This paper is part of the following report:

TITLE: Annual Research Briefs - 2003 [Center for Turbulence Research]

To order the complete compilation report, use: ADA420749

Constrained aeroacoustic shape optimization using the surrogate management framework

By Alison L. Marsden, Meng Wang, and John E. Dennis, Jr.

1. Motivation and objectives

Reduction of noise generated by turbulent flow past the trailing-edge of a lifting surface is a challenge in many aeronautical and naval applications. Numerical predictions of trailing-edge noise necessitate the use of advanced simulation techniques such as large-eddy simulation (LES) in order to capture a wide range of turbulence scales which are the source of broadband noise. Aeroacoustic calculations of the flow over a model airfoil trailing edge using LES and aeroacoustic theory have been presented in Wang & Moin (2000) and were shown to agree favorably with experiments. The goal of the present work is to apply shape optimization to the trailing edge flow previously studied, in order to control aerodynamic noise.

There are several considerations in choosing a tractable optimization method for the trailing-edge problem. The primary concern is the computational expense of the function evaluations, and additional considerations include availability of gradient information and robustness of the optimization method. Although adjoint solvers have been successfully applied for gradient-based optimization in aeronautics problems (for example in Jameson *et al.* (1998)), they present difficulties with implementation, portability, and data storage for unsteady problems. Approximation modeling was used for trailing-edge optimization in Marsden *et al.* (2002), and results showed significant reduction in acoustic power with reasonable computational cost. In these methods, optimization is performed not on the expensive actual function, but on a surrogate function, which is cheap to evaluate. Although the approximation method presented in Marsden *et al.* (2002) was effective, it lacks rigorous convergence properties.

The surrogate management framework (SMF), developed by Booker *et al.* (1999), incorporates the use of surrogate functions into a pattern search framework, hence providing a theoretical basis for convergence. The convergence of pattern search methods has been studied extensively by Audet & Dennis (2003, 2000). Use of the SMF method has been demonstrated, among others, by Booker *et al.* (1999), where the method was successfully applied to a helicopter rotor blade design problem with 31 design variables.

The SMF method provides a robust and efficient alternative to traditional gradient-based optimization methods. In this work, the SMF method is applied for trailing-edge optimization in a time-dependent laminar flow problem with and without constraints on lift and drag. Several interesting optimal shapes have been identified, all of which result in significant reduction of trailing-edge noise. In particular, the development of a trailing-edge bump in the constrained case is an unexpected result which illustrates the trade-off between noise reduction and loss of lift.

FIGURE 1. Blake airfoil used in model problem. Right half section of upper surface is allowed to deform

2. Problem formulation and cost function

The general optimization problem may be formulated with bound constraints as follows

$$\text{minimize} \quad J(x)$$
$$\text{subject to} \quad x \in \Omega. \tag{2.1}$$

In the above problem statement, $J : \mathbb{R}^n \to \mathbb{R}$ is the cost function, and x is the vector of design parameters. The bounds on the parameter space are defined by $\Omega = \{x \in \mathbb{R}^n | l \leq x \leq u\}$ where $l \in \mathbb{R}^n$ is a vector of lower bounds on x and $u \in \mathbb{R}^n$ is a vector of upper bounds on x.

In this work, the surrogate management framework is implemented and validated for optimization of a time-dependent flow problem. The airfoil geometry for the model problem is shown in Figure 1 and is a shortened version of the airfoil used in experiments of Blake (1975). The airfoil chord is 10 times its thickness, and the right half of the upper surface is allowed to deform. The flow is from left to right and results presented in this work are at a chord Reynolds number of $Re = 10,000$.

The cost function is defined in terms of noise radiation from an acoustically compact airfoil, calculated using Curle's extension to the Lighthill theory (Curle 1955). The compactness assumption is valid for unsteady laminar flow past an airfoil at low Mach number since the acoustic wavelength associated with the vortex shedding is typically long relative to the airfoil chord. Details of the cost function derivation are given in Marsden *et al.* (2003), and the final cost function expression is

$$\bar{J} = \overline{\left(\frac{\partial}{\partial t} \int_S n_j p_{1j}(\mathbf{y}, t) d^2 \mathbf{y} \right)^2} + \overline{\left(\frac{\partial}{\partial t} \int_S n_j p_{2j}(\mathbf{y}, t) d^2 \mathbf{y} \right)^2}, \tag{2.2}$$

which is directly proportional to the radiated acoustic power.

The cost function, $\bar{J}$, depends on control parameters for the airfoil surface deformation. Each parameter corresponds to a deformation point on the airfoil surface, and its value must be within prescribed allowable bounds. The value of each parameter is defined as the displacement of the control point relative to the original airfoil shape, in the direction normal to the surface. A positive parameter value corresponds to displacement in the outward normal direction, and a negative value corresponds to the inward normal direction. A spline connects all the deformation points to the trailing edge point and the left (un-deformed) region to give a continuous airfoil surface. Both ends of the spline are fixed. While the surface must be continuous and smooth on the left side, the trailing edge angle is free to change.

