JOURNAL OF AIRCRAFT
Vol. 35, No. 6, November–December 1998

Aerodynamic Optimization Using Analytic Descriptions of the Design Space

C. A. Toomer,* M. E. Topliss,† and D. P. Hills‡
British Aerospace (Operations), Ltd., Bristol, England BS12 7QW, United Kingdom

This paper discusses a strategy for performing aerodynamic single-point and multipoint optimization studies. The approach uses a combination of three methods: an Euler computational fluid dynamics (CFD) code that provides accurate information at one chosen point in the design space, a linear approximation method of the quasianalytic type that rapidly evaluates data at many design points, and a function-fitting algorithm that maps aerodynamic and geometric data across the whole design space. As the method proceeds, a good quality database of aerodynamic and geometric quantities is rapidly constructed at a fraction of the cost of using conventional methods. The availability of this database and the analytic representation of the quantities over the whole design space permits optimizations to be performed using constraints on specific aerodynamic parameters, e.g., maximum surface Mach number and pressure gradients, over chosen regions of the aerofoil. The approach is also highly flexible, enabling optimization problems of increasing complexity to be built up quickly, and is easily integrated within modern engineering methods. Results are presented for a typical transonic aerofoil giving comparisons with individual CFD analyses.

Nomenclature

C_D = coefficient of drag
C_L = coefficient of lift
C_M = pitching moment
C_p = pressure coefficient
k = constant
M = Mach number
N_s = number of vertices in the defining B-spline representation
p = pressure
s = distance along aerofoil surface
w = weighting coefficient, used in objective function
X = distance from leading edge along chord
Y = distance normal to chord
α = angle of attack, deg

Superscripts
opt = optimization result
orig = results on original configuration

Introduction

THE area of design optimization is of great relevance to the aerospace industry and is currently advancing rapidly. The numerical tools under development should provide the designer with the flexibility to address different types of problems, while providing rapid and accurate solutions. One type of problem is to understand the characteristic parametric variations and interrelations in a multidimensional design domain. A suitable numerical tool would enable the designer to investigate the design space, identify regions in this parameter space of specific interest, and locate as many solutions to a given optimization problem as possible. Another type of problem is to select a particular configuration of interest and use numerical tools to modify this design to a more exacting standard.

This paper discusses some of the methods currently being developed at the Sowerby Research Centre to address the first type of problem—forward rather than inverse design. This work is part of a program that aims to provide an integrated software environment for the design engineer. In its maturity, the system will comprise a wide variety of techniques that will facilitate aerodynamic, structural, and manufacturing optimization for realistic three-dimensional vehicles in representative flight envelopes.

One particular strategy is described here for exploring the design space for use in aerodynamic optimization. The value of this approach is in the large time and cost savings that can be made compared with the more traditional methods of data generation using experimental techniques and computational fluid dynamics (CFD) codes. The use of nonlinear CFD programs to fully explore the product design space for anything but simple shapes and flow regimes is still prohibitively expensive, even with the use of powerful computing facilities.

This paper is structured into three main sections. In the Methodology section, the three complementary methods used to generate the data are described. These methods are used in sequence with the aim of providing analytic descriptions of aerodynamic and geometric quantities over the design space. The next two sections contain single-point and multipoint optimization results for a typical transonic aerofoil. Initially, constraints are applied to the aerofoil area and the force and moment coefficients alone. Then the number of constraints is increased by adding conditions on other aerodynamic quantities on the aerofoil surface. These examples show the flexibility of this approach where further constraints can be added into an optimization problem based on results of previous optimizations.

Methodology

Design Space

The design space of interest is that associated with typical two-dimensional, transonic aerofoils in an inviscid, isenthalpic

Presented as Paper 96-4141 at the AIAA/USAF/NASA/ISSMO 6th Symposium on Multidisciplinary Analysis and Optimization, Bellevue, WA, Sept. 4–6, 1996; received Feb. 16, 1997; revision received May 12, 1998; accepted for publication June 30, 1998.

*Research Scientist, Sowerby Research Centre, CFD Group, FPC 267, P.O. Box 5, Filton; currently at NASA Langley Research Center, Institute for Computer Applications in Science and Engineering, M/S 403, Hampton, VA 23681-0001. E-mail: chris.toomer@src.bae.co.uk.

†Currently Banker, SBC Warburg, London, England, UK.

‡Group Leader, Sowerby Research Centre, CFD Group, FPC 267, P.O. Box 5, Filton. E-mail: david.hills@src.bae.co.uk.

flow. The domain is characterized by a unique set of design variables that reflect the operating conditions (freestream Mach number and angle of attack) and the sectional shape of the aerofoil. The aerofoil configuration is defined by a set of rational B-spline curves,[1] from which the resulting N_s vertices are identified as the set of shape variables. Thus, with the angle of attack and freestream Mach number, a total of $(N_s + 2)$ design variables fully describe the product. The design variables are independent of each other. For the aerofoil configurations considered, the surface of the aerofoil is described using 24 B-spline curves of tenth order. Note that the length of the chord line is always scaled to one unit. A C-grid is then generated around the aerofoil surface with (33×257) points in the normal and circumferential directions, respectively, and 129 gridpoints on the surface.

