

Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 2

James J. Reuther*

NASA Ames Research Center, Moffett Field, California 94035

Antony Jameson† and Juan J. Alonso‡

Stanford University, Stanford, California 94305

and

Mark J. Rimlinger§ and David Saunders§

NASA Ames Research Center, Moffett Field, California 94035

This is the second part of a two-paper series that describes the derivation, development, and practical application of a shape optimization method for aerodynamic design problems involving complex aircraft configurations including multiple design points, subject to geometric constraints. The aerodynamic performance is evaluated using the Euler equations discretized on body-conforming multiblock meshes. The design process is greatly accelerated by the use of control theory for the calculation of sensitivity information and the implementation of the method on distributed memory parallel computers. Whereas the first paper of the series (Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 1," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 51–60) placed emphasis on the foundations of the control theory approach and presented a general framework for the design of complex configurations, this second paper focuses on the performance of the overall design method and its practical application to realistic design problems. More specifically, the accuracy of the adjoint-based sensitivity information is examined and compared with the sensitivities provided by the finite difference method. The parallel scalability of the design method is analyzed and demonstrated. Aerodynamic performance improvements are obtained for example wing design problems using both multipoint drag minimization and single-point inverse techniques. Design examples include both transonic and supersonic cases and involve wing aerodynamic design in the presence of complex aircraft configurations with all resulting configurations satisfying geometric constraints.

Introduction

THE development of aerodynamic shape optimization methods based upon computational fluid dynamics (CFD) has become an important topic of research in recent years. However, the true potential value of the technology resides in its direct applicability to industrial design problems. Unlike topics often examined by academia, these real-world design challenges are frequently ill-defined and fraught with constraints that do not lend themselves to elegant mathematical solutions. It is thus important that the overall flexibility and robustness of the proposed shape design method are given due consideration. The first part of this paper¹ discusses the convergence properties of the multiblock flow solver, followed by a similar discussion of the convergence of the adjoint solver. In this second part of the paper, a section is devoted to the speed-ups attainable by the parallel design algorithm. A section attending to an analysis of the accuracy achieved by adjoint-based sensitivities follows. Finally, three test cases are presented that explore the overall ability of the design method to treat real-world design challenges.

Analysis of Flow Solver Convergence

Fast and robust flow solver convergence is an essential ingredient of a successful automatic design method. If the flow solver cannot quickly produce accurate steady-state solutions, the design method becomes unusable because of the ensuing computational cost. Moreover, if the flow solver is unable to produce converged solutions robustly for a variety of design points (Mach numbers, angles of attack, and coefficients of lift) and design shapes, the design process is bound to fail because the accuracy in the calculation of the cost function will be compromised and the numerical optimization algorithm will be unable to produce meaningful improvements in the existing design.

An additional requirement for the success of an automatic design method is the ability of the flow solver to produce solutions at a fixed coefficient of lift C_L , robustly and efficiently. Problems involving aerodynamic shape design at cruise conditions often require that the computed lift remain fixed during the design process. To preserve a fixed value of C_L , it is commonplace to adjust the angle of attack of the configuration periodically during the convergence of the flow solution. The magnitude of the angle-of-attack adjustment is proportional to the difference between the target C_L and its current value. For this iterative algorithm to converge quickly, it is important to have a flow solver that rapidly develops a good approximation to the total lift of the configuration. In our method, this property is satisfied with Jameson's implementation of the multi-grid method.²

The convergence properties of the parallel multiblock flow solver are examined in this section. The geometry chosen for this analysis is a full-configuration business jet composed of

Received April 8, 1998; revision received June 15, 1998; accepted for publication June 22, 1998. Copyright © 1998 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

*Research Scientist, Research Institute for Advanced Computer Science, M/S 227-6. Member AIAA.

†T. V. Jones Professor of Engineering, Department of Aeronautics and Astronautics. Fellow AIAA.

‡Assistant Professor, Department of Aeronautics and Astronautics. Member AIAA.

§Research Scientist, Sterling Software, M/S 227-6.

wing, fuselage, pylon, nacelle, and empennage. The inviscid multiblock mesh around this configuration follows a general C-O topology with special blocking to capture the geometric details of the nacelles, pylons, and empennage. A total of 240 point-to-point matched blocks with 4,157,440 cells (including halos) are used to grid the complete configuration. This mesh allows the use of four multigrid levels obtained through recursive coarsening of the initial fine mesh. The upstream, downstream, upper, and lower far-field boundaries are located at an approximate distance of 15 wing semispans, whereas the far-field boundary beyond the wing tip is located at a distance approximately equal to 5 semispans. An engineering-accuracy solution (with a decrease of 4 orders of magnitude in the average density residual) can be obtained in 100 multigrid cycles. This kind of solution can be routinely accomplished in under 20 min of wall clock time using 32 processors of an SGI Origin2000 computer. Figure 1 shows a close-up view of the configuration in question; the block boundaries appear as dark lines.

The convergence properties of the flow solver are clearly demonstrated in Fig. 2. This figure shows the convergence history of a typical fixed angle-of-attack flow calculation. The average rms density residual decreases by 6 orders of magnitude in approximately 280 cycles. As discussed in Part 1¹ of this series, this fast convergence is achieved using a combi-



Fig. 1 Closeup view of typical business jet configuration, black lines show block boundaries.




Fig. 2 FLO87-MB convergence history of the flow solver for fixed angle of attack and fixed coefficient of lift calculations.

nation of multigrid and the damping properties provided by a modified Runge-Kutta time-stepping scheme. The figure also shows a typical convergence history for the case in which the angle of attack is periodically adjusted to obtain a specified coefficient of lift. At the points in the convergence history where the angle of attack is changed (every 20 multigrid iterations after iteration 60), a small increase in the average residual of the calculation is observed. Despite these periodic convergence setbacks, the residual of the calculation still decreases by six orders of magnitude in the same number of cycles as for the fixed angle-of-attack case. The convergence histories in Fig. 2 show that the current multiblock flow solver, FLO87-MB, meets the strict requirements necessary for a practically useful computational design method.

Analysis of Adjoint Solver Convergence

Because the design method presented here uses the solution of an adjoint equation to minimize the computational expense of calculating the sensitivities of the cost function, the numerical procedures used to discretize and solve the adjoint equation must also meet certain requirements of accuracy, convergence, and robustness.

The adjoint equation is a linear partial differential equation with constant coefficients that are derived from the values of the steady-state flow solution. Because of the difference in character between the flow and adjoint equations (nonlinear vs linear), different numerical methods could be used to solve each set efficiently. However, for simplicity of implementation, it is convenient to use the same numerical techniques for both the flow and adjoint solvers. In this work, a Runge-Kutta time-marching multigrid strategy is also used for the solution of the adjoint equation. This approach has been shown to provide satisfactory convergence rates in previous works by the authors.³⁻⁵

In the present study, the same mesh used for the flow solver convergence analysis is employed to examine the convergence properties of the parallel multiblock adjoint solver. Figure 3 shows a typical convergence history of the average residual in the first adjoint variable. As we can see, multigriding provides for rapid convergence of the solution. The robustness of the adjoint solution procedure is demonstrated by the fact that the magnitude of the average residual typically decreases by two orders in approximately 100 multigrid cycles. As will be shown later, to obtain accurate gradients through the adjoint-based approach, only a moderate level of convergence in the adjoint equations is required. Therefore, Fig. 3 shows that the current discretization of the adjoint equations fulfills the convergence requirements that the design method imposes on the




Fig. 3 Typical convergence history of the adjoint solver, ADJ87-MB.

numerical procedure. The issue of numerical accuracy will be treated later in this paper but has also been addressed in previous publications. The reader is encouraged to examine Refs. 6–8 for more details.

Parallel Performance and Scalability

Despite the fact that adjoint-based design methods drastically decrease the computational cost of aerodynamic shape optimization, the ability to perform useful design work still relies on the use of supercomputing resources. The need for these resources is a direct consequence of the very fine computational meshes required to obtain accurate flow and adjoint solutions when using high-fidelity mathematical models. Moreover, design calculations must have short enough turnaround times such that they can comfortably fit within an industrial design cycle.

A cost-effective option for this kind of design calculation is to use distributed computing platforms such as an integrated parallel computer or a network of workstations. For the large




Fig. 4 Scalability study for multiblock design method. Parallel speed-up of SYN87-MB.




