AIAA 2000–2545

Airfoil Design Optimization Using Reduced Order Models Based on Proper Orthogonal Decomposition

Patrick A. LeGresley and Juan J. Alonso
Stanford University, Stanford, CA 94305

FLUIDS 2000 Conference and Exhibit
June 19–22, 2000/Denver, CO

AIAA 2000–2545

Airfoil Design Optimization Using Reduced Order Models Based on Proper Orthogonal Decomposition

Patrick A. LeGresley* and Juan J. Alonso†
Stanford University, Stanford, CA 94305

This paper presents a method for inviscid airfoil analysis and design optimization that uses reduced order models to reduce the cost of computation. Strong emphasis is placed on obtaining reasonably accurate solutions to the Euler equations with computational costs which are far lower than those required by traditional Computational Fluid Dynamics (CFD) techniques. The design procedure presented here begins by computing a series of flow solutions (*snapshots*) in which the design variables of interest are perturbed using a Design of Experiments approach. Proper Orthogonal Decomposition (POD) is then used to produce the optimal linear representation of these *snapshots* using a finite series of basis functions or modes. These basis modes are then used to construct arbitrary solutions to the Euler equations about modified airfoil geometries with very small computational expense. The flow solution problem is reduced in this way to a non-linear least squares fit problem with a small number of variables that can be solved efficiently. For design purposes, a gradient-based optimization procedure is used with the information supplied by the reduced order model. Results for both direct airfoil analysis and for an inverse design optimization problem are presented. Observations regarding the useability of this technique in a design environment are also discussed.

Nomenclature

a_j	generic coefficient of the j-th POD mode
E	total energy (internal plus kinetic)
$\mathbf{f}, \mathbf{g}$	Euler flux vectors
H	total enthalpy
M	number of modes used in approximation
p	static pressure
$R(x, x')$	autocorrelation function
$\mathbf{R}$	autocorrelation tensor, finite-volume residual
$\mathcal{R}$	autocorrelation matrix for method of snapshots
u	x-component of velocity
v	y-component of velocity
$\mathbf{u}$	arbitrary function to be generated
x	vector of independent variables
λ	Lagrange multiplier, an eigenvalue
η_i	coefficient of the ith mode in a function expansion
Ω	domain of interest
ρ	density
$\varphi^j(x)$	j-th POD basis mode
$\langle \cdot \rangle$	averaging operator
$\| \cdot \|$	L^2-norm

*Graduate Student, Student Member AIAA
†Assistant Professor, AIAA Member

Introduction

RECENT efforts in the field of aerodynamic shape optimization (ASO) have yielded remarkable improvements in our ability to design shapes with certain types of optimal behavior. By leveraging concepts from control theory, the computational cost of traditional gradient-based optimization schemes can be substantially decreased via solution of an adjoint equation. This procedure effectively yields the gradient with respect to an arbitrary number of design variables with the cost of a single flow and adjoint solution.[1–3] The adjoint procedure has undergone intensive scrutiny during the last few years and has matured to the point where both complex geometries and viscous flows are starting to be treated with reasonable confidence .[4–6]

In spite of these great contributions, the application of the adjoint procedure has been limited to the computations confined to the fields of aerodynamic and structural design. The fundamental reason being that the adjoint equations, boundary conditions, and gradient calculation formulae are cost function dependent, and therefore need to be re-derived every time the cost function changes. Moreover, it is not possible to treat arbitrary forms of the cost function, thus limiting the applicability of this procedure. Although a library of adjoint formulations for some of the more relevant cost functions of interest has been developed, the use of this procedure in the midst of more realistic design en-

vironments involving multiple disciplines and a large number of design constraints has been somewhat limited. In order to make the results of direct ASO valid, artificial constraints were imposed to include the effects of trade-offs with other disciplines. For example, in the case of aerodynamic wing design, planform and thickness constraints have often been imposed so that structural weight, fuel volume, and takeoff/landing requirements would not be adversely affected by changes in the wing shape dictated by the ASO procedure.

However, the importance of using high-fidelity modeling in the design of new aircraft configurations cannot be underestimated.[7–9] ASO methods have carved a niche in the transonic aircraft design area, and the focus of the attention must switch to the problem of integrating additional disciplines at the same level of fidelity. This objective poses fundamental challenges to the organization of the design procedure; a new paradigm of design must be generated to ensure that our goal of developing a truly high-fidelity multidisciplinary design environment can be attained. Work is required in many areas: new reorganizations of the design optimization process, development of software integration environments, enhanced optimization techniques, and novel low-cost, low-order approximations to high fidelity models will be required to minimize the formidable computational cost that this type of approach will generate. In addition, creative thinking in the area of computation of coupled sensitivities of responses that involve the interaction between two or more disciplines will be necessary.

The Proper Orthogonal Decomposition (POD) appears to be an excellent match to our need for low-cost, low-order approximations. Given a series of *snapshots* of a system (aerodynamics alone or a more realistic multi-disciplinary system) a simple algorithm produces a series of basis functions that are guaranteed to be optimal for the description of the system snapshots provided. However, the POD is nothing but a linear basis of the space that is being described, and therefore, it is an approximation to this design space and a certain amount of error will result when we use the POD to describe systems that fall outside of the trust region of the basis. Using the modes that form the POD basis, the governing equations of the flow (the Euler equations in our case) can be suitably projected to produce a much smaller set of non-linear equations that can be very efficiently solved using non-linear least squares solution methodologies. The full partial differential equations or their semi-discrete form no longer need to be solved.

Our long-term objective is to develop the various components of a truly high-fidelity multidisciplinary design environment where aerodynamics, structures, mission performance, and a realistic description of the aircraft/spacecraft model are included. To facilitate this goal, the work in this paper focuses on the low-order approximation of inviscid aerodynamic models for use in sample ASO problems. However, the intent here is broader; we are aiming to understand more clearly the ability of these low-order models to play a role in our proposed design environment, and, more importantly, to investigate the possibility of using low-order modeling to obtain coupled sensitivities which are extremely costly to derive without any sort of approximation. Additional numerical experiments are still necessary to determine the range of capabilities of these methods. Since POD-based reduced order models admittedly can only accurately represent the solution of problems that are *similar* to the information provided in the form of snapshots, a certain level of error is expected. However, they may find their place in variable fidelity design models with a multigrid-like structure where the *coarser* level information can be extracted from these models, while the *finest* levels are still provided by some of the high-fidelity models described above.

The paper presents a detailed description of our use of the POD approximation for purposes of both analysis and design of inviscid airfoils. Several examples of both subsonic and transonic flow analysis are shown with the basis modes obtained from a variety of different snapshots. The dependence of the quality of the solution on the number of modes used in its representation is also shown. Finally, some preliminary results of the application of this procedure to an airfoil inverse design test case are presented. Conclusions from these results are drawn and an outline of our current and future research efforts in this area is given.

Proper Orthogonal Decomposition (POD)

The POD procedure has been used in a wide range of disciplines, including fluid mechanics, random variables, image processing, signal analysis, data compression, process identification and control in chemical engineering, and oceanography.[10] In Ref.[11] the use of POD to construct reduced order models of synthetic jet actuators for flow control is described. Flow separation control experiments in which the control formulation was based on reduced order models from POD showed the low-order models to be quite effective. Reduced order modeling of the unsteady aerodynamic and aeroelastic behavior of a transonic airfoil using basis modes from POD is described in Ref.[12]

The details of the POD procedure and its properties have been well presented in Ref.[10] A brief description of the fundamental points is transcribed below. Needless to say, the analysis of the properties of the POD can be lengthy and have already been the subject of separate papers.

