

A Gradient Accuracy Study for the Adjoint-Based Navier-Stokes Design Method

*Sangho Kim**, *Juan J. Alonso†* and *Antony Jameson‡*

Department of Aeronautics and Astronautics

Stanford University

Stanford, California 94305 U.S.A.

1 Abstract

A continuous adjoint method for Aerodynamic Shape Optimization (ASO) using the compressible Reynolds-Averaged Navier-Stokes (RANS) equations and the Baldwin-Lomax turbulence model was implemented and tested. The resulting implementation was used to determine the accuracy in the calculation of aerodynamic gradient information for use in ASO problems. For completeness, the formulation and discretization of the Navier-Stokes equations and the resulting adjoint equations are discussed. However, the reader is referred to previous work for details of the derivations. The accuracy of the resulting derivative information is investigated by direct comparison with finite-difference gradients. In the process, shortcomings of the finite difference method for the calculation of derivative information are pointed out and discussed. The advantages of the use of an adjoint method become apparent because of the strict requirements that the finite difference method imposes on the level of flow solver convergence and the sensitivity of the value of the gradients with respect to the choice of step size. Design examples for both inverse and drag minimization problems are presented. A parallel implementation using a domain decomposition approach and the MPI (Message Passing Interface) standard is used to reduce the computational cost of automatic design involving viscous flows.

*Doctoral Candidate

†AIAA Member, Assistant Professor, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, U.S.A

‡AIAA Fellow, Thomas V. Jones Professor of Engineering, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, U.S.A

2 Introduction

Since its introduction for the study of fluid flow, Computational Fluid Dynamics (CFD) has played an important role in the analysis of aerodynamic configurations of varying complexity. Successive improvements to our ability to predict complex flow fields have addressed both the physics of the flow model (potential, Euler, Navier-Stokes equations) and the level of geometric complexity present (airfoils, wings, wing-bodies, complete configurations). Although several attempts were made in the past to use CFD as a design tool [1, 2, 3], it has not been until recently, with the advent of high speed computers and the development of accurate numerical algorithms, that the focus of CFD applications has shifted to aerodynamic design [4, 5, 6, 7, 8, 9].

Typical CFD design methods use gradient-based optimization techniques in which a control function to be optimized (an airfoil shape for example) is parameterized with a set of design variables, and a suitable cost function to be minimized is defined (drag coefficient, lift/drag ratio, difference from a specified pressure distribution, etc). Then, the sensitivity derivatives of the cost function with respect to the design variables are calculated in order to get a direction of improvement. A step is taken in this direction and the procedure is repeated until convergence to a minimum is achieved. Finding a fast and accurate way of calculating the necessary gradient information is essential to developing an effective design method, since this can be the most time consuming portion of the design algorithm.

Gradients can be computed using finite differences in a straight-forward manner. In the finite difference method, it is customary to take small steps in each

and every one of the design variables independently, in order to find the sensitivity of the cost function with respect to those design variables. Since each of these steps requires a complete flow solution, the computational cost of this method is proportional to the number of design variables, and, consequently, it cannot be afforded for problems with large dimensionality.

The control theory approach to optimal aerodynamic design, in which gradient information is obtained via the solution of an adjoint equation, was first applied to transonic flow by Jameson [4, 5, 10] and has become a popular choice for design problems involving fluid flow [11, 7, 12, 13]. In fact, the method has even been successfully used for the aerodynamic design of complete aircraft configurations [14, 6]. The adjoint method is extremely efficient since the computational expense incurred in the calculation of the complete gradient with respect to an arbitrary number of design variables is effectively *independent* of the number of design variables. The only cost involved is the calculation of *one* flow solution and *one* adjoint solution whose complexity is similar to that of the flow solution.

Most of the early work in the formulation of the adjoint-based design framework was done using the potential and Euler equations as models of the fluid flow. The development of the adjoint system for the Euler equations, as well as the assessment of the accuracy of the gradient information that can be obtained from these adjoint equations have been presented in the past [5, 15]. Aerodynamic design calculations using the Reynolds Averaged Navier-Stokes equations as the flow model have only recently been tackled. Most of the early work was done directly in three dimensions using an inviscid adjoint equation [16, 10]. In order to validate this procedure, it was considered appropriate to conduct a series of numerical experiments in two dimensions that would establish the soundness of the three-dimensional implementations. The extension of adjoint methods for optimal aerodynamic design to viscous flows is necessary to provide the increased level of modeling which is crucial for certain types of flows. This cannot only be considered an academic exercise, but also a very important issue for the design of viscous dominated applications such as the flow in high lift systems.

It is interesting to point out that the different existing approaches to the adjoint method can be classified into two categories: the continuous-adjoint and discrete-adjoint methods. If the adjoint equations are directly derived from the governing equations and then discretized, they are termed *continuous*, and if they are derived from the discretized form of the governing equations then they are referred to as *discrete*.

In this paper a *continuous* adjoint formulation for aerodynamic shape optimization using the compressible Reynolds-Averaged Navier-Stokes equations and the Baldwin-Lomax turbulence model is presented. This work derives from that published in Refs. [16] and [17]. In particular, this paper focuses on a study of the accuracy of the gradient information resulting from the application of a viscous adjoint to the problems of inverse design and drag minimization. The comparisons made here are meant to validate the use of a viscous adjoint formulation in aerodynamic design and to answer questions regarding the level of convergence in the flow and adjoint solutions necessary to obtain sufficiently accurate gradient information.

Section 3 presents the governing equations of the flow in the specific form used in the derivation of the adjoint equations. Section 4 presents a general description of the adjoint design method and its formulation for the Navier-Stokes equations. In Sec. 5, the convergence properties of the flow and adjoint solvers are briefly discussed. Section 6 summarizes the derivation of the viscous adjoint terms and the gradient formulae. Finally, Sec. 7 discusses the results of the gradient accuracy study by direct comparison with finite-difference gradients. The study focuses on the effects of mesh resolution, flow solver convergence, adjoint solver convergence, and step size, for the calculation of finite-difference and adjoint gradients. In addition, Sec. 7 includes the results of sample design calculations.

Because of the increased mesh point count needed to resolve the boundary layer, the additional cost of computing the viscous terms, the need to use a turbulence model, and the slower convergence of the Navier-Stokes equations on stretched meshes, the computational costs for viscous design are at least an order of magnitude greater than for inviscid design. Therefore, the reduction of computational costs is

crucial to the further development of practically effective viscous design methods. The design method, which is greatly accelerated by the use of control theory, can be further enhanced by the use of parallel computing. In this study a parallel implementation using a domain decomposition approach and the MPI standard for communication was used. Section 7 contains information regarding the speedup obtained with the use of parallel computing.

3 The Navier-Stokes Equations

For the rest of this paper, three-dimensional expressions are derived for formulational completeness, although only two-dimensional computations were performed during the course of this study. For the derivations that follow, it is convenient to adopt the convention of indicial notation where a repeated index “ i ” implies summation over $i = 1$ to 3. The three-dimensional Navier-Stokes equations in Cartesian coordinates (x_1, x_2, x_3) take the form

$$\frac{\partial w}{\partial t} + \frac{\partial f_i}{\partial x_i} = \frac{\partial f_{vi}}{\partial x_i} \quad \text{in } \mathcal{D}, \quad (1)$$

where the state vector w , inviscid flux vector f , and viscous flux vector f_v in are described respectively by

$$w = \begin{Bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ \rho u_3 \\ \rho E \end{Bmatrix}, \quad (2)$$

$$f_i = \begin{Bmatrix} \rho u_i \\ \rho u_i u_1 + p \delta_{i1} \\ \rho u_i u_2 + p \delta_{i2} \\ \rho u_i u_3 + p \delta_{i3} \\ \rho u_i H \end{Bmatrix}, \quad (3)$$

$$f_{vi} = \begin{Bmatrix} 0 \\ \sigma_{ij} \delta_{j1} \\ \sigma_{ij} \delta_{j2} \\ \sigma_{ij} \delta_{j3} \\ u_j \sigma_{ij} + k \frac{\partial T}{\partial x_i} \end{Bmatrix}. \quad (4)$$

In these definitions, ρ is the density, u_1, u_2, u_3 are the Cartesian velocity components, E is the total

energy and δ_{ij} is the Kronecker delta function. The pressure is determined by the equation of state

$$p = (\gamma - 1) \rho \left\{ E - \frac{1}{2} (u_i u_i) \right\},$$

and the stagnation enthalpy is given by

$$H = E + \frac{p}{\rho},$$

where γ is the ratio of the specific heats. The viscous stresses may be written as

$$\sigma_{ij} = \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \lambda \delta_{ij} \frac{\partial u_k}{\partial x_k},$$

where μ and λ are the first and second coefficients of viscosity. The coefficient of thermal conductivity and the temperature are defined by

$$k = \frac{\gamma \mu}{Pr}, \quad T = \frac{p}{(\gamma - 1) \rho}. \quad (5)$$

For discussion of real applications using a discretization on a body conforming structured mesh, it is also useful to consider a transformation to the computational coordinates (ξ_1, ξ_2, ξ_3) defined by the metrics

$$K_{ij} = \left[\frac{\partial x_i}{\partial \xi_j} \right], \quad J = \det(K), \quad K_{ij}^{-1} = \left[\frac{\partial \xi_i}{\partial x_j} \right].$$

The Navier-Stokes equations can then be written in computational space as

$$\frac{\partial (Jw)}{\partial t} + \frac{\partial (F_i - F_{vi})}{\partial \xi_i} = 0 \quad \text{in } \mathcal{D}, \quad (6)$$

where the inviscid and viscous flux contributions are now defined with respect to the computational cell faces by $F_i = S_{ij} f_j$ and $F_{vi} = S_{ij} f_{vj}$, and the quantity $S_{ij} = JK_{ij}^{-1}$ is used to represent the projection of the ξ_i cell face along the x_j axis.