For a given set of parameter values, there is a unique corresponding airfoil shape. To calculate the cost function value for a given shape, a mesh is generated and the flow simulation is performed until the solution is statistically converged. A finite difference code discussed in Wang & Moin (2000) is used to solve the time-dependent incompressible two-dimensional Navier-Stokes equations in generalized curvilinear coordinates. Because

the flow has unsteady vortex shedding, the cost function is oscillatory. In the optimization procedure, the mean cost function $\bar{J}$ (cf. (2.2)) is used, which is obtained by averaging in time until convergence. With each shape modification, the flow field is allowed to evolve for sufficiently long time to establish a new quasi-steady state before the time averaging is taken.

3. Outline of the surrogate management framework

In this section, we outline the steps used for trailing-edge shape optimization with the surrogate management framework. The SMF, introduced in Booker *et al.* (1999), is a pattern search method which incorporates surrogate functions to make the optimization cost effective. The main idea behind the SMF method is to use a surrogate function as a predictive tool, while retaining the robust convergence properties of pattern search methods. Like pattern search methods, SMF is a mesh based algorithm, so that all points evaluated are restricted to lie on a mesh.

The first step in the optimization is to choose a set of initial data. Latin hypercube sampling (LHS) McKay *et al.* (1979), is commonly used to find a well distributed set of initial data in the parameter space. Latin hypercube sampling ensures that each input variable has all portions of its range represented in the chosen data set. Once the initial data set, $x_1 \ldots x_m$, has been chosen, the cost function, $J(x)$ is evaluated at these points, and an initial surrogate model is constructed.

A Kriging surrogate model is used to interpolate the data, and to predict the value of the function at a particular location in the parameter space. As the optimization progresses, the surrogate model should be updated to include new data. Kriging, also called DACE, is a statistical method based on the use of spatial correlation functions. It is easily extended to multiple dimensions, making it attractive for optimization problems with several parameters. A detailed derivation of the Kriging approximation is given in Marsden *et al.* (2003), following Lophaven *et al.* (2002). After constructing an initial surrogate, all points subsequently evaluated by the algorithm are restricted to lie on a mesh in the parameter space. The mesh definition is flexible so long as the vectors connecting a point x to any $2n$ points adjacent to x form a positive basis for $\mathbb{R}^n$. The mesh may be refined or rotated during the optimization as long as it satisfies this definition.

The SMF algorithm consists of two steps, SEARCH and POLL. The exploratory SEARCH step uses the surrogate to aid in the selection of points which are likely to improve the cost function. The SEARCH step provides means for local and global exploration of the parameter space, but is not strictly required for convergence. Because the SEARCH step is not integral to convergence, it affords the user a great deal of flexibility and may be adapted to a particular engineering problem.

Convergence of the SMF algorithm is guaranteed by a POLL step, in which points neighboring the current best point on the mesh are evaluated to check whether the current best point is a mesh local optimizer. A set of POLL points are required to generate an $n + 1$ positive basis. An example of such a basis is constructed in $\mathbb{R}^n$ as follows. We let V be the matrix whose columns are the basis elements. Then construct $D = [V, -V \cdot e]$, where e is the vector of ones and $-V \cdot e$ is the negative sum of the columns of V. The columns of D form an $n + 1$ positive basis for $\mathbb{R}^n$. For example, in three dimensions such a basis could be given by $(1, 0, 0), (0, 1, 0), (0, 0, 1), (-1, -1, -1)$.

Following evaluation of the initial data, the first step in the optimization is a SEARCH step. In the SEARCH step, optimization is performed on the surrogate in order to predict

the location of one or more minimum points, and the function is evaluated at these points. If an improved cost function value is found, the search is considered successful, the surrogate is updated, and another search step is performed. If the SEARCH fails to find an improved point, then it is considered unsuccessful and a POLL step is performed, in which the set of POLL points are evaluated. If the POLL produces an improved point, then a SEARCH step is performed on the current mesh. Otherwise, if no improved points are found, then the current best point is a local minimum of the function on the mesh. For greater accuracy, the mesh may be refined, at which point the algorithm continues with a SEARCH. Convergence is reached when a local minimum on the mesh is found, and the mesh has been refined to the desired accuracy. Each time new data points are found in a SEARCH or POLL step, the data is added to the surrogate model and the surrogate is updated. The steps in the algorithm are summarized below, where the set of points in the initial mesh is M_0, the mesh at iteration k is M_k, and the current best point is x_k.

1. SEARCH
(*a*) Identify finite set of trial points T_k on the mesh M_k.
(*b*) Evaluate $J(x_{trial})$ for trial points $T_k \in M_k$.
(*c*) If for any point in T_k, $J(x_{trial}) < J(x_k)$, a lower cost function value has been found, the *SEARCH is successful.* Increment k and go back to SEARCH.
(*d*) Else if none of T_k improves the cost function, *SEARCH is unsuccessful.* Increment k and go to POLL.