The POD is simply a procedure that provides an optimal linear basis for the reconstruction of multidimensional data (such as the result of a partial differential

equation) which allows a reduction in the order of the system under consideration from very large numbers (in the tens of millions for typical full-configuration Navier-Stokes fluid flow calculations) to very small ones (in the tens to hundreds). This order reduction is accomplished using a set of modes, or a modal decomposition, that will be explained below.

The POD has its roots in statistical analysis and has appeared in multiple disciplines with various names. Among others, the terms principal component analysis, empirical eigenfunctions, Karhunen-Loéve decomposition, and empirical orthogonal eigenfunctions have been used in the literature to describe the same concept. The nature of the POD is similar to the usual Fourier decomposition that is taught at the freshman calculus level and in multiple courses in approximation theory: a function of interest (i.e. the solution of a partial differential equation in a large mesh) is projected onto a set of basis functions (or modes) thus providing a finite set of scalar coefficients that *represent* the function in question. The POD provides a particular set of basis functions (or vectors in the finite dimensional sense) which are guaranteed to make up the optimal linear basis for the description of an ensemble of empirical (or numerical) observations of finite size.

The key idea is that, if the behavior of the system in question is such that a certain structure is present in the response, one may only need a small number of modes to represent the behavior of the overall system. If the number of modes required is of order 100, while the size of a typical mesh is of order $1,000,000$ then the description of the system of interest can be achieved at a much reduced cost while suffering only a small accuracy penalty.

We are then seeking finite dimensional representations of a function, $\mathbf{u}(\mathbf{x})$ in terms of a basis $\{\varphi_j(x)\}_{j=1}^{\infty}$ which allows an approximation to $\mathbf{u}$ to be constructed as follows:

$$\mathbf{u}_M = \sum_{j=1}^{M} a_j \varphi_j(x). \tag{1}$$

In order to construct this basis, assume that we have an ensemble of N *empirical* (these will be the outcome of several computational runs) solutions that we will denote by $\{\mathbf{u}^k\}$. We would like to choose $\{\varphi_j(x)\}_{j=1}^{\infty}$ so that these basis functions describe the functions in the ensemble $\{\mathbf{u}^k\}$ better than any other linear basis available. In order to formulate this problem mathematically, let's introduce the averaging operator, which is assumed to commute with the spatial integral used in the inner product operation.

The mathematical statement of optimality is that we should choose our basis functions so that they maximize the averaged projection of our ensemble of functions $\{\mathbf{u}^k\}$ onto φ. In other words,

$$\max_{\varphi} \frac{\langle |(\mathbf{u}, \varphi)|^2 \rangle}{\|\varphi\|^2}, \tag{2}$$

where $|\cdot|$ denotes the modulus and $\|\cdot\|$ is the L^2-norm given by

$$\|f\| = (f, f)^{\frac{1}{2}},$$

and the notation $(\cdot, \cdot)$ simply expresses the inner product of two functions over a pre-defined interval or domain.

Note that the solution of (2) would yield the best approximation to the ensemble of functions by a *single* function φ. However, the functional in (2) can have multiple local maxima, which would provide additional basis functions for the decomposition in (1). We, therefore, have a problem of calculus of variations in which we would like to maximize $\langle |(\mathbf{u}, \varphi)|^2 \rangle$ subject to the constraint that $\|\varphi\|^2 = 1$. This is a constrained optimization problem where the function to be maximized is given by:

$$J[\varphi] = \langle |(\mathbf{u}, \varphi)|^2 \rangle - \lambda(\|\varphi\|^2 - 1), \tag{3}$$

where λ is a Lagrange multiplier. A necessary condition for an extremum of this cost function is that for all variations $\varphi + \delta\psi, \delta \in R$, the following expression must hold:

$$\frac{d}{d\delta} J[\varphi + \delta\psi]|_{\delta=0} = 0.$$

From Eq. 3 and for real functions u, φ, and ψ, we have that

$$\frac{d}{d\delta} J[\varphi + \delta\psi]|_{\delta=0} = 2[\langle (u, \psi)(\varphi, u) \rangle - \lambda(\varphi, \psi)] = 0.$$

With some amount of algebra (shown in Ref[10]), and since the function ψ can be chosen arbitrarily, it can be shown that the basis functions we are seeking must satisfy:

$$\int_{\Omega} \langle u(x)u(x') \rangle \varphi(x') dx' = \lambda\varphi(x). \tag{4}$$

Therefore, the optimal POD basis that we were seeking is composed of the eigenfunctions $\{\varphi_j\}$ of the integral equation 4 above, whose kernel is the averaged autocorrelation function $\langle u(x)u(x') \rangle = R(x, x')$.

In the finite dimensional case that we encounter in numerical computations, the ensemble of functions $\mathbf{u}^k$ become a group of *N-dimensional* vectors instead of functions (or infinitely dimensional vectors), and the autocorrelation function becomes the autocorrelation tensor given by

$$\mathbf{R} = \langle \mathbf{u} \otimes \mathbf{u} \rangle.$$

Therefore, in finite-dimensional spaces, the integral eigenvalue problem mentioned above becomes

$$\mathbf{R}\varphi = \lambda\varphi.$$

Once the basis functions/vectors are computed from this eigenvalue problem (under most situations, $\mathbf{R}$ is guaranteed to have a complete set of eigenvectors), the member functions/vectors of the ensemble $\mathbf{u}^k$ can be decomposed as follows:

$$u(x) = \sum_{j=1}^{\infty} a_j \varphi^j(x). \tag{5}$$

The meaning of the corresponding eigenvalues in the problem depends on the function that is being approximated. For fluid flow problems, where the ensemble of functions represents various snapshots of the density, velocity, and pressure fields obtained via either experiment or computation, the values of the eigenvalues of the velocity modes, for example, represent twice the average kinetic energy contained in each mode. Therefore, an ordering of the eigenvectors according to the magnitude of their correponding eigenvalue allows for truncated models which contain the maximum amount of energy.

Note, however, that even in the event in which only a small number of modes is necessary to represent a new function which is not a member of the original ensemble, we still have to solve an eigenvalue problem of order equal to that of the original problem (that is the size of the autocorrelation matrix). The computational cost involved in this solution can be ameliorated by using iterative techniques to find the eigenvectors corresponding only to the largest eigenvalues with highest energy content. However, a more elegant procedure is available that reduces the cost of the solution of the eigenvalue problem to an essentially trivial amount. This method has been named the *method of snapshots* and is due to Sirovich.[13]

Assume that the numerical simulation from which we derive our ensemble of vectors is performed on a grid with a large number N of points. Furthermore, assume that the ensemble of functions contains n snapshots and is deemed to provide a reasonable description of the typical solutions to the system that we are seeking. According to the derivations previously presented, we would have to solve an $N \times N$ eigenvalue problem. It can be shown[10, 11, 13] that using the method of snapshots, the problem can be easily reduced to an $n \times n$ eigenvalue/eigenvector problem, which, under the premises of reduced order modeling should be quite a bit more manageable.