For high Reynolds numbers, the Navier-Stokes equations are time averaged, and a suitable turbulence model needs to be included to account for the appearance of the Reynolds stresses. In this work, the Baldwin-Lomax turbulence model is used.

4 General Formulation of the Optimum Design Problem for the Navier-Stokes equations

In this section, the general formulation of the optimum design method using the Navier-Stokes equations is briefly described. A more thorough derivation can be found in Refs. [16, 17].

For the adjoint-based optimal design of airfoils and wings, a control function which represents the physical location of the boundary, \mathcal{F} , is defined and parameterized using a fixed set of design variables. A cost function, I , which represents an aerodynamic performance measure such as the drag coefficient, C_d , the lift-to-drag ratio, L/D , or the difference from a specified pressure distribution, is constructed. The fact that the flow solution must satisfy the governing equations of the flow (Navier-Stokes equations) can be introduced as a constraint in the problem. With all this in mind, the design problem is treated as a control problem where the control function, \mathcal{F} , is altered to minimize the cost function, I , subject to the above constraint, which expresses the interdependence of the cost function and the control function in the domain of interest. The process of minimization of the control function is based on the direction of steepest descent provided by the sensitivity derivatives of the cost function with respect to the design variables. As mentioned in Sec. 2 the fundamental advantage of the adjoint-based optimum design method lies in the procedure followed to obtain these sensitivity derivatives.

Suppose that the performance of the design is measured by a cost function of the form

$$I = \int_{\mathcal{B}} \mathcal{M}(w, S) dB_{\xi} + \int_{\mathcal{D}} \mathcal{P}(w, S) dD_{\xi},$$

containing both boundary and field contributions, where dB_{ξ} and dD_{ξ} are the surface and volume elements in the computational domain. In general, \mathcal{M} and \mathcal{P} will depend on both the flow variables w and the metrics S defining the computational space.

A shape change produces a variation in the flow solution δw and the metrics δS which in turn produce a variation in the cost function

$$\delta I = \int_{\mathcal{B}} \delta \mathcal{M}(w, S) dB_{\xi} + \int_{\mathcal{D}} \delta \mathcal{P}(w, S) dD_{\xi}, \quad (7)$$

with

$$\begin{aligned} \delta \mathcal{M} &= [\mathcal{M}_w]_I \delta w + \delta \mathcal{M}_{II}, \\ \delta \mathcal{P} &= [\mathcal{P}_w]_I \delta w + \delta \mathcal{P}_{II}, \end{aligned} \quad (8)$$

where we use the subscripts I and II to distinguish between the contributions associated with the variation of the flow solution δw and those associated with the metric variations δS . Thus $[\mathcal{M}_w]_I$ and $[\mathcal{P}_w]_I$ represent $\frac{\partial \mathcal{M}}{\partial w}$ and $\frac{\partial \mathcal{P}}{\partial w}$ with the metrics fixed,

while $\delta \mathcal{M}_{II}$ and $\delta \mathcal{P}_{II}$ represent the contribution of the metric variations δS to $\delta \mathcal{M}$ and $\delta \mathcal{P}$.

In the steady state, the constraint equation (6) specifies the variation of the state vector δw by

$$\frac{\partial}{\partial \xi_i} \delta (F_i - F_{vi}) = 0. \quad (9)$$

Here δF_i and δF_{vi} can also be split into contributions associated with δw and δS using the notation

$$\begin{aligned} \delta F_i &= [F_{iw}]_I \delta w + \delta F_{iII} \\ \delta F_{vi} &= [F_{v iw}]_I \delta w + \delta F_{viII}. \end{aligned} \quad (10)$$

Multiplying by a co-state vector ψ and integrating over the domain produces

$$\int_{\mathcal{D}} \psi^T \frac{\partial}{\partial \xi_i} \delta (F_i - F_{vi}) dD_{\xi} = 0. \quad (11)$$

If ψ is differentiable this may be integrated by parts to give

$$\int_{\mathcal{B}} n_i \psi^T \delta (F_i - F_{vi}) dB_{\xi} \quad (12)$$

$$- \int_{\mathcal{D}} \frac{\partial \psi^T}{\partial \xi_i} \delta (F_i - F_{vi}) dD_{\xi} = 0. \quad (13)$$

Since the left hand expression equals zero, it may be subtracted from the variation in the cost function (7) to give

$$\begin{aligned} \delta I &= \int_{\mathcal{B}} [\delta \mathcal{M} - n_i \psi^T \delta (F_i - F_{vi})] dB_{\xi} \\ &+ \int_{\mathcal{D}} \left[\delta \mathcal{P} + \frac{\partial \psi^T}{\partial \xi_i} \delta (F_i - F_{vi}) \right] dD_{\xi}. \end{aligned} \quad (14)$$

Now, since ψ is an arbitrary differentiable function, it may be chosen in such a way that δI no longer depends explicitly on the variation of the state vector δw . The gradient of the cost function can then be evaluated directly from the metric variations without having to recompute the variation δw resulting from the perturbation of each design variable.

Comparing equations (8) and (10), the variation δw may be eliminated from (14) by equating all field terms with subscript “ I ” to produce a differential adjoint system governing ψ

$$\frac{\partial \psi^T}{\partial \xi_i} [F_{jw} - F_{vjw}]_I + [\mathcal{P}_w]_I = 0 \quad \text{in } \mathcal{D}. \quad (15)$$

The corresponding adjoint boundary condition is produced by equating the subscript “ I ” boundary terms in equation (14) to produce

$$n_i \psi^T [F_{jw} - F_{vjw}]_I = [\mathcal{M}_w]_I \quad \text{on } \mathcal{B}. \quad (16)$$

The remaining terms from equation (14) then yield a simplified expression for the variation of the cost function which defines the gradient

$$\begin{aligned} \delta I &= \int_{\mathcal{B}} \left\{ \delta \mathcal{M}_{II} - n_i \psi^T [\delta F_i - \delta F_{vi}]_{II} \right\} d\mathcal{B}_\xi \\ &+ \int_{\mathcal{D}} \left\{ \delta \mathcal{P}_{II} + \frac{\partial \psi^T}{\partial \xi_i} [\delta F_i - \delta F_{vi}]_{II} \right\} d\mathcal{D}_\xi, \end{aligned} \quad (17)$$

The details of the formula for the gradient depend on the way in which the boundary shape is parameterized as a function of the design variables, and the way in which the mesh is deformed as the boundary is modified. Using the relationship between the mesh deformation and the surface modification, the field integral is reduced to a surface integral by integrating along the coordinate lines emanating from the surface. Thus the expression for δI is finally reduced to

$$\delta I = \int_{\mathcal{B}} G \delta S d\mathcal{B}_\xi = \mathcal{G} \delta \mathcal{F}, \quad (18)$$

where \mathcal{F} represents the design variables, and \mathcal{G} is the gradient, which is a function defined over the boundary surface. The advantage is that (18) is independent of δw , with the result that the gradient of I with respect to an arbitrary number of design variables can be determined with only a single flow-field evaluation and a single adjoint evaluation in a design cycle. Then, the computational cost of a single design cycle is roughly equivalent to the cost of two flow solutions since the adjoint problem has similar complexity. When the number of design variables becomes large, the computational efficiency of the control theory approach over the traditional approach (which requires direct evaluation of the gradients by individually varying each design variable and recomputing the flow field) becomes compelling.

Once equation (18) is established, an improvement can be made with a shape change

$$\delta \mathcal{F} = -\lambda \mathcal{G},$$

where λ is positive and small enough that the first variation is an accurate estimate of δI . Then

$$\delta I = -\lambda \mathcal{G}^T \mathcal{G} < 0.$$

After making such a modification, the gradient can be recalculated and the process repeated to follow a path of steepest descent until a minimum is reached.

The boundary conditions satisfied by the flow equations restrict the form of the left hand side of

the adjoint boundary condition (16). Consequently, the boundary contribution to the cost function \mathcal{M} cannot be specified arbitrarily. Instead, it must be chosen from the class of functions which allow cancellation of all terms containing δw in the boundary integral of equation (14). On the other hand, there is no such restriction on the specification of the field contribution to the cost function \mathcal{P} , since these terms may always be absorbed into the adjoint field equation (15) as source terms.