2. POLL
(*a*) Find a set of poll points X_k around x_k which are neighboring mesh points that generate a positive basis.
(*b*) If for any point in X_k , $J(x_{poll}) < J(x_k)$, a lower cost function has been found and the *POLL is successful.* Increment k and go to SEARCH.
(*c*) Else if none of X_k improves the cost function, *POLL is unsuccessful.*
  (*a*) If convergence criteria are satisfied, a converged solution has been found. STOP.
  (*b*) Else if convergence criteria are not met, refine mesh. Increment k and go to SEARCH.

Because the method has distinct SEARCH and POLL steps, convergence theory for the SMF method reduces to convergence of pattern search methods. Convergence of the SMF method is discussed at length by Booker *et al.* (1999). Pattern search convergence theory is presented by Audet & Dennis (2003).

4. Unconstrained two parameter results

In this section, the full SMF method is validated using two shape parameters. The control points, a and b, are evenly distributed on the upper surface of the airfoil. For each set of parameters, the airfoil surface is interpolated using a hermite spline. The lower bound on the parameters is chosen to correspond to a straight line connecting the left edge of the deformation region and the trailing edge, and the upper bound is an equal deformation in the outward normal direction.

To give a basis for comparison, results using two parameters without the POLL step are presented in Table 1. We term this method the "strawman method." The case shown produces a cost function reduction of 54% with 18 function evaluations. The optimal shape obtained in this case is shown on the left of Figure 2. The corresponding normalized cost function reduction is given on the right of Figure 2, and the initial and final Kriging surrogates are shown in Figure 3. To make the "strawman" case consistent with cases
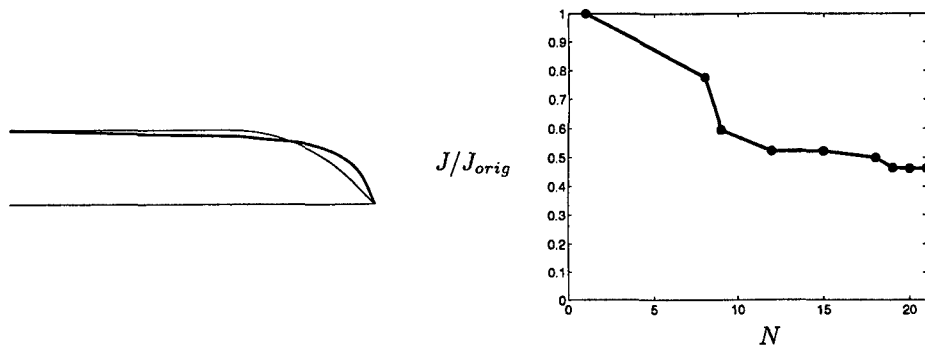
FIGURE 2. Left: initial (thin line) and final (thick line) airfoil shapes using two parameters with no poll step ("strawman method"). Right: normalized cost function (acoustic power) vs. total function evaluations
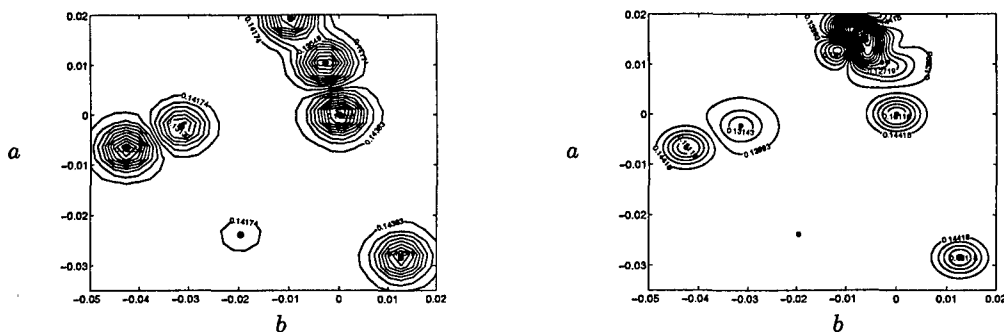

FIGURE 3. Initial and final Kriging surrogate functions for "strawman" case. Left plot shows initial data obtained with latin hypercube sampling; right side shows final surrogate fit.

using the SMF method, all points evaluated were restricted to lie on a mesh of the same size, and the mesh was refined twice. Convergence was reached when surrogate minimum point was the same as the previous iteration.

For purpose of comparison, all full SMF method cases use the same set of LHS initial data (7 points) as the "strawman" case shown on the left side of Figure 3. The two-parameter results obtained using the full SMF method are given in Table 1. The table lists the total number of function evaluations as well as the number of iterations, where one iteration is a complete SEARCH or POLL step. In all cases, the number of function evaluations includes the number of initial data points.

The second line of table 1 shows a case in which one point, the surrogate minimum on the mesh, is evaluated in each SEARCH step. With two mesh refinements, the total cost function reduction is 72%. Comparing with the "strawman" case, the POLL step adds some computational expense, but also results in a significantly lower cost function value. As discussed in Section 3, the POLL step ensures convergence to a local minimum on the mesh in the parameter space.