Using the method of snapshots, the resulting elements of the modified autocorrelation matrix for each of the density, velocity, and pressure fields is given by

$$\mathcal{R}_{ij} = \frac{1}{M} \int_{\Omega} u_i(x,y) u_j(x,y) dxdy \tag{6}$$

where u_i corresponds to the i-th snapshot of the solution, $i, j = 1, 2, \ldots, M$, and M is the number of snapshots provided. The integral simply represents the inner product of two of the snapshots, i and j, resulting in $\mathcal{R}$, a non-negative definite, symmetric matrix which therefore has non-negative eigenvalues and a corresponding full set of orthogonal eigenvectors. The eigenvectors of $\mathcal{R}$ are computed, as an intermediate step, to determine the actual POD modes

$$\mathcal{R}a = \lambda a \tag{7}$$

The POD basis functions can now be calculated as

$$\varphi^K = \sum_{i=1}^{M} a_i^K u_i(x,y) \quad K = 1, 2, \ldots, M \tag{8}$$

where a_i^K is the ith element of eigenvector a corresponding to the eigenvalue λ_K. The resulting POD modes are fully orthogonal, and are normalized to simplify future operations:

$$(\varphi^K, \varphi^{K'}) = \begin{cases} 1 & : \quad K = K' \\ 0 & : \quad K \neq K' \end{cases} \tag{9}$$

In the case of the Euler equations, the procedure is repeated for each of the primitive variables (ρ, u, v, p) that are needed to compute the conservative state vector (ρ, ρu, ρv, ρE) required in our Euler flow solver. Once these basis modes have been obtained, we can, to a certain degree of accuracy, expand the flow solution about an arbitrary airfoil shape. For example, the density field of these solutions will be expanded in the form

$$\rho(x,y) = \sum_{i=1}^{M} \eta_{\rho_i} \varphi_\rho^i, \tag{10}$$

where the subscripts indicate the fact that the coefficients of the expansion, η_{ρ_i} are those particular to the expansion of the density field.

After this quick outline of the POD procedure, the first step in the POD process is the preparation of an ensemble of snapshots from which a set of basis modes will be derived. Because of the relationship between the snapshots and the resulting modes described above, it is important to use snapshots that, as closely as possible, resemble the solutions that we will try to obtain. In this work, the snapshots are produced by computing a series of flow solutions using FLO82, the two-dimensional, cell-centered Euler solver of Jameson[14] for the analysis of arbitrary airfoil geometries. The different solutions are obtained by perturbing a baseline airfoil using a parameterization of its geometry and a design of experiments approach.

It must be finally mentioned that modifications to the POD procedure described above and in Ref.[10] can be included to produce the optimal basis to represent a set of function snapshots and their derivatives with respect to a series of design parameters. The resulting eigenvalue problem is similar to the one in Eq. (7) and will be studied in the future with gradient information provided by efficient adjoint-based methodologies.

Flow Analysis Procedure

Traditional uses of POD expansions in fluid dynamics have focused on the development of a series of POD modes which are later used to project the full incompressible Navier-Stokes equations. This procedure transforms a complicated non-linear partial differential equation (the Navier-Stokes equations) into a number (equal to the number of modes used) of non-linear ordinary differential equations that can be integrated in time to describe the evolution of the fluid system. This type of approach is currently being used to analyze the behavior of turbulent flows with a small number of modes.

In our case, however, we are not particularly interested in the time evolution of the flow solution at hand. Our flow solutions are *steady* and, for airfoil analysis and design purposes, the *evolution* parameters, rather than time, will be identified with the coefficients of a surface parameterization that allows us to change the shape of the geometry of interest. In a nutshell, we are seeking to use POD modes to describe the flow solutions of arbitrary airfoil shapes using the information contained in the POD modes, and the fact that the resulting solution must satisfy, as closely as possible, both the governing equations of the flow, and its wall and far-field boundary conditions.

The approach we have chosen to take is based on the well-known *finite-volume* procedure which is often used to discretize the governing equations of the flow. Let p, ρ, u, v, H, and E denote the pressure, density, cartesian velocity components, total enthalpy, and total energy respectively. Consider an arbitrary control volume Ω with boundary $\partial\Omega$. The equations of motion of the fluid can then be written in integral form as

$$\frac{d}{dt}\iint_{\Omega} \mathbf{w}\; dx\, dy + \oint_{\partial\Omega} (\mathbf{f}\, dy - \mathbf{g}\, dx) = \mathbf{0}, \qquad (11)$$

where $\mathbf{w}$ is the vector of conserved flow variables

$$\mathbf{w} = \begin{Bmatrix} \rho \\ \rho u \\ \rho v \\ \rho E \end{Bmatrix},$$

and $\mathbf{f}$, $\mathbf{g}$ are the Euler flux vectors

$$\mathbf{f} = \begin{Bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uH \end{Bmatrix}, \mathbf{g} = \begin{Bmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho vH \end{Bmatrix}.$$

Also, for an ideal gas, the equation of state may be written as

$$p = (\gamma - 1)\ \rho\ \left[E - \frac{1}{2}(u^2 + v^2)\right].$$

Applying Eq. 11 independently to each cell in the mesh we obtain a set of ordinary differential equations of the form

$$\frac{d}{dt}(\mathbf{w}_{ij}\, V_{ij}) + \mathbf{R}(\mathbf{w}_{ij}) = \mathbf{0}, \qquad (12)$$

where V_{ij} is the volume of the i, j cell and the residual $\mathbf{R}(\mathbf{w}_{ij})$ is obtained by evaluating the flux integral in Eq. 11. In the steady state, the time derivative term drops out and we are left with

$$\mathbf{R}(\mathbf{w}_{ij}) = \mathbf{0},$$

which already incorporates the wall and far-field boundary conditions in the calcuation of the boundary fluxes at the edges of the domain.

Since, using a POD expansion of the kind of Eq. 10, both the primitive and conservative flow variables can be considered to be functions of the expansion coefficients, η_i, we can develop the residual of the Euler equations into a local non-linear function of each and every one of the expansion coefficients, η_l, for each of the primitive fields of interest

$$\mathbf{R}(\eta_l) = \mathbf{0}, \;\; l = 1, \ldots, M. \qquad (13)$$

Notice the drastic reduction in the complexity of the problem: from a total number of unknowns equal to $4 \times N$, where N is the number of nodes in the grid, to $4 \times M$, where M is the number of modes used for the expansion of each of the solution fields. This is where the drastic complexity reduction derived from POD can be seen. For a typical two-dimensional Euler calculation, $4 \times 161 \times 33 = 21,252$ unknowns are solved for, while for a calculation based on POD using 15 modes, the total number of unknowns is only $4 \times 15 = 60$. When this procedure is used in the solution of three-dimensional flows, the reduction in computational cost becomes compelling.

Given that an exact solution to Eq. 13 will typically not be possible since we have drastically decreased the number of free parameters in the problem, we have chosen to define the following cost function

$$I_{POD} = \sum_{l} \mathbf{R}^2(\eta_l), \qquad (14)$$

where the summation, l, is over all the cells in the domain. Notice that the residuals for the four governing equations (continuity, momenta, and energy) are weighed appropriately so that the residuals of any equation are not given uneven weighting. This approach renders our problem well-posed and defines the solution that, with a given series of modes, most closely satisfies the equations of motion. The definition of I_{POD} guarantees that both the governing equations and the solid wall boundary conditions are satisfied in the least-squares sense. Forcing terms that guarantee compliance with the far-field boundary conditions for this problem are added to the residual, but omitted here for simplicity.

The problem of finding the solution to our problem has now been reduced to obtaining the least-squares minimizer of Eq. 14 which can be done readily with a variety of methods. For this work we have chosen to use the method of Levenberg-Marquardt which combines elements of both quasi-Newton and steepest descent optimization procedures. Using this method, typical solutions for problems involving $15-30$ modes can be found inexpensively in only a few iterations (typically less then 10.)

This projection procedure dramatically reduces the cost of flow calculation from having to solve a complete set of hyperbolic partial differential equations (Euler equations) to the solution of a small set of nonlinear equations. Using this reduced order model, a gradient-based approach may be used for optimization of an arbitrary cost function with respect to a variety of design variables that can be used to parametrize the changes in geometry of an airfoil. Possible design variables that can be used include camber and thickness distributions, angle of attack, and series of distributed Hicks-Henne bump functions for full aerodynamic shape optimization.