Fig. 5 Finite difference gradients on medium mesh for flow solver convergence of 4, 5, 6, and 7 orders of magnitude in the density residual. Accuracy for different levels of flow solver convergence. Step size = 0.0001.

mesh sizes required in this research, high parallel efficiencies can be achieved for large numbers of processors using a domain decomposition strategy. The computational mesh is subdivided into a number of smaller subdomains that are each assigned to a different processor in the parallel computer. Because of the physical distribution of the data, the need for communication among processors arises. Nevertheless, provided that the amount of work performed by each processor is large compared with the time spent communicating between processors, the parallel performance of the method will be high.

The multiblock strategy adopted in this work, and discussed in detail in Part 1,¹ allows the independent update of the in-




Fig. 6 Finite difference gradients on medium mesh for step sizes = 0.001 and 0.0001. Seven orders of magnitude in the density residual. Accuracy for different step sizes.




Fig. 7 Continuous adjoint gradients on medium mesh for flow solver convergence of 3, 4, 5, 6, and 7 orders of magnitude in the density residual. Fixed adjoint solver convergence level: 2.5 orders of magnitude. Accuracy for different levels of flow solver convergence.

ternal cells of every block in the mesh by using *halo* or *ghost* cells. The vast majority of the communication cost is associated with the transfer of the flowfield values between blocks to fill these halos. The update procedure is repeated at several points in each flow or adjoint solver iteration using message passing.

Efficient parallel computation on a set of distributed processors is achieved by a combination of minimizing the overhead of communication between these processors and balancing the partitioned work load among them. The first obvious choice to improve parallel performance is then to minimize the amount of communication required between processors. This includes minimizing the number of messages sent and received

(latency) as well as the total amount of data to be transferred (bandwidth). Further improvements may be achieved by using other techniques such as latency hiding and scheduled communication, if either the problem or the specific architecture of the distributed platform allows it.⁹ The second approach to improve the parallel performance of the implementation is to use a load-balancing strategy that correctly accounts for the total amount of work performed by the processors.

For the explicit multistage multigrid time-stepping algorithm used in the present implementation, the computational work performed by a processor is proportional to the number of cells contained in its mesh partition. The time spent communicating with other processors depends on the exposed surface area of the mesh partition in a given processor. As the number of




Fig. 8 Continuous adjoint gradients on medium mesh for adjoint solver convergence of 1, 1.5, 2, 2.5, and 3 orders of magnitude. Fixed flow solver convergence level: seven orders of magnitude. Accuracy for different levels of adjoint solver convergence.




Fig. 10 Comparison of adjoint and finite difference gradients on coarse mesh. Adjoint gradient convergence: flow solver/adjoint solver = 6.2/2.5 orders of magnitude. Finite difference gradients: flow solver convergence/step size = 6.2 orders/0.0001.




Fig. 9 Comparison of adjoint and finite difference gradients on medium mesh. Adjoint gradient convergence: flow solver/adjoint solver = 7/3.5 orders of magnitude. Finite difference gradients: flow solver convergence/step size = 7 orders/0.0001.




Fig. 11 Comparison of adjoint and finite difference gradients on fine mesh. Adjoint gradient convergence flow: solver/adjoint solver = 6/3.5 orders of magnitude. Finite difference gradients: flow solver convergence/step size = 6 orders/0.0001.

processors in the calculation grows, the individual mesh partitions get progressively smaller. The ratio of surface area to volume increases, and the parallel performance decreases. A load-balancing scheme that accounts for not only the computational load, but also the communication time will produce parallel algorithms with better scalability. Both avenues of improvement were explored in the parallel implementation of our design method. Details can be found in Ref. 9.

Figure 4 shows the result of a scalability study performed using the complete configuration mesh described earlier. Because the current parallel implementation is such that the convergence history of the flow and adjoint solvers is identical regardless of the number of processors used, the results presented translate directly into wall clock time improvements. The figure shows the excellent scalability properties of the flow solver. For 32 processors, the achieved speed-up is 27.7. It must be mentioned that, for this particular mesh that contains blocks of varying sizes, it was impossible to distribute the computational load perfectly among 32 processors. The highest speed-up that could be hoped for was 29.5 in this case (if the communication time is assumed to be nonexistent).

Accuracy and Efficiency of Adjoint Sensitivities

The most critical element in determining the validity of adjoint-based sensitivities is the establishment of their accuracy level. As discussed in Part 1¹ of this series, two main alternatives for the development of the adjoint equations exist: the continuous and the discrete approaches. The discrete approach, if rigorously implemented, will produce gradients that exactly match those obtained by using accurate finite differences. Thus, comparing such gradients serves as a consistency check for the method. In the case of the continuous approach that we use throughout this work, gradient comparisons have a different interpretation. The adjoint equations derived in the discrete approach may be thought of as one of the possible discretizations of the continuous adjoint equations. Therefore, unless this particular discretization of the continuous adjoint equations is used, the gradients produced using the continuous approach will differ from those obtained by finite differencing in proportion to the truncation error of the approximation. For example, if the schemes used to discretize the flow and adjoint equations are both second-order accurate, the differences between the discrete and continuous adjoint solutions should also

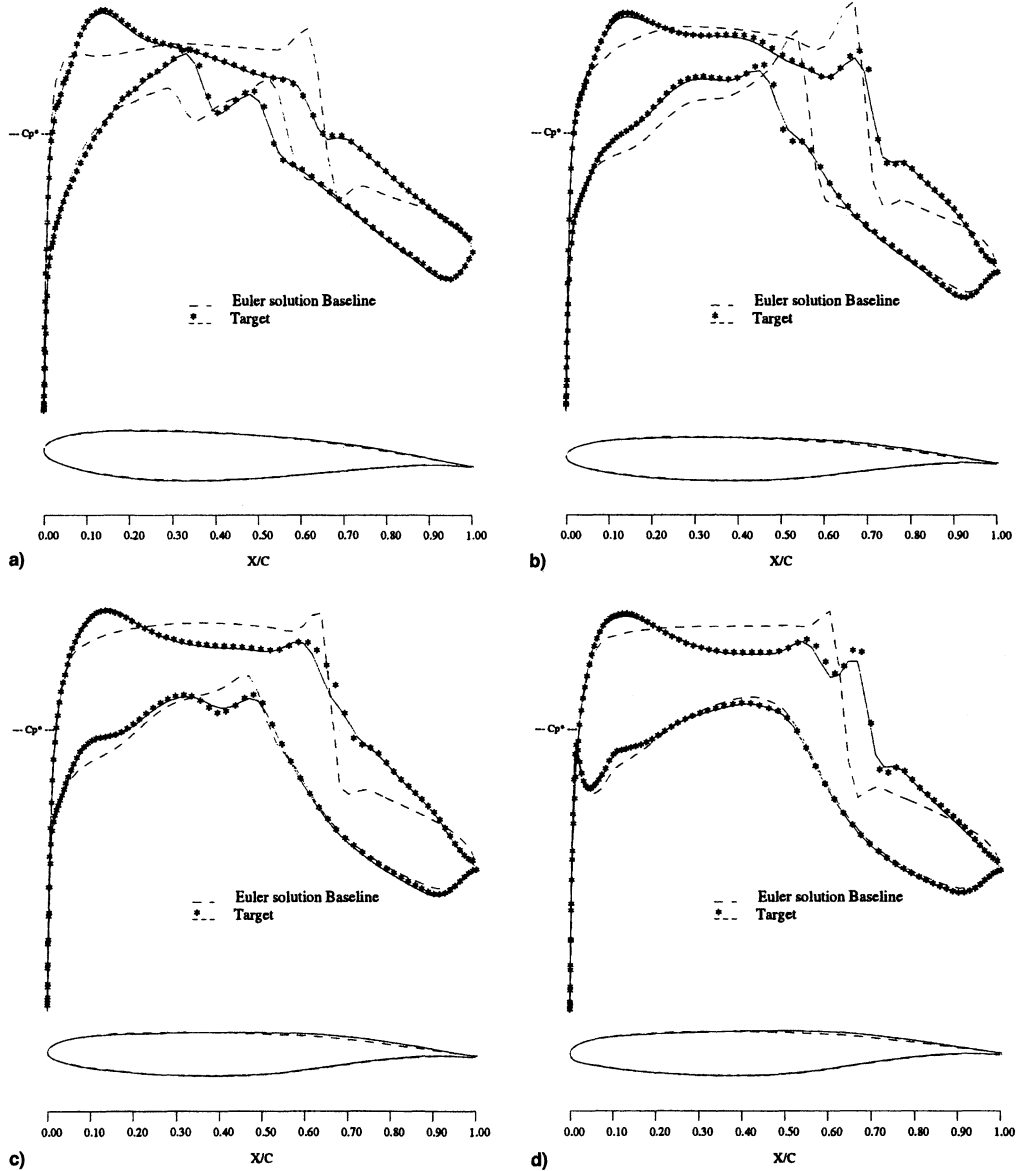


Fig. 12 Business jet configuration. Inverse design wing design to target pressure distribution. $M = 0.82$, $C_L = 0.30$, 108 Hicks-Henne variables and no active constraints. Span station $z =$ a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, initial pressures; *---, targets; and —, pressures after 18 design cycles.

be of second order. Consequently, the differences between the two formulations should vanish in the limit of zero mesh width. This conclusion, which has been demonstrated in the recent work of Anderson and Venkatakrishnan,¹⁰ can be helpful in determining the proper mesh resolution required for accuracy in the adjoint solution; if a properly implemented continuous adjoint formulation fails to provide sensitivity information that closely matches the result of finite differencing, the probable cause is that the cost function being evaluated is poorly resolved on the mesh in question.

