Hybrid Inverse Design Method for Nonlifting Bodies in Incompressible Flow

Benjamin A. Broughton* and Michael S. Selig[†]
University of Illinois at Urbana–Champaign, Urbana, Illinois 61801

A methodology for the inverse design of nonlifting axisymmetric and nonaxisymmetric bodies in incompressible flow is presented. In this method, an inverse design approach based on conformal mapping is used to design a set of airfoils in isolation. These airfoils are then assembled into a three-dimensional body and the flow over the body is calculated using a panel method. The inverse design parameters for the isolated airfoils are adjusted by a multidimensional nonlinear solver to achieve the desired aerodynamic properties on the three-dimensional body. The method can be used with fairly complex geometries, such as bodies in the presence of a wing or keel. The suitability and performance of several numerical schemes are compared in the paper. Several examples are presented that demonstrate the flexibility of the design method when applied to various representative design problems and they also show the ability of the method to match a known velocity distribution.

Nomenclature

A	=	area
$\boldsymbol{B}$	=	approximate Jacobian
c	=	chord length
c_i	=	control airfoil
$\boldsymbol{F}$	=	vector of functional relations
f_i	=	functional relations to be zeroed
h	=	surface displacement
$\boldsymbol{J}$	=	Jacobian matrix
$\dot{m}$	=	mass-flux
n	=	dimension of nonlinear system
$\boldsymbol{n}$	=	surface normal
Re	=	Reynolds number
s	=	arc length
s_i	=	control section
$\tilde{s}$	=	arc length relative to beginning of section
u	=	normal velocity
u_t	=	transpiration velocity
V_∞	=	freestream velocity
$\dot{v}$	=	volume flux
$\tilde{v}_i$	=	relative velocity along a splined segment
x_i	=	unknown variables
$\boldsymbol{x}$	=	vector of unknown variables
$\delta\boldsymbol{x}$	=	correction vector
α	=	angle of attack, deg
α^*	=	segment design angle of attack
ϕ	=	arc limit
$\tilde{\phi}$	=	arc limit relative to beginning of segment
σ	=	source strength
ΔV	=	velocity difference over a segment normalized by the freestream velocity

Subscripts

i	=	segment number on an airfoil or the equivalent segment of the body cross-section
p	=	unperturbed surface panel
t	=	transpiration

Introduction

EXTENSIVE research has been conducted throughout the past century towards the design of more economical and efficient aircraft by improving their aerodynamic performance. Much of this research focused on improved wing design, which in turn led to powerful inverse design methods for airfoils in isolation and lifting surfaces. Fuselages and other nonlifting bodies received somewhat less attention, to some extent due to the more complex three-dimensional nature of the flow over these bodies and the fact that it is more difficult to create a generalized scheme that can model a large variety of geometries. Today, there are relatively few inverse design schemes available that are suitable for the design of these types of bodies. In the last few decades, various studies on the aerodynamics of nonlifting bodies have, however, underlined the need for more general inverse design schemes.

Under certain flow conditions and for simple streamlined bodies or forebody shapes, there actually exist optimal shapes that can be directly derived from linearized flow theories. One example is that of axisymmetric bodies in supersonic flow, where the total drag is dominated by wave drag. Examples of such optimal bodies include those derived in the work of von Karman [1], Sears [2] and Haack [3]. The drag of streamlined bodies in shock-free subsonic flow, the topic of this paper, is mainly due to viscous effects that are in turn strongly related to the boundary-layer development. The complexities of this type of flow makes it difficult to derive generalized equations that describe the shapes of bodies with minimum drag. Furthermore, the designer often requires a more complex shape than a simple body of revolution. More general inverse design schemes are therefore required.

In contrast to the limited number of methods available for the inverse design of nonlifting bodies such as fuselages, a number of highly refined methods have been developed for the design of isolated airfoils and even multi-element airfoil geometries in two-dimensions. In fact, inverse aerodynamic design of airfoils has been a subject of investigation for a long time, dating as far back as the early airfoil research of the 1930s and 40s [4,5]. Since the 1970s, some highly specialized airfoils have been developed using more modern inverse aerodynamic design techniques [6], to a large extent made possible by the rapid advances in computer technology. Early methods were limited to single design points [7–9]. Over the last few

*Graduate Research Assistant, Department of Aerospace Engineering, 306 Talbot Laboratory, 104 S. Wright Street; e-mail: bbrought@uiuc.edu. Student Member AIAA.
†Associate Professor, Department of Aerospace Engineering, 306 Talbot Laboratory, 104 S. Wright Street; e-mail: m-selig@uiuc.edu. Senior Member AIAA.

decades, however, these methods have been refined to allow inverse design with multiple design points [10], and even the direct prescription of boundary-layer growth parameters [11,12].

Eppler [10] developed an airfoil design theory based on conformal mapping that is today still a popular practical airfoil design method. His method gives the designer the ability to divide the airfoil into segments. A design angle of attack relative to the zero-lift line is prescribed for each segment, which corresponds to an angle of attack at which the velocity along that segment will be constant. This formulation provides the designer with a genuine multipoint design tool, which provides the ability to prescribe multipoint objectives directly at the design stage. Selig and Maughmer [11,12] developed Eppler's [10] original theory further by adding the ability to prescribe a nonconstant velocity along a segment at the segment's design angle of attack. The method was further generalized to allow for both finite trailing edges and cusped trailing edges. In addition, the method can be linked to a multidimensional Newton iteration procedure to satisfy additional design constraints such as airfoil pitching moment coefficient and maximum thickness. The Selig and Maughmer [11,12] method was implemented in the PROFOIL inverse design code. PROFOIL was developed further to provide a multi-element airfoil design capability [13,14], and it has also been used for the design of finite wings [15,16]. When used as a multi-element scheme or for the design of three-dimensional wing geometries, several airfoils are designed in isolation and then assembled to represent the actual flowfield. A custom Newton iteration scheme is used to adjust the conformal mapping parameters to match desired aerodynamic and geometric properties on the three-dimensional wing or the multi-element airfoil system. In the case of a wing, the airfoils are used to define the cross sections of the wing at various spanwise locations.

Work on the inverse design of fuselages and other nonlifting bodies has been more limited, but the development of such a method would have great significance. For example, Dodbele, et al. [17] noted that the fuselage of a typical commercial aircraft may produce slightly less than 50% of the total parasite drag of the complete aircraft when the surfaces (including the wing and empennage) have all-turbulent flow. When extended laminar flow is present on the lifting surfaces, which is typical of modern designs, the fuselage can produce as much as 70% or more of the total parasite drag. These numbers demonstrate that a reduction in the fuselage drag contribution will have a significant impact on the total drag of an aircraft. To improve the extent of the laminar boundary-layer, while limiting any areas of separation, more direct control over the velocity and pressure distribution on the fuselage is needed. The possible applications of a general inverse design method also extend beyond fuselage design: work done by Lutz and Wagner [18] on shape optimization of airship bodies demonstrated the usefulness of a custom inverse design method to decrease drag through delayed laminar-to-turbulent transition. Unfortunately, the scheme used in their study was limited to axisymmetric bodies at zero angle of attack. The purpose of the current research is to extend the capability of existing inverse design methods to the design of nonlifting bodies and, ultimately, to complex three-dimensional geometries that include both lifting and nonlifting components.