Results

In this paper, the POD methodology is used to evaluate the ability of reduced order models to accurately represent the inviscid aerodynamics of arbitrary airfoil shapes. In particular, we are interested in the ability of POD-based models to provide solutions of sufficient accurary that can be used to *guide* the design process, even if occasional solutions of the full Euler equations are necessary. The projection / flow solution technique was then used with an inverse design procedure where a pre-specified target pressure distribution is achieved by variations in the airfoil surface geometry.

Flow Computations

For purposes of evaluating our projection algorithm, the ability of the method to accurately compute flow solutions was validated. Because any of the snapshots used in constructing a set of modes can be *exactly* represented by the modes, the projection algorithm should be able to compute the exact solution for any of the geometries represented in the snapshots.

We parameterize the airfoil surface with a series of Hicks-Henne bump functions,[15] which make smooth changes in the geometry,

$$b(x) = \{sin[\pi x^{log(1/2)/log(t_1)}]\}^{t_2} \quad 0 \le x \le 1, \qquad (15)$$

where the maximum of the bump is located at $x = t_1$ and the parameter t_2 controls the width of the bump. Airfoil geometries for the snapshots were computed by taking two baseline airfoils (the RAE 2822 and NACA 1413) and adding 14 bump functions, 7 each on the upper and lower surfaces. The bumps were distributed evenly along the chord with an amplitude of 0.1% of the chord. Flow solutions for the baseline airfoils, plus 28 flow solutions for each of the modified airfoils, were computed using FLO82 and used to form the basis modes for ρ, u, v, and p, each with a total of 30 modes.

Using the projection technique previously described, the flow solution for the RAE 2822 airfoil was projected using the first n of these 30 modes, where n was varied from 1 to 30, and the modes are numbered in order of decreasing eigenvalue. In Fig. 1, the results for pressure from the projected solution made using all 30 modes is shown. As expected, the exact solution for the RAE 2822 airfoil is recovered. Also shown is the convergence history of the lift and drag coefficients as more modes are added to the computation. Note that the lift coefficient quickly converges with relatively few modes, while the drag coefficient requires significantly more modes to reach the exact solution.

These results validate the projection technique, but not the ability of POD based models to be used for design, since we have only reconstructed information we already knew. For design we must be able to compute, to a reasonable degree of accuracy, the flow solutions for new geometries for which we have calculations for similar geometries.

A set of basis modes was computed from a NACA 4412 airfoil by placing a total of 30 of the Hicks-Henne bump functions with a maximum amplitude of 0.1% of the chord over the upper and lower surfaces. The results of projecting the flow solution for a NACA 3413 section, which has different maximum camber and thickness ratios from the NACA 4412, are shown in Fig. 2. Because the projected geometry was not in the snapshots used to compute the modes, we cannot expect to get an exact solution. However, the pressure contours and surface pressure distributions in Fig. 2 indicate the projected solution is quite good, and certainly sufficient for our goal of guiding the design process. Again, the lift coefficient convergence is quite rapid, while the drag coefficient takes significantly more modes. Note that nothing about POD implies that using more of the modes that are available will result in a more accurate pressure distribution at the surface of the airfoil. Additional modes will add to the overall accuracy when the entire domain is considered, of which the airfoil is only a boundary.

A partial set of basis modes for pressure are shown in Fig. 3. These are the first 8 modes (corresponding to the 8 largest eigenvalues) computed from the RAE 2822 and NACA 1413 airfoils, and used in the computations shown in Fig. 1. The first mode looks very much like the pressure contour itself, while subsequent modes are concentrated at various locations around the surface. It was mentioned that the errors in Fig. 1 when not all 30 modes were used in the projection, particularly for the drag coefficient, were due to the fact that the predominant error was located in

in places with greater projected area in the drag direction. Notice that the modes beyond the first are highly concentrated, with large magnitudes in a small area at the airfoil surface with a significant projected area in the drag direction, but are relatively flat elsewhere. It is this characteristic of the modal decomposition, which occurs when the snapshots represent changes in geometry, that allows us to project changes in geometry.

At a Mach number of $M = 0.50$ the flows are shock free and relatively linear. However, for transonic flow the presence of shocks, and the fact that they move slightly as the airfoil geometry is perturbed to create the snapshots, may lead to difficulties in projecting a solution. A set of 30 modes was computed by using snapshots of the RAE 2822 and NACA 1413, as well as the perturbed variations, at $M = 0.75$. The original RAE 2822 airfoil was then projected, and the results are shown in Fig. 4. Again, as expected, we recover the exact solution by using all 30 modes. Also, we again see quite rapid convergence for the lift coefficient, while the drag coefficient converges more slowly. The fact that the lift coefficient is almost exact with just a single mode is a coincidence - note that the corresponding drag coefficient shows a significant error.

The sensitivity of the lift and drag coefficients to the number of modes utilized in the projection is an important piece of information if these models are to be used efficiently in design. For example, if one is interested in a target lift coefficient, the results shown indicate that relatively few of the modes can be used, and one will still get a satisfactory solution. However, if drag plays a role in the design criteria, such as drag minimization or lift-to-drag ratio, more of the modes will be needed to get a satisfactory projected solution. Further investigation into the numbers of modes required for each of the individual flow quantities (ρ, u, v, and p) may lead to even greater efficiencies for these reduced order models.

Inverse Design

The ability of POD based design models to provide gradients for design optimization was investigated by performing an inverse design in which a target pressure distribution was specified, and the geometry is modified to achieve such a distribution. To parameterize the airfoil surface the same Hicks-Henne bumps previously described were used, where the amplitude of the bumps are the design variables. An inverse design cost function is defined as

$$I_{ID} = \int_S (p - p_T)^2 ds, \tag{16}$$

where p_T is the target or specified pressured distribution. Gradients of the cost function in Eq. 16 with respect to the design variables were obtained by finite differencing of the POD model.

A set of basis modes using the RAE 2822 and NACA 1413, and perturbed variations, with a total of 30 modes were used for the inverse design problem (the same modes used for the projection results shown in Fig. 1). The pressure distribution for the RAE 2822 airfoil was specified as the target. A parameterization of the surface was made using a total of 10 bump functions, 5 each on the upper and lower surface. The starting shape for the design was produced by making arbitrary changes to the RAE 2822 airfoil, as shown in Fig. 5a. The initial pressure distribution, computed by projecting the initial geometry, is shown along with the target pressure in Fig. 5b.

After one design iteration, the airfoil cross section appears to have been changed in the wrong direction - away from the known solution. However, the cost function has in fact been reduced. The second iteration moves the lower surface towards the target geometry again. At this point the pressure distribution is still quite rough.

At the end of the third iteration, the cross section is moving out toward the target solution, the aft camber at the trailing edge is more pronounced and the pressure distribution has smoothed out significantly. The cost function after three design iterations has been reduced by a factor of five from the initial design.

After four design iterations (Fig. 6a-b), the design is quite close to the target in terms of geometry and pressure. Almost an order of magnitude decrease in the cost function has been achieved from the previous iteration, and almost two orders of magnitude from the initial design.

In subsequent iterations, the cross section in the aft portion of the airfoil is actually moving away from the true solution as the pressure distribution is smoothed out. The final design at iteration 40 shows the pressure distribution was matched fairly well, although there appear to be difficulties in resolving the pressure distribution near the trailing edge. The overall reduction in the inverse design cost function was approximately two orders of magnitude.