Lines three and four in Table 1 explore the effect of using multiple points in the SEARCH step of the algorithm. In each case, one SEARCH point is found by a direct search for the minimum of the surrogate function on the mesh. Additional "space-filling" points are added in an effort to keep the data set well distributed and prevent degradation of surrogate accuracy. The addition of "space-filling" points does not necessarily increase the overall cost since the SEARCH points may be evaluated in parallel. Results using two
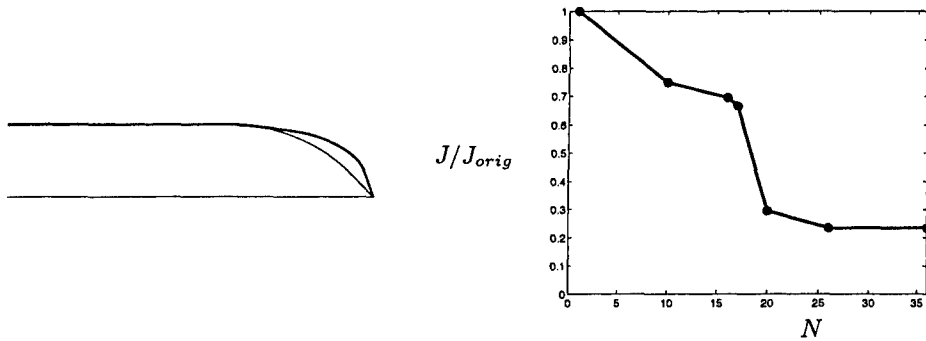
FIGURE 4. Left: initial (thin line) and final (thick line) airfoil shapes using two parameters with SMF method. Right: normalized cost function (acoustic power) vs. total function evaluations. Each SEARCH step used three function evaluations.

optimization method	parameters	search points	% reduction	evaluations	iterations
strawman	2	1	54%	18	11
SMF	2	1	72%	33	15
SMF	2	2	72%	42	15
SMF	2	3	77%	46	13

TABLE 1. Two parameter SMF - Comparison of cases using one, two, and three search points.

points (the surrogate minimum and one "space-filling" point) in each SEARCH step are given in line three of Table 1. The cost function reduction and the optimal shape are identical to the case with only one search point. In both cases, 15 iterations were required using two mesh refinements, a slight increase in cost over the "strawman" case.

Using three SEARCH points (one surrogate minimizer and two 'space-filling' points) a larger cost function reduction of 77% was achieved, as shown in line four of Table 1. Figure 4 shows the optimized airfoil shape (left) and the cost function reduction (right) for this case. Comparing with the other cases in Table 1, the case with three search points required fewer iterations. This savings is explained by a higher surrogate quality in the final iterations, resulting in fewer polling steps. A comparison of surrogate quality is made by evaluating the mean squared error.

We have shown that efforts to reduce surrogate degradation can pay off, resulting in a lower cost function solution. In addition, the two parameter cases, shown in Figures 2 and 4, resulted in airfoil shapes with a blunt trailing-edge, which was at first sight counter-intuitive. Figure 5 illustrates the time-dependent nature of the problem and compares cost function vs. time for the original airfoil, the "strawman" case, and the best SMF case. The magnitude of acoustic power has decreased significantly for the optimized shapes compared to the original. However, with the 77% acoustic power reduction in the best two parameter case, the optimized airfoil has 20% lower lift than the original, which is often not allowed in engineering practice. This emphasizes the need for addition of constraints on lift and drag, which is addressed in Section 5.1.

FIGURE 5. Comparison of cost functions for original airfoil (upper plot), "strawman" with 54% reduction (middle plot) and SMF with 77% reduction (lower plot). Thin line is instantaneous cost function (acoustic power), and thick line is averaged cost function (acoustic power).

5. Constrained optimization with multiple parameters

5.1. *Filtering method for constrained optimization*

Lift and drag constraints can be enforced using the filtering method of Audet & Dennis (2000) and Audet *et al.* (2000), an extension of the pattern search method, and hence of the SMF method. We consider the general constrained optimization problem

$$\text{minimize} \quad J(x),$$
$$\text{subject to} \quad x \in \Omega, \; C(x) \leq 0. \tag{5.1}$$

In the above problem statement, $J : \mathbb{R}^n \to \mathbb{R}$ is the cost function, and x is the vector of parameters. The constraints are given by m functions contained in $C(x)$ so that $C = (c_1(x) \ldots c_m(x))^T$. The bounds on the parameter space are defined by a polyhedron in $\mathbb{R}^n$ denoted by Ω.

We begin by defining a constraint violation function, H, the value of which indicates how closely the problem constraints are being met. With multiple constraints, H may be the sum of several constraint functions which are weighted according to relative importance. The goal of the optimization problem is to find solutions which have a small cost function value, together with a small (or zero) value of H.