In this section, the gradients resulting from our adjoint implementation are compared with those obtained from finite differencing. Finite difference gradients are subject to two classes of error: numerical cancellation error (error caused by inaccuracies in cost function because of the lack of solver convergence and machine precision) and gradient discretization (truncation) error (error caused by the discrete finite difference approximation used to estimate the gradient). To address these issues, a short study is performed to establish the level of convergence and discrete step size necessary to produce accurate

gradients. These finite difference gradients are then compared with those obtained by using the current adjoint implementation.

To perform this gradient check, an inverse wing design test problem is defined. Although only the pressures on the surface of the wing are driven toward a target, the design is carried out in the presence of a fuselage and nacelle. The same gradient calculation is carried out on a sequence of coarse, medium, and fine meshes containing 157,676, 734,976, and 4,280,320 cells, respectively. These cell counts include all halo cells. The medium and coarse meshes are obtained by recursively taking alternate points in the original fine mesh in each coordinate direction. The medium mesh is used to study the effect of flow solver convergence and step size, whereas all three meshes are used to determine the mesh sensitivity of the gradients. Figures 5 and 6 show the values of gradient components obtained using one-sided differencing for different values of flow solver convergence and step size. Only the first 30 out of a total of 108 Hicks–Henne design variables are compared because of the prohibitive computational cost of calculating all gradient components using finite differences. It can

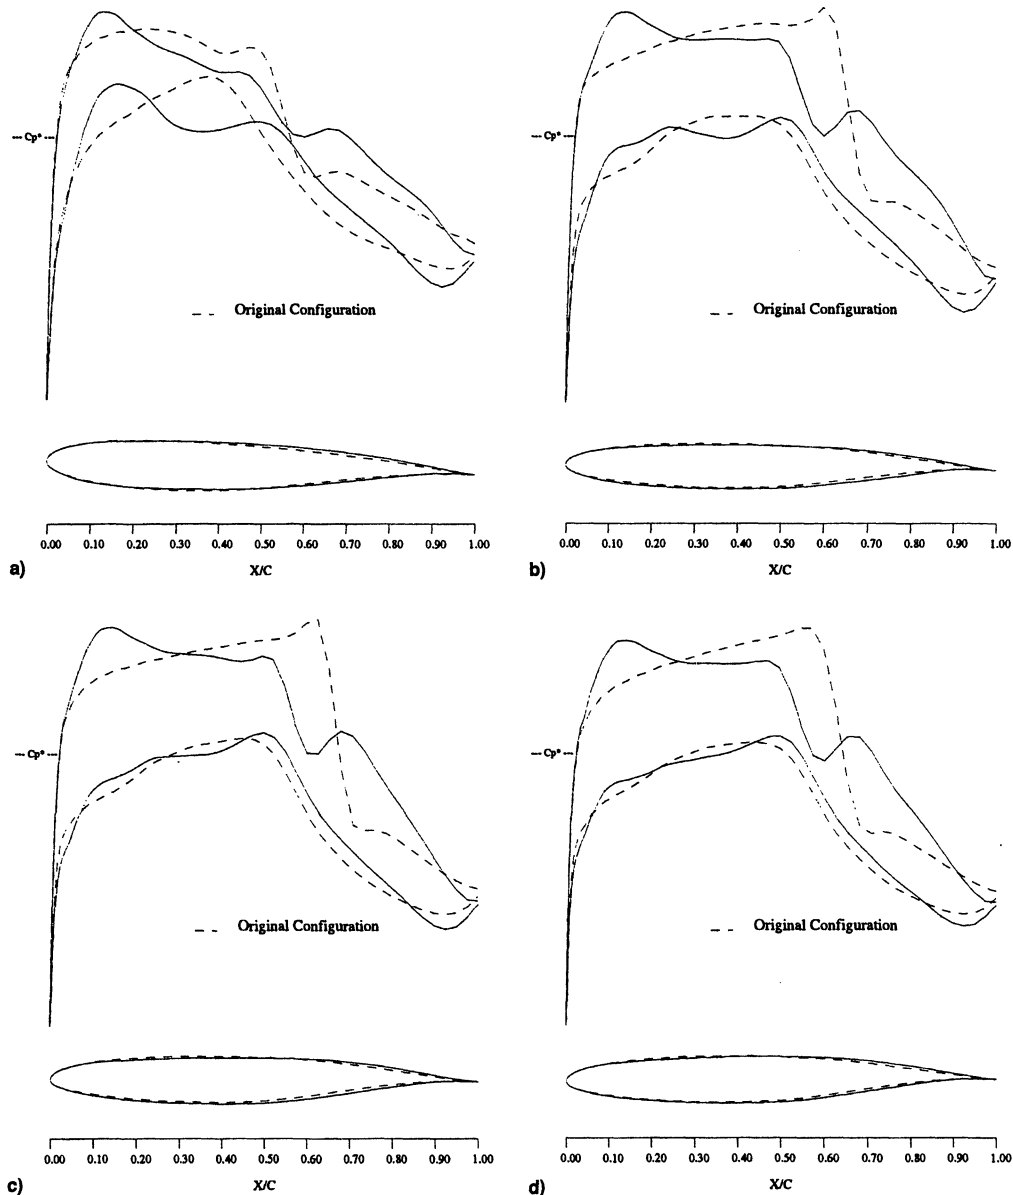


Fig. 13 Business jet configuration with the multipoint drag minimization at fixed lift. Design point 1, $M = 0.81$, $C_L = 0.35$, 90 Hicks–Henne variables, and spar constraints active. Span station $z =$ a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, initial pressures; and —, pressures after five design cycles.