There are a number of different approaches to inverse design. The curvature method [19–21] uses the curvature to obtain the shape in three-dimensions. This method uses the change in velocity to drive the curvature. An alternative method is to use the surface of an airfoil, whereby this surface is defined by specifying the velocity distribution in two-dimensions. The latter method is adopted in the current work to assist in the design of general bulbous bodies and is a continuation of previous work done on axisymmetric bodies [22]. Several airfoils are designed in an inverse manner in isolation through conformal mapping and then used to define a three-dimensional body such as a fuselage. A surface is wrapped around the resulting wire frame and then discretized for analysis. The three-dimensional velocity distribution over the body is calculated with a low-order panel method and compared with the distribution over the isolated airfoils. The velocity distribution on the isolated airfoils are then adjusted through a nonlinear solver to match the three-dimensional velocity distributions to the target distributions. Two

suitable numerical schemes for solving the nonlinear system will be discussed, as well as a method to increase the speed of Jacobian evaluation at a slight cost in accuracy and convergence rate.

Methodology

Inverse Airfoil Design

The Selig and Maughmer [11,12] method, as briefly described in the introduction, provides the capability to design airfoils through direct manipulation of the velocity distribution over the airfoil within specified geometry constraints. The method, as implemented in the PROFOIL airfoil design code, has been used successfully for the design of a range of custom airfoils. The goal of the present study is to develop a method with the same level of control over the velocity distributions on three-dimensional bodies that PROFOIL offers over the velocity distributions of airfoils in isolation, while also maintaining the capability to add geometric constraints such as maximum thickness.

Various velocity distributions can be specified at the design angle of attack of an airfoil segment when using PROFOIL for the inverse design, including 1) Constant velocity along a segment (similar to Eppler's [10] design code); 2) Linear velocity distribution along a segment; and 3) Splined velocity distribution with control over the intermediate points and end conditions. Figure 1 shows how one segment is selected in the ζ-plane and its physical representation in the z-plane. Figure 2 gives some examples of the different distributions listed. The velocity distributions in these figures are shown relative to the value at the start of the segment and, therefore, denoted $\tilde{v}_i$ for each segment i. Correspondingly, $\tilde{\phi}_i$ and $\tilde{s}_i$ are, respectively, used for the relative arc limit and the arc length along the airfoil surface for segment i.

The notation ΔV_{2D} will be used to denote the velocity difference over an airfoil segment. The subscript 2D is used to differentiate between the velocity distribution over the isolated element and the element when assembled into a three-dimensional body (subscript 3D), which will be described later. The design angle of attack α^* relative to the zero-lift line on a given airfoil segment can be specified directly, or adjusted by the Newton iteration to satisfy one of the additional constraints such as the design pitching moment coefficient or the airfoil thickness. Because lift coefficient is approximately equal to $0.11 \times \alpha$, with α given in degrees relative to the zero-lift line, it is also possible to use this parameter to prescribe the velocity distribution over a segment for a design lift coefficient rather than a design angle of attack. The concept of these design angles of attack was pioneered by Eppler [10,23,24]. As mentioned in the introduction, Selig and Maughmer [11,12] rewrote the governing equations so that nonconstant velocities at the design angle of attack can also be used. The present work deals with nonlifting bodies and, therefore, the design angle of attack is usually set to zero. Nonzero values are, however, used in cases where the design angle of attack is used as an iteration variable, such as for satisfying the thickness constraints.

Three-Dimensional Inverse Design

The current method uses the airfoils designed in isolation with PROFOIL to define a three-dimensional wire frame. A nonuniform rational B-spline surface is then wrapped over this frame and discretized for analysis. A low-order panel method is used to calculate the three-dimensional velocity distribution, which in turn is
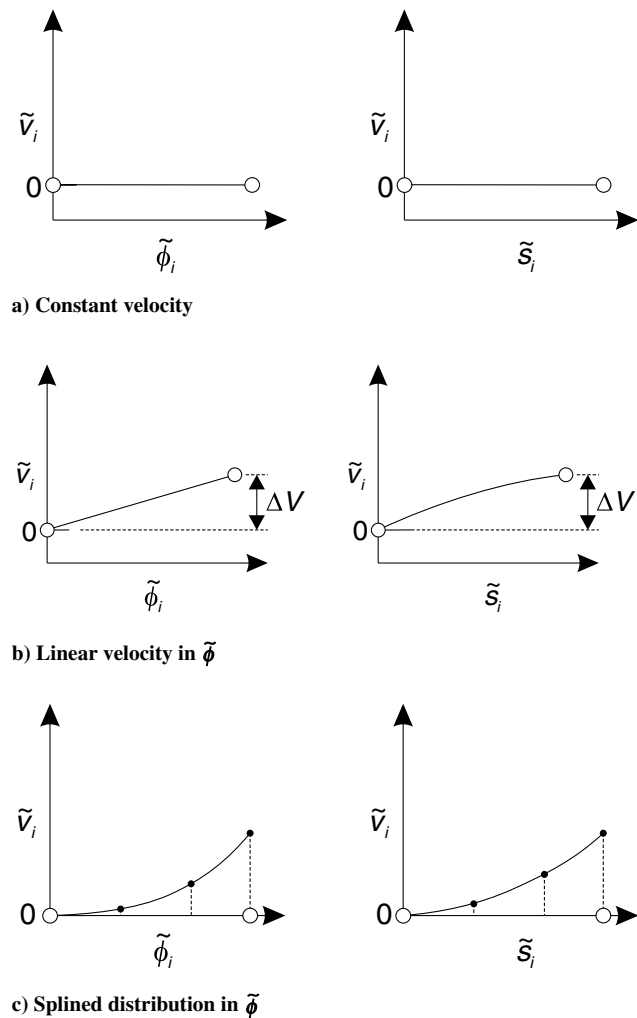

Fig. 1   Airfoil segment in ζ-plane and z-plane.

a) Constant velocity

b) Linear velocity in $\tilde{\phi}$

c) Splined distribution in $\tilde{\phi}$

Fig. 2   Most common prescribed velocity distributions.

Fig. 3   Axisymmetric body resulting from an airfoil designed via the inverse method.

used in the Newton iteration scheme. The method is outlined in Fig. 3.