The detailed inviscid adjoint terms have been previously derived in [18, 4] and a detailed derivation of the viscous adjoint terms and the corresponding viscous adjoint boundary conditions can also be found in [16, 17].

5 Flow Solvers

Although, as will be shown later, the adjoint method is fairly insensitive to the convergence level of either the flow or the adjoint systems, an accurate, robust and fast solver for both the flow and the adjoint equations is required for an usable design method. In this study FLO103, a RANS solver developed by Martinelli [19, 20], is used to satisfy the requirements of accuracy, convergence, and robustness. FLO103 solves the steady two-dimensional RANS equations using a modified explicit multistage Runge-Kutta time-stepping scheme. A finite volume technique and second order central differencing in space are applied to the integral form of the Navier-Stokes equations.

The Jameson-Schmidt-Turkel (JST) scheme with adaptive coefficients for artificial dissipation is used to prevent odd-even oscillations and to allow the clean capture of shock waves and contact discontinuities. In addition, local time stepping, implicit residual smoothing, and the multi-grid method are applied to accelerate convergence to steady-state solutions. The algebraic turbulent model of Baldwin and Lomax is used to calculate the eddy-viscosity necessary to estimate the effect of the Reynolds stress. Because of the need to achieve high levels of convergence in the flow solver, the turbulence model is *frozen* after the average density residual has converged 4 orders of magnitude. Figure 1 shows a typical convergence history of the average density residual for both calculations at fixed angle of at-

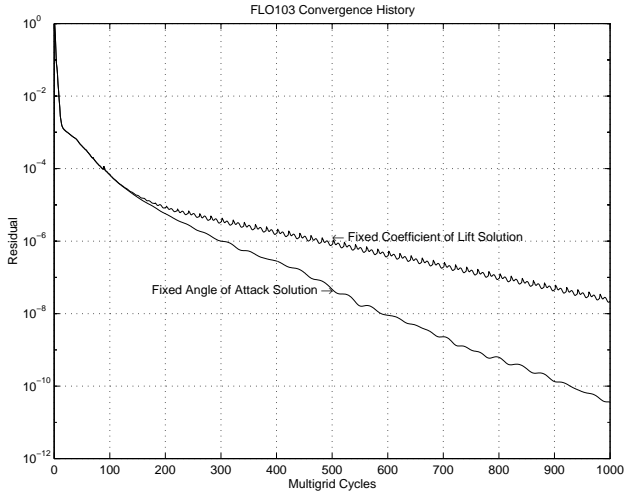


Figure 1: Convergence History for the Average Density Residual of a Typical FLO103 Calculation.

tack and fixed coefficient of lift (obtained by periodic updates of the angle of attack). As can be seen, engineering accuracy calculations can be obtained in under 100 multigrid cycles. In addition, at the expense of larger computational cost, the flow solver can be made to converge at least 10 orders of magnitude. This ability to obtain highly converged RANS solutions will be necessary in the results of Sec. 7. The adjoint solution is obtained with the exact same numerical techniques used for the flow solution. The implementation exactly mirrors the flow solution modules inside FLO103, except for the boundary conditions which are imposed on the co-state variables. Acceptable convergence rates are obtained for the solution of the adjoint equation. It is clear from Fig. 2 that additional work is needed to make sure that the convergence of the adjoint solver achieves the same asymptotic rates as the flow solution. However, as will be shown in Sec. 7, only mild levels of convergence in the adjoint equation are necessary to obtain highly accurate gradient information. Therefore, the current implementation of the adjoint solution is fully adequate for all practical purposes.

6 Derivation of the Viscous Adjoint Terms

Following the steps described in the previous section, the derivations of the adjoint equations, the adjoint boundary conditions, and the gradient terms

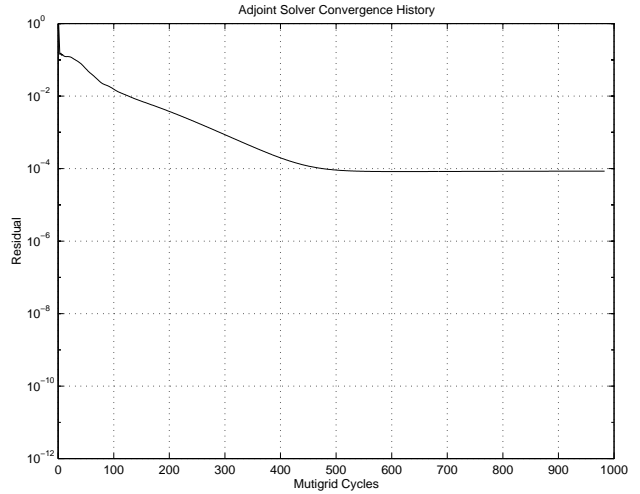


Figure 2: Convergence History of the Adjoint Solver for the Average Residual in the First Co-State Variable.

are briefly presented here. The reader is referred to Refs. [16, 17] for some of the missing details omitted for brevity.

Adjoint equations

The resulting adjoint equations are given by

$$C_i^T \frac{\partial \psi}{\partial \xi_i} - M^{-1T} \tilde{L} \psi = 0 \quad \text{in } \mathcal{D}. \quad (19)$$

The first and the second terms come from the convective and diffusive terms of the Navier-Stokes equations respectively. C_i are the inviscid Jacobian matrices in the transformed space which are given by

$$C_i = S_{ij} \frac{\partial f_j}{\partial w}.$$

The derivation of the viscous adjoint terms is simplified by transforming to primitive variables

$$\tilde{w}^T = (\rho, u_1, u_2, u_3, p)^T,$$

because the viscous stresses depend on the velocity derivatives $\frac{\partial u_i}{\partial x_j}$, while the heat flux can be expressed as

$$\kappa \frac{\partial}{\partial x_i} \left(\frac{p}{\rho} \right),$$

where $\kappa = \frac{k}{R} = \frac{\gamma \mu}{Pr(\gamma-1)}$. The relationship between the conservative and primitive variations is defined by the expressions

$$M = \frac{\partial w}{\partial \tilde{w}}, M^{-1} = \frac{\partial \tilde{w}}{\partial w}.$$

The transpose of M^{-1} is expressed as

$$M^{-1T} = \begin{bmatrix} 1 & -\frac{u_1}{\rho} & -\frac{u_2}{\rho} & -\frac{u_3}{\rho} & \frac{(\gamma-1)u_i u_i}{2} \\ 0 & \frac{1}{\rho} & 0 & 0 & -(\gamma-1)u_1 \\ 0 & 0 & \frac{1}{\rho} & 0 & -(\gamma-1)u_2 \\ 0 & 0 & 0 & \frac{1}{\rho} & -(\gamma-1)u_3 \\ 0 & 0 & 0 & 0 & \gamma-1 \end{bmatrix}.$$

The viscous adjoint operator in primitive variables, \tilde{L} , can be expressed as

$$\begin{aligned} (\tilde{L}\psi)_1 &= -\frac{p}{\rho^2} \frac{\partial}{\partial \xi_l} \left(S_{lj} \kappa \frac{\partial \theta}{\partial x_j} \right) \\ (\tilde{L}\psi)_{i+1} &= \frac{\partial}{\partial \xi_l} \left\{ S_{lj} \left[\mu \left(\frac{\partial \phi_i}{\partial x_j} + \frac{\partial \phi_j}{\partial x_i} \right) + \lambda \delta_{ij} \frac{\partial \phi_k}{\partial x_k} \right] \right\} \quad i = 1, 2, 3 \\ &+ \frac{\partial}{\partial \xi_l} \left\{ S_{lj} \left[\mu \left(u_i \frac{\partial \theta}{\partial x_j} + u_j \frac{\partial \theta}{\partial x_i} \right) + \lambda \delta_{ij} u_k \frac{\partial \theta}{\partial x_k} \right] \right\} \\ &- \sigma_{ij} S_{lj} \frac{\partial \theta}{\partial \xi_l} \\ (\tilde{L}\psi)_5 &= \rho \frac{\partial}{\partial \xi_l} \left(S_{lj} \kappa \frac{\partial \theta}{\partial x_j} \right), \end{aligned}$$

where ψ_{j+1} is set to ϕ_j for $j = 1, 2, 3$ and ψ_5 is set to θ in order to make use of the summation convention.

In the derivation of the viscous adjoint terms, the viscosity and heat conduction coefficients μ and k are assumed independent of the flow so that their variations may be neglected. In the case of turbulent flow, if the flow variations are found to result in significant changes in the turbulent viscosity, it may eventually be necessary to include its variation in the calculations.

Adjoint Boundary Conditions

It was recognized in Sec.4 that the boundary conditions satisfied by the flow equations restrict the form of the performance measure that may be chosen for the cost function. There must be a direct correspondence between the flow variables for which variations appear in the variation of the cost function, and those variables for which variations appear in the boundary terms arising during the derivation of the adjoint field equations. Otherwise it would be impossible to eliminate the dependence of δI on δw through proper specification of the adjoint boundary condition.