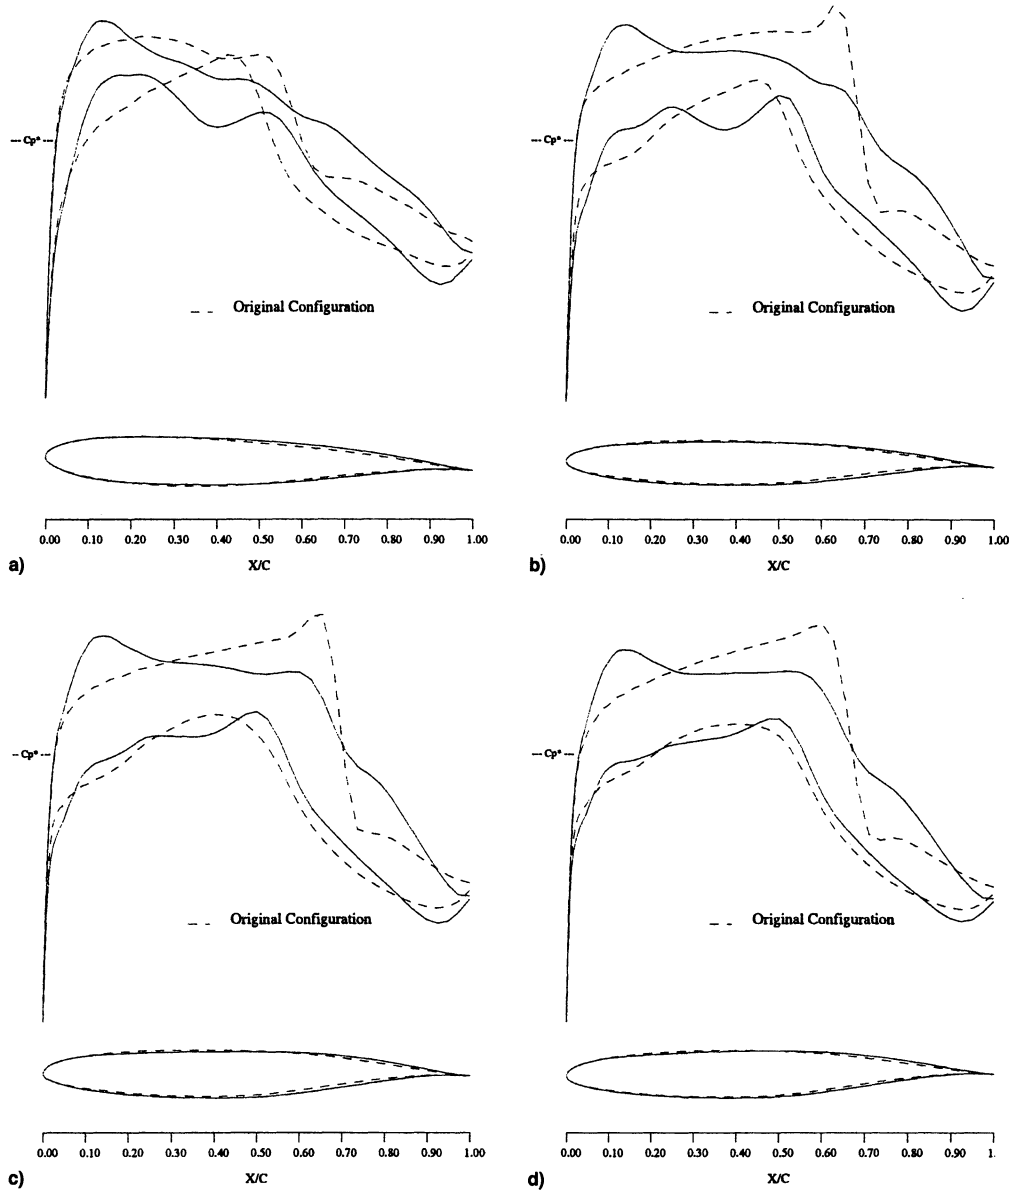


Fig. 14 Business jet configuration. Multipoint drag minimization at fixed lift. Design point 2, $M = 0.82$, $C_L = 0.30$, 90 Hicks-Henne variables, and spar constraints active. Span station $z =$ a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, initial pressures; and —, pressures after five design cycles.

be seen from Fig. 5 that, for a step size of 0.0001, a flow solver convergence level of at least 6 orders of magnitude is necessary to achieve accurate gradients. For a flow solver convergence of 7 orders of magnitude in the average residual, Fig. 6 indicates that the discretization error is acceptably small for step sizes smaller than 0.001. From the results in these figures, one can be certain that if a convergence level of 7 orders of magnitude and step sizes of 0.0001 are used for the calculation of finite difference gradients, the sensitivities will be accurate and provide a good reference for comparison with sensitivities obtained using the continuous adjoint method.

Adjoint-based gradients must also be resolved properly to ensure realistic comparisons. In contrast with using the finite difference method, gradients obtained via the adjoint method do not contain a gradient discretization error because they are independent of step size. Instead, they have sources of numerical cancellation error resulting from inadequate convergence of the flow and/or adjoint solvers. Thus, by examining the levels of convergence in both systems, it is possible to obtain adjoint sensitivities that are accurate given the choice of discretization for the overall adjoint system and the choice

of mesh resolution. Figure 7 shows the dependence of the continuous adjoint gradients on the flow solver convergence, whereas Fig. 8 shows their dependence on the level of convergence of the adjoint solver. From these figures we can conclude that accurate gradient information can be obtained using the continuous adjoint formulation provided that the flow and adjoint solvers are converged by at least four and two orders of magnitude, respectively. It is thus apparent that the level of convergence required to obtain accurate gradient information with the continuous adjoint formulation is drastically lower than that necessary in the finite difference method. This relaxation of the convergence requirement makes the use of an adjoint formulation even more attractive. Reference 6 provides an explanation of the mechanisms by which accurate gradient information can be obtained with moderate convergence in both the flow and adjoint systems.

Figure 9 compares adjoint and finite difference sensitivities for the medium mesh case. The flow solver was converged 7 orders of magnitude in both approaches. The costate system was converged 3 orders of magnitude, and a constant step size of 0.0001 was chosen for the calculation of the finite difference

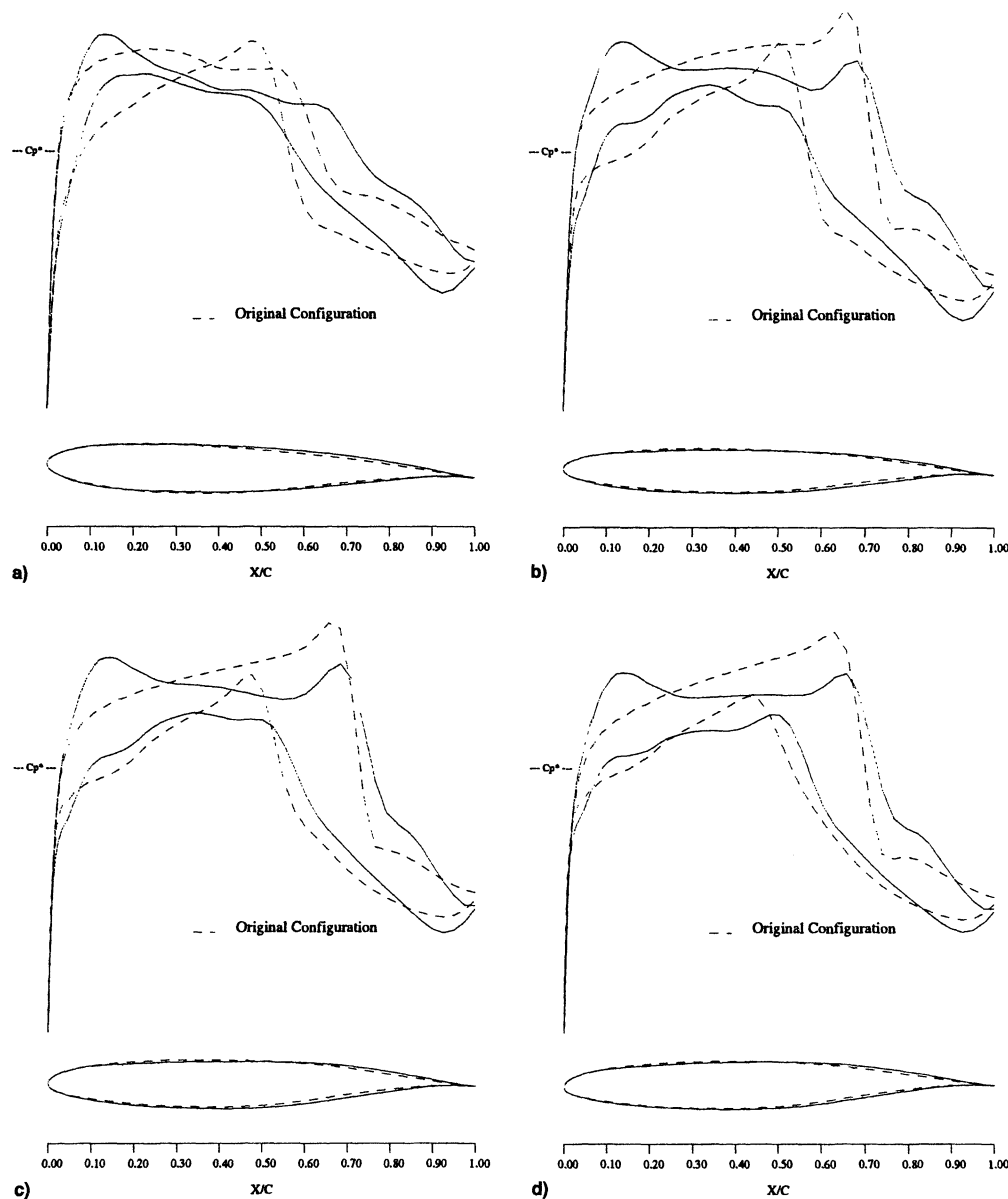


Fig. 15 Business jet configuration. Multipoint drag minimization at fixed lift. Design point 3, $M = 0.83$, $C_L = 0.25$, 90 Hicks-Henne variables, and spar constraints active. Span station $z =$ a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, initial pressures; and —, pressures after five design cycles.

gradients. Because these choices of convergence tolerances and step sizes should eliminate all sources of numerical cancellation and gradient discretization errors from the comparisons, the explanation of the small discrepancies evident in Fig. 9 must be a consequence of the difference in discretization between the adjoint system used in this work and the discretization that would have been obtained via a discrete adjoint implementation. It should now be possible to show that the differences between the adjoint and finite difference sensitivities will diminish as the mesh is refined. Figures 10 and 11 show comparisons of the same gradients for both the coarse and fine meshes with convergence tolerances adequately adjusted to provide low numerical errors on each mesh. It is apparent from these comparisons that, while the general character of the gradients remains the same in all three meshes, there is a tendency for the relative magnitudes of the individual gradient components to increase as the mesh is refined. This suggests that the cost function of the design problem may not be properly resolved on any of the meshes explored here. However, the figures do show that the discrepancy between the adjoint and finite difference gradients decreases as the

mesh is refined. The \mathcal{L}^2 norm of the difference between the most accurate finite difference gradients and the sensitivities obtained using the adjoint formulation on each of the three meshes is computed to be 0.010883, 0.009506, and 0.002684 for the coarse, medium, and fine meshes, respectively. The reduction in the norm of the error as the mesh is refined supports the assertion that the continuous adjoint method is consistent with other techniques of obtaining sensitivities in the limit of infinite mesh density.