The feasible region in a plot of J vs. H is defined as the set of points that exactly satisfy $H(x) = 0$. Thus, a point x is infeasible if $H(x) > 0$. A point x' is considered filtered, or dominated, if there is a point x belonging to the filter for which $H(x) \leq H(x')$ and $J(x) \leq J(x')$. A filter, $\mathcal{F}$, is defined here to be the finite set of non-dominated points found so far. The above concepts are exemplified in figure 8, which depicts the final filter corresponding to the constrained trailing-edge optimization (to be discussed later). The points in the filter are connected with vertical and horizontal lines to form a dividing line between filtered and unfiltered regions. The best feasible point, marked with a square, is the point with the lowest cost function value, which exactly satisfies the constraint

(i.e. $H = 0$). The least infeasible point, marked with a triangle, is the filter point with the lowest non-zero constraint function value. Other points in the filter are marked with circles.

The steps in the filtering optimization algorithm fit within the framework of the SMF method, and the basic structure is the same as presented in Section 3. The differences in implementation between unconstrained SMF and the filter method lie in the criteria which make SEARCH and POLL steps successful or unsuccessful. In this implementation of the filter method, a SEARCH step is formally considered successful if it improves the filter, which means that a new non-dominated point was identified. A POLL step is considered successful if it improves either the the best feasible point or the least infeasible point. Convergence theory of this method is also based on pattern search theory, and is discussed at length in Audet & Dennis (2000). The set of points in the initial mesh is M_0, the mesh at iteration k is M_k, the current least infeasible point is LF_k and the current best feasible point is BF_k. The SEARCH and POLL steps are as follows:

1. SEARCH
(*a*) Identify finite set of trial points T_k on the mesh M_k.
(*b*) Evaluate $J(x_{trial})$ and $H(x_{trial})$ for trial points $T_k \in M_k$.
(*c*) If any point in T_k, is an unfiltered point, the *SEARCH is successful*. Increment k and go back to SEARCH.
(*d*) Else if none of T_k is an unfiltered point, *SEARCH is unsuccessful*. Increment k and go to POLL.

2. POLL
(*a*) Find a set of POLL points X_k around LF_k or BF_k, which are neighboring mesh points that form a positive basis.
(*b*) If any point in X_k dominates LF_k or BF_k, the *POLL is successful*. Increment k and go to SEARCH.
(*c*) Else if none of X_k dominates LF_k or BF_k, the *POLL is unsuccessful*.
  (*a*) If convergence criteria are satisfied, a converged solution has been found. STOP.
  (*b*) Else if convergence criteria are not met, refine mesh. Increment k and go to SEARCH.

5.2. *Incorporation of a penalty function*

Many optimization methods rely on the use of a penalty function to enforce constraints. Penalty functions are attractive due to ease of implementation into existing optimization frameworks. Challenges usually involve the choice of arbitrary weighting of the constraint function relative to the cost function. Penalty functions can be easily incorporated into the SMF filtering method, and can be extremely useful in aiding the selection of SEARCH points. Here, we present a systematic approach for choosing the penalty constant by making use of the filtering framework.

In this approach, a penalty is added to the cost function as follows,

$$\hat{J} = J_{orig} + \alpha H. \tag{5.2}$$

A surrogate is constructed to model the function $\hat{J}$, so that it is used to predict areas of the function which satisfy the constraint. The minimum of the modified surrogate function can then be evaluated in the SEARCH step.

The parameter α is chosen based on the current set of filter points (including the best feasible point). We wish to choose α so as to bias the surrogate towards points with low values of H, which will improve the filter. Let us first consider a filter with two points,

a and b, for which $H(a) < H(b)$ and $J(a) > J(b)$. We wish to choose α so that point a is favored by the surrogate model because it has a smaller constraint violation. We therefore require

$$J(a) + \alpha H(a) < J(b) + \alpha H(b).$$

For this pair of points, α must be at least as large as the slope of the line connecting them. When considering the set of points making up a filter, α should be at least as large as the slope of the steepest line connecting any two points in the filter. This choice guarantees domination of the point with the smallest value of H in the filter. Since the filtered points are not of interest in the optimization, they need not be considered in the choice of α. If there are less than two points in the filter set, $\alpha = 0$. As points are evaluated in the optimization, the filter evolves and the value of α is updated in each iteration. Values of J and H for all data points should be saved so that previous data points may be updated in the surrogate model as the value of α changes.

5.3. *Issues with multiple parameters*

The constrained optimization is performed using five design parameters. An increase in the number of parameters gives greater flexibility in the airfoil geometry and will also demonstrate feasibility and cost of the SMF method for more realistic applications. Increasing the number of parameters also presents challenges for searching on the surrogate model, and these are first discussed briefly.