The ΔV_{2D} value over a segment on an isolated airfoil, as well as its design angle of attack α^*, are specified directly in PROFOIL. Figure 4a shows an example of this specification. In the example, the isolated airfoil is divided into three segments on the upper surface. In the figure, three different ΔV_{2D} values are compared over the second segment of three symmetrical airfoils. The design angle of attack for the second segment is set to zero and, therefore, the velocity distributions are also compared with the airfoil set to zero angle of attack. The three resulting airfoils are shown in Fig. 4b.

The three test airfoils shown in Fig. 4b are used to define three different axisymmetric bodies. Figure 4c was generated using CMARC [25] to calculate the potential flow velocity distributions on the axisymmetric bodies. CMARC is a C + + port of the Fortran code PMARC [26], and was used for all the three-dimensional flow computations in the current work. It can be seen in Fig. 4c that changing ΔV_{2D} over a segment on the isolated airfoil produces a similar trend on ΔV_{3D} on the axisymmetric body, even though the absolute velocities are quite different, and the linear distributions for the 2D case become slightly curved for the 3D case. Nevertheless, the similarity in the trends between ΔV_{2D} and ΔV_{3D} make the isolated airfoil ΔV_{2D} values ideal design parameters and much easier to work with than the entire set of surface nodes on the airfoils. A nonlinear solver is used to adjust the ΔV_{2D} values until the required ΔV_{3D} is satisfied for that segment.

Another important trend seen in Fig. 4c is that the resulting change on the three-dimensional velocity distribution is localized to the part of the distribution that corresponds with the segment being controlled in the two-dimensional distribution. This is of particular
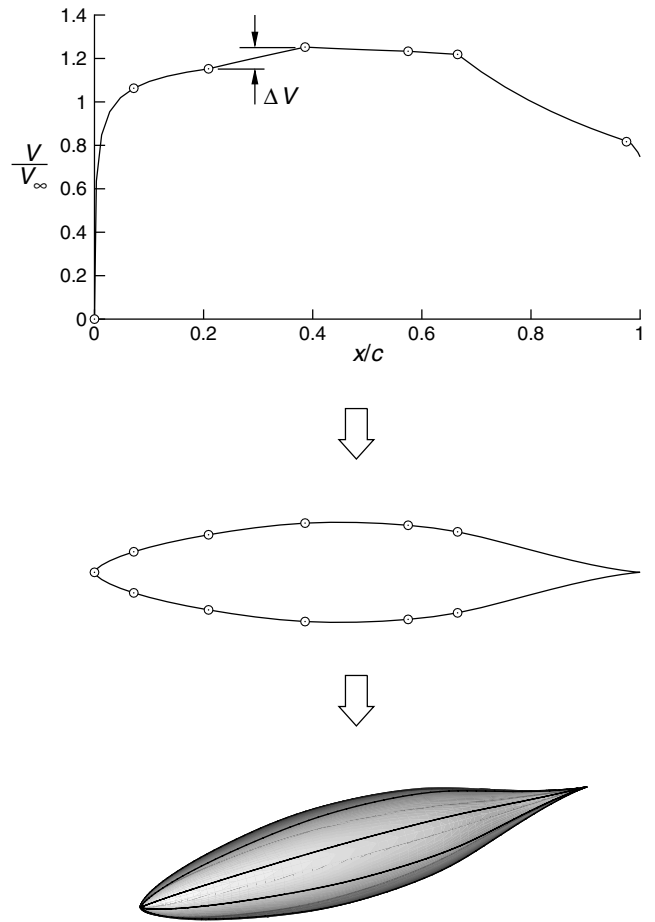
significance as it helps to decouple the three-dimensional velocity distributions on the individual segments. Although flow over nonaxisymmetric bodies with some crossflow components are inherently more complex, the current approach still succeeds as long as the control airfoils are aligned with the streamlines in an approximate way. The validity of the approach will be demonstrated in the examples to follow.

The iteration procedure as explained earlier was originally implemented into a Newton solver and was able to produce the required pressure distribution on axisymmetric three-dimensional bodies [22] reliably. One equation is added to the multidimensional iteration scheme for each segment on the three-dimensional body where a specified ΔV_{3D} is required. The iteration parameter in the scheme is the isolated ΔV_{2D} whereas the value to be set equal to zero (the root) is the difference between the required ΔV_{3D} and the ΔV_{3D} calculated from the three-dimensional panel method analysis. Additional design parameters may be added to control thickness, absolute velocity at a point on the airfoil, or other geometric or aerodynamic requirements.

In the case of nonaxisymmetric bodies, each segment of interest on the body corresponds to a unique control airfoil. In cases where geometric and flow symmetry is required, one control airfoil can be "shared" by more than one body segment. The numbering scheme in Fig. 5 will be used in the current paper unless noted otherwise. In this figure, a body that is symmetrical in both the xy- and xz-planes is shown, demonstrating how control airfoils may be shared by different control segments. In the example in the figure, control airfoil c_1 is assigned to sections s_1 and s_5. Control airfoil c_2 is assigned to sections s_2 and s_4. Finally, control airfoil c_3 is assigned to section s_3. Symmetry through the xz-plane is enforced by the panel code to minimize the computational time for symmetrical flow problems.

a) 3 airfoils with different ΔV_{2D} values on the 2nd segment

b) Symmetric airfoils resulting from the ΔV_{2D} prescriptions in Fig. 4a

c) On corresponding axisymmetric bodies

Fig. 4   Comparison of two- and three-dimensional velocity distributions.

Numerical Schemes

Newton–Raphson Method

Newton–Raphson iteration is frequently used to solve nonlinear systems of equations due to its quadratic convergence properties [27,28]. This method is also already used internally in the PROFOIL airfoil design code. The method solves a system of nonlinear equations of the form

$$f_i(x_1, x_2, \ldots, x_n) = 0 \quad \text{for } i = 1, 2, \ldots, n \tag{1}$$

which can also be written as

Fig. 5   Numbering of body sections and corresponding control airfoils.

$$F(x) = 0 \tag{2}$$

The Jacobian of the system is defined as

$$J(x) = \begin{bmatrix} \dfrac{\partial f_1(x)}{\partial x_1} & \dfrac{\partial f_1(x)}{\partial x_2} & \cdots & \dfrac{\partial f_1(x)}{\partial x_n} \\ \dfrac{\partial f_2(x)}{\partial x_1} & \dfrac{\partial f_2(x)}{\partial x_2} & \cdots & \dfrac{\partial f_2(x)}{\partial x_n} \\ \vdots & \vdots & & \vdots \\ \dfrac{\partial f_n(x)}{\partial x_1} & \dfrac{\partial f_n(x)}{\partial x_2} & \cdots & \dfrac{\partial f_n(x)}{\partial x_n} \end{bmatrix} \tag{3}$$