Boundary Conditions Arising from the Momentum Equations

In the inverse design case, in order to design a shape which will lead to a desired pressure distribution, a

natural choice for the cost function is to set

$$I = \frac{1}{2} \int_{\mathcal{B}} (p - p_d)^2 dS,$$

and the corresponding boundary condition is

$$\phi_k = n_k (p - p_d). \quad (20)$$

For high Reynolds number flows, the pressure gradient in the direction normal to the boundary layer is very small, and the specification of a target pressure in inverse design should work in a similar fashion to inviscid flow.

For inviscid drag minimization, the cost function reflects the pressure drag alone, which can be expressed as

$$D_{nf} = \int_{\mathcal{B}} n_i S_{2j} p d\mathcal{B}_{\xi},$$

and the boundary condition is similar to that of inviscid inverse design in Eq. 20.

For viscous drag minimization, the cost function should measure the skin friction drag. The friction force is given by

$$D_{nf} = \int_{\mathcal{B}} n_i S_{2j} \sigma_{ij} d\mathcal{B}_{\xi}.$$

Expressed in terms of the surface stress τ_i , this corresponds to

$$D_{nf} = \int_{\mathcal{B}} n_i \tau_i dS,$$

and the resulting adjoint boundary condition is given by,

$$\phi_k = n_k.$$

Note that this choice of boundary condition eliminates one of the terms in the variation in the cost function so that it need not be included in the gradient calculation. This is also a special case of the boundary condition for inviscid drag minimization, and, therefore, it is used for the total drag minimization cases.

Boundary Conditions Arising from the Energy Equation

The form of the boundary terms arising from the energy equation depends on the choice of temperature boundary condition at the wall. For the adiabatic case, the boundary contribution is

$$\int_{\mathcal{B}} k \delta T \frac{\partial \theta}{\partial n} d\mathcal{B}_{\xi},$$

while for the constant temperature case the boundary term is slightly different and can be found in Ref. [17]. One possibility is to introduce a contribution into the cost function which is dependent on T or $\frac{\partial T}{\partial n}$ so that the appropriate cancelation would occur. Since there is little physical intuition to guide the choice of such a cost function for aerodynamic design, a more natural solution is to set

$$\theta = 0$$

in the constant temperature case or

$$\frac{\partial \theta}{\partial n} = 0$$

in the adiabatic case.

Gradient Formulae

The specific form of the gradient formulae for the aerodynamic cost functions used in this work are omitted here for the sake of brevity. However, they have been presented in Refs. [16, 17] and the reader is referred to them for more details.

7 Results

This section presents the results of the gradient accuracy study for both the Euler and Navier-Stokes equations, as well as some simple examples of the use of the resulting gradient information in inverse design and drag minimization problems. Gradient accuracy is assessed by comparison with finite-difference gradients and by examination of the changes in the magnitude of the gradients as various parameters are changed.

Two main considerations need to be taken into account when discussing the characteristics of finite difference gradients: the level of numerical convergence in the flow solution, and the step size chosen to discretize the gradient formula. The level of numerical convergence in the flow solution gives an indication of the error in the calculation of the aerodynamic figure of merit that is being differentiated: the lower the size of the flow solution residual, the more accurate the aerodynamic information will be. On the other hand, for simple one-sided finite difference approximations of the first derivative of a given aerodynamic quantity, the truncation error in the approximation is proportional to the step size

chosen. Therefore, in order to minimize the truncation error, a small step size is preferred. In theory, very small step sizes and high levels of convergence would provide the most accurate approximation to the true value of the gradient. However, one cannot arbitrarily decrease the step size, because, due to the finite precision in the representation of floating point quantities, the error due to roundoff can become significant. In addition, the choice of step size and convergence level are intimately related if one intends to obtain accurate gradient information.

Suppose that the gradient, G , we are interested in is defined as

$$G = \frac{dI}{d\mathcal{F}},$$

with appropriate definitions of I and \mathcal{F} given earlier. If a small step in a given design variable produces a change in the aerodynamic cost function, ΔI , and the error in our numerical solution of the flow equations is proportional to ϵ , we have, approximately, that

$$G = \frac{(I + \Delta I \pm \epsilon) - (I \pm \epsilon)}{\Delta \mathcal{F}} = \frac{\Delta I}{\Delta \mathcal{F}} \left(1 \pm \frac{\epsilon}{\Delta I}\right).$$

That is, the error in the approximation of the derivative depends on the value of $\frac{\epsilon}{\Delta I}$. The error in the numerical approximation, ϵ , is proportional to the residual in the calculation, R_a , while the change in the aerodynamic cost function, ΔI , is proportional to the step size. Suppose, for the sake of argument that $I = C_d$. Then, for a good approximation, $\frac{\epsilon}{\Delta I} < 10^{-2}$, and if the change in the cost function is measurable, say one drag count, $\Delta I = 0.0001$, the requirement on the level of convergence of the flow solver is

$$R_a < 10^{-6},$$

which places a heavy burden on the CFD procedure both in terms of computational cost and in the requirement that the flow solver converge for at least 6 orders of magnitude.

Experience shows, and we will see this confirmed by numerical experiments later, that, in order to obtain accurate gradient information, the step size for the finite difference approximation needs to be chosen to be $\Delta \mathcal{F} < 10^{-4}$. Since the gradients that are typically produced are of $\mathcal{O}(1)$, this implies that

$$\Delta \mathcal{F} \sim \Delta I > 10^{-4},$$

which leads to a contradiction, since earlier, we had required that the change in the shape should be

$\Delta\mathcal{F} < 10^{-4}$. Since the aerodynamic cost function will typically be an integral over a surface, the requirements on step size and convergence level are even more severe. Despite the fact that this analysis states some simple assumptions, it is representative of the fact that, in the finite difference method, an *a priori* choice of step size and convergence level that leads to accurate gradient information is at the very least problematic.

The adjoint method is only mildly affected by these issues. The formula for the gradient in general terms is given by

$$\delta I = \left\{ \frac{\partial I}{\partial \mathcal{F}}^T - \psi^T \left[\frac{\partial R}{\partial \mathcal{F}} \right] \right\} \delta \mathcal{F}$$

which is completely analytic, and therefore, the adjoint gradients should be *independent* of the step size chosen to modify the geometry. Moreover, since the expression for the gradient above is a function of both the flow solution, w , and the adjoint solution, ψ , its evaluation will have a numerical error, $\epsilon_{adjoint} \sim \max(R_{flow}, R_{costate})$, where R_{flow} and $R_{costate}$ are the residuals of the flow and costate solutions. This form of the dependence of the accuracy of the gradient on the residuals of both solutions places less stringent requirements on convergence than the finite difference method, making the adjoint approach even more computationally affordable.

The results that follow focus on the variations in the finite difference and adjoint gradients obtained when each of the above parameters is varied in turn while keeping the others constant. In addition, the issue of mesh resolution will be studied. The remainder of this section is organized into three subsections for the different problems for which the gradient study was performed.

Euler Inverse Design Problem

In this section we present the results of the gradient accuracy study for the case of inverse design using the Euler equations. These results are meant to serve as a benchmark for later comparison with the viscous data. The aerodynamic cost function chosen is given by:

$$I = \frac{1}{2} \int_{\mathcal{B}} (p - p_d)^2 dS, \quad (21)$$

which is simply the Euclidean norm of the difference between the current pressure distribution and a desired target, p_d , at constant angle of attack, α . The gradient of the above cost function is obtained with respect to variations in 54 Hicks-Henne sine ‘‘bump’’ functions centered at various locations along the upper and lower surfaces of a NACA 64A410 airfoil. The locations of these geometry perturbations are ordered sequentially such that they start at the lower surface of the trailing edge, proceed forward to the leading edge, and then back around to the upper surface of the trailing edge. The target pressure used is that produced by the same flow solver, FLO103, for a Korn airfoil at $M = 0.75$ and at an angle of attack $\alpha = 0.124$. The mesh used is an C-mesh of size 192×32 , whose resolution has been shown to be appropriate for inviscid calculations.

Figures 3 and 4 show the values of the 54 components of the computed gradients for different levels of flow solver convergence in the finite difference and adjoint methods. As we can see, for Euler calculations, the finite difference information requires a minimum of 4 to 5 orders of magnitude of convergence before the gradients are fully converged. In contrast, the adjoint information is essentially unchanged if the level of convergence in the flow solver is at least 2.5 orders of magnitude. Figures 5 and 6 show the variation in the values of the gradients for both the finite difference and adjoint techniques when the step size is varied. As predicted by the analysis at the beginning of this section, the results from the finite difference method are much more sensitive to the value of the step size than those obtained with the adjoint method. In fact, although convergence to a set of values is achieved for a range of step sizes, we see how, for a fixed convergence level in the flow solver of 7 orders of magnitude, the gradient information starts to deviate when the step size is chosen to be 10^{-6} . In contrast, the adjoint method produces gradients which are completely independent of the the step size chosen as predicted earlier. Figure 7 shows the dependence of the gradients produced by the adjoint method on the convergence tolerance of the adjoint solver. As mentioned earlier in Sec. 5 the adjoint solver is only required to converge a small amount before the gradients stabilize; in this case, the threshold appears to be at a convergence level of 1 order of magnitude. This allows the

computation of the adjoint solution with substantially decreased computational cost. Figure 8 shows a comparison between the most accurate gradients obtained using both techniques. It is clear that small discrepancies exist on the trailing edge area (first and last few design variables) and around the location of the upper surface shock (design variable no. 38), but it is clear that the gradient information resulting from the adjoint method can be used with confidence in design problems of this kind. Finally, Figs. 9 and 10 show the outcome of mesh resolution studies on the accuracy of the finite difference and adjoint gradients. Both figures show that the results converge nicely as the mesh is refined, except in the neighborhood of the trailing edge, where the mesh resolution is perhaps still not fully adequate. When the results of both plots are superimposed, we find, as expected, that the finite difference and adjoint gradient discrepancies disappear as the mesh is refined. This is also to be expected because the continuous adjoint formulation should reduce to the discrete one in the limit of infinite mesh resolution.