The use of reduced order models introduces errors into the calculation of the gradients used in the inverse design. Despite the errors that will inevitably exist, as long as sensitivities and gradients are generally correct, a significant advantage has been obtained in inexpensively obtaining their values. However, the more accurate we can compute gradients while remaining within the context of reduced order models, the better the final solution will reflect the true solution. When using finite differencing, there is a compromise between taking a small enough step to obtain more accurate gradients and taking a large enough step to be captured by the model. Therefore, a partial or fully analytical derivative based on the models presented here, or by supplementing the function models with models of the derivatives with respect to design

parameters as computed from adjoint methodologies, could significantly increase the usability of these types of reduced order models for design.

Although this problem was setup so that the exact solution of the inverse design problem could be respresented exactly by the modes used, a true design process would not have the benefit of such *a priori* knowledge of the outcome. However, a coupled optimizer could be used to refine the POD based model by adding additional snapshots, which would come from flow solutions that might be computed during the optimization itself, to refine the model as the design evolved. In this case, arbitrary initial conditions and target pressure distributions could be treated.

Conclusions

The theory of POD based models and their application to projecting arbitrary airfoils from snapshots of other similar airfoils was presented. The cost of computation of a flow solution is reduced from that of solving a set of partial differential equations to the solution of a coupled set of non-linear equations for the steady flows in this work. The number of unknowns to be solved for was reduced by approximately three orders of magnitude for a typical two-dimensional inviscid flow calculation. A simple solution method based on a combination of quasi-Newton and steepest descents was used to compute approximate steady solutions for arbitrary airfoil geometries.

An inverse design problem in which a target pressure distribution is specified and the corresponding airfoil shape is modified to produce that distribution was performed using information only from a reduced order model. Although this reduced order model based design did not attain the exact solution, it moved quite close to the target solution in only a few design steps. Although this method may require more function evaluations than the adjoint method, it potentially offers advantages in robustness, handling of constraints, and can be used for coupled sensitivity analysis of systems where multiple disciplines interact.

Future work will address means to accelerate the projection solution process, perhaps by the same multigrid techniques used to speed convergence in higher fidelity flow solutions. Also, the selection of snapshots to form the best set of modes as possible - in terms of the accuracy of solutions computed using them, as well as for possible re-use in other design work - needs to be addressed. Further integration of the design optimization routine with the POD method will allow arbitrary design problems to be posed, with no prior knowledge of the solution, and allow the POD model to evolve during the design process.

Ultimately, this work will be extended to 3-D aerodynamic flows, both invsicid Euler flows and the full Navier-Stokes equations. By then applying POD to other disciplines, such as structures, mission performance, and ultimately the entire system, a truly multidisciplinary design environment can exist with an acceptable computational cost and at a higher level of fidelity than is currently possible.

References

[1]A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3:233–260, 1988.

[2]A. Jameson and J.J. Alonso. Automatic aerodynamic optimization on distributed memory architectures. *AIAA paper 96-0409*, AIAA 34th Aerospace Sciences Meeting, Reno, Nevada, January 1996.

[3]J. Reuther, A. Jameson, J.J. Alonso, M.J. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers. *AIAA paper 97-0103*, AIAA 35th Aerospace Sciences Meeting, Reno, Nevada, January 1997.

[4]A. Jameson, N. Pierce, and L. Martinelli. Optimum aerodynamic design using the Navier-Stokes equations. *Theoretical Computational Fluid Dynamics*, 10:213–237, 1998.

[5]J. Reuther, A. Jameson, J.J. Alonso, M.J. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part i. *Journal of Aircraft*, 36:51–60, 1999.

[6]J. Reuther, A. Jameson, J.J. Alonso, M.J. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part ii. *Journal of Aircraft*, 36:61–74, 1999.

[7]R. Braun, P. Gage, I. Kroo, and I. Sobieski. Implementation and performance issues in collaborative optimization. *AIAA paper 96-4017*, AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Bellevue, Washington, September 1996.

[8]J. Reuther, J.J. Alonso, J.R.R.A Martins, and S.C. Smith. A coupled aero-structural optimization method for complete aircraft configurations. *AIAA paper 99-0187*, AIAA 37th Aerospace Sciences Meeting, Reno, Nevada, January 1999.

[9]J. Sobieszczanski-Sobieski and R.T. Haftka. Multidisciplinary aerosapce design optimization: Survey of recent developments. *AIAA paper 96-0711*, AIAA 34th Aerospace Sciences Meeting, Reno, Nevada, January 1996.

[10]P. Holmes, J.L. Lumley, and G. Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, 1998.

[11]O.K. Rediniotis, J. Ko, X. Yue, and A.J. Kurdila. Synthetic jets, their reduced order modeling and applications to flow control. *AIAA paper 99-1000*, AIAA 37th Aerospace Sciences Meeting, Reno, Nevada, January 1999.

[12]K.C. Hall, J.P. Thomas, and E.H. Dowell. Reduced order-modelling of unsteady small-disturbance flows using a frequency-domain proper orthogonal decomposition technique. *AIAA paper 99-0655*, AIAA 37th Aerospace Sciences Meeting, Reno, Nevada, January 1999.

[13]L. Sirovich. Turbulence and the dynamics of coherent structures. I - coherent structures. II - symmetries and transformations. III - dynamics and scaling. *Quarterly of Applied Mathematics*, 45:561–571,573–590, 1987.

[14]A. Jameson. Multigrid algorithms for compressible flow calculations. In W. Hackbusch and U. Trottenberg, editors, *Lecture Notes in Mathematics, Vol. 1228*, pages 166–201. Proceedings of the 2nd European Conference on Multigrid Methods, Cologne, 1985, Springer-Verlag, 1986.

[15]R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15:407–412, 1978.