The Newton–Raphson iteration process is started with an initial approximation of the unknown variables x. In the current implementation, the initial approximation is usually selected by assuming the aerodynamic properties of the flow along the design sections on the three-dimensional body are similar to the two-dimensional properties on the control airfoils as designed in PROFOIL. Once the initial values of x are known, it is possible to evaluate the function values F and the Jacobian J. The correction vector δx is now computed by solving the linear system

$$J(x)\delta x = -F(x) \tag{4}$$

Once the Newton step is known, the unknown variables x are updated via

$$x_{\text{new}} = x_{\text{old}} + \delta x \tag{5}$$

The likelihood of convergence for most solvers dedicated to the solution of nonlinear problems is highly dependent on the nature of the nonlinear system and the initial estimate of the unknown variables x. To improve the likelihood of convergence from almost any initial approximation, various methods such as line searches and backtracking can be used [27,28]. The current implementation was modified from the algorithm presented in [27] and automatically uses a line search and backtracking if the full Newton step does not decrease the residue sufficiently. The line search and backtracking are usually only required when the initial approximation is far from solution.

Broyden's Method

A disadvantage of Newton's method is that the Jacobian has to be computed before each Newton step can be taken. In cases where the function evaluation is performed via a numerical scheme, such as the panel method analysis used in the current design procedure, the exact

Jacobian is generally not known. Derivatives have to be evaluated numerically, typically via a forward difference scheme. Although a numerical procedure may be acceptable for relatively small systems, the numerical evaluation of the Jacobian can quickly become the most expensive component of the solution process. A family of methods known as quasi-Newton or secant methods exist that provide rapid approximations to the Jacobian for root finding. One of the most popular of these methods is Broyden's method [29].

At a given time step k, it is assumed that an approximation for the Jacobian $\boldsymbol{B}$ is known. In a process similar to Newton's method, it is possible to solve $\delta\boldsymbol{x}_k$ via the system

$$\boldsymbol{B}_k\delta\boldsymbol{x}_k = -\boldsymbol{F}_k \qquad (6)$$

The unknown values $\boldsymbol{x}$ can now be computed for the next iteration step $(k+1)$ via

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \delta\boldsymbol{x}_k \qquad (7)$$

Once $\boldsymbol{x}_{k+1}$ is known, $\boldsymbol{F}_{k+1}$ can be evaluated. Broyden's method now consists of finding $\boldsymbol{B}_{k+1}$ so that it satisfies

$$\boldsymbol{B}_{k+1}\delta\boldsymbol{x}_k = \delta\boldsymbol{F}_k \qquad (8)$$

with $\delta\boldsymbol{F}_k = \boldsymbol{F}_{k+1} - \boldsymbol{F}_k$. Broyden showed that $\boldsymbol{B}_{k+1}$ can be computed via

$$\boldsymbol{B}_{k+1} = \boldsymbol{B}_k + \frac{(\delta\boldsymbol{F}_k - \boldsymbol{B}_k\delta\boldsymbol{x}_k)}{\delta\boldsymbol{x}_k \cdot \delta\boldsymbol{x}_k}\delta\boldsymbol{x}_k^t \qquad (9)$$

The iteration process can now proceed without computing the Jacobian directly through function evaluations. The method still needs an initial approximation for the Jacobian, which is provided in the current implementation through a finite difference analysis. The same line search and backtracking procedure as used in the Newton–Raphson method is used to improve the likelihood of convergence. It is also possible to reevaluate the Jacobian through direct function evaluations in instances where the backtracking scheme fails. Although Broyden's scheme exhibits superlinear convergence instead of the quadratic convergence rate of Newton–Raphson iteration, the savings in the Jacobian evaluation is significant.

Sensitivity Analysis

Both Newton–Raphson iteration and Broyden's method require a Jacobian evaluation. Although Broyden's method only requires an initial Jacobian, even this one Jacobian evaluation can take a significant amount of time, as a minimum of one full flow evaluation is required for each design variable. One method to accelerate the sensitivity analysis is to simulate small geometric perturbations by adding a transpiration velocity to each panel. The transpiration method is often used to model viscous/inviscid interaction when combining a boundary-layer model to a panel method. These methods use a transpiration velocity through each panel to displace the surface streamlines the same distance as given by the boundary-layer displacement thickness. In the current methodology, the purpose of the transpiration velocity is to displace the surface streamlines the same distance that a small geometric perturbation of the surface would displace them.

The advantage of the transpiration approach over physically adjusting the surface of the body is that the influence coefficient matrix of the flow analysis method can be kept unaltered, whereas only the boundary conditions have to be adjusted. In the case of a panel method such as the one used in the current design procedure, the transpiration values are incorporated into the solution by adjusting the known source strengths, which in turn are used to compute the right-hand side of the linear system so that only the doublet strengths remain as unknowns.

The amount of transpiration, or blowing, required can be determined by considering a control volume between the original, unchanged surface and the new surface. The lower face of the control volume represents the original surface and the surface to which the transpiration velocity u_t will be applied. The upper face of the control volume represents the new body and, because this new surface will become a stream surface, there is no flow through the upper face. It can be shown that the transpiration velocity through the original surface is given by [30]

$$u_t = \frac{\partial uh}{\partial x} + \frac{\partial vh}{\partial y} \qquad (10)$$

Equation (10) relates the required local transpiration velocity through a surface to the required deflection of the off-surface streamlines to simulate a small perturbation to the surface shape. The average transpiration velocity through each surface panel can be approximated using this equation when a panel method analysis is used for the flow evaluation. However, the calculation of the local partial derivatives $\partial uh/\partial x$ and $\partial vh/\partial y$ can become complicated in practice. In addition, the equation only gives the local required transpiration velocity at one point on the panel. Because the velocity distribution over one panel can be quite complex, the local transpiration velocity required at, for instance, the control point of the panel may be significantly different to the average transpiration velocity required over the entire panel. A more practical approach is to consider the volume between the original surface and the perturbed surface. This volume is then divided into a number of control volumes, each with its lower face represented by a panel on the original surface, and its upper face represented by an equivalent panel on the perturbed surface. When using this approach, it is necessary to move the nodes approximately normal to the original surface when discretizing the perturbed surface. An example of a control volume formed by this method is shown in Fig. 6. It should be noted that the surfaces in this control volume may be slightly deformed and that the original and perturbed panel may also not be completely flat surfaces. Some accuracy is lost when the faces of the control volume are not completely flat.

