified range of Mach numbers, an automated technique is proposed involving two steps. In the first step, the goal is to equalize the drag coefficients at the specified sampling points. This is accomplished by updating the weights according to the following formula:

$$w_i^{new} = w_i^{old} + c\left(\frac{C_{Di}}{\sum_{i=1}^{N} C_{Di}} - \frac{1}{N}\right) \tag{4}$$

where N is the number of sampling points, and c is a user-specified constant. The weights are updated according to the above formula after a complete multipoint optimization with the previous set of weights. Once the drag coefficients at the specified Mach numbers are sufficiently equalized, the drag coefficients of the designed airfoil are evaluated over the entire range of Mach numbers, i.e. in between the specified sampling points. If there exists a significant local maximum between the specified Mach numbers, then an additional sampling point is added at the Mach number where the maximum exists. The new point is given an initial weight of zero, and the weights then evolve according to Eq. 4 until the drag coefficients at the sampling points, including the newly introduced Mach number, are sufficiently equalized. These two steps can be repeated until no significant local maxima occur between sampling points.

The results of this automated procedure are presented in Tables 2 and 3 and Fig. 3. Table 2 shows the evolution of the weights, Table 3 shows the evolution of the drag coefficients at the specified Mach numbers, and Fig. 3 displays the evolution of the drag coefficients over the entire range of Mach numbers. A value of $c = 15$ is used in Eq. 4. Each iteration shown in the tables is a complete multipoint optimization. After eight such iterations, the difference between the drag coefficient at $M = 0.76$ and that at $M = 0.68$ has been reduced from over ten counts to roughly one count. The weight on the objective function at $M = 0.76$ has more than doubled, while the others have decreased. However, evaluation of the drag coefficient over the complete range of Mach numbers from 0.68 to 0.76 reveals a significant local maximum at $M = 0.753$, where the drag is over six counts higher than that at $M = 0.76$, as shown in Fig. 3, where the result is labelled "8 iterations". Consequently, an additional sampling point is added at $M = 0.753$ with an initial weight of zero. Eight further iterations are performed with the five sampling points until the difference between the largest drag coefficient and the smallest is reduced to one-quarter of a count. Fig. 3 shows that the resulting airfoil (labelled "16 iterations") produces no significant local maxima between operating points, although there are two local minima. The final airfoil produces a maximum drag coefficient of 0.014787 for Mach numbers ranging from 0.68 to 0.76.

Fig. 4 displays the Mach number contours for the final optimized airfoil at $C_L = 0.733$, $M = 0.76$. The surface pressure coefficient distributions at the five specified Mach numbers are compared with the initial RAE 2822 airfoil in Fig. 5. The high weighting on the two highest Mach numbers is evident from these figures. The performance at $M = 0.70667$ is actually inferior to that of the RAE 2822 airfoil.

The automated procedure converged such that the drag coefficients at the sampling points are equalized, and there are no significant local maxima between sampling points. Nevertheless, due to the presence of two local minima, constant drag has not been achieved throughout the Mach number range. In order to achieve constant drag, sampling points would have to be introduced at the Mach numbers where the local

Table 2. Evolution of weights.

Iteration	c	Weights (Mach No.) 0.68	0.70667	0.73333	0.753	0.76
1		0.25000	0.25000	0.25000		0.25000
2	15	0.14736	0.20276	0.23624		0.41365
3	15	0.07167	0.17371	0.24624		0.50837
4	15	0.01497	0.15141	0.24749		0.58613
5	15	0.01598	0.14456	0.23031		0.60915
6	15	0.02188	0.13996	0.21372		0.62444
7	15	0.03808	0.13543	0.19683		0.62966
8	15	0.04999	0.13010	0.18127	0.00000	0.63865
9	15	0.01311	0.09592	0.15396	0.11384	0.62317
10	15	0.05062	0.10483	0.12189	0.12695	0.59571
11	15	0.04409	0.10294	0.10973	0.15193	0.59131
12	15	0.04441	0.10647	0.09879	0.16495	0.58539
13	15	0.04430	0.11093	0.09259	0.17228	0.57990
14	15	0.04529	0.11475	0.08648	0.17772	0.57576
15	15	0.04622	0.11815	0.08234	0.18116	0.57214
16	**15**	**0.04637**	**0.12041**	**0.07901**	**0.18481**	**0.56940**

Table 3. Evolution of drag coefficients and standard deviation.

