6.0 Results and Discussion

6.1 Two-Dimensional Aerodynamic Optimization

In the present work the two-dimensional configurations studied are an initially symmetric NACA-0012 airfoil and a four-element airfoil. For the NACA-0012 airfoil shape optimization several solution strategies, which vary the spatial accuracy of the nonlinear fluid and linear shape sensitivity equations, are investigated. This study examines the effect of inconsistency in the order of accuracy between the fluid and sensitivity equations. This optimization was also used to demonstrate the importance of proper problem formulation by reformulating the optimization problem to yield a more desirable airfoil design. The multielement airfoil configuration was used to demonstrate the advantage of using unstructured grids for shape optimization.

6.1.1 Transonic Airfoil Design

The unstructured mesh for the NACA-0012 airfoil, shown in Fig. 6.1a, consists of 3,044 triangular cells and 1,578 nodes. The design variables, which control the shape of the airfoil, are chosen as the y-locations of the interior Bezier control points and are depicted in Fig. 6.1b. This parameterization consists of two Bezier curves; one for the upper surface and one for the lower surface, with each containing 7 interior control points. Thus, this optimization has a total 14 design variables. To verify that the ADIFOR generated grid sensitivities were correct, they were compared with central finite-difference. This comparison is shown in Fig. 6.2 for design variable 10 on the upper surface of the

airfoil. It is obvious from the figure that the results of the two methods look the same; in fact, ADIFOR grid sensitivities match central difference sensitivities to approximately 8 significant figures.

The flow condition assumed in this design problem is a transonic Mach number of 0.8 at a zero degree angle-of attack. The objective function is lift, with constraints that the lift and drag be above and below specified values C_L and C_D, respectively. For cases (1)-(4) presented below, C_L was 0.70 and C_D was 0.03. An additional geometric constraint is placed on the thickness of the trailing edge to keep it from collapsing or becoming excessively thick. This geometric constraint required the trailing edge included angle to be no less that 8° and no greater than 25°.

In an attempt to study the influence of spatial accuracy on the sensitivity analysis, several solution strategies were investigated. The solution strategies are distinguished by the order of accuracy of the nonlinear fluid and the aerodynamic shape sensitivity analyses. These strategies are as follows: (1) *Fully 1st order*, in which both the fluid and sensitivity analysis are spatially first-order accurate; (2) *Inconsistent 1st order*, where the higher-order spatially accurate fluid equations are solved and the resulting state vector used in the first-order sensitivity equations (hence, there is an inconsistency in the order of accuracy between the nonlinear fluid and linear sensitivity analysis); (3) *Fully higher-order*, where both systems of equations have higher-order spatial accuracy, with the aerodynamic shape sensitivity equations being solved in the pre-eliminated form given in equation 3.25. The fourth case is also fully higher-order; but the sensitivity equations are now solved in the incremental iterative form given in equation 4.11a.

Before any optimization procedure is deemed complete or correct, analytic sensitivity derivatives should be compared with finite-difference sensitivities. Table 6.1 presents this comparison for the four strategies listed above. As should be expected, the analytically

obtained sensitivity derivatives match their central finite-difference equilivants to approximately 5 significant digits for strategies (1), (3) and (4). It is interesting to note that the inconsistent strategy (2) produced an estimate of the sensitivity derivative which nearly splits the first-order and higher-order computations; however, strategy (2) does not require the evaluation of the higher-order Jacobians for either the fluid or the linear sensitivity analysis. (*Note*: the process of obtaining sensitivity derivatives in an inconsistent manner is not encouraged due to reasons which will be discussed below.)

For this design problem, Fig. 6.3 depicts the initial NACA-0012 and the final optimized airfoil surfaces for each strategy. Note that since strategy (3) and (4) produce the same optimized shape as expected, only one is shown. As can be seen, near supercritical airfoil shapes (characterized by an undercutting of the forward lower-surface, a reduced curvature in the mid-chord region, and an aft lower-surface concavity or cusp [169,170]) are obtained. Figures 6.4a and 6.4b illustrate the coefficient of pressure distributions on these airfoils for the first- and higher-order spatially accurate solutions, respectively. In all cases, the lower-surface shock has been completely eliminated and the upper-surface shock moved aft. Unfortunately the drag constraint, which was arbitrarily specified to an unrealistically high value, allowed the upper-surface shock to locate itself in very close proximity to the trailing edge. This undesirable feature, however, is easily circumvented by maximizing the lift to drag ratio instead of just lift, or by simply reducing the drag constraint, as will be seen.