Thickness constraints are defined by drawing a straight line from the left side of the upper surface deformation region to the trailing-edge point. The maximum airfoil thickness is defined by an equal displacement from the surface in the outward normal direction. Use of a hermite cubic spline as the interpolating function for the airfoil surface guarantees that no point on the surface will be displaced more than the maximum allowable displacement distance. This method can be easily generalized for any prescribed function which defines the thickness constraint.

Using a surrogate of multiple dimensions requires use of an optimization method to search for the surrogate minimum. To search the surrogate in the five parameter cases, a standard covariance matrix adaptation evolutionary strategy, or CMA-ES (Hansen & Ostermeier 1996) is employed. Accuracy is increased by running the CMA-ES optimization several times and taking the minimum value.

Based on experience with two parameters, three points are used in each SEARCH step for the five-parameter cases. In this case, the three SEARCH points are chosen to meet three goals: (1) global search, (2) local search around current best point and (3) model improvement.

The first point is chosen using the surrogate function as a predictor. The minimum of the surrogate is found using the CMA-ES, and the nearest mesh point is evaluated in the SEARCH. The second SEARCH point takes advantage of the surrogate to do a local search around the current best point. In order to pre-empt the need for a POLL step, the surrogate is used to predict the values of the POLL points neighboring the current best point. In the constrained case, the current best point may be either the best feasible point or the least infeasible point. The POLL point with the smallest surrogate value is then evaluated. In the event that the SEARCH step fails, one of the POLL points has already been evaluated, reducing the cost of a POLL step to N evaluations. The third point is for both model improvement and global search. The CMA-ES is used to search the Kriging surrogate for the point of maximum mean squared error (MSE), and the nearest mesh

parameters	constraints	% reduction	%change lift	% change drag	evaluations	iterations
5	no	77%	-17%	-12%	88	22
5	yes	43%	+0.2%	-9%	92	22

TABLE 2. Five parameter cases with SMF method. The first line is unconstrained and the second line is with constraints on lift and drag.
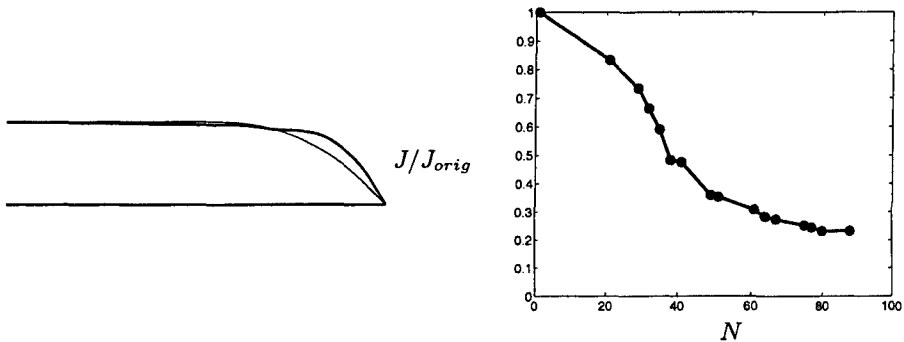

FIGURE 6. Left: initial (thin line) and final (thick line) airfoil shape with 5 shape parameters, unconstrained case. Right: corresponding normalized cost function reduction vs. number of function evaluations.

point is used in the SEARCH. This is done in an attempt to find areas with the fewest data points, as well as mitigate surrogate degradation.

5.4. Comparison of constrained and unconstrained results

Unconstrained results using five shape parameters are shown in the first line of Table 2. The number of evaluations in the table includes the initial data set of 15 points, found with LHS. The number of iterations includes all search and poll steps after one mesh refinement. The converged airfoil shape for the unconstrained case is shown on the left side of Figure 6, and the cost function reduction for this case is shown on the right. The blunt trailing-edge shape is qualitatively similar to the shapes obtained with the two parameter optimization, confirming robustness of the SMF method. The maximum cost function reduction for this case is 77% using 22 iterations after one mesh refinement, which agrees with the two parameter results. As in the two parameter case, the blunt trailing-edge results in a significant (nearly 20%) loss in lift. We observe that the blunt trailing-edge reduces the surface area affected by the separation region at the trailing-edge, and that pressure fluctuations are reduced in this region due to the smaller separation area. This in turn results in less vorticity generation, and a smaller vorticity magnitude in the wake.

Results for the constrained case are presented in the second line of Table 2. The surrogate for this case is constructed with an L_1 constraint violation function, aimed at keeping lift and drag at desirable levels:

$$H = max\left(0, \frac{L^* - L}{L^*}\right) + max\left(0, \frac{D - D^*}{D^*}\right), \tag{5.3}$$