	C_D (C_L=**0.733**) 0.68	0.70667	0.7333	0.753	0.76	St Dev
1	0.014176	0.014391	0.014521		0.015211	0.0004473
2	0.014254	0.014435	0.014586		0.014915	0.0002801
3	0.014351	0.014484	0.014576		0.014873	0.0002216
4	0.014683	0.014652	0.014611		0.014769	0.0000668
5	0.014720	0.014679	0.014632		0.014757	0.0000538
6	0.014779	0.014698	0.014649		0.014736	0.0000553
7	0.014750	0.014682	0.014642		0.014738	0.0000503
8	0.014634	0.014647	0.014681	0.015378	0.014740	0.0003169
9	0.015039	0.014897	0.014694	0.014918	0.014717	0.0001452
10	0.014745	0.014768	0.014718	0.014901	0.014756	0.0000713
11	0.014784	0.014800	0.014729	0.014847	0.014754	0.0000452
12	0.014781	0.014804	0.014751	0.014818	0.014755	0.0000294
13	0.014784	0.014798	0.014749	0.014806	0.014759	0.0000246
14	0.014782	0.014794	0.014757	0.014795	0.014760	0.0000182
15	0.014777	0.014787	0.014760	0.014794	0.014763	0.0000150
16	**0.014770**	**0.014787**	**0.014762**	**0.014786**	**0.014764**	**0.0000119**

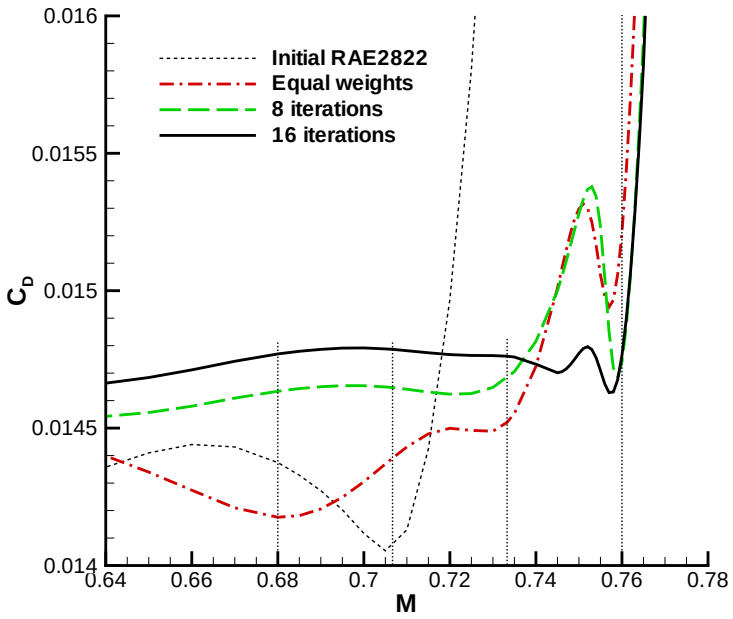

Figure 3. Variation of drag coefficient with Mach number at $C_L = 0.733$.

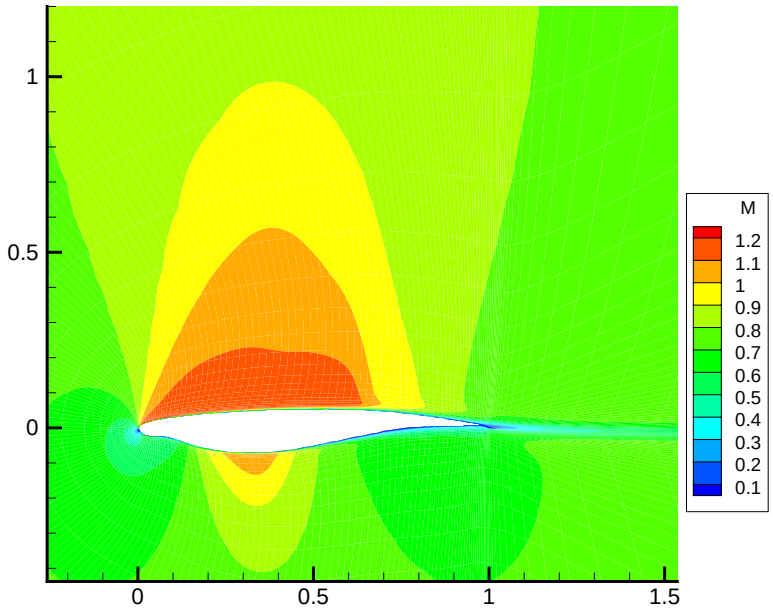

Figure 4. Mach number contours for the final airfoil at $C_L = 0.733$, $M = 0.76$.