Design Test Cases

Beyond the discussions of convergence, parallel scalability, and gradient accuracy, the true test of a design method is its success in treating problems of interest. In this section, three illustrative design examples are presented to demonstrate the versatility and capabilities of our aerodynamic shape optimization method.

Wing Inverse Design with a Complete Aircraft Configuration

In the first test case, inverse design is performed on the wing of a business jet in transonic flow. The problem is essentially

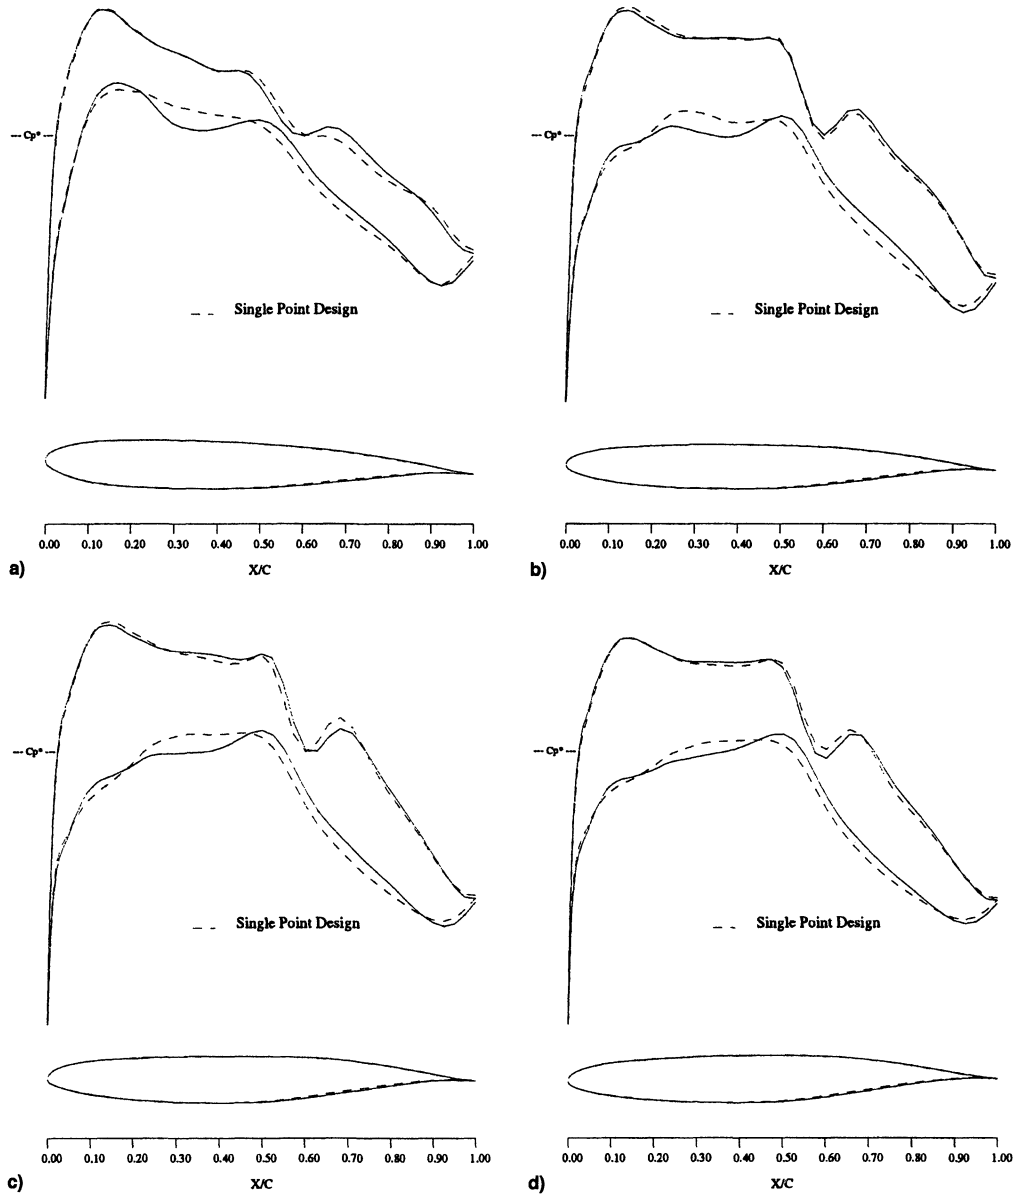


Fig. 16 Business jet configuration. Single point vs multipoint drag minimization at fixed lift. Design point 1, $M = 0.81$, $C_L = 0.35$, 90 Hicks–Henne variables, and spar constraints active. Span station $z =$ a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, single-point design pressures; and —, multipoint design pressures.

the same as that outlined earlier for examining the gradient accuracies. Again, wing target pressure distributions are specified and design variables are chosen to modify the wing. The main difference is that, to ensure the accuracy of the resulting design, the very fine complete configuration mesh used earlier to demonstrate the flow solver convergence properties replaces the wing/body/nacelle mesh employed in the gradient studies. For demonstration purposes, the target pressure distributions are taken from a wing geometry that has shown to have lower drag at $M = 0.82$ and $C_L = 0.30$.¹¹ This choice of the target, which would not normally be available in a production-design environment, serves as an excellent test case because it ensures that the design is realizable and that it will result in an improvement in performance. The initial wing was designed for $M = 0.80$ and $C_L = 0.30$ and, thus, experiences a significant drag rise at Mach = 0.82.

The design process is driven by the modification of 108 Hicks–Henne design variables chosen to perturb an underlying wing geometry entity. The wing geometry is defined by 14 airfoil sections, with 18 design variables applied to six of these sections. Linear lofting between the sections where the design

variables are applied governs the changes at the remaining sections. A fuselage geometry entity is also present but serves only in the calculation of new wing–body intersection locations as the design algorithm proceeds. To permit the design to converge fully on the realizable target, no constraints are imposed on the design process.

Figure 12 shows the initial, final, and target airfoil geometries and their associated coefficient of pressure C_p distributions. It is noted that in a true design environment the target shape would not be known as it is here. The design process was run for 18 design cycles with the resulting design almost exactly matching the target. Further improvements were possible, but did not warrant the added computational costs. Indeed 95% of the reduction in the cost function [$I = \frac{1}{2} \int (p - p_d)^2 ds$] seen in Fig. 12 was obtained in the first nine design cycles.