Similar results to the ones presented here were first published by Reuther [15]. The intent here is to contrast the inviscid results with the viscous ones in the following sections. The reader is referred to Reuther’s work for more detailed comparisons.

Navier-Stokes Inverse Design Problem

The results in this section exactly mimic those in the previous one, with the main difference being that the Reynolds Averaged Navier-Stokes equations are used as a model of the flow. For inverse design, the same aerodynamic cost function in Eq. 21 is used with 50 Hicks-Henne sine “bump” functions distributed in a similar fashion. The mesh used is a Navier-Stokes C-mesh of size 512×64 around a Korn airfoil. The target pressure specified is that of NACA 64A410 airfoil at $M = 0.75$ and $\alpha = 0.0$. The Reynolds number of all calculations was $Re = 6.5$ million.

Figures 11 and 12 show the values of the 50 components of the computed gradients for different levels of flow solver convergence in the finite difference and adjoint methods. For Navier-Stokes calculations, the finite difference information requires a minimum of 6 orders of magnitude of convergence before the gradients are fully converged. In con-

trast, once again, the adjoint information is essentially unchanged if the level of convergence in the flow solver is at least 4 orders of magnitude. This requirement of lower levels of convergence is even more relevant in viscous calculations since the convergence rates are slower than in the inviscid case, and, in addition, not all flow solvers can converge to that large extent. Figures 13 and 14 show the variation for both the finite difference and adjoint gradients when the step size is changed. Once again, the variation of step size is reflected much more heavily on the finite difference information. In fact, as opposed to the inviscid case, there doesn’t seem to be a range of step sizes for which the gradient information is converged. In particular, as the step size is reduced beyond 10^{-5} , the error in the calculation of the gradient increases changing drastically both the magnitude and the sign of the values, which can lead to erroneous design decisions. As predicted by the analysis at the beginning of this section, the results from the adjoint method are effectively *independent* of the step size chosen to calculate the gradients. Figure 15 shows the dependence of the gradients produced by the adjoint method on the convergence tolerance of the adjoint solver. As was the case in the inviscid results, the adjoint solver is only required to converge 1.5 orders of magnitude to produce stable gradients. This allows the computation of the adjoint solution with substantially decreased computational cost. Figure 16 shows a comparison between the most accurate gradients obtained using both techniques. Slightly larger discrepancies exist when compared with the inviscid case, which are not surprising due to the degree of uncertainty in the choice of the finite difference gradients that one should be comparing to. However, there is a general agreement on all trends that validates the implementation of our adjoint method. Finally, Figs. 17 and 18 show the outcome of mesh resolution studies on the accuracy of the finite difference and adjoint gradients. The second figure shows that, as the mesh is refined, subtle changes are seen on the upper surface of the airfoil, particularly at the locations of the initial double shock wave. Otherwise the gradients have achieved grid convergence. The finite difference result, although at first sight reasonable, under closer scrutiny shows some disturbing features. The gradient solution that departs more clearly from the

others is in fact the one corresponding to the finest mesh in the sequence. We had expected that the differences in the gradients would decrease as the mesh resolution was increased, but this is clearly not the case. In fact, when the finite difference and adjoint gradients are over-plotted in Figs. 19 and 20 we see that there is closer agreement between the two methods for the coarser mesh, rather than the finest. This issue is one whose resolution will require further study. Figures 21 and 22 show the progress in two typical viscous inverse design calculations. The norm of the pressure error is decreased from 0.0573 to 0.0056 in 100 design iterations in the first case, and from 0.0504 to 0.0043 in the second. Both target pressure distributions are achieved around the majority of the airfoil. The areas around the trailing edge lower surface are very sensitive to small changes in the geometry and require more geometry control to be exactly matched.

Navier-Stokes Drag Minimization Problem

Finally, the results of the gradient accuracy study for the drag minimization problem are discussed in this section. The aerodynamic figure of merit is the total drag of the airfoil (pressure + skin friction). The presentation of the results exactly follows that of the previous two sections. Computations were performed with an RAE 2822 airfoil at a fixed coefficient of lift, $C_l = 0.84$, and $M = 0.73$. The Reynolds number of all calculations was set to $Re = 6.5$ million. A C-mesh of size 512×64 was also used for all calculations. As before, 50 Hicks-Henne sine “bump” functions are placed along the surface of the airfoil.

Figures 23 and 24 show the values of the 50 components of the computed gradients for different levels of flow solver convergence in the finite difference and adjoint methods. For viscous drag minimization cases, the finite difference gradients are also very sensitive to the level of flow solver convergence to the point that 6 to 7 orders of magnitude are necessary to assure converged gradients. On the other hand, the adjoint gradients are effectively independent of the flow solver convergence after 4 orders of magnitude. As mentioned before, less severe convergence requirements result in large computational savings because typical Navier-Stokes solvers have

much slower asymptotic convergence rates than inviscid ones. In addition, the asymptotic convergence rate slows down after the solution has converged 2 orders of magnitude as can be seen in Fig. 1. Figures 25 and 26 paint a similar picture to the one presented before: finite difference gradients are highly dependent on the step size chosen, to the extreme that, even as the step size is refined, and for a high level of convergence in the flow solver, we are unable to find values of the gradient that are step size independent. Once again, as predicted by the analysis at the beginning of this section, the results from the adjoint method are effectively *independent* of the step size chosen to calculate the gradients. Figure 27 shows the dependence of the gradients produced by the adjoint method on the convergence tolerance of the adjoint solver. In the case of drag minimization it appears that the requirement of adjoint solver convergence is a bit more severe. From the Figure, one can say that 2.5 orders of magnitude of convergence are necessary to produce accurate gradients. Finally, Figure 28 shows a comparison between the gradients obtained using both techniques. Some discrepancies exist which are attributed to several sources. First, there is the issue that the finite difference gradients never achieved step size independence, and, therefore, a step size had to be selected for this plot which may not be the correct one. Second, there is the issue of mesh resolution. It is possible that finer resolution is necessary to obtain closer agreement between the gradients from finite differences and those from the adjoint method when small changes in pressure and viscous drag need to be resolved. The source of the discrepancies is likely to be a combination of these two influences. However, we believe that the gradient information resulting from the adjoint method can be used with confidence in aerodynamic design problems. To substantiate this claim, the following figures show the result of two separate drag minimization studies using the RAE 2822 airfoil at $M = 0.83$, and fixed $C_l = 0.84$. Figure 29 shows the initial and final results of a drag minimization study where the cost was chosen to be the *inviscid* pressure drag alone. As we can see, after 14 design iterations, the *total* (pressure + skin friction) coefficient of drag has been reduced from $C_d = 0.0168$ to $C_d = 0.0108$. Notice that in this design case, the gradients used for minimization are

those of the *inviscid* drag *only*. Figure 30 shows the initial and final results of a viscous drag minimization study for the same airfoil, mesh, and coefficient of lift. In this case, the aerodynamic cost function was the *total* coefficient of drag (pressure + skin friction) and therefore, the gradient terms include information regarding the *viscous* effects present. In this case, after 17 design cycles, the total coefficient of drag was reduced from $C_d = 0.0168$ to $C_d = 0.0096$. This result is 12 counts lower than the previous one which emphasizes the fact that the additional gradient information coming from the viscous adjoint can lead to substantial improvements in performance of the design.

Parallel Implementation

Details of the parallel implementation of this program can be found in Refs. [21, 22]. The parallel implementation uses a domain decomposition approach, a SPMD (Single Program Multiple Data) structure, and the MPI (Message Passing Interface) standard for message passing. Figure 31 shows the parallel speedups obtained with up to 8 processors of an SGI Origin 2000 which can further enhance the usability of the design method. Additional processors can be added to the calculation, but, due to the relatively small size of two-dimensional calculations, it would be inefficient to do so.

8 Conclusions

The work presented in this paper addresses the accuracy of the sensitivity derivatives of aerodynamic cost functions that can be obtained using an adjoint method for flows governed by the Reynolds Averaged Navier-Stokes equations. The accuracy of these gradients is assessed by comparison with gradients obtained via the finite difference method, and by mesh and parameter refinement studies. The following conclusions are reached:

- The flow solver convergence requirements for the finite difference method to produce accurate gradient information are quite severe and substantially increase the computational cost of this method.
- As opposed to what may be expected from direct truncation error analysis, the choice of step

size for the calculation of derivatives using the finite difference method can cause a large error in the approximation of the gradients. Smaller step sizes (even disregarding the effect of round-off error) are not necessarily better; the error in the resulting gradients depends on both the convergence level of the flow solver and the step size chosen.