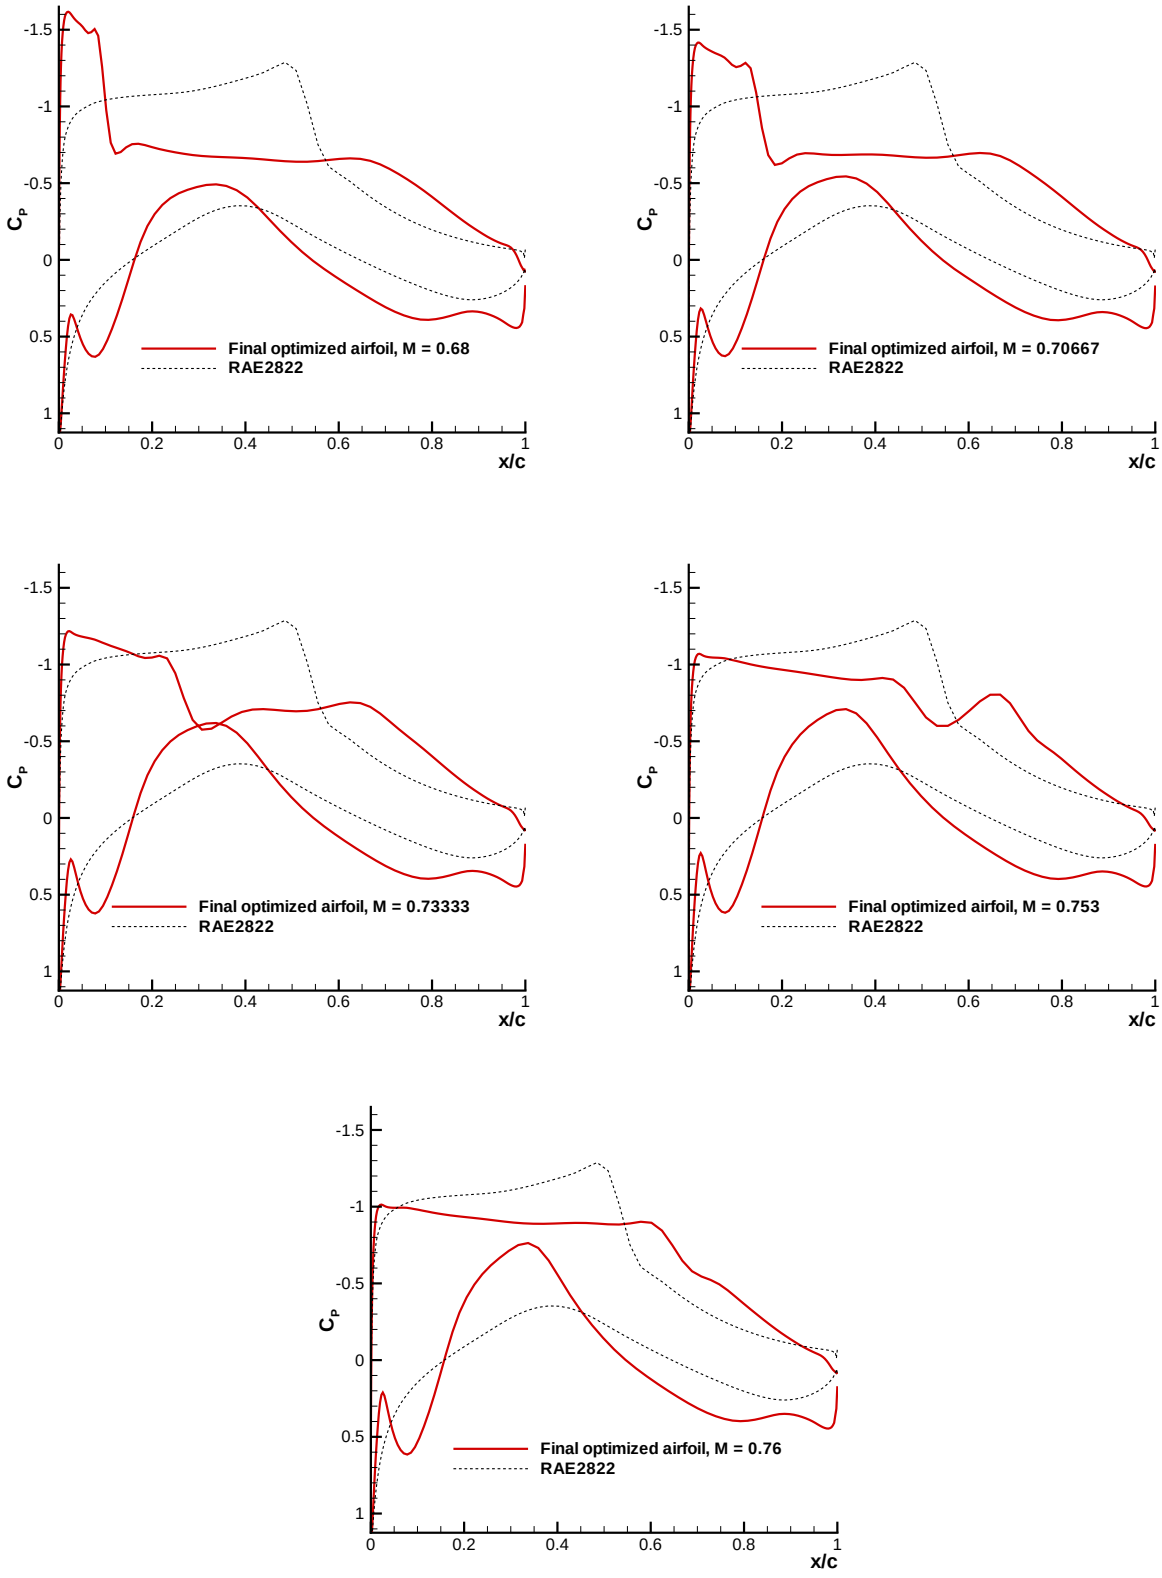

Figure 5. Final airfoil (after 16 iterations): C_p graphs for all five design points with $C_L = 0.733$.

minima occur, and the drag forced to *increase* at these points. This would not appear to be a worthwhile exercise. In some examples, the weight assigned to a given sampling point can be reduced to zero[c] yet the drag coefficient remains lower at this sampling point than at the others. This indicates that this sampling point is redundant and can be omitted, and furthermore that constant drag cannot be achieved without purposely increasing the drag at this sampling point. This suggests that constant drag over the entire Mach number range is not a particularly good goal, unless there is sufficient benefit to such performance to justify an increase in drag at some points in the operating range.

Rather than producing constant drag in the interval, the automated procedure has minimized the maximum drag over the specified range of Mach numbers. This is generally preferable to constant drag. However, the reduction in drag at the highest Mach numbers, $M = 0.753$ and $M = 0.76$, is achieved at the expense of a significant increase in drag between $M = 0.68$ and $M = 0.73$; compare the solution with equal weights to the converged solution after sixteen iterations in Figure 3. It is up to the designer to decide whether the minimax solution obtained using the automated weight update formula is preferred over other possibilities such as the optimized airfoil obtained using equal weights. The feedback from the optimizer can aid the designer in making this decision.

VI. Off-Design Performance

Next we consider an example in which performance at off-design conditions is considered, taken from Driver and Zingg.[10] The objective function is the endurance factor, given by:

$$J = \frac{C_L^{3/2}}{C_D} \tag{5}$$

The reciprocal of the endurance factor is minimized. The following parameters are used:

- M = 0.25, Re = 2 million;
- fifteen B-spline control points of which six are used as design variables; the angle of attack is also a design variable;
- thickness constraints: $t/c \geq 0.01$ at $x/c = 0.15$, $t/c \geq 0.164$ at $x/c = 0.35$, $t/c \geq 0.07$ at $x/c = 0.60$, $t/c \geq 0.01$ at $x/c = 0.92$, $t/c \geq 0.001$ at $x/c = 0.99$.