Transonic Multipoint Constrained Aircraft Design

As a second test case of the aerodynamic shape optimization method, the transonic business jet configuration is re-examined in the context of drag minimization. The medium mesh seen

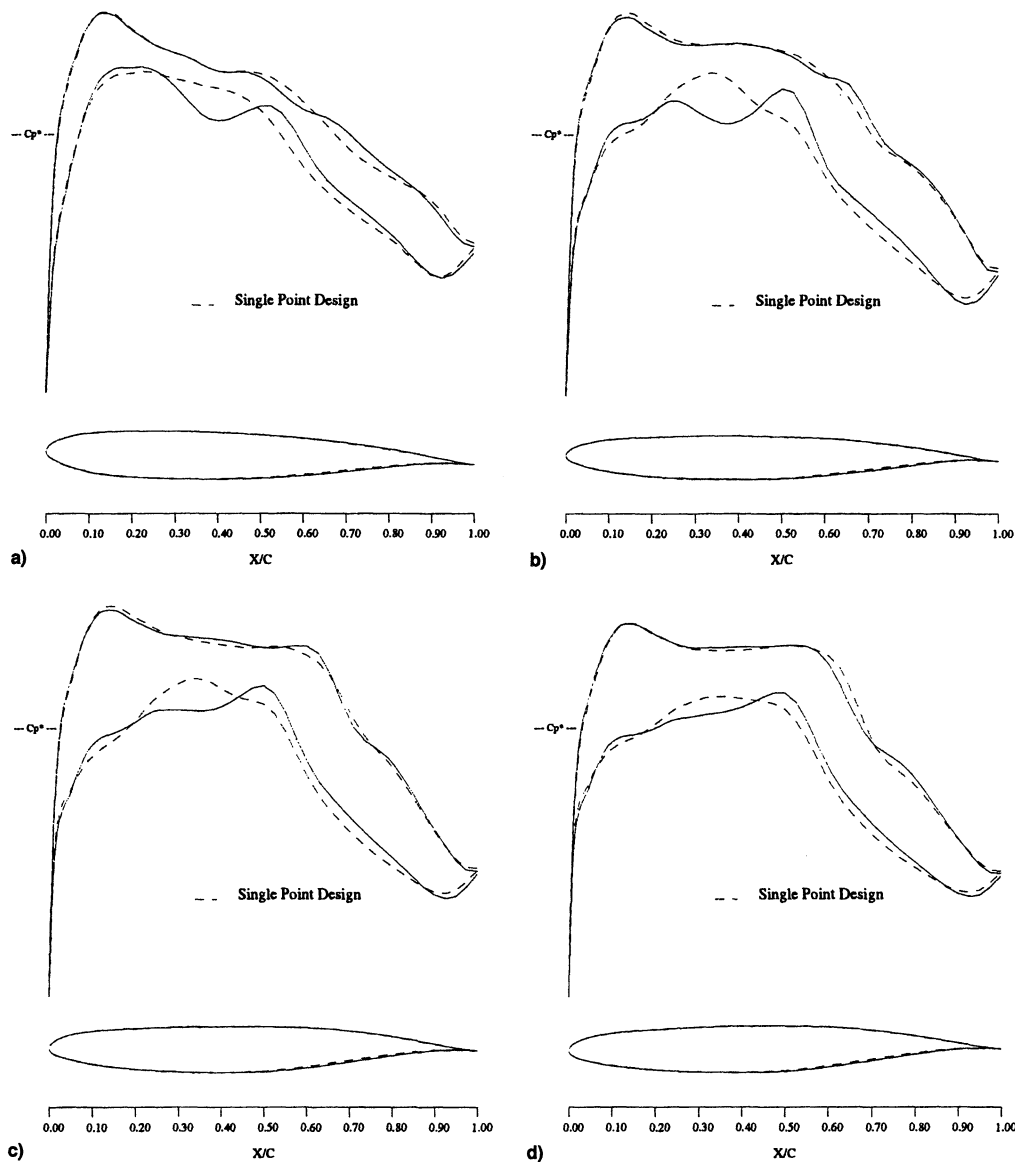


Fig. 17 Business jet configuration. Single point vs multipoint drag minimization at fixed lift. Design point 2, $M = 0.82$, $C_L = 0.30$, 90 Hicks-Henne variables, and spar constraints active. Span station z = a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, single-point design pressure; —, multipoint-design pressures.

in the gradient accuracy tests is used for this design example. The objective of the design is to alter the geometry of the wing to minimize the configuration inviscid drag at three different flight conditions simultaneously. Realistic geometric spar thickness constraints are enforced.

The initial configuration was designed for $M = 0.8$ and $C_L = 0.3$. The three operating points chosen for this design are $M = 0.81$ with $C_L = 0.35$, $M = 0.82$ with $C_L = 0.30$, and $M = 0.83$ with $C_L = 0.25$. For each of the design points, both Mach number and lift coefficient are held fixed. To demonstrate the advantage of a multipoint design approach, the final solution at the middle design point will be compared with a single point design at the same conditions. The underlying geometry entities chosen for this example consist of separate wing and fuselage components. As the geometry of the wing is modified, the design algorithm computes new wing-fuselage intersections. The wing component is made up of six airfoil defining sections. Eighteen Hicks-Henne design variables are applied to five of these sections for a total of 90 design variables. The sixth section at the symmetry plane is not modified. Spar thickness constraints were also enforced on each defining station at the $x/c = 0.2$ and $= 0.8$ locations. Maximum thickness was

forced to be preserved at $x/c = 0.4$ for all six defining sections. To ensure an adequate included angle at the trailing edge, each section was also constrained to preserve thickness at $x/c = 0.95$. Finally, to preserve leading-edge bluntness, the upper surface of each section was forced to maintain its height above the camber line at $x/c = 0.02$. Combined, a total of 30 linear geometric constraints were imposed on the configuration.

Figures 13–15 show the initial and final airfoil geometries and C_p distributions after five NPSOL design iterations. It is evident that the new design has significantly reduced the shock strengths on both upper and lower wing surfaces at all design points. The transitions between design points are also quite smooth. For comparison purposes, a single-point drag minimization study ($M = 0.81$ and $C_L = 0.25$) is carried out starting from the same initial configuration and using the same design variables and geometric constraints.

Figures 16–18 show comparisons of the solutions from the three-point design with those of the single-point design. Interestingly, the upper surface shapes for both final designs are very similar. However, in the case of the single-point design, a strong lower surface shock appears at the $M = 0.83$, $C_L = 0.25$ design point. The three-point design is able to suppress

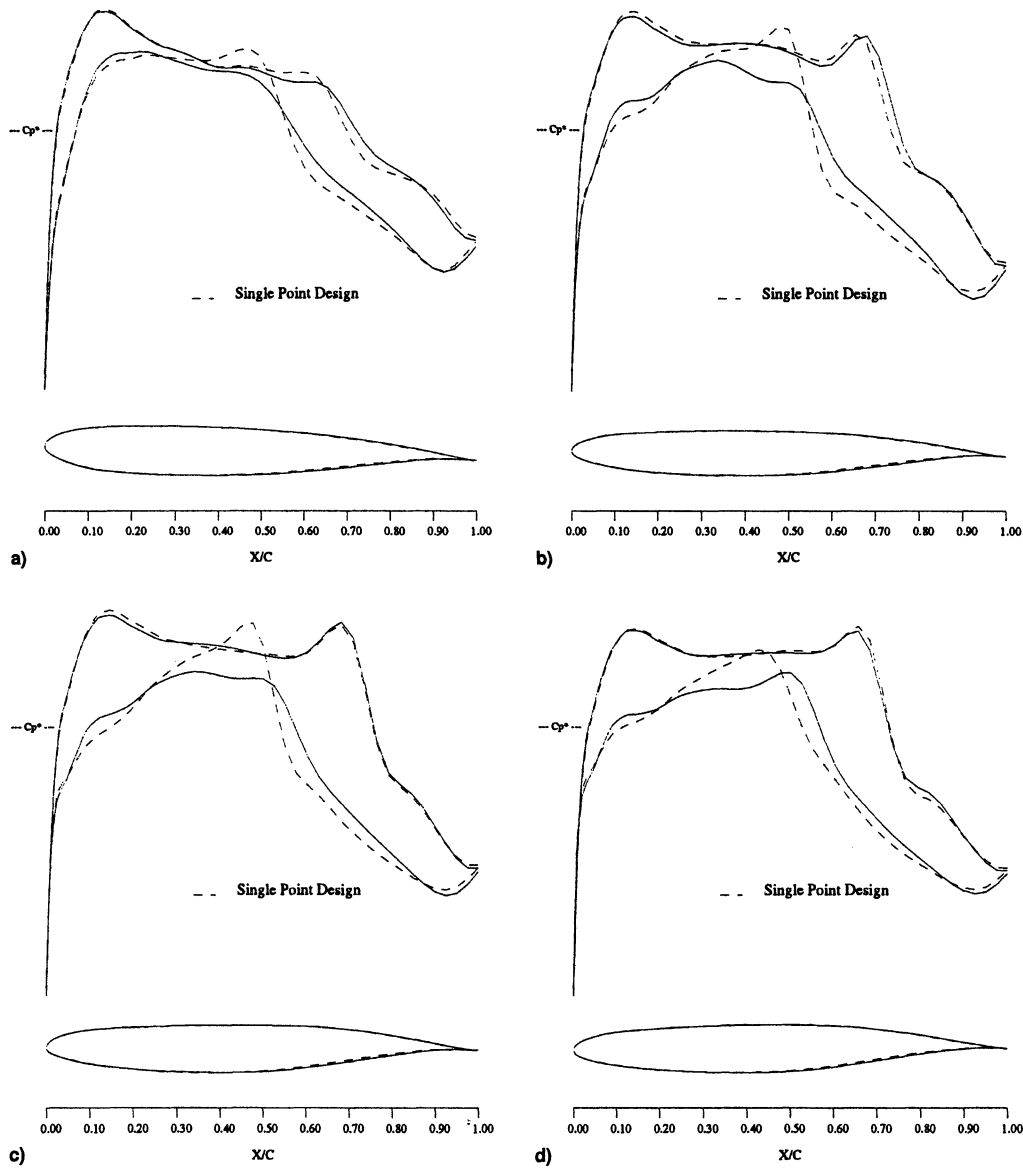


Fig. 18 Business jet configuration. Single point vs multipoint drag minimization at fixed lift. Design point 3, $M = 0.83$, $C_L = 0.25$, 90 Hicks-Henne variables, and spar constraints active. Span station $z =$ a) 0.190, b) 0.475, c) 0.665, and d) 0.856. ---, single-point design pressures; and —, multipoint design pressure.