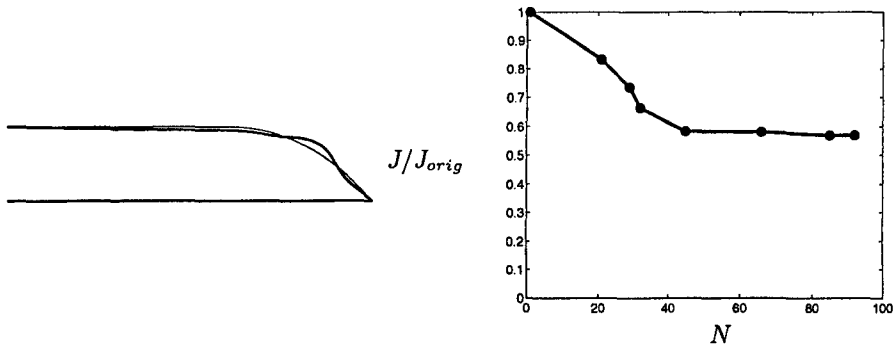where L^* and D^* are the original airfoil lift and drag. The surrogate is constructed using

FIGURE 7. Left: initial (thin line) and final (thick line) airfoil shape with 5 shape parameters, and constraints on lift and drag. Right: corresponding normalized cost function reduction vs. number of function evaluations.

equation (5.2), so that a penalty is added if *either* the lift decreases or the drag increases. In this way, we allow the lift to increase and/or the drag to decrease.

The optimized airfoil shape for the constrained case is shown on the left side of Figure 7. The constraints are effective in keeping the lift at the target value, and the drag for this case has fortuitously decreased by 9%. The cost function reduction for this case is 43%, requiring 22 iterations for convergence on a once-refined mesh. The bump near the trailing-edge reduces the magnitude of the unsteady vortex shedding by reducing the size of the separation region. However, when compared to the unconstrained case, the bump size has been compromised to maintain lift, and the trailing-edge shape is closer to a cusp. Comparison of the shapes for the constrained and unconstrained cases also illustrates the sensitivity of the flow to very small changes in the shape of the airfoil.

The final filter is shown in Figure 8. The left side shows the entire filter domain, and the right side shows a magnified view of the filter region. The filter shows the trade-off between cost function reduction and constraint violation. The cluster of points around the filter, and on the $H = 0$ axis verifies that the algorithm is expending much effort in the relevant region of the function. For comparison, the rightmost filter point corresponds to a shape with a 64% cost function reduction and a 13% loss in lift. The other filter points show the range of possible airfoil designs between this point and the optimal point.

Reduction in acoustic power can be caused by reduction in the amplitude, or a decrease in frequency of lift and drag oscillations (see (2.2)). Results do not show a change in frequency when comparing optimized cases with the original, and the influence of the unsteady drag was found to be small. The reduction of unsteady lift amplitude can be illustrated by closer examination of the flow field. Instantaneous vorticity contours are shown in Figure 9 for the original (upper), unconstrained (middle) and constrained (lower) cases. In this plot, we verify that the magnitude of vortex shedding has decreased for both optimized shapes compared to the original. In addition, we observe that vortices are shed much closer to the trailing-edge in the original case compared with the optimized cases. Movement of the unsteady region away from the trailing-edge results in smaller pressure fluctuations on the airfoil surface in this region, and explains the reduction in unsteady lift, and therefore of acoustic power.

FIGURE 8. Final filter for constrained 5 parameter optimization problem. Cost function J vs. constraint violation function H. The best feasible point is the square, the least infeasible point is the triangle, the filter points are the circles, and filtered points are stars. The original airfoil cost function is marked with a diamond. Right figure is close-up of filter region in left figure.


FIGURE 9. Comparison of vorticity contours for original (upper), unconstrained (middle) and constrained (lower) optimal shapes. Each plot shows instantaneous vorticity contours with minimum -25, maximum 25 and 20 contour levels.

6. Summary and future work

Application of the SMF method to optimize a model airfoil trailing-edge with laminar vortex shedding has resulted in significant reduction in acoustic power, as well as several interesting and previously unexpected airfoil shapes. The SMF method is robust and efficient for several design parameters with and without constraints. A filtering method has been applied to enforce constraints on airfoil lift and drag. It was implemented to include use of a penalty function, and a systematic method for choosing the penalty parameter. Comparison between the constrained and unconstrained cases using five parameters clearly showed a trade-off between noise reduction and loss of lift.

Theoretical analysis of trailing-edge noise for the Blake airfoil geometry in turbulent flow is presented by Howe (1988). Results from this work show that the lift dipole is much more significant in contributing to the noise spectra than the thickness (drag)

dipole, which was confirmed by our simulations. Although Howe's work is an analysis of turbulent trailing-edge flow, it is worth noting that his analysis predicted a decrease in trailing-edge noise with an increase in trailing-edge angle. This result agrees qualitatively with the blunt shapes found by the optimization method in the two and five parameter cases, both of which resulted in a dramatic reduction in trailing-edge noise.

The similarity between optimal shapes obtained in the unconstrained two and five parameter cases demonstrates the robustness of the SMF method. Comparison of the full SMF method with the "strawman" approach showed that the POLL step can lead to a greater cost function reduction with minimal additional cost. In general, the number of iterations required by the SMF method was modest. However, it may be possible to further reduce the cost of the method through surrogate quality improvement. The use of a second Gaussian process in Kriging, as in Audet *et al.* (2000), can prevent surrogate degradation and has been shown to reduce the number of POLL steps in several test cases. This is an area for future study.