Figure 6 also shows an additional complication when calculating the transpiration volume, namely that even if the nodes are moved normal to the local surface when perturbed, this movement is not necessarily normal to the panel to which that corner point belongs because the surface panels only approximate the actual surface of the body that is analyzed. Once each "control volume" is established between the original surface and perturbed surface, a process very similar to the method used for deriving Eq. (10) can be used to determine the average transpiration velocity required through the original surface panel. The mass-flux through each face of the control volume is required to satisfy continuity:

$$\dot{m}_1 + \dot{m}_2 + \dot{m}_3 + \dot{m}_4 + \dot{m}_t = 0 \qquad (11)$$

In the preceding continuity equation shown, all the mass-flux values are taken to be positive into the control volume. If the flow is


Fig. 6   Control volume between original and perturbed surface panel.

assumed to be incompressible, the mass-flux values can be replaced with volume flux values:

$$\dot{v}_1 + \dot{v}_2 + \dot{v}_3 + \dot{v}_4 + \dot{v}_t = 0 \quad (12)$$

The volume flux through faces 1–4 is computed by first calculating the average velocity through each face, based on a flow analysis over the original surface, and then multiplying it by the area of each face. The volume flux through face 1 is, for instance, given by

$$\dot{v}_1 = \bar{u}_1 A_1 \quad (13)$$

where $\bar{u}_1$ is the average normal velocity through face 1 and A_1 is the area of face 1. The transpiration velocity is then calculated using Eq. (12) that, when rewritten to isolate u_t, becomes

$$u_t = -\frac{\dot{v}_1 + \dot{v}_2 + \dot{v}_3 + \dot{v}_4}{A_p} \quad (14)$$

where A_p is the area of the original panel.

In the current design methodology, one flow analysis is performed with the original, unperturbed shape. Once the velocities at all the nodes are determined, the geometry is perturbed slightly and care is taken to move the nodes approximately normal to the local surface tangents. Equation (14) is used together with the procedure previously outlined to determine the necessary transpiration velocity through each panel. The source strengths, which are incorporated into the right-hand side of the linear system in the CMARC formulation, are then adjusted to account for the normal velocities through each panel surface via the relation

$$\sigma = \frac{1}{4\pi}(u_t - \mathbf{n} \cdot \mathbf{V}_\infty) \quad (15)$$

Once the source strengths are known, the right-hand side of the linear system can be recomputed and the system can then be solved for the doublet strengths. The solver built into CMARC/PMARC is an iterative solver that was not designed for multiple right-hand sides. Instead of using the solution from CMARC, the current design methodology uses an LU-factorization on the influence coefficient matrix originally computed by CMARC. Back substitution is then performed on the multiple right-hand side vectors to evaluate the Jacobian after a single LU-factorization.

Performance of Numerical Schemes

The preceding numerical schemes described were evaluated by designing a fairly complex bulbous body. The example used to obtain the figures below was chosen so that no backtracking was required during the iteration process. A total of 37 design variables
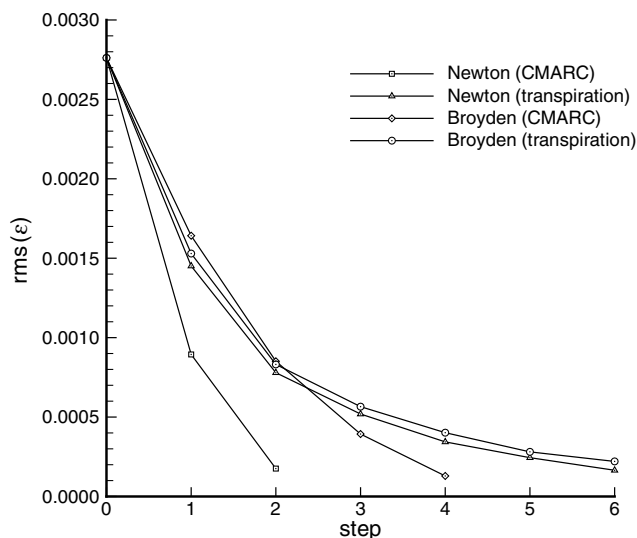
Fig. 7   Convergence history of 37-variable test case.

Table 1   Performance of different numerical schemes for a system with 37 design variables

Scheme	Steps	t
Newton (CMARC)	2	1536 sec
Newton (transpiration)	6	808 sec
Broyden (CMARC)	4	760 sec
Broyden (transpiration)	6	392 sec

were used over three control airfoils, which results in 37 unknowns in the nonlinear system. The example uses 2400 surface panels, which is representative of the number of surface panels typically used in more complex problems. Four different combinations of solvers were compared. In the following list, all Jacobian evaluations were performed via a forward difference procedure: 1) Newton–Raphson iteration combined with direct CMARC evaluations for the Jacobian; 2) Newton–Raphson iteration combined with Jacobian evaluations through the transpiration method; 3) Broyden's method combined with direct CMARC evaluations for the initial Jacobian; and 4) Broyden's method combined with an initial Jacobian evaluated through the transpiration method.

Figure 7 shows the convergence history using each of the preceding methods listed. Table 1 shows the computational time for the entire design process on a 1.8 GHz Pentium 4 processor. As expected, the Newton–Raphson method combined with direct flow analysis for the Jacobian evaluations displays the highest convergence rate (fewest iteration steps). However, this method is also the most costly in computational time due to the slow Jacobian evaluations. The remaining methods all use a more approximate Jacobian and, therefore, display lower convergence rates (more iteration steps) than the first method. Computational time varies considerably between the methods. Broyden's method combined with a Jacobian evaluation through the transpiration method is almost 75% faster than the first method. Experience with the design method has, however, shown that the less accurate Jacobian resulting from Broyden's method, as well as the approximate transpiration method, tend to be less robust in terms of global convergence, so that the slower method is sometimes preferred to improve the likelihood of finding a solution.

Examples

Three examples are shown in this section to demonstrate the flexibility of the new design method. The first demonstrates the ability of the current method to exactly match a target velocity distribution. Only one control airfoil is used to produce an axisymmetric shape. The second example shows two cases of "squashed" bodies. In this example set, the α^* value over the first segment behind the leading edge is used as a design variable to control the maximum thickness of each airfoil. The maximum thickness of each airfoil in turn is chosen to produce the squashed shape. A slightly favorable pressure distribution with a constant slope over most of the forward 60% of the chord is prescribed. Finally, a complex example showing the influence of a second component in the flowfield is shown. In this last example, a constant-chord symmetrical keel at zero angle of attack is added to the body to simulate the bulb and keel of a competition sailing yacht.

Design of an Airship Body

The goal of the first example is to demonstrate the capability of the current method to match a known velocity distribution. The example verifies that the method will converge from an arbitrary (but physical) initial velocity distribution to another known velocity distribution. The reader should note that when a designer is faced with a new design, he/she will usually not know exactly what the velocity distribution should look like over the entire body because the influence of the geometric constraints on the velocity distribution are not known. Later examples will demonstrate how the designer can specify velocity gradients over only part of the body while

allowing those on other parts of the body to be adjusted by the method to match additional constraints.