Off-surface pressure contours about the initial nonlifting NACA-0012 airfoil, for the first- and higher-order spatially accurate solutions, are depicted in Fig. 6.5a and 6.5b, respectively. As expected, a smearing of the shock, for the first-order solution, is observed. Presented in Figs. 6.6a, 6.6b, and 6.6c are the off-surface pressure contours for the optimized airfoils from strategies (1) through (3), respectively. Once again, it can be seen that the upper-surface shock has moved as aft as the drag constraint will allow, thus,

locating it at the trailing edge. A summary of the optimized lift and drag values obtained from each strategy is illustrated in Table 6.2. As noted earlier, the incremental iterative strategy (4) allows higher-order spatially accurate optimization while only utilizing the memory of a first-order method, but suffers from a CPU time increase as shown in the Table. It should also be noted that the maximum lift of strategy (1) was 0.5931 with a drag of 0.03517. This design violated the specified constraints due to the fact that the first-order spatially accurate assumption underpredicted the shock strength and location. A higher-order solution yielded a lift of 0.9327 and a drag of 0.04019 for this same design.

At this time a potential hazard from the use of an inconsistent method should be noted. It was seen in Fig. 6.3 above that the optimized shape for the inconsistent strategy (2) significantly deviated from the airfoils obtained with consistent methods. The difference in airfoil shapes between the consistent strategies (1) and (3) can be directly attributed to the difference in the order of spatial accuracy of the fluid equations (i.e., the "physics"). However, since the order of the spatial accuracy of the aerodynamic analysis is the same for both strategies (2) and (3), the difference between the optimized shapes must now be attributed to the inconsistency of the sensitivity derivative calculation within strategy (2). Due to the fact that the sensitivity derivatives act as the *directions* or *road map* for the optimizer, it is thus reasonable to assume that given *wrong directions* it might not find the *right location*. This is not to say that the *right location* (i.e., that particular isolated local minimum) is always the best location. It is conceivable that the inconsistent method may find (due to the *wrong directions*) a region of the design space that yields a design which is an improvement over that obtained with a consistent method. In the present airfoil optimization this is not the case, with strategy (3) producing a slightly better design than strategy (2). Furthermore, it is interesting to note that when the optimized airfoil from strategy (2) is used as the initial shape in the consistent fully higher-order scheme, the optimized airfoil from strategy (3) is recaptured. Thus it is concluded that the optimized

shape produced from the fully consistent, higher-order strategy (3) is the isolated local minimum, and that strategy (2) failed to reach it due to the inconsistency in the sensitivity derivative calculation.

As an indication of the importance of problem formulation, the above aerodynamic design optimization was performed again with the drag constraint changed to 0.015 (half of its previous value). The result of the new formulation is shown in Fig. 6.7. Now, as illustrated by the coefficient of pressure distributions on the surfaces of the initial and optimized airfoil shapes, the shock is located at approximately 85% chord. Thus, proper formulation of the design problem is required to avoid unrealistic or undesirable features.

6.1.2 Subsonic Multielement Airfoil Design

The multielement airfoil configuration considered in the current research has four elements. These elements consists of a leading edge slat up-stream of the main airfoil, and a double slotted flap down-stream of the main airfoil. The double slotted flap contains a vane in front of a main flap. The unstructured mesh for the multielement airfoil contains 14,919 triangular cells and 7,614 nodes and is shown in Fig. 6.8a. In this design case the shape of the vane is the surface being optimized. The vane is parameterized with two Bezier curves; once again, one for the upper surface and one for the lower surface, with 5 interior control points for each. This parameterization is illustrated in Fig. 6.8b, and has a total of 10 design variables. As before the ADIFOR grid sensitivities were verified by comparison with finite-difference values. Figure 6.8c depicts the sensitivity of the interior mesh with respect to a design variable on the lower surface of the vane.