For this example, the location of laminar-turbulent transition is free and is predicted using the e^n method. Hence the optimizer exploits this by designing an airfoil for which transition occurs relatively far aft. The final airfoil has transition points located at 56 and 65% chord on the upper and lower surfaces respectively at an angle of attack of 4.36 degrees. The resulting endurance factor is 115.

Whenever an airfoil is designed with a region of natural laminar flow, one must consider the possibility that transition will occur earlier than expected for some reason, such as frost, roughness, damage, etc. When an analysis is performed with the assumption of fully turbulent flow of the airfoil optimized under the assumption of free transition, the endurance factor drops below 50. This is a drastic reduction in performance and may not be acceptable. A conservative approach is to optimize the design under fully turbulent conditions. This leads to an endurance factor of 58 under fully turbulent conditions and roughly 75 if analyzed with free transition. This is a substantial penalty under free transition conditions compared to the value of 118 achieved by the previous design.

The trade-offs between good performance under fully turbulent conditions and free transition can be assessed using a Pareto front. The composite objective function is:

$$J = w_{ft} J_{ft} + (1 - w_{ft}) J_{lt} \tag{6}$$

where the subscript ft denotes fully turbulent conditions, and lt denotes free transition. The Pareto front computed by varying w_{ft} is shown in Figure 6. Note that the front is plotted as a maximization problem. The extreme ends of the front are not good designs. For example, increasing w_{ft} from 0.1 to 0.2 leads to a significant increase in the endurance factor under fully turbulent conditions with little penalty in the free-transition endurance factor. Similarly, reducing w_{ft} from 0.9 to 0.7 produces a large increase in the

[c]Negative weights are not permitted.

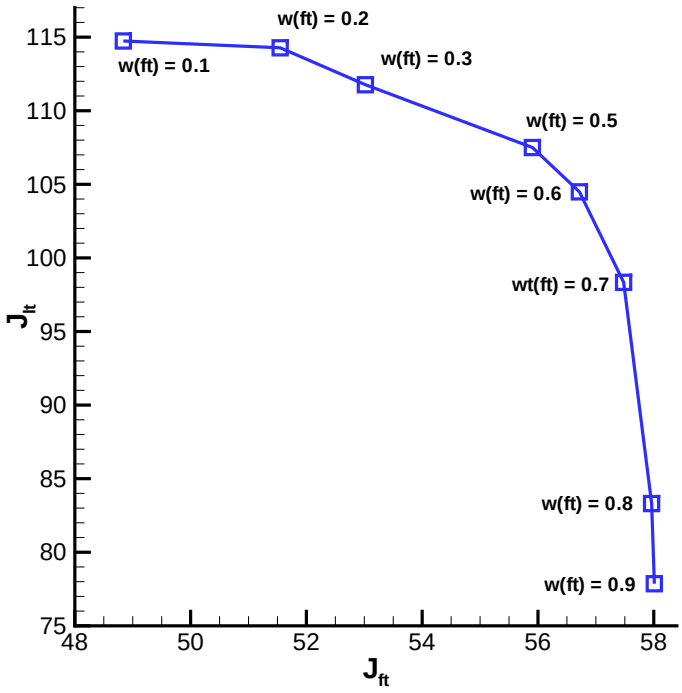

Figure 6. Pareto front arising from maximization of endurance factor. J_{lt} and J_{ft} are the endurance factors computed with free transition and fully turbulent conditions, respectively.

endurance factor with free transition without significantly penalizing performance under fully-turbulent conditions. Hence, the Pareto front provides feedback needed by the designer to select the value of w_{ft} that produces the airfoil which best meets the designer's priorities. In this case, a value of w_{ft} between 0.2 and 0.7 should be selected.

VII. Discussion

The above examples have shown that formulation of a multipoint aerodynamic optimization problem is not a straightforward task. The designer must have a deep understanding of the desired performance characteristics of the aerodynamic component and their implications on the performance of the aircraft. Furthermore, the designer must be aware of what is feasible before the problem can be fully posed. Well-designed aerodynamic optimization software can be of great use in understanding the trade-offs between various competing performance objectives, especially through the use of Pareto fronts.