Lutz and Wagner [18] investigated possibilities for drag reduction on airship bodies in various design regimes through shape optimization. Their method used an inverse design scheme based on a source distribution along the body axis. One of the bodies designed by them had a velocity distribution optimized for a volumetric Reynolds number of 1×10^6. The velocity distribution produced by their method has a favorable pressure gradient over the first $0.75c$, followed by a fairly steep recovery. In the current example, it will be demonstrated how the new method can be used to reproduce the velocity distribution over the entire airship body. The source-based method produced a body with a blunt trailing edge with a diameter of approximately 1% of the body length. PROFOIL's ability to design airfoils with blunt trailing edges was used to reproduce the same blunt trailing edge geometry for the current example.

The design process was initiated by dividing the control airfoil into 13 segments. In the design of this airfoil, the segment endpoints are moved to specific x/c positions where the local velocities were known from the original design. Across each of these segments, the required velocity difference for the final design was entered as a ΔV_{2D} value. Although this will not produce the target velocity distribution exactly on the axisymmetric body, it does result in a fairly good initial condition for the nonlinear solver. A splined segment was used over the recovery region between $0.70c$ and $0.90c$. In addition, a splined segment was also used over the pressure peak near the leading edge. Some freedom was allowed over the last $0.10c$ so that a reasonable airfoil could be produced for each step in the iteration process. The resulting initial three-dimensional velocity distribution is compared with the target distribution in Fig. 8. The distribution of segments and the numbering scheme is also shown in the figure. Although there are some differences at this stage, the general trend is already very similar to that of the target distribution. The resulting geometry is compared to the geometry produced by Lutz and Wagner's method [18] in Fig. 9. The most noticeable difference between the two geometries at this stage is the difference in thickness. It is possible to directly prescribe a thickness in the new method, but for the purposes of this demonstration only the velocity distribution is controlled with no constraint on the thickness.

The velocity differences over all the segments on the axisymmetric body were specified as target velocity differences for the three-dimensional Newton iteration scheme. Segment 2, which models the recovery region, was divided into four subsections and the velocity difference at each node of the resulting spline was specified. The spline used over the pressure peak only consisted of two segments, and it was used to lower the peak level slightly to closer approximate the required distribution. Finally, the absolute


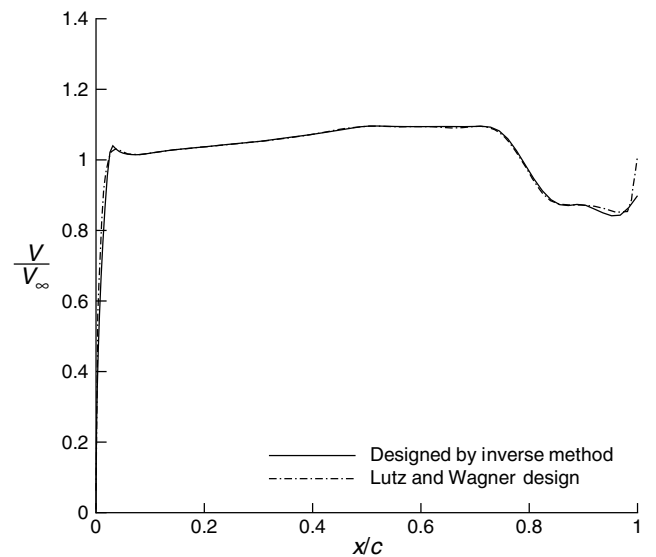**Fig. 9   Comparison of target and initial design geometries.**

Fig. 10   Comparison of target and final design velocity distributions.


Fig. 11   Comparison of target and final design geometries.

velocity at the $0.10c$ location was specified by iterating on the two-dimensional velocity difference on the leading edge segment.

The resulting velocity distribution after completion of the iteration process is compared with the target distribution in Fig. 10. It can be seen that the current method was able to closely match the specified velocity distribution, except for a slight difference over the last $0.10c$ of the airfoil. The shapes are compared in Fig. 11. Note that the designed and target shapes are very similar, and that the initial difference in thickness between the designed and the target geometry had been eliminated.

Fig. 8   Comparison of target and initial design velocity distributions.

a) Section s_1

b) Section s_2

c) Section s_3

Fig. 12 Velocity distributions over the control sections of a squashed body.

The number of segments selected was arbitrary. Once the inverse design process is complete and the resulting velocity distribution has been compared with the target distribution, the designer has the option to increase the number of segments in areas where he requires a closer match to the target distribution. An alternative approach is to convert the velocity distributions on some segments to splined velocity distributions, which give more control over the shape of the distribution on those particular segments.

Squashed Bodies

The following example shows two different designs that are related in that they both have identical velocity gradient requirements. In the first case, a body is squashed by prescribing a different thickness in the xz-plane than in the xy-plane. The cross-section at each station of the body is defined by four quarter-ellipses, implemented in the design code via quadratic rational B-splines with the weight values selected to exactly represent the quarter-ellipses. In the second case, the control airfoil describing the lower surface of the body is constrained to be thinner than the control airfoil describing the upper surface, thereby producing a slight droop or camber in the body. Each body uses three control airfoils: an upper vertical, a horizontal, and a lower vertical control airfoil. Symmetry in the xy-plane is assumed. This scheme is the same as shown in Fig. 5, except that control airfoils s_2 and s_4 shown in that figure are not required.

A target velocity distribution with a constant, favorable gradient over the forward portion of the body is prescribed. This gradient is prescribed from approximately $0.08c$ to $0.60c$. Eleven segments are used along the airfoils. Figure 12a shows the numbering of these segments. Splined velocity distributions are used on segment 2 (the recovery region) and segments 3, 8, 9, and 10.

Case 1

In case 1, which is symmetrical in both the xz-plane and the xy-plane, it was found that no further iteration is required in the recovery regions of the control airfoils. However, the shape of the recovery region of the airfoil in the horizontal plane was specified slightly differently to that of the other two control airfoils. The splined velocity distributions used on segments 3, 8, 9, and 10 were only added as a refinement to produce more linear velocity distributions on the three-dimensional body. The thicknesses of the three control airfoils were selected to produce a body that is 18% thick in width and 12% thick in height. The α^* value of segment 11 was used as the iteration variable to achieve the respective thicknesses of the control airfoils. This does result in slightly different velocity distributions in the leading edge region of the control sections, as well as a slight difference in absolute velocities.


Fig. 13 Comparison of three-dimensional velocity distributions over sections s_1–s_3 of a squashed body.