In this work, we have demonstrated successful use of the SMF method for an expensive, time-dependent cost function. The methodology described here is not restricted to the laminar flow problem, but can be applied to a wide range of fluids problems with complex geometries, unsteadiness and turbulence. Because of the portability of the method, it can be coupled to turbulent flow solvers based on LES, unsteady RANS, or DES (detatched-eddy simulation) for high Reynolds number flows. Use of the SMF method for time-dependent fluid dynamics problems avoids significant difficulties with the addition of constraints, implementation and data storage that arise with adjoint solvers. Even in problems in which gradients are available, the SMF method has many desirable properties. Using only the sign of the gradient, polling directions can be 'pruned' to reduce cost as in Abramson *et al.* (2003). The SMF method has proven to reduce the risk of quickly converging to a shallow local minimum, as is often the case in standard gradient methods.

Constrained optimization of both the upper and lower surfaces of the trailing-edge in laminar flow is currently underway, and initial results are very promising. Deformation of both sides of the airfoil allows for greater flexibility in the trailing-edge shape. In this case, the airfoil thickness is used as an optimization parameter, and the trailing-edge point is free to move in the vertical direction.

In future work, the SMF method will be applied for constrained optimization of the upper and lower surface of trailing-edge in fully turbulent flow using LES. Considerations of computational expense may lead us to incorporate a wall model (Wang & Moin 2002). In the turbulent case, the airfoil is not acoustically compact for all the frequencies of interest, and the cost function may need to be reconsidered. Alternatively, an approximation of the cost function can be used so long as it is well correlated with the true acoustic source function.

rum for collaboration, as well as Charles Audet and Petros Koumoutsakos for valuable discussions.

REFERENCES

ABRAMSON, M. A., AUDET, C. & DENNIS, JR., J. E. 2003 Generalized pattern searches with derivative information. *Tech. Rep.* TR02-10. Department of Computational and Applied Mathematics, Rice University, Houston, TX, to appear in *Mathematical Programming Series B.*

AUDET, C. & DENNIS, JR., J. E. 2000 A pattern search filter method for nonlinear programming without derivatives. *Tech. Rep.* TR00-09. Department of Computational and Applied Mathematics, Rice University, Houston TX.

AUDET, C. & DENNIS, JR., J. E. 2003 Analysis of generalized pattern searches. *SIAM Journal on Optimization* **13**, 889–903.

AUDET, C., DENNIS, JR., J. E. & MOORE, D. W. 2000 A surrogate-model-based method for constrained optimization. *AIAA Paper* 00-4891.

BLAKE, W. K. 1975 A statistical description of pressure and velocity fields at the trailing edge of a flat strut. DTNSRDC Report 4241. David Taylor Naval Ship R & D Center, Bethesda, Maryland.

BOOKER, A. J., DENNIS, JR., J. E., FRANK, P. D., SERAFINI, D. B., TORCZON, V. & TROSSET, M. W. 1999 A rigorous framework for optimization of expensive functions by surrogates. *Structural Optimization* **17**, 1–13.

CURLE, N. 1955 The influence of solid boundary upon aerodynmaic sound. *Proc. Royal Soc. Lond. A* **231**, 505–514.

HANSEN, N. & OSTERMEIER, A. 1996 Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. *Proc. of the 1996 IEEE Int'l. Conf. on Evolutionary Computation* pp. 312–317.

HOWE, M. S. 1988 The influence of surface rounding on trailing edge noise. *J. Sound & Vib.* **126**, 503–523.

JAMESON, A., MARTINELLI, L. & PIERCE, N. A. 1998 Optimum aerodynamic design using the navier-stokes equations. *Theoret. Comp. Fluid Dynamics* **10**, 213–237.

LOPHAVEN, S. N., NIELSEN, H. B. & SONDERGAARD, J. 2002 Dace: A MATLAB kriging toolbox version 2.0. *Tech. Rep.* IMM-TR-2002-12. Technical University of Denmark, Copenhagen.

MARSDEN, A. L., WANG, M., DENNIS, JR., J. E. & MOIN, P. 2003 Optimal aeroacoustic shape design using the surrogate management framework. Submitted for review.

MARSDEN, A. L., WANG, M. & KOUMOUTSAKOS, P. 2002 Optimal aeroacoustic shape design using approximation modeling. Annual Research Briefs-2001, Center for Turbulence Research,Stanford Univ./NASA Ames.

McKAY, M. D., CONOVER, W. J. & BECKMAN, R. J. 1979 A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**, 239–245.

WANG, M. & MOIN, P. 2000 Computation of trailing-edge flow and noise using large-eddy simulation. *AIAA J.* **38**, 2201–2209.

WANG, M. & MOIN, P. 2002 Dynamic wall modeling for large-eddy simulation of complex turbulent flows. *Physics of Fluids* **14**, 2043–2051.