Table 1 Drag reduction of single and multipoint designs

Mach	C_L	Relative C_D		
		Initial	Single-point design	Three-point design
0.81	0.35	1.00257	0.85003	0.85413
0.82	0.30	1.00000	0.77350	0.77915
0.83	0.25	1.08731	0.81407	0.76836

the formation of this lower surface shock and achieves a nine-count drag benefit over the single-point design at this condition. However, it has a one-count penalty at the single-point design condition. The three-point design features a weak single shock for one of the three design points and a very weak double shock at another design point. Table 1 summarizes the drag results for the two designs. The C_D values have been normalized by the drag of the initial configuration at the second design point.

Before proceeding to the next section, it should be noted that these business jet design examples are only representative of the potential for automated design, and are not intended to

provide designs for actual construction. Firstly, in each case, only five NPSOL design steps were taken when a larger number could have yielded designs with further improvements. More importantly, for the case of transonic designs, the inclusion of viscous effects may prove to have an important impact on the optimized shape. Preliminary use of the Navier-Stokes equations has already started for transonic design studies, in place of the Euler equations solved here.

The results of the multipoint design example illustrate a well-established fact: single-point designs typically achieve high-cost function improvements at the expense of off-design performance. Multipoint designs achieve more moderate improvements that are spread over a larger range of operating conditions. It is this latter attribute that is more desirable in real-life configurations.

Supersonic Constrained Aircraft Design

For supersonic design, provided that turbulent flow is assumed over the entire configuration, the inviscid Euler equations suffice for aerodynamic design because the pressure drag is not greatly affected by the inclusion of viscous effects. Moreover, flat-plate skin friction estimates of viscous drag are




Fig. 19 Supersonic transport configuration, Iso- C_p contours on upper and lower surfaces, baseline and optimized designs, and $M = 2.2$, $C_L = 0.105$: a) lower and b) upper surface.

often very good approximations. In this study, the generic supersonic transport configuration used in Ref. 12 is revisited.

The baseline supersonic transport configuration was sized to accommodate 300 passengers with a gross takeoff weight of 750,000 lb. The supersonic cruise point is Mach 2.2 with a C_L of 0.105. Figure 19 shows that the planform is a cranked-delta configuration with a break in the leading-edge sweep. The inboard leading-edge sweep is 68.5 deg, whereas the outboard is 49.5 deg. Because the Mach angle at $M = 2.2$ is 63 deg, it is clear that some leading-edge bluntness may be used inboard without a significant wave drag penalty. Blunt leading-edge airfoils were created with thickness ranging from 4% at the root to 2.5% at the leading-edge break point. These symmetric airfoils were chosen to accommodate thick spars at roughly the 5 and 80% chord locations over the span up to the leading-edge break. Outboard of the leading-edge break, where the wing sweep is ahead of the Mach cone, a sharp leading edge was used to avoid unnecessary wave drag. The airfoils were chosen to be symmetric, biconvex shapes modified to have a region of constant thickness over the midchord. The four-engine configuration features axisymmetric nacelles tucked close to the wing lower surface. This layout favors reduced wave drag by minimizing the exposed boundary-layer diverter area. However, in practice, it may be problematic because of the channel flows occurring in the juncture region of the diverter, wing, and nacelle at the wing trailing edge. The leading-edge heights of the diverters are determined by the local boundary-layer displacement thickness such that entrainment of boundary-layer flow into the engines is avoided. Because the distances from the wing leading edge to the diverter leading edge are different for the two nacelles, this results in a corresponding diverter height difference.

The computational mesh on which the design is run has 180 blocks and 1,500,000 mesh cells (including halos), whereas the underlying geometry entities define the wing with 16 sectional cuts and the body with 200 sectional cuts. In this case, where we hope to optimize the shape of the wing, care must be taken to ensure that the nacelles remain properly attached with diverter heights being maintained. To accomplish this without the inclusion of additional geometry entities, the portions of the nacelles and diverters that are actually below the wing planform outline take their associated surface mesh point motion from their projected locations on the lower surface of the wing.

The objective of the design is to reduce the total drag of the configuration at a single design point ($M = 2.2$, $C_L = 0.105$) by modifying the wing shape. Just as in the transonic cases, 18 design variables of the Hicks-Henne type are chosen for each wing defining section. Similarly, instead of applying them to all 16 sections, they are applied to eight of the sections and then lofted linearly to the neighboring sections. Spar thickness constraints are imposed for all wing defining sections at $x/c = 0.05$ and 0.8. An additional maximum thickness constraint is specified along the span at $x/c = 0.5$. A final thickness constraint is enforced at $x/c = 0.95$ to ensure a reasonable trailing-edge included angle. An Iso- C_p representation of the initial and final designs is depicted in Fig. 19 for both the upper and lower surfaces.

It is noted that the strong oblique shock evident near the leading edge of the upper surface on the initial configuration is largely eliminated in the final design after five NPSOL design iterations. Also, it is seen that the upper surface pressure distribution in the vicinity of the nacelles has formed an unexpected pattern. However, recalling that thickness constraints abound in this design, these upper surface pressure patterns