To perform an optimization, information on the design space is required and a suitable optimization algorithm must be chosen. The computational approach to generating the required aerodynamic data is to produce a suitable mesh for each design and employ a CFD tool to obtain the flowfields and, hence, aerodynamic coefficients, at selected points within the design space.

Topliss et al.[2] developed an alternative approach for investigating the design space. This methodology involved the use of three subprograms, which when called in sequence enabled rapid generation of information about the whole design space. The choice of applying the codes in sequence is not unique, other options are available and form part of an ongoing development program. The data were produced using only one CFD result at a user-defined point in the design space. At other points in the domain, the information was provided by two other codes: Projected Implicit Reconstruction (PIR) and Design Spare Approximation Technique (DSAT). The algorithms used in each of these codes are briefly described next. Further information is in Ref. 2.

CFD Code

The CFD code solves the unsteady, two-dimensional Euler equations for a given angle of attack and freestream Mach number on the computational mesh. A cell-centered, finite volume approximate Riemann solver approach has been employed and an explicit time-integration scheme is used to march the solution to convergence (taken as at least a three-order-of-magnitude reduction in the residuals and in the variation of the aerodynamic coefficients). The basic flux discretization is based upon a second-order total variation-diminishing version of the AUSMD flux splitting scheme,[3] where the residual of each cell is written as a balance of fluxes across cell interfaces. Closure of the system is attained by using the ideal gas law, with a constant ratio of specific heats. The total enthalpy of the system is constant for a constant freestream Mach number.

Projected Implicit Reconstruction

Once the flowfield for the conserved field variables (density and momentum) at a point in the design space is obtained, this is used as the starting solution for the PIR method. PIR provides a linearization around a design point of interest and is based on the fundamental sensitivity equations cited in the literature.[4,5] The application of sensitivity analysis to aerodynamic optimization has been widely reported.[6,7] From the initial design point, the PIR calculation involves stepping through the design space by perturbing one or more components of the design vector (composed of the design variables). At each step a linearized system of equations is solved to provide the necessary changes to the flowfield that was calculated at the previous step. In addition to the reconstruction of the flowfield, aerodynamic and geometric data are also extracted and stored. The sweep through the design space then continues, accumulating data at every step, until the user-defined final position is reached. The process is illustrated in Fig. 1 for the case where the only change in the design vector is because of the angle of attack being incremented at each step by 1 deg from an initial value of 3 deg.

The changes to the flowfield are found by evaluating the sensitivity of each conserved field variable with respect to changes in the design vector at every cell center. An approach based on the quasianalytic method is applied, where a large sparse system of linear algebraic sensitivity equations (incorporating the gridpoints, the residual flux vector, the conserved field variables, and the design vector) is solved using a quasi-Newton iterative method. All of the boundary conditions applied are consistent with those in the CFD code. Required