Figure 12 shows the two- and three-dimensional velocity distributions at each of the control sections. Both the initial and converged three-dimensional distributions are shown. Notice how close the initial distribution is to the target distribution. The main reason for the small difference before and after the iteration process is that the maximum thickness of each control airfoil is held constant throughout the design process. The absolute velocity, therefore, changes very little throughout the design process. The final velocity distributions for all three control sections are coplotted on Fig. 13. Notice that section s_2, which is thicker than sections s_1 and s_3, has a slightly different absolute velocity over segments 2–9, but that it is very close to the other two. Section s_2 does, however, have a very different velocity distribution in the leading edge region and a slightly different distribution in the recovery region. These differences are a result of the different thickness constraints and is, therefore, a compromise between aerodynamic and geometric requirements. The fact that it is possible to isolate these differences to specific regions of the velocity distribution is, however, very useful to the designer. Figure 14 shows the locations of the control sections on the final body. Figure 15 shows an isometric view and front view of the final geometry.

Case 2

Case 2 is similar to case 1 except that different thickness constraints were used on the three control airfoils. The upper surface of 18% thick symmetrical airfoils were used to define sections s_1 and s_2, whereas the upper surface of a 12% thick symmetrical airfoil was used to define section s_3. The velocity distribution requirements were the same as in case 1.

The final velocity distributions for the three control sections are coplotted on Fig. 16. Notice that all three sections differ slightly in the velocity levels as well as in the leading edge and recovery regions. However, the velocity gradient over segments 2–9 is virtually identical and is required for all three sections. Figure 17 shows the locations of the control sections on the final body. Figure 18 shows an isometric view and front view of the final geometry.

Racing Yacht Bulb with Keel

It is a common practice for designers to divide a complex geometry, such as that of an aircraft, into smaller and more simple components and then to design these in isolation. Once these components are assembled, the flow may be considerably altered from the conditions assumed during the design of the isolated components. Usually a system of fairings are used to alleviate some of the problems that arise, whereas other problems may need a cut-and-try method where the isolated components are redesigned until the flow over the assembled geometry is deemed acceptable. The design process would be considerably more efficient if the designer could design the components in their final assembled form from beginning. In the current example, the design of a typical underwater bulb of the type used by some racing yachts is shown. The keel is placed in the flowfield throughout the design process. The size and geometry of the keel is kept constant for this example, and it is assumed that the keel is nonlifting.

a) Isometric view

b) Front view

Fig. 15   Final squashed body geometry.

Fig. 16   Comparison of three-dimensional velocity distributions over sections s_1–s_3 for case 2.


Fig. 14   Location of control airfoils on final squashed body.


Fig. 17   Location of control airfoils on final body for case 2.

a) Isometric view

b) Front view

**Fig. 18   Final squashed body geometry of case 2.**

The section numbering is identical to that shown in Fig. 5, with the vertical keel intersecting section s_1 on the upper surface of the bulb. A velocity distribution with a constant gradient similar to those in the previous examples was required on the bulb. There are, however, significant complications due to the presence of the keel. Most notably, the presence of the leading edge stagnation point makes it unrealistic to expect a constant gradient on the upper surface of the bulb in the direct vicinity of this point. It is, however, possible to achieve some control over the velocity distribution slightly outside this region (for instance in the vicinity of section s_2). A keel with a chord of 20% of the length of the bulb is used. Because this keel is relatively small, it is also possible to control the velocity distribution in front of and behind it with some flexibility allowed for about 10% of the length of the bulb in front of and behind the keel.

The initial and final three-dimensional velocity distributions, as well as the final two-dimensional velocity distributions over the upper three control sections of the bulb, are shown in Fig. 19. The numbering scheme used on the individual segments is shown in Fig. 19a. Notice that the recovery region is labeled as segment 1 in this case, as it was not defined with a splined segment as in the previous example. The three-dimensional distribution shown over segment 3 of control section s_1 as shown in Fig. 19a is actually the distribution as calculated around the root of the keel, as no velocity is defined inside the keel where segment 3 is blanked out. The two stagnation points are located in front of and behind segment 3. The panel method applies boundary conditions at the panel centroids rather than at panel nodes, and some interpolation and/or extrapolation between neighboring panels is, therefore, required to calculate velocities at nodes. Because the two stagnation points are located at panel corners, the velocities shown at these nodes are approximate and the stagnation points are not fully captured; the velocities shown in the figures are, therefore, nonzero at these points.

Control section s_2 is particularly interesting. Figure 19b shows the initial and final three-dimensional distributions over this section. Notice the fact that the initial distribution, shown by the dashed line in the figure, displays some unwanted bumps and oscillations. The "bump" in the region of segment 3 is particularly noticeable. To "straighten out" the velocity distribution in this region, splined velocity distributions were used on segments 2, 3, and 4 of section s_2. A splined distribution was used on segment 9, near the leading edge,
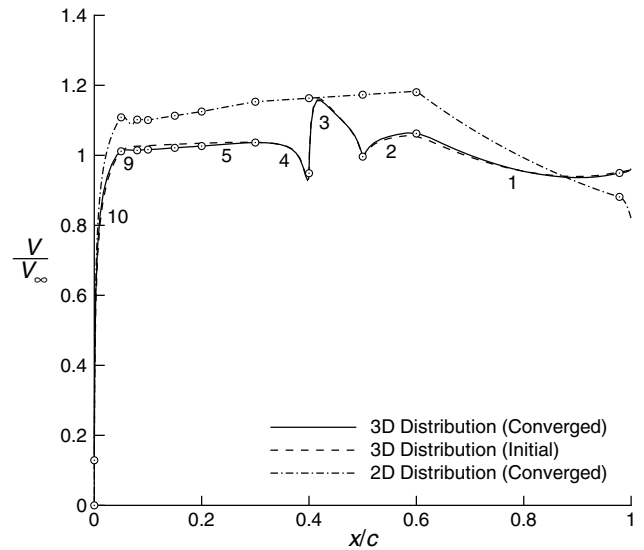
a) Section s_1

b) Section s_2

c) Section s_3

Fig. 19   Velocity distributions over the upper three control sections of the bulb.

on all control sections including section s_2. As can be seen in the figure, the new method was able to reduce the bumps in the velocity distribution to a point where their effects will have been eliminated. The result is a long, constant favorable gradient extending from approximately 8% to 60% of the bulb length.

The velocity distribution over the control section in the horizontal plane (section s_3), is shown in Fig. 19c. It can be seen in this figure that the effect of the keel is barely noticeable in the initial velocity distribution and all remaining effects had been removed by the three-dimensional iteration process. Because of the small effect of the keel on section s_3, it was not necessary to use a splined velocity distribution on segments 2, 3, and 4 of this control airfoil. The resulting velocity distributions on the two remaining control sections were very similar to that on section s_3. The final distributions on all five control sections are coplotted in Fig. 20. Note that the differences are almost completely limited to the leading edge region and to a small extent to the recovery region. As explained in the preceding paragraph, the distribution on section s_1 is different over segments 2, 3, and 4 due to the presence of the keel. Figure 21 shows the final geometry with the locations of the control sections. The mesh used by the panel method is shown in Fig. 22.