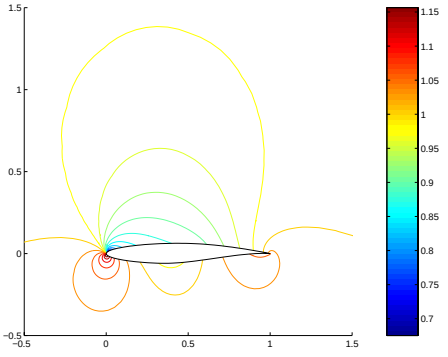

a) Exact Pressure Contours

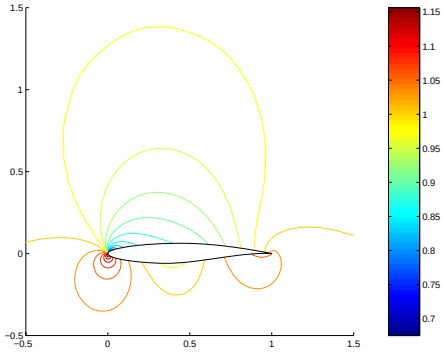

b) Projected Pressure Contours

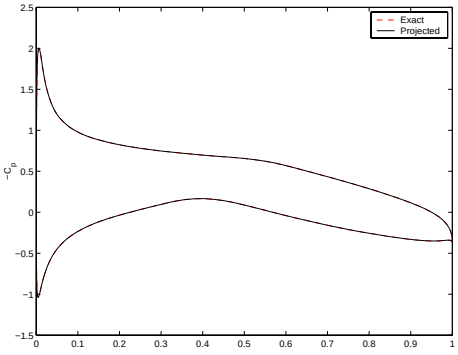

c) Surface Pressure

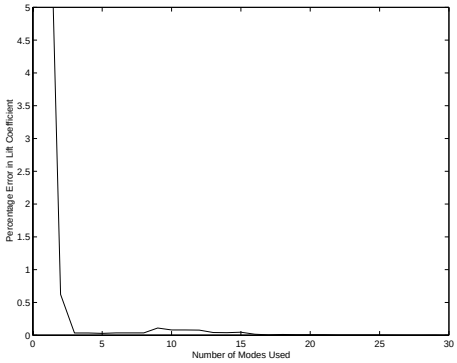

d) Error in Lift Coefficient

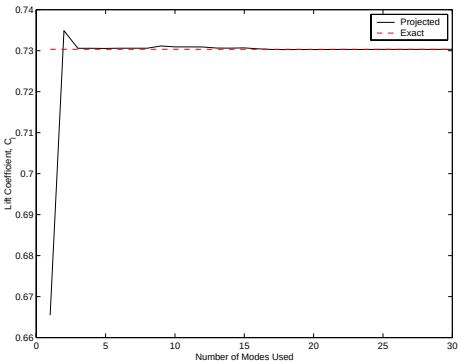

e) Lift Coefficient Convergence

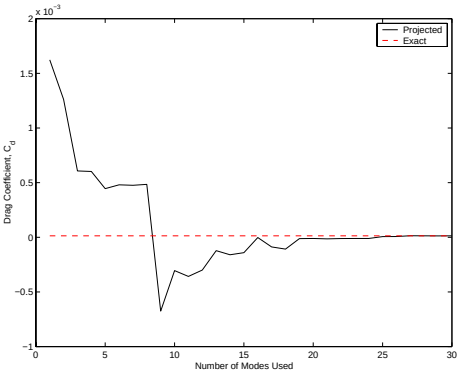

f) Drag Coefficient Convergence

Fig. 1 Projection Results for RAE 2822 Airfoil, $M = 0.50$**.**

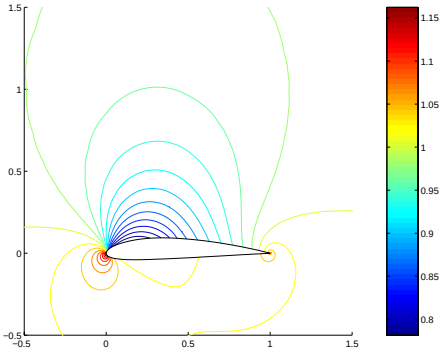

a) Exact Pressure Contours

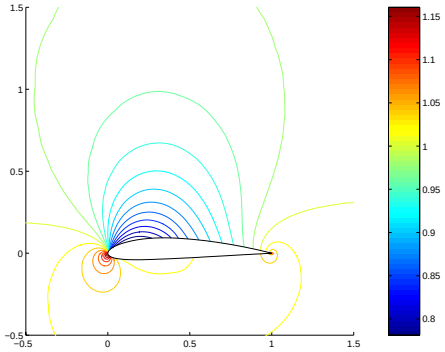

b) Projected Pressure Contours

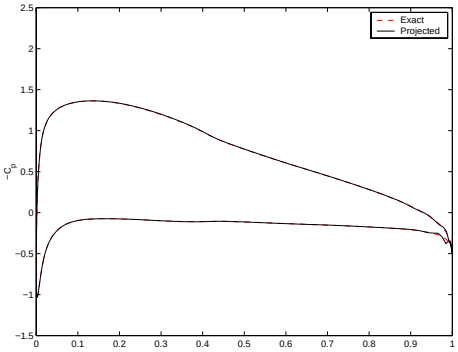

c) Surface Pressure

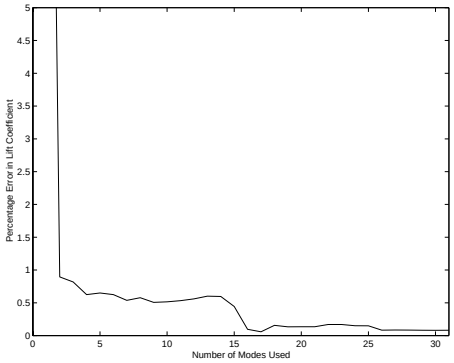

d) Error in Lift Coefficient

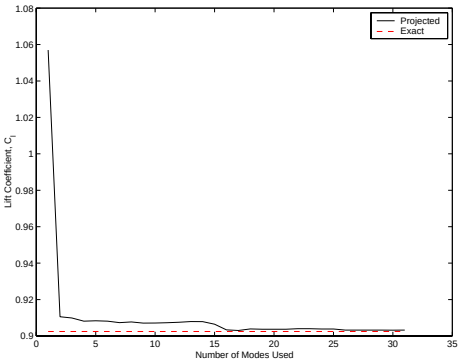

e) Lift Coefficient Convergence

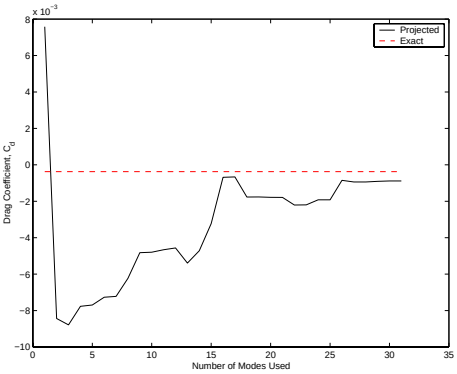

f) Drag Coefficient Convergence

Fig. 2 Projection Results for NACA 3413 Airfoil, $M = 0.50$**.**

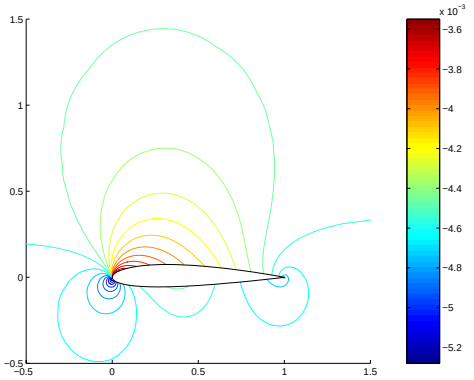

a) Pressure Mode 1

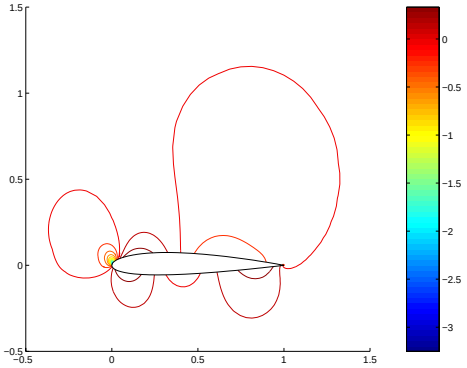

b) Pressure Mode 2

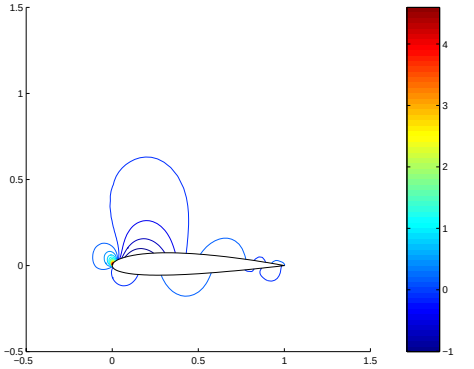

c) Pressure Mode 3

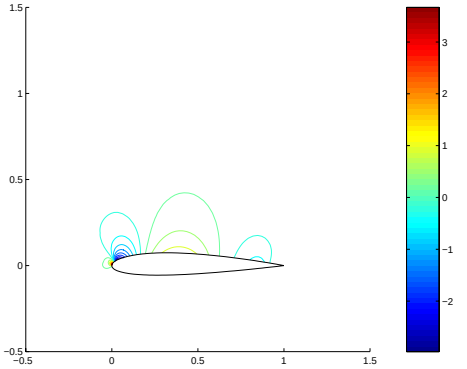

d) Pressure Mode 4

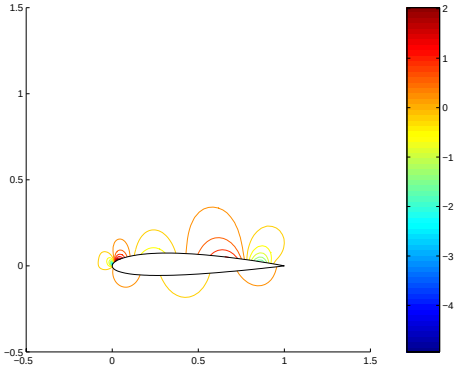

e) Pressure Mode 5

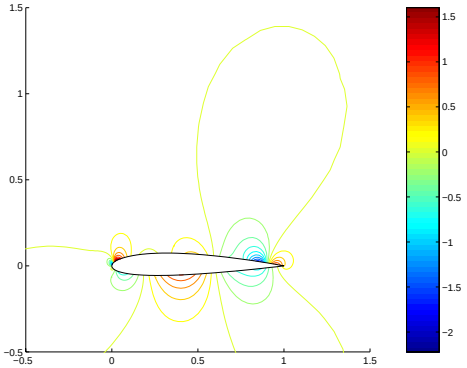

f) Pressure Mode 6

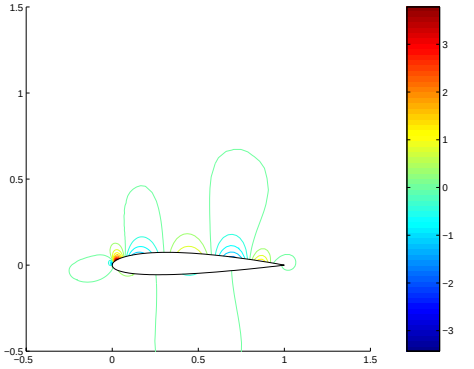

g) Pressure Mode 7

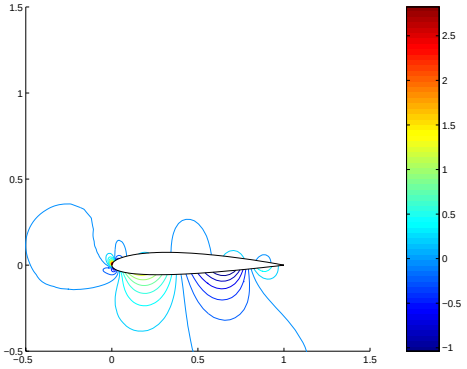

h) Pressure Mode 8

Fig. 3 Sample Pressure Modes for a Series of Snapshots from a NACA 1413 Airfoil.