- Gradient information obtained using the adjoint method is much less dependent on the level of convergence of the flow solver. Typically a convergence level 2 orders of magnitude larger than in the finite difference method can be tolerated.
- Gradient information obtained using the adjoint method is insensitive to the step size chosen in the deformation of the aerodynamic configuration.
- The adjoint method requires only modest levels of convergence (1.5 – 2.5 orders) of the adjoint solver, thus reducing even further the computational cost of this procedure.
- Accurate gradient information can be obtained for flow governed by the RANS equations using the procedure outlined in Ref. [17].

Finally, several design examples were presented which show the ability to use the resulting gradient information in useful viscous aerodynamic design problems.

9 Acknowledgments

This research has been made possible by the generous support of the David and Lucille Packard Foundation in the form of a Stanford University School of Engineering Terman fellowship.

References

- [1] F. Bauer, P. Garabedian, D. Korn, and A. Jameson. *Supercritical Wing Sections II*. Springer Verlag, New York, 1975.
- [2] P. R. Garabedian and D. G. Korn. Numerical design of transonic airfoils. In B. Hubbard,

- editor, *Proceedings of SYNSPADE 1970*, pages 253–271, Academic Press, New York, 1971.
- [3] R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15:407–412, 1978.
- [4] A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3:233–260, 1988.
- [5] A. Jameson. Optimum aerodynamic design using CFD and control theory. *AIAA paper 95-1729*, AIAA 12th Computational Fluid Dynamics Conference, San Diego, CA, June 1995.
- [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. *AIAA paper 97-0103*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.
- [7] J. Reuther, J. J. Alonso, J. C. Vassberg, A. Jameson, and L. Martinelli. An efficient multiblock method for aerodynamic analysis and design on distributed memory systems. *AIAA paper 97-1893*, June 1997.
- [8] J. J. Reuther, A. Jameson, J. J. Alonso, M. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers: Part I. *Journal of Aircraft*, 1998. Accepted for publication.
- [9] J. J. Reuther, A. Jameson, J. J. Alonso, M. Rimlinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers: Part II. *Journal of Aircraft*, 1998. Accepted for publication.
- [10] A. Jameson. Re-engineering the design process through computation. *AIAA paper 97-0641*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.
- [11] J. Reuther, J.J. Alonso, M.J. Rimlinger, and A. Jameson. Aerodynamic shape optimization of supersonic aircraft configurations via an adjoint formulation on parallel computers. *AIAA paper 96-4045*, 6th AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Bellevue, WA, September 1996.
- [12] O. Baysal and M. E. Eleshaky. Aerodynamic design optimization using sensitivity analysis and computational fluid dynamics. *AIAA paper 91-0471*, 29th Aerospace Sciences Meeting, Reno, Nevada, January 1991.
- [13] W. K. Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *AIAA paper 97-0643*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.
- [14] J. Reuther, A. Jameson, J. Farmer, L. Martinelli, and D. Saunders. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. *AIAA paper 96-0094*, 34th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1996.
- [15] J. Reuther. Aerodynamic shape optimization using control theory. *Ph. D. Dissertation*, University of California, Davis, Davis, CA, June 1996.
- [16] A. Jameson, N. Pierce, and L. Martinelli. Optimum aerodynamic design using the Navier-Stokes equations. *AIAA paper 97-0101*, 35th Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1997.
- [17] A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamic design using the Navier-Stokes equations. *Theoretical and Computational Fluid Dynamics*, 10:213–237, 1998.
- [18] A. Jameson and J. Reuther. Control theory based airfoil design using the Euler equations. *AIAA paper 94-4272*, 5th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Panama City Beach, FL, September 1994.
- [19] L. Martinelli and A. Jameson. Validation of a multigrid method for the Reynolds averaged equations. *AIAA paper 88-0414*, 1988.
- [20] S. Tatsumi, L. Martinelli, and A. Jameson. A new high resolution scheme for compressible

viscous flows with shocks. *AIAA paper To Appear*, AIAA 33rd Aerospace Sciences Meeting, Reno, Nevada, January 1995.

- [21] J. J. Alonso, T. J. Mitty, L. Martinelli, and A. Jameson. A two-dimensional multigrid Navier-Stokes solver for multiprocessor architectures. In Satofuka, Periaux, and Ecer, editors, *Parallel Computational Fluid Dynamics: New Algorithms and Applications*, pages 435–442, Kyoto, Japan, May 1994. Elsevier Science B. V. Parallel CFD '94 Conference.
- [22] J. J. Alonso, L. Martinelli, and A. Jameson. Multigrid unsteady Navier-Stokes calculations with aeroelastic applications. *AIAA paper 95-0048*, AIAA 33rd Aerospace Sciences Meeting and Exhibit, Reno, NV, January 1995.

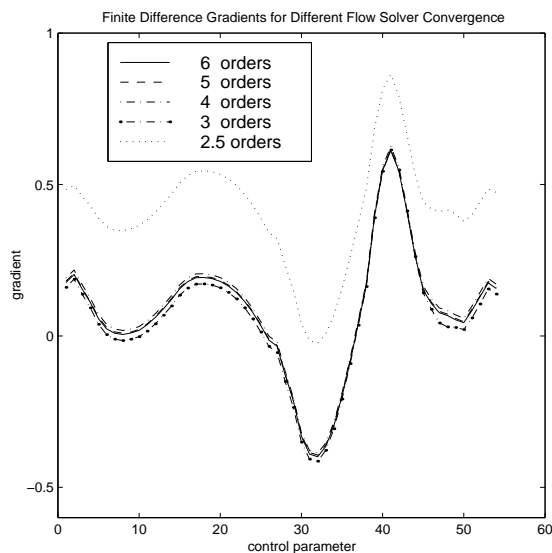


Figure 3: Inviscid Inverse Design: Finite Difference Gradients for Varying Flow Solver Convergence.




Figure 4: Inviscid Inverse Design: Adjoint Gradients for Varying Flow Solver Convergence.

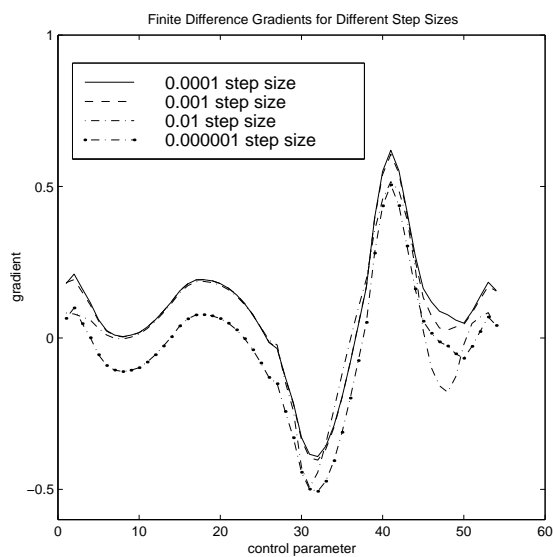


Figure 5: Inviscid Inverse Design: Finite Difference Gradients for Varying Step Sizes.

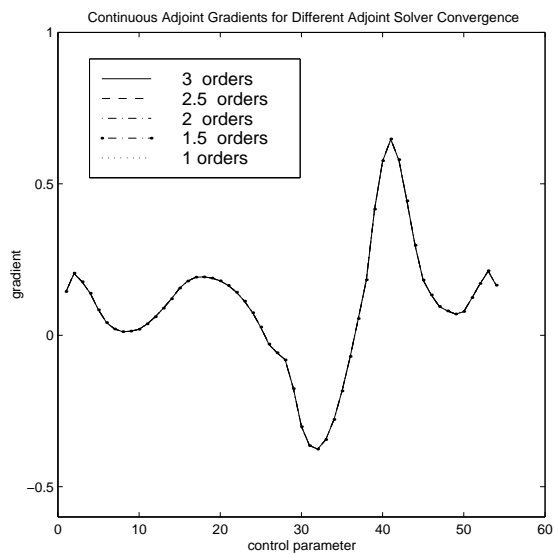


Figure 7: Inviscid Inverse Design: Adjoint Gradients for Varying Adjoint Solver Convergence Levels.

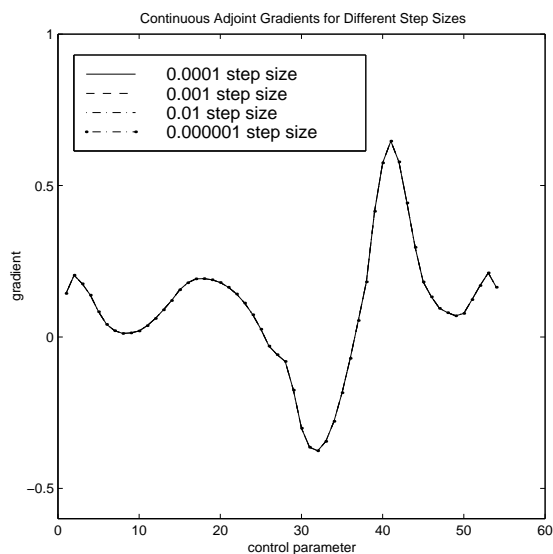


Figure 6: Inviscid Inverse Design: Adjoint Gradients for Varying Step Sizes.