Fig. 22   Isometric view of panel method mesh.

Conclusions

A hybrid approach has been developed for the design of complex nonlifting bodies. A number of options for different numerical schemes were compared. Some of the capabilities of this new method were demonstrated by designing three different configurations. The first demonstrated the capability of the current method to reproduce a known velocity distribution on an axisymmetric body. The second example showed the capability of the method to produce specified velocity distributions over some segments of a body, while also satisfying certain geometric constraints, such as different thicknesses in this particular example. The third example demonstrated the capability of the current method to satisfy velocity distribution requirements in a more complex flowfield, such as in the vicinity of a second body or component. The speed of the numerical schemes allow the designer to quickly and efficiently adjust the design parameters as needed. Work is currently underway to expand the method as demonstrated here to design more complex configurations that consist of a combination of lifting and nonlifting components.

Acknowledgments

Fig. 20   Comparison of three-dimensional velocity distributions over sections s_1–s_3 of a squashed body.

**Fig. 21   Location of control airfoils on final squashed body.**

References

[1] von Karman, T., "The Problem of Resistance in Compressible Fluids," *Fifth Volta Congress*, Roma Reale Accademia D'Italia, Rome, Italy, 1936.

[2] Sears, W. R., "On Projectiles of Minimum Wave Drag," *Quarterly of Applied Mathematics*, Vol. 4, No. 4, 1947, pp. 361–366.

[3] Haack, W., "Projectile Shapes for Smallest Wave Drag," Translation A9-T-3, Contract W33-038-ac-15004(16351), Air Technical Index Rept. 27736, Air Material Command, U.S. Air Force, Brown Univ., Providence, RI, 1948.

[4] Theodorsen, T., "Theory of Wing Sections of Arbitrary Shape," NACA TR 411, 1931.

[5] Lighthill, M. J., "A New Method of Two-Dimensional Aerodynamic Design," Aeronautical Research Council R&M 2112, April 1945.

[6] Liebeck, R. H., "A Class of Airfoils Designed for High Lift in Incompressible Flow," *Journal of Aircraft*, Vol. 10, No. 10, Oct. 1973, pp. 610–617.

[7] Glauert, M. B., "The Application of the Exact Method of Aerofoil Design," Aeronautical Research Council R&M 2683, Oct. 1947.

[8] Timman, R., "The Direct and Inverse Problem of Airfoil Theory," National Aeronautical Research Inst. Rept. F.16, The Netherlands, 1951.

[9] Strand, T., "Exact Method of Designing Airfoils with Given Velocity Distribution in Incompressible Flow," *Journal of Aircraft*, Vol. 10, No. 11, 1973, pp. 651–659.

[10] Eppler, R., "Direkte Berechnung von Tragügelprofilen aus der Druckverteilung," *Ingenieur-Archive*, Vol. 25, No. 1, 1957, pp. 32–37; also "Direct Calculation of Airfoil from Pressure Distribution," NASA TT F-15, 417, 1974 (in English).

[11] Selig, M. S., and Maughmer, M. D., "A Multi-Point Inverse Airfoil Design Method Based on Conformal Mapping," *AIAA Journal*, Vol. 30, No. 5, May 1992, pp. 1162–1170.

[12] Selig, M. S., and Maughmer, M. D., "Generalized Multipoint Inverse Airfoil Design," *AIAA Journal*, Vol. 30, No. 11, Nov. 1992, pp. 2618–2625.

[13] Gopalarathnam, A., and Selig, M. S., "Multipoint Inverse Method for multi-element Airfoil Design," *Journal of Aircraft*, Vol. 35, No. 3, May-June 1998, pp. 398–404.

[14] Gopalarathnam, A., and Selig, M. S., "A Multipoint Viscous Design Method for Multi-Element Airfoils," AIAA Paper 98-2404, June 1998.

[15] Gopalarathnam, A., and Selig, M. S., "A Hybrid Approach to Inverse Design of Complex Aerodynamic Systems," AIAA Paper 2000-0784, Jan. 2000.

[16] Gopalarathnam, A., and Selig, M. S., "Hybrid Inverse Airfoil Design Method for Complex Three-Dimensional Surfaces," *Journal of Aircraft*, Vol. 39, No. 3, May 2002, pp. 409–417.

[17] Dodbele, S. S., van Dam, C. P., and Vijgen, P. M. H. W., "Design of Fuselage Shapes for Natural Laminar Flow," NASA CR 3970, 1986.

[18] Lutz, T., and Wagner, S., "Drag Reduction and Shape Optimization of Airship Bodies," *Journal of Aircraft*, Vol. 35, No. 3, 1998, pp. 345–351.

[19] Barger, R. L., and Brooks, C. W., Jr., "A Streamline Curvature Method for Design of Supercritical and Subcritical Airfoils," NASA TN D-7770, Sept. 1974.

[20] Campbell, R. L., "Efficient Viscous Design of Realistic Aircraft Configurations," AIAA Paper 98-2539, June 1998.

[21] Milholen, W. E., II, "An Efficient Inverse Aerodynamic Design Method for Subsonic Flows," AIAA Paper 2000-0780, Jan. 2000.

[22] Broughton, B. A., and Selig, M. S., "A Hybrid Inverse Design Method for Axisymmetric Bodies in Incompressible Flow," AIAA Paper 2002-3142, June 2002.

[23] Eppler, R., and Somers, D. M., "A Computer Program for the Design and Analysis of Low-Speed Airfoils," NASA TM 80210, Aug. 1980.

[24] Eppler, R., *Airfoil Design and Data*, Springer–Verlag, New York, 1990, pp. 514–521.

[25] Pinella, D. F., and Garrison, P., *CMARC: Three-Dimensional Panel Code*, Aero-Logic, Los Angeles, 1996.

[26] Ashby, D. L., Dudley, M. R., and Iguchi, S. K., "Development and Validation of an Advanced Low-Order Panel Method," NASA TM 101024, Oct. 1988.

[27] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., *Numerical Recipes in Fortran*, 2nd ed., Cambridge Univ. Press, New York, 1992.

[28] Burden, R. L., and Faires, J. D., *Numerical Analysis*, 5th ed., PWS Publishing Co., Boston, MA, 1993.

[29] Broyden, C., "A Class of Methods for Solving Nonlinear Simultaneous Equations," *Mathematics of Computation*, Vol. 19, No. 92, Oct. 1965, pp. 577–593.

[30] Katz, J., and Plotkin, A., *Low-Speed Aerodynamics*, 2nd ed., Cambridge Aerospace Series, Cambridge Univ. Press, New York, 2001.