The flow conditions for the multielement airfoil calculation were a subsonic Mach number of 0.20 at 16.02° angle-of-attack. To verify the accuracy of the sensitivity derivative computations for the multielement airfoil, the geometric design variable selected was the same Bezier control point used above in Fig. 6.8c for the grid sensitivity illustration. For

this configuration the analytically obtained derivative $dC_L/d\ _7$, for higher-order spatial accuracy, was computed to be 0.210754 as compared to the central finite-difference value, 0.210733. As a qualitative illustration of the influence of shape changes on the flow field solution, Fig. 6.9a and 6.9b depict the pressure and the sensitivity of the pressure with respect to the above discussed design variable, respectively.

Before presenting the design results, some brief words about the importance of design problem formulation need to be re-asserted. Sensitivity analysis is merely an extra level of computation that provides additional information to the designer. When the sensitivity analysis routines are coupled with the fluid solver, a mesh movement strategy, and a numerical optimizer, a functional design tool is produced. The eventual designs created with this tool will be only as good as the formulated design problem. If improperly formulated, designs can be produced that will violate constraints such as those needed for manufacturability or structural feasibility, or produce a design that has superb performance at one operating condition, but is unacceptable off the design point. Thus, experience is required in formulating meaningful design problems.

For the multielement airfoil case, it is recognized that the goal of a flap system is to maintain the highest possible lift-to-drag ratio at the maximum lift coefficient [171]. With this in mind, the problem formulation consisted of maximizing the lift coefficient. It is not asserted that inviscid flow analysis and sensitivity analysis are capable of modeling the physics or properly designing such a configuration (especially at high angle-of-attack situations such as take off and landing), but rather that an unstructured grid method easily discretizes the domain and that it is possible to carry out this type of design study. A geometric constraint has been placed on the trailing edge angle to keep it from becoming excessively thick or thin. This constraint required the trailing edge included angle to be within $5^\circ\quad _{TE}\quad 45^\circ$.

Results of the multielement airfoil design study are summarized in Table 6.3. This optimization required 7 design cycles with 83 CFD analyses along the line searches and a total run time of a little over 40.5 min on a CRAY Y/MP. The objective function was increased by 6.2%. Figure 6.10a depicts the initial and final optimized vane produced by this shape optimization procedure. Illustrated in Fig. 6.10b are the corresponding pressure coefficient distributions about the multielement airfoil. It should be noted that the horizontal distances between the vane, the main airfoil, and the flap have been exagerated in the figure so that their C_P distributions may be easily distinguished. As seen, the pressure distributions about the leading edge slat and the main element remain roughly unchanged, but those on the vane and flap have been greatly altered. Another interesting design problem, which may have produced more dramatic improvements, could have been the design optimization of the shape, location, and orientation of the leading edge slat.

6.2 Three-Dimensional Aerodynamic Optimization

The three-dimensional configurations studied were an initially rectangular wing with uniform cross-sections and a complete Boeing 747-200 aircraft. For the rectangular wing, the planform shape of the wing was optimized using both fully consistent first- and higher-order spatial accuracy. This optimization study permitted the investigation of the effect of order of accuracy on the final planform designs. The second case studied was the wing re-design of a Boeing 747-200 aircraft in transonic flow. Higher-order spatial accuracy was used for the flow ananlysis and the shape sensitivity analysis. This case was used to demonstrate the advantage of using unstructured grid schemes for the analysis and shape optimization of geometrically complex configurations.

6.2.1 Wing Planform Design

The wing planform parameterization used in the present work has previously been shown in Fig. 5.2. These design variables consisted of local chord multipliers and setback distances at the cranks and tip. For the rectangular wing planform the root chord and semi-span of the wing remain constant throughout the design and are 1.0 and 2.0, respectively. Furthermore, a total of two interior cranks are selected, with their spanwise locations also fixed throughout the design. The distance between the cranks and tip (see Fig. 5.2) are B_1 = 0.48, B_2 = 0.18, and B_3 = 1.34. Thus, this design problem had a total of 6 design variables. The unstructured surface mesh for this initially rectangular wing, which has uniform NACA-0012 cross-sections and a rounded tip, is shown in Fig. 6.11a. This mesh contained 18,856 tetrahedral cells and 3,867 nodes. As usual, to verify the ADIFOR grid sensitivities a design variable is selected and comparisons are made with central finite-difference values. The surface grid sensitivity for the first chord multiplier is depicted in Fig. 6.11b and 6.11c for the finite-difference and ADIFOR grid sensitivity, respectively. Once again, identical results are obtained, with the grid sensitivities matching to approximately 8 significant between the two methods.