Figure 8: Inviscid Inverse Design: Comparison of Finite Difference and Adjoint Gradients for Medium Mesh.

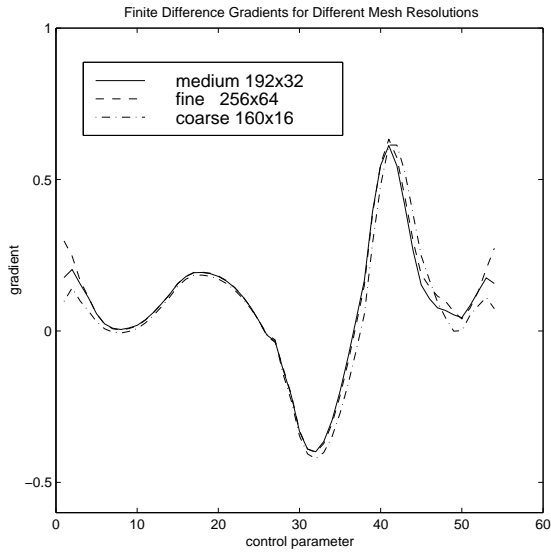


Figure 9: Inviscid Inverse Design: Comparison of Finite Difference Gradients for Varying Mesh Sizes.

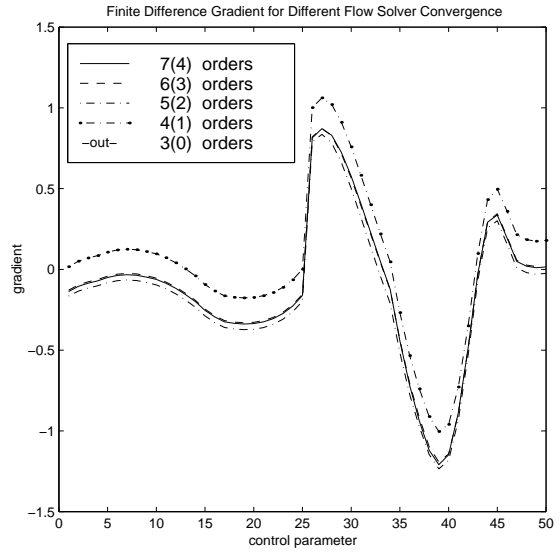


Figure 11: Navier-Stokes Inverse Design: Finite Difference Gradients for Varying Flow Solver Convergence.

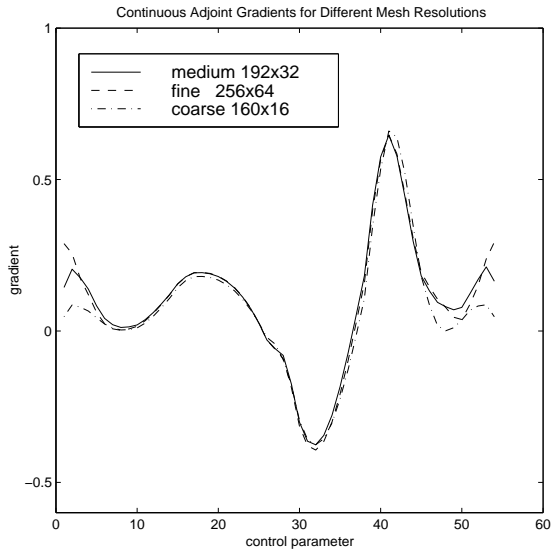


Figure 10: Inviscid Inverse Design: Comparison of Adjoint Gradients for Varying Mesh Sizes.




Figure 12: Navier-Stokes Inverse Design: Adjoint Gradients for Varying Flow Solver Convergence.

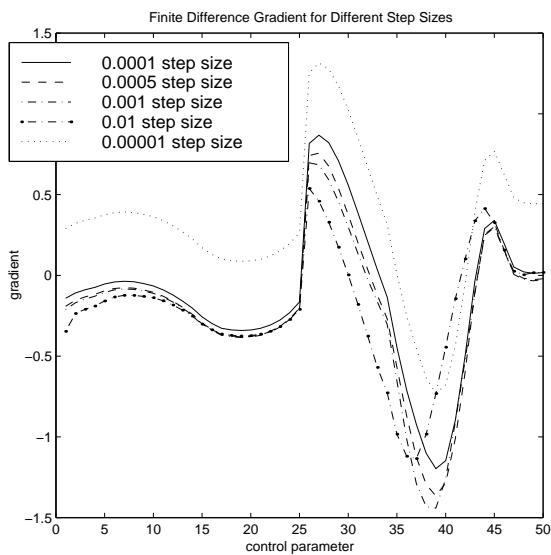


Figure 13: Navier-Stokes Inverse Design: Finite Difference Gradients for Varying Step Sizes.

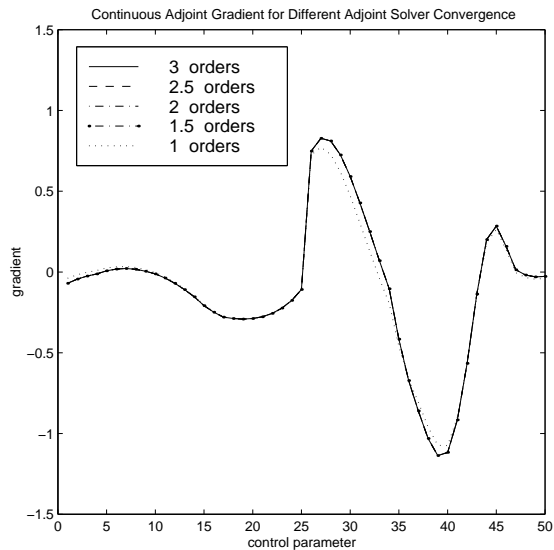


Figure 15: Navier-Stokes Inverse Design: Adjoint Gradients for Varying Adjoint Solver Convergence Levels.

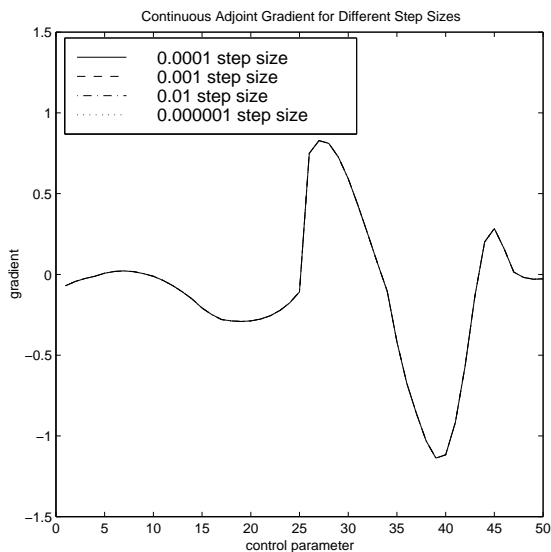


Figure 14: Navier-Stokes Inverse Design: Adjoint Gradients for Varying Step Sizes.

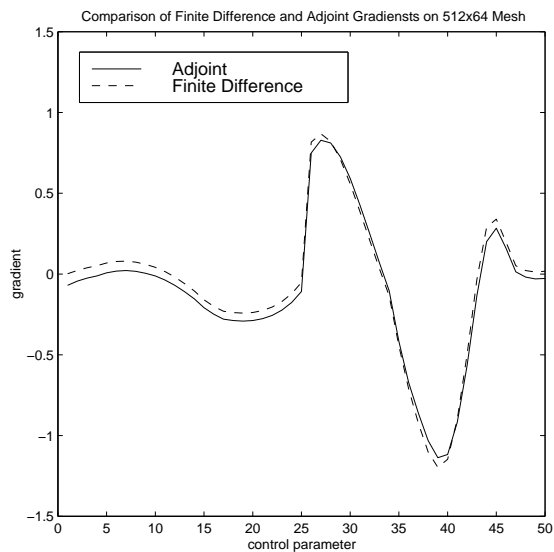


Figure 16: Navier-Stokes Inverse Design: Comparison of Finite Difference and Adjoint Gradients for Medium Mesh.




Figure 17: Navier-Stokes Inverse Design: Comparison of Finite Difference Gradients for Varying Mesh Sizes.