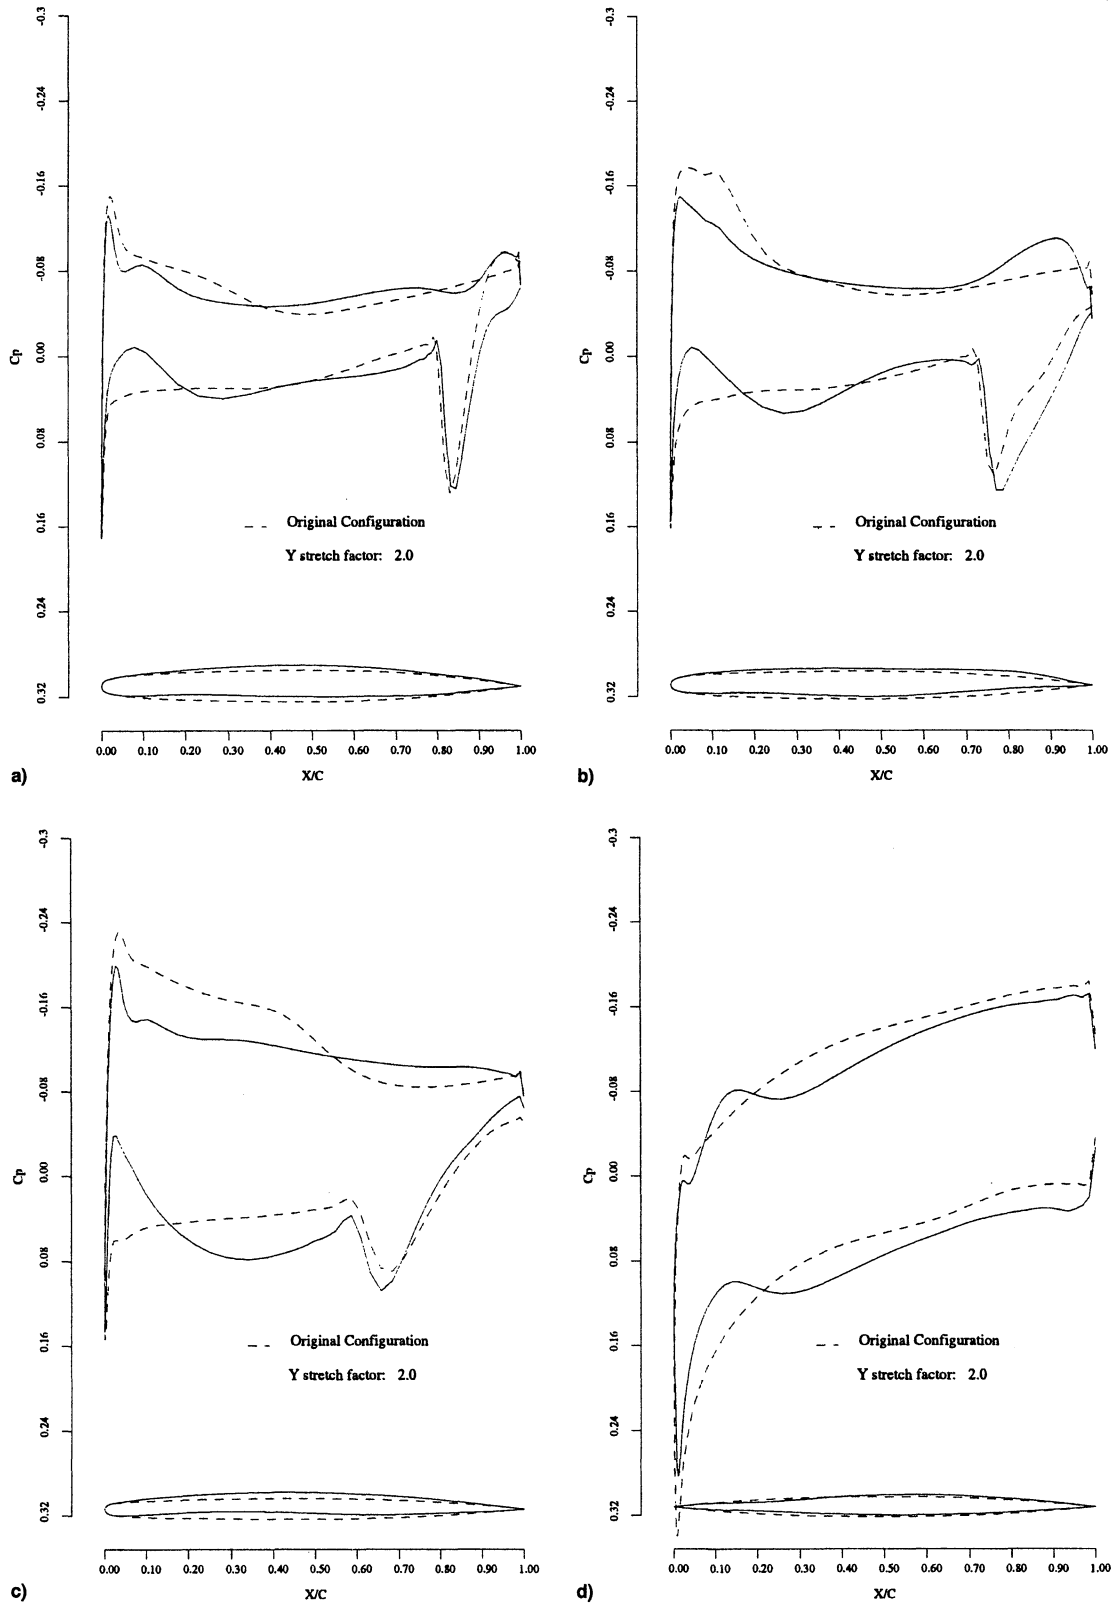


Fig. 20 Supersonic transport configuration. Drag minimization at fixed lift. $M = 2.20$, $C_L = 0.105$, 144 Hicks-Henne variables, and spar constraints active. Span station $z =$ a) 0.194, b) 0.387, c) 0.581, and d) 0.775. ---, initial pressures; and —, pressures after five design cycles.

are assumed to be the result of sculpting of the lower surface near the nacelles that affects the upper surface shape via the thickness constraints. For the lower surface, the leading edge has developed a suction region while the shocks and expansions around the nacelles have been somewhat reduced. Figure 20 shows the pressure coefficients and (scaled) airfoil sections for four sectional cuts along the wing. These cuts further dem-

onstrate the removal of the oblique shock on the upper surface and the addition of a suction region on the leading edge of the lower surface. The airfoil sections have been scaled by a factor of 2 so that shape changes may be seen more easily. Most notably, the section at 38.7% span has had the lower surface drastically modified such that a large region of the aft airfoil has a forward-facing portion near where the pressure spike

from the nacelle shock impinges on the surface. The final overall pressure drag was reduced by 8%, from $C_D = 0.0088$ to 0.0081.

Conclusions

In this paper, we have shown how the detailed wing design of both transonic and supersonic aircraft configurations including airframe/nacelle integration effects can be accomplished in a routine fashion for multiple design points and with the inclusion of geometric constraints. While the results presented in this paper have been restricted to the inviscid Euler equations, this limitation will soon be overcome in our future work. The focus here was to demonstrate the feasibility of realistic designs governed by simple geometry entities that are assembled and treated during the design process. Furthermore, with the proven coupling to NPSOL and the addition of multiple design-point capability, the door is truly open for the method to act as a crucial element of a high-fidelity multidisciplinary optimization technique capable of revolutionizing aircraft design.

In the period since the current approach to optimal shape design was first proposed, the method has been verified by numerical implementation for both potential flow and flows governed by the Euler equations. The extension to treat the Navier–Stokes equations has also been carried out and is currently under study. Furthermore, results have been presented for three-dimensional calculations using both analytic mappings and general finite volume implementations for the calculation of the gradients. Recently, the technique has been adopted by some industry participants to perform the aerodynamic design of future configurations. With the parallel implementation of the multiblock design algorithm now in hand, the technology has advanced to the degree that aerodynamic shape design of complete aircraft configurations with very rapid turnaround is possible.

All analysis and design cases in this study were performed on parallel architectures using <1 day of computer time. This demonstrates that complete configuration designs may be achieved with rapid turnaround even with the most conservative estimates of available computational resources. In future efforts, additional disciplines will be coupled into the techniques presented here while work continues on the more systematic treatment of complex geometries and optimization

strategies, on unstructured grid approaches, and on the inclusion of viscous effects.

Acknowledgments

This research has benefited greatly from the generous support of the U.S. Air Force Office of Scientific Research under Grant AFOSR-91-0391, ARPA under grant number N00014-92-J-1976, USRA through RIACS, the High Speed Research branch of NASA Ames Research Center, and IBM.

References

- ¹Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 1," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 51–60.
- ²Jameson, A., "Multigrid Algorithms for Compressible Flow Calculations," *Lecture Notes in Mathematics*, edited by W. Hackbusch and V. Trottenberg, Vol. 1228, Springer-Verlag, 1986, pp. 166–201.
- ³Reuther, J., and Jameson, A., "Control Theory Based Airfoil Design for Potential Flow and a Finite Volume Discretization," AIAA Paper 94-0499, Jan. 1994.
- ⁴Jameson, A., "Optimum aerodynamic design via boundary control," *AGARD-VKI Lecture Series, Optimum Design Methods in Aerodynamics*, von Kármán Inst. for Fluid Dynamics, 1994.
- ⁵Jameson, A., and Alonso, J. J., "Automatic Aerodynamic Optimization on Distributed Memory Architectures," AIAA Paper 96-0409, Jan. 1996.
- ⁶Reuther, J., "Aerodynamic Shape Optimization Using Control Theory," Ph.D. Dissertation, Univ. of California, Davis, CA, June 1996.
- ⁷Jameson, A., "Optimum Aerodynamic Design Using CFD and Control Theory," AIAA Paper 95-1729, June 1995.
- ⁸Jameson, A., "Automatic Design of Transonic Airfoils to Reduce the Shock Induced Pressure Drag," *Proceedings of the 31st Israel Annual Conference on Aviation and Aeronautics* (Tel Aviv), 1990, pp. 5–17.
- ⁹Reuther, J., Alonso, J. J., Vassberg, J. C., Jameson, A., and Martinelli, L., "An Efficient Multiblock Method for Aerodynamic Analysis and Design on Distributed Memory Systems," AIAA Paper 97-1893, June 1997.
- ¹⁰Anderson, W. K., and Venkatakrishnan, V., "Aerodynamic Design Optimization on Unstructured Grids with a Continuous Adjoint Formulation," AIAA Paper 97-0643, Jan. 1997.
- ¹¹Reuther, J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers," AIAA Paper 97-0103, Jan. 1997.