The flow condition assumed in this design problem is a subcritical Mach number of 0.7 at a 2.5˚ angle-of-attack. The objective function is the lift-to-drag ratio, with constraints that the lift and drag coefficients be above and below their initial values, respectively. It should be noted that the reference area used in the lift and drag coefficient calculations was the actual wetted surface area of the wing. This made it possible to observe an increase in the lift coefficient during the planform design at fixed angle-of-attack. Additional side constraints are placed on the design variables to ensure a *structurally feasible* geometry. Note that for actual aircraft design, a more meaningful design problem would be to minimize drag at fixed lift, with strict geometric constraints on the aspect ratio of the wing.

The present formulation, however, has been successfully implemented in the past for proof-of-concept optimization studies [162,164] with structured-grid algorithms and is therefore utilized here.

Regardless of problem formulation, the most vital part of any optimization procedure is the evaluation of accurate sensitivity derivatives. Once the code has been built to evaluate these derivatives, the designer may then formulate meaningful optimization problems at his discretion. Thus, before any optimization procedure is deemed complete or correct, analytic sensitivity derivatives should be compared with finite-difference values. Table 6.4 presents this comparison for the planform design variables. As expected, the *consistent* analytically obtained sensitivity derivatives, for spatially first- and higher-order accurate computations, compare well with finite-difference values. A discussion on the requirement of consistency to obtain accurate sensitivity derivatives has been presented in section 6.1.1 above. On a more qualitative note, the surface pressure contours and the sensitivity of the pressure with respect to the first design variable are illustrated in Fig. 6.12 and 6.13 for the first- and higher-order computations, respectively, for the initial geometry.

For the above formulated design problem, Fig. 6.14 depicts the initial and final wing planforms for both the first- and higher-order optimizations. As seen, aft-swept wings were obtained for both cases, with the final designs having a much higher aspect ratio than the initial rectangular wing (which was expected for the current problem formulation). A summary of the planform optimization results are illustrated in Table 6.5. From this table it is observed that the objective function was improved by 13.3% and 12.1% for the first- and higher-order accurate optimizations, respectively. Note that even though both optimization cases produced similar designs, accurate lift-to-drag ratio prediction requires the use of a spatially higher-order accurate scheme. This is due to the inability of a spatially first-order scheme to accurate predict the drag. However, the cost of this accuracy is realized as a factor of 3 increase in CPU time over the first-order optimization. As for the

memory quoted, it is that needed in addition to the memory required to run the CFD solver. Note that this assumes that the flux Jacobians at the cell faces are stored within the solver. This itself is an additional memory if a faced-based data structured is adopted in the unstructured grid solution algorithm.

6.2.2 Boeing 747-200 Wing Redesign

The unstructured surface mesh for the twin-engine Boeing 747-200 aircraft configuration is shown in Fig. 6.15a and is derived from the model tested in the NASA Ames 11 foot Transonic Pressure Tunnel (Test AR0502). The volume grid around this surface contains 352,547 tetrahedral cells and 63,828 nodes. Outboard of the outermost engine nacelle, the design variables for this configuration fit the twist and dihedral schedules as a function of the span, with cubic polynomials. At the point where the engine strut meets the wing, point and slope continuity for both the twist and dihedral are enforced. Thus, only the coefficients of the quadratic and cubic terms are free and are therefore chosen as the design variables. This constitutes 2 design variables for the twist and 2 for the dihedral, giving a total of 4 design variables for this optimization problem. Figure 6.15b depicts the surface grid sensitivity for a twist design variable; specifically the coefficient of the quadratic term in the spanwise twist schedule.