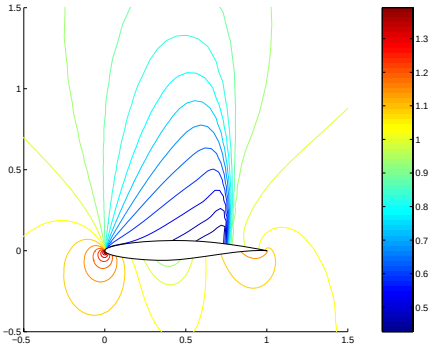

a) Exact Pressure Contours

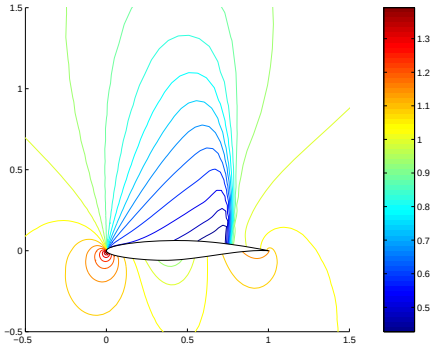

b) Projected Pressure Contours

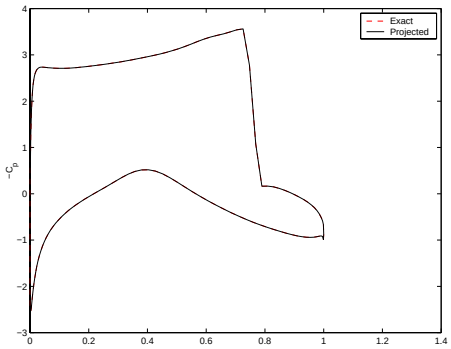

c) Surface Pressure

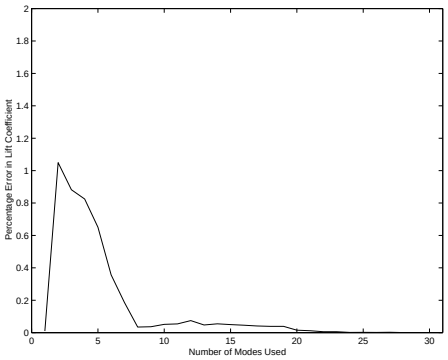

d) Error in Lift Coefficient

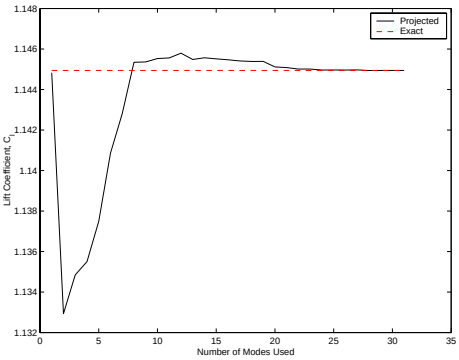

e) Lift Coefficient Convergence

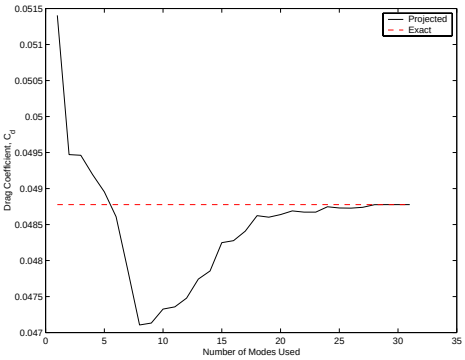

f) Drag Coefficient Convergence

Fig. 4 Projection Results for RAE 2822 Airfoil, $M = 0.75$**.**

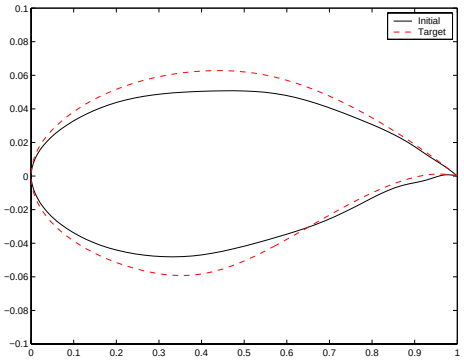

a) Geometry for Initial Design

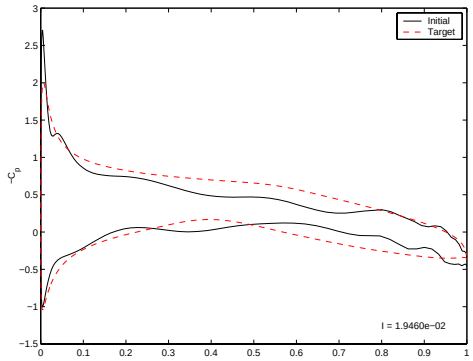

b) Pressure Distribution for Initial Design

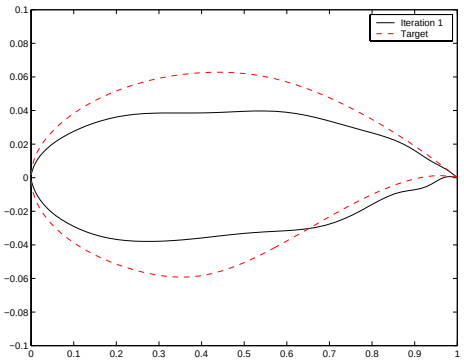

c) Geometry at Design Iteration 1

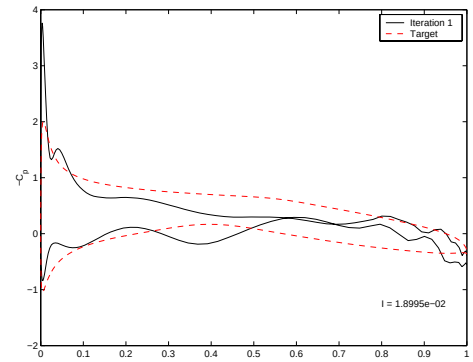

d) Pressure at Design Iteration 1

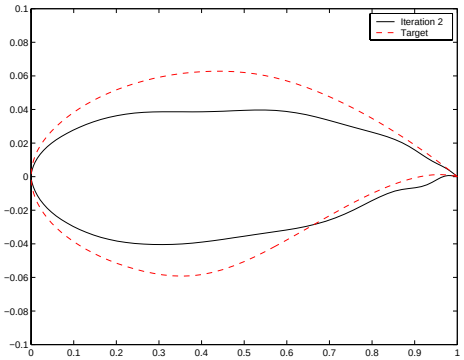

e) Geometry at Design Iteration 2

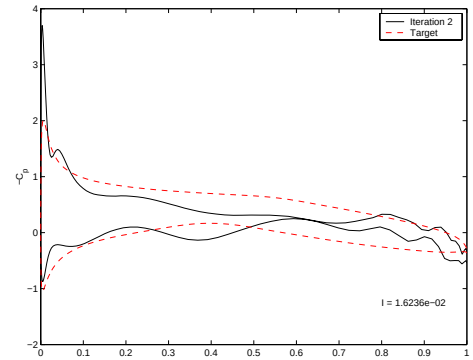