A complete multipoint optimization will involve consideration of various lift coefficients as well as a range of Mach numbers plus some off-design requirements. While one can achieve nearly constant drag over a range of Mach numbers, it does not seem reasonable to design for constant drag over a range of lift coefficients. A constant lift-to-drag ratio may be more appropriate. Minimizing a weighted integral over the on-design operating envelope will likely be the most effective approach if a rational strategy based on the aircraft mission requirements can be used to determine the appropriate weighted integral. Using the mid-point, trapezoidal, or Simpson integration formulas, the weighted integral can be converted to a multipoint optimization problem with specified sampling points and weights. The automated procedure presented in this paper can be used for such a problem, not to determine the weights, since they are fixed by the specified weighted integral, but to determine the sampling points needed to provide a good approximation to the weighted integral. The above numerical integration formulas are accurate only if the function is well-behaved between sampling

points. Hence new sampling points can be added automatically if significant local maxima or minima are detected between existing sampling points.

The ability to optimize aerodynamic components for specific performance over a range of operating conditions also increases the importance of careful specification of off-design requirements. Off-design performance is usually imposed as a constraint. For example, a certain maximum lift coefficient may be required for low-speed performance, and constraints may be placed on performance under dive conditions as well. With formal multipoint aerodynamic optimization, the constraints associated with off-design performance requirements can have a significant detrimental impact on the performance under on-design conditions. Consequently, the off-design performance requirements must be specified as carefully as the on-design requirements, which in turn requires a good understanding of uncertainty and risk under off-design conditions.

VIII. Conclusions

Issues in aerodynamic optimization under variable operating conditions have been discussed and addressed. The examples show that it is difficult to pose a multipoint optimization problem a priori and that the feedback from the optimization can lead to better problem specification. Pareto fronts are shown to be particularly useful in revealing the trade-offs associated with different weightings in composite objective functions. A method is presented to automatically select sampling points and their weights in order to achieve desired performance over a range of operating conditions, in this case constant drag over a range of Mach numbers. The results presented provide insight both in formulating and in solving multipoint aerodynamic optimization problems.

Acknowledgments

The funding of the first author by the Natural Sciences and Engineering Research Council of Canada and the Canada Research Chairs program is gratefully acknowledged.

References

[1]Nielsen, E.J., and Anderson, W.K., "Aerodynamic Design Optimization on Unstructured Meshes Using the Navier-Stokes Equations," *AIAA J.*, Vol. 37, No. 11, 1999, pp. 1411-1419.

[2]Reuther, J.J., Jameson, A., Alonso, J.J., Rimlinger, M.J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part I," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 51-60.

[3]Elliott, J., and Peraire, J., "Constrained Multipoint Shape Optimization for Complex 3D Configurations," *Aeronautical Journal*, Vol. 102, No. 1017, 1998, pp. 365-376.

[4]Drela, M., "Pros & Cons of Airfoil Optimization," *Frontiers of Computational Fluid Dynamics*, edited by D.A. Caughey and M.M. Hafez, World Scientific, Singapore, 1998, pp. 363-381.

[5]Nemec, M., and Zingg, D.W., "Newton-Krylov Algorithm for Aerodynamic Design Using the Navier-Stokes Equations," *AIAA J.*, Vol. 40, No. 6, 2002, pp. 1146-1154.

[6]Nemec, M., Zingg, D.W., and Pulliam, T.H., "Multipoint and Multi-Objective Aerodynamic Shape Optimization," *AIAA J.*, Vol. 42, No. 6, 2004, pp. 1057-1065.

[7]Li, W., Huyse, L., and Padula, S., "Robust Airfoil Optimization to Achieve Consistent Drag Reduction Over a Mach Range," NASA/CR-2001-211042, Aug. 2001.

[8]Nocedal, J. and Wright, S.J., *Numerical Optimization*, Springer-Verlag, New York, 1999.

[9]Nemec, M., "Optimal Shape Design of Aerodynamic Configurations: A Newton-Krylov Approach," Ph.D. Thesis, University of Toronto Institute for Aerospace Studies, 2002.

[10]Driver, J., and Zingg, D.W., "Optimized Natural-Laminar-Flow Airfoils," AIAA Paper 2006-247, Reno, Jan. 2006.