For the Boeing 747-200 optimization study, a transonic Mach number of 0.84 was chosen with a free-stream angle-of-attack of 2.73 degrees. To verify the accuracy of present discrete sensitivity analysis approach, the sensitivity derivatives of lift-to-drag ratio, with respect to the geometric design variables previously discussed, are computed and compared with finite-difference values. It should be noted that the work associated with computing sensitivity derivatives via the direct differentiation method does not scale with the number of output functions. Hence, for the computations shown, any number of output function sensitivities (i.e., for lift, drag, pitching moments...etc.) can be computed with

little effort. To evaluate the corresponding finite-difference derivatives, however, requires two analysis runs per design variable for central differences. Thus, due to the expense of finite-difference derivatives, only those derivative comparisons with respect to one of the geometric design variables were performed. The computed derivative $d(L/D)/d_1$ yielded 1.3231, and the central finite-difference produced 1.3229. As should be expected with a consistent, discrete sensitivity analysis, favorable agreement between analytic derivatives and finite-difference derivatives is observed.

The design problem formulation for the Boeing 747-200 aircraft was to maximize the lift-to-drag ratio. Once again, constraints that the lift coefficient at the final design be greater than the initial value and that the drag coefficient be reduced for the optimized shape have been incorporated. The results of this optimization are shown in Table 6.6. Since this design study was for purely demonstrative purposes, the optimization was halted after 3 design cycles and not allowed to continue until an isolated local minimum was found. Observe that the objective function has been improved by 2.7%, but at the cost of 23.4 CRAY Y/MP hours. This represents about 6 converged CFD analyses. As noted in the introduction, however, techniques are now being studied to reduce these excessive CPU times. The initial and final twist and dihedral schedules are shown in Fig. 6.16a. Note that a positive twist angle is defined herein as leading edge up. As seen, this optimized wing has a greater twist at the tip station and an altered dihedral distribution. The initial and optimized surface meshes are viewed from an upstream vantage point in Fig 6.16b to show the dihedral distributions of each wing. Surface pressure contours for both the initial and final designs are illustrated in Fig. 6.17 and 6.18 for the upper and lower surfaces, respectively. It can be observed from close inspection that the upper surface of the optimized wing has a greater region of lower pressure than the initial wing and that the lower surface has a slightly higher pressure. Once again, a design problem which possibly

could have produced more significant increases in the lift-to-drag ratio would have been to perform the shape optimization of the wing airfoil sections. However, this current procedure has demonstrated the ease with which an unstructured grid approach to aerodynamic shape sensitivity analysis and optimization may be used to analyze and design geometrically complex configurations.

6.3 Efficient Static Aeroelastic Analysis

The aerodynamic surface mesh and the structural model have been previously shown in Fig. 4.1. The aerodynamic configuration used for the aeroelastic analysis is the wing that resulted from the above planform optimization study. For the structural model skins, webs, and ribs are all modeled as constant strain triangle (CST) membrane elements and the spar caps are modeled as truss elements. The final model contains 190 elements (162 CST and 28 truss elements) with 159 degrees-of-freedom. Note that for the present results the surface structural nodes for the finite-element model were taken to be a subset of the aerodynamic surface grid. However, with the unstructured-grid CFD capability, one can start from the structural layout (or design) and require that the aerodynamic surface triangulation conform to the structural surface nodes. That is, an unstructured CFD surface grid can be constructed that contains the structural surface nodes as a subset.

The rigid wing and aeroelastic wing analyses performed herein have a free-stream angle of attack of 2.5°, and the trimmed aeroelastic wing analyses are initiated at this value. In addition, free-stream data at an altitude of 35,000 ft ($= 7.382 \times 10^{-4}$ slug/ft^3 and a $=$ 973.14 ft/s) have been adopted to dimensionalize the aerodynamic forces. The dimensions of the wing have been dimensionalized to have a semi-span of 65 ft and a root chord of 32.5 ft. For the structural analysis, all members are considered to be made of Aluminum

6061-T6 with a modulus of elasticity, a Poisson's ratio, and a density of 10^7 lb/in^2, 0.33, and 0.097 lb/in^3, respectively. Truss elements are considered to have a cross-sectional area of 0.05 in^2, and CST elements are considered to have a thickness of 0.3 in.