f) Pressure at Design Iteration 2

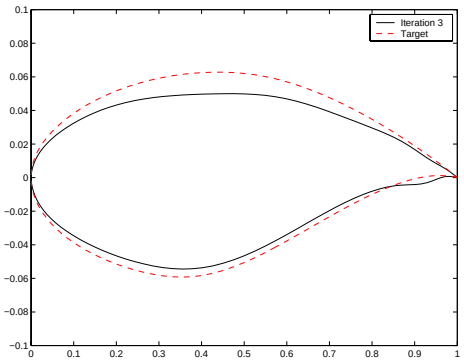

g) Geometry at Design Iteration 3

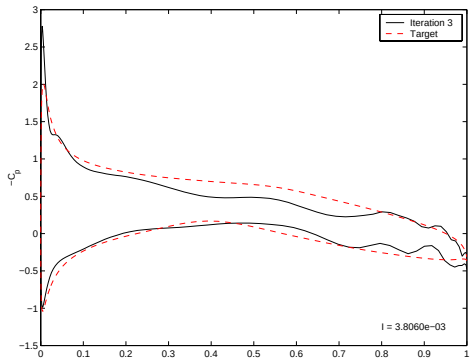

h) Pressure at Design Iteration 3

Fig. 5 Inverse Design Iterations for RAE 2822 Target Pressure, $M = 0.50$**.**

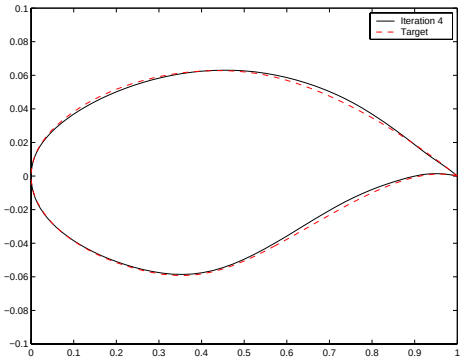

a) Geometry at Design Iteration 4

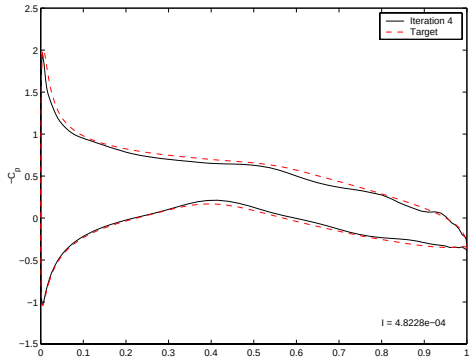

b) Pressure at Design Iteration 4

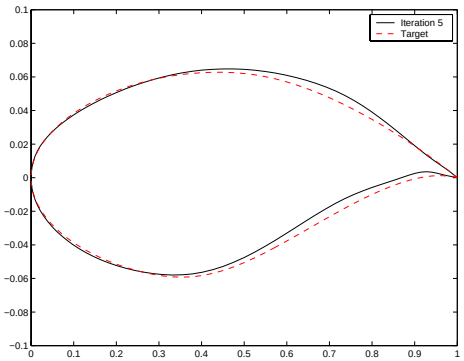

c) Geometry at Design Iteration 5

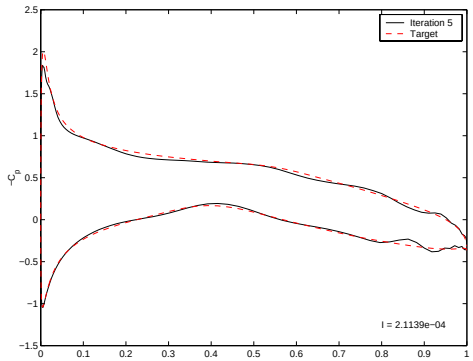

d) Pressure at Design Iteration 5

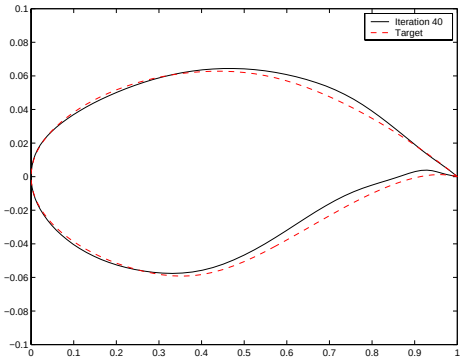

e) Geometry after Last Design Iteration

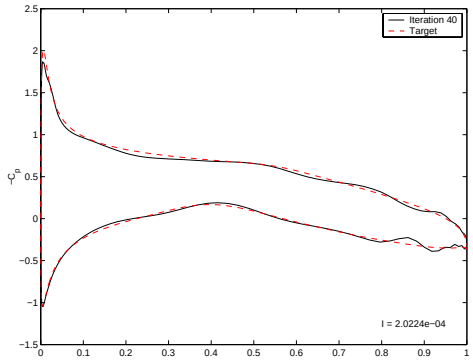

f) Pressure after Last Design Iteration

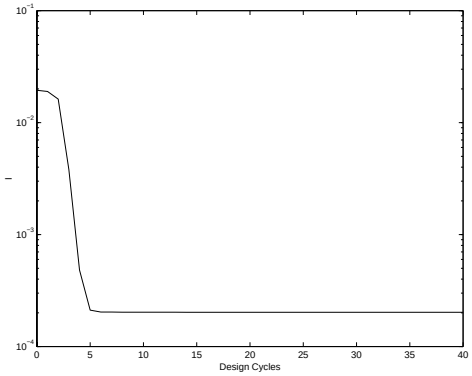

g) Cost Function Convergence

Fig. 6 Inverse Design Iterations and Convergence for RAE 2822 Target Pressure, $M = 0.50$**.**