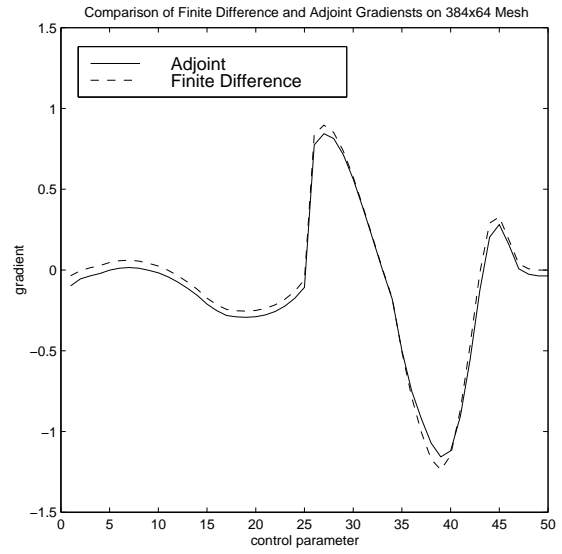


Figure 19: Navier-Stokes Inverse Design: Comparison of Finite Difference and Adjoint Gradients on Coarse Mesh.




Figure 18: Navier-Stokes Inverse Design: Comparison of Adjoint Gradients for Varying Mesh Sizes.






Figure 20: Navier-Stokes Inverse Design: Comparison of Finite Difference and Adjoint Gradients on Fine Mesh.




21a: Initial, $P_{error} = 0.0573$




21b: 100 Design Iterations, $P_{error} = 0.0056$

Figure 21: Typical Navier-Stokes Inverse Design Calculation, Korn airfoil to NACA 64A410



22a: Initial, $P_{error} = 0.0504$



22b: 100 Design Iterations, $P_{error} = 0.0043$

Figure 22: Typical Navier-Stokes Inverse Design Calculation, RAE 2822 airfoil to NACA 64A410

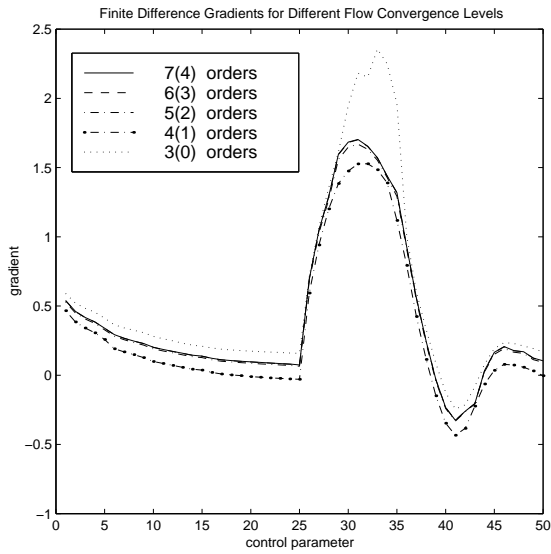


Figure 23: Navier-Stokes Drag Minimization: Finite Difference Gradients for Varying Flow Solver Convergence.

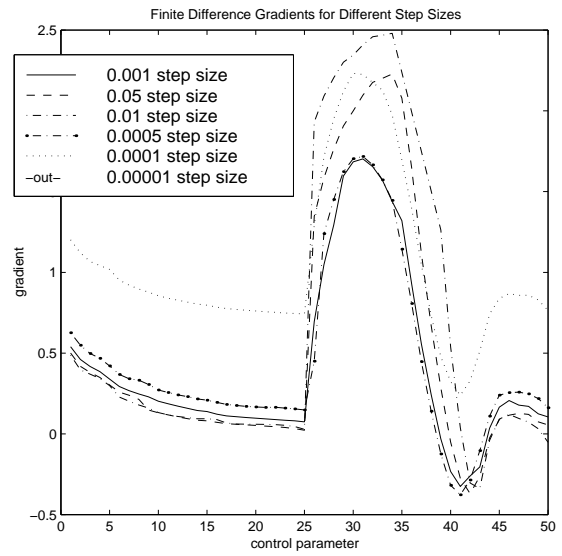


Figure 25: Navier-Stokes Drag Minimization: Finite Difference Gradients for Varying Step Sizes.

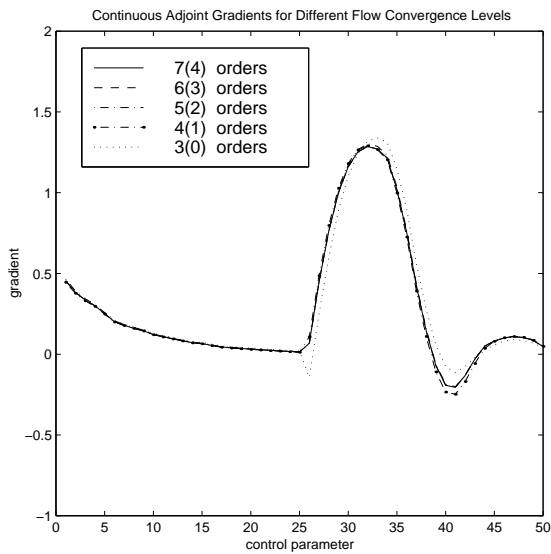


Figure 24: Navier-Stokes Drag Minimization: Adjoint Gradients for Varying Flow Solver Convergence.




Figure 26: Navier-Stokes Drag Minimization: Adjoint Gradients for Varying Step Sizes.

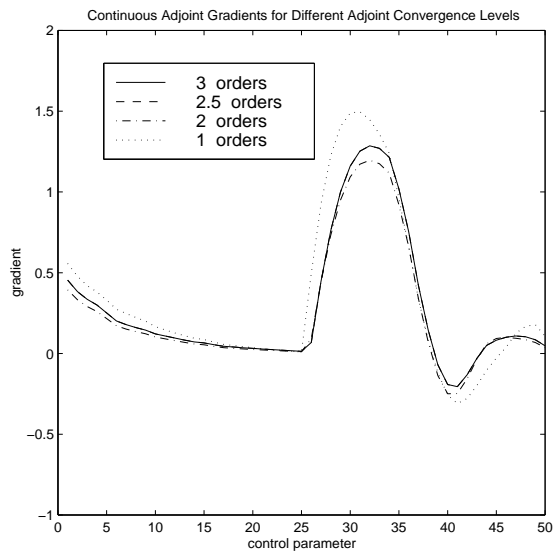


Figure 27: Navier-Stokes Drag Minimization: Adjoint Gradients for Varying Adjoint Solver Convergence Levels.

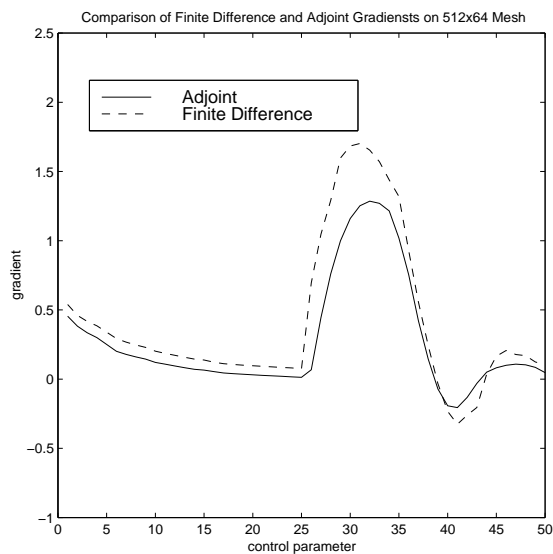




Figure 28: Navier-Stokes Drag Minimization: Comparison of Finite Difference and Adjoint Gradients for Medium Mesh.




29a: Initial, $CD_t=0.0168$
 $M = 0.73$, $\alpha = 2.756$, $CL_t = 0.8363$




29b: 14 Design Iteration, $CD_t=0.0108$
 $M = 0.73$, $\alpha = 2.665$, $CL_t = 0.8463$

Figure 29: Typical Navier-Stokes Drag Minimization Calculation, RAE 2822 Airfoil



30a: Initial, $CD_t=0.0168$
 $M = 0.73$, $\alpha = 2.756$, $CL_t = 0.8363$



30b: 17 Design Iterations, $CD_t=0.0096$
 $M = 0.73$, $\alpha = 2.565$, $CL_t = 0.8519$

Figure 30: Typical Navier-Stokes Drag Minimization Calculation, RAE 2822 Airfoil

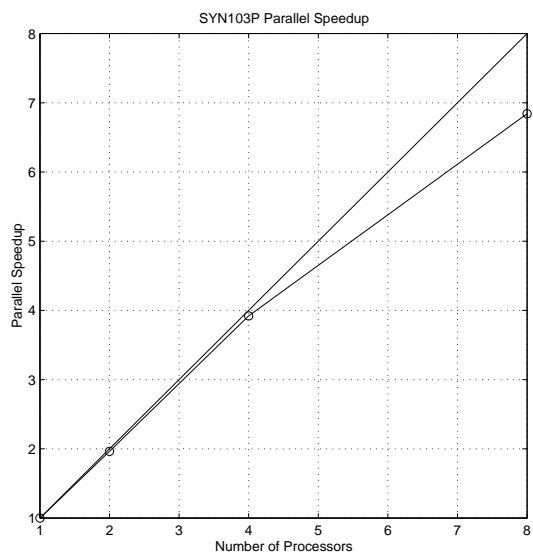


Figure 31: Parallel Speedup of SYN103P Viscous Design Code.