Some system convergence results are depicted in Figs. 6.19(a) and 6.19(b) for the subcritical (M = 0.70) and supercritical (M = 0.85) flow cases, respectively. The two methods have been previously discussed in section 4.3.2; method 1 interacts the disciplines when a CFD criterion is satisfied and method 2 interacts them at a prescribed constant rate. In these figures, the ordinate has been normalized by the CPU time of the rigid-wing analysis and the abscissa is the value of the interaction control parameter. As observed for both flow cases, neither interaction of the disciplines with fully converged solutions nor interaction at every iteration is the most efficient means for performing the static aeroelastic analysis. The most efficient approach, as seen in the figures, requires only 10-percent more CPU time than that required for the rigid-wing CFD analysis. Proper specification of the interaction control parameters, moreover, is seen to depend on the flow conditions and is difficult to determine a priori. In fact, for small values of F_n, highly converged CFD solutions are obtained before structural interaction, and the CPU time required to converge the coupled system is off the scale. However, intermediate values of F_n exist at which the interaction diverged for the supercritical flow case. Based on the current study, however, a compromise would be to allow interaction of the disciplines at every iteration; this level of interaction was found to converge within 20-percent of the CPU time required for the rigid-wing at all flow conditions. This extreme is more robust as well. Thus, interaction of the disciplines at every iteration is used in the following aeroelastic computations.

Static aeroelastic analyses of the wing are made for flow conditions that range from low subsonic to supersonic. In particular, four cases are examined. For each case the rigid-wing CFD analysis, the static aeroelastic wing analysis, and the static aeroelastic wing analysis trimmed to the rigid-wing lift are presented. These cases consist of low subsonic

($M_\infty = 0.30$), high subsonic but still subcritical ($M_\infty = 0.70$), transonic or supercritical ($M_\infty = 0.85$), and supersonic ($M_\infty = 1.2$) wing flows.

Computational results for the four flow cases are given in Figs. 6.20 through 6.23, respectively. Each of these figures consists of three iteration history plots as well as three sketches of the rigid-wing, aeroelastic wing, and trimmed aeroelastic wing structural sections at convergence. These three sketches, which plot the ordinate and abscissa at the same scale, are at the lower right and in each serve as the key for the labeled curves. At the upper left of each figure is the CFD residual history plot (Log(R/R0) versus CPU time in seconds on the CRAY Y/MP); at the upper right are the corresponding lift coefficient histories. System convergence is monitored by the rms value of the surface deflections, and interaction history plots are shown at the lower left for the aeroelastic wing solutions.

As seen in Figs. 6.20 through 6.23, the residual histories of the coupled aeroelastic analyses closely mimic those of the rigid-wing computations; the extra CPU time required for each iteration is attributed to the additional structural analysis involved. Also note that for the low subsonic case presented in Fig. 6.20 the compressible equations should utilize preconditioning to more efficiently model the nearly incompressible flow. For the trimmed aeroelastic analyses, a greater flow-field disturbance is observed as the fluid and the structure interact. Consequently, the trimmed aeroelastic analysis requires more CPU time to converge. The double hump in all convergence histories, which is particularly noticeable in the trimmed aeroelastic cases, is caused by the flow solver using a first-order solution until the CFD residual drops one order of magnitude (i.e., -1) before switching to higher-order spatial accuracy. This procedure was also adopted from USM3D [71]. Thus, the solver effectively trims the aeroelastic wing based on the first-order spatially accurate solution and, when the residual switching criteria are met, re-trims the wing based on the higher-order solution. Some CPU time may potentially be saved for the trimmed aeroelastic

cases by relaxing these criteria; however, for consistency all cases herein use the same residual tolerance of –1.

As the Mach number is increased, loading on the wing also increases; thus, greater wing deflections are observed. The supercritical case shown in Fig. 6.22 demonstrates the applicability of the present approach to flows in which shock waves are present. This case also illustrates the dramatic loss of lift associated with the flexible wing, the high angle of attack required to trim the wing, and the correspondingly large wing deflections and twist that can occur. This loss of lift is caused by the washout effect of the aeroelastic wing and by the sensitivity of the lift to the shock strength and location. Figure 6.23 presents an application of this method to supersonic flow, although this particular wing is not ideally suited for this flow regime.

As a final indication of how the aeroelastic wing deforms across the Mach number range studied herein, the rigid wing and aeroelastic wing for each flow case are shown in Fig. 6.24. Once again, the large structural deflections which occur at supercritical Mach numbers are observed.