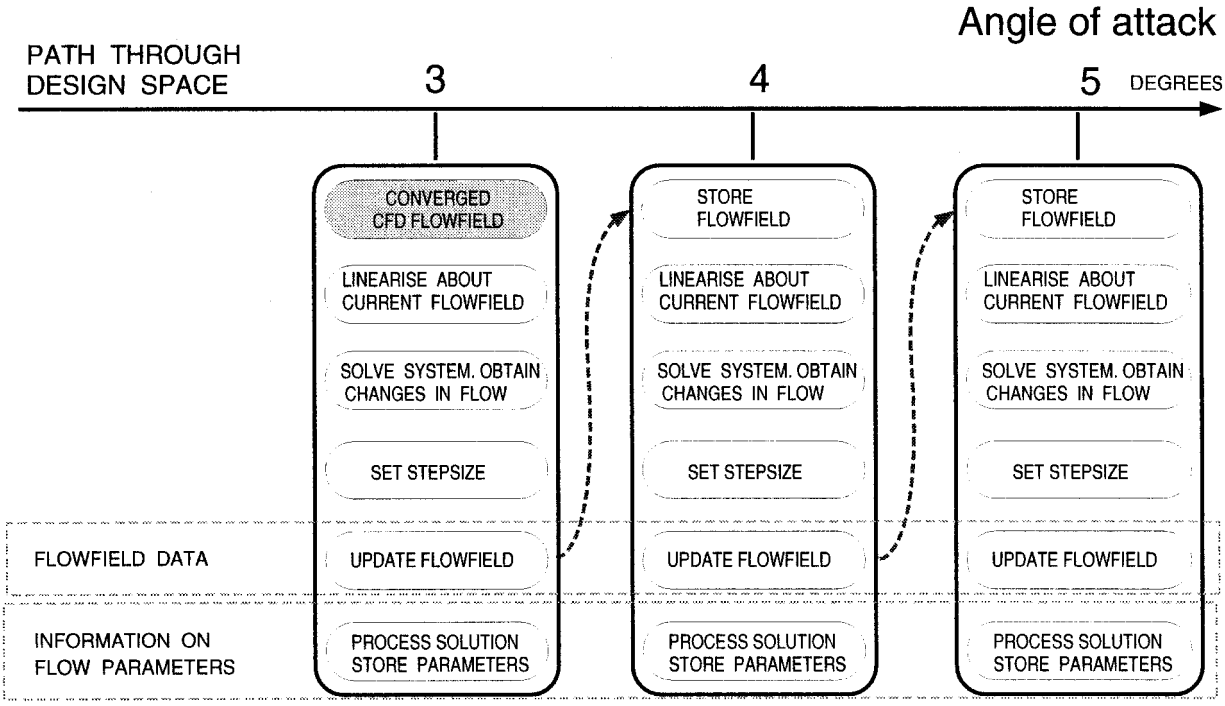

Fig. 1 Progression through the design space using the PIR method. In this example the angle of attack is the only design variable perturbed.

solutions have a residual reduction of at least six orders of magnitude.

The aerodynamic coefficients are evaluated as functions of the normal and axial force coefficients. These are calculated in the standard way as the sums of the contributions at the surface of each cell face on the aerofoil.

The parameters of interest recorded at each step in the PIR marching process are the overall section forces and moments (lift coefficient, drag coefficient, pitching moment), flowfield quantities, and geometric quantities. The flowfield quantities are the pressure, density, Mach number, entropy, and pressure loss, and three options are available (of which all, some, or none of them can be chosen by the user). The first option is for the maximum and minimum values over the field, the second option is on the aerofoil surface, and the third option is a user-defined section of the surface. For the surface options, the flowfield gradient values are available. The user can specify any combination of upper, lower, or both surfaces. The geometric quantities include the aerofoil area and areas of sections defined by the user, the radius of curvature, curvature and curvature derivatives, maximum thickness, mean thickness, spar thickness(es), and aerofoil trailing-edge angle.

Note that although not requested by users of the code, it is straightforward to add in constraints on the leading-edge radius. This would be achieved by adding the leading-edge radius in as a parameter in the PIR code and in similar fashion to the other geometric parameters, storing the leading-edge radius values as a function of the B-spline positions. When time permits, further tests will be performed and reported using the geometric quantities.

PIR produces accurate data at a substantial saving in time and cost compared with the CFD method. However, both approaches produce results at discrete points in the design domain. Hence, a third technique is employed that uses the CFD and PIR data to map each chosen parameter across the whole design space in terms of analytic functions.

Design Space Approximation Technique

To economically produce an approximation to the function, the smallest number of PIR iterations possible must be performed without unduly compromising the accuracy. Hence, the direction taken through the design space during the PIR runs is controlled. The CFD analysis used to provide the initial solution for the PIR method is chosen to be at the center of the design domain. A small number of PIR analyses are then undertaken at specific points in the domain radiating out from the center, to capture the variations in the fields caused by perturbations in the design variables. Each design variable is perturbed separately.

For example, for each shape variable, PIR analyses are performed in each direction, allowing the surface to be lowered and raised. The data from four of these analyses are used to find the coefficient values of the function. Hence, for two shape variables, data from eight PIR steps are required. With the initial CFD analysis, this provides information at nine points and their distribution is shown in Fig. 2, where the center point (CFD analysis) is the starting point for all of the PIR runs. Data are taken at four points near the center and at four points on the boundary of the domain. Large stepsizes can lead to a loss in accuracy and, thus, extra points (the results of which are discarded) are evaluated in between to ensure the accuracy of the results at the points on the boundary.

The typical change in the position of the vertices is approximately $\pm 1.92 \times 10^{-2}$ in steps of about 6.4×10^{-3}. The stepsize used for perturbing the angle of attack is a half-a-degree over the range $\alpha = [-4, 8]$ deg. As expected from a first-order scheme, there are limitations on the extent to which each design variable can be perturbed about any design point.

A curve-fitting technique is applied to each of the quantities in the PIR database. The function used is the sum of a fourth-order polynomial in each design parameter, so there are no cross terms. Using the data supplied by the PIR and CFD methods, the coefficients of the function are obtained by Gauss–Jordan elimination.[8]

The time taken to perturb the angle of attack and 23 shape vertices about a chosen design point, to record all of the required flow quantities in the regions specified by the user, and to function fit these aerodynamic quantities over the whole design domain requires 2 to 3 CPU hours on a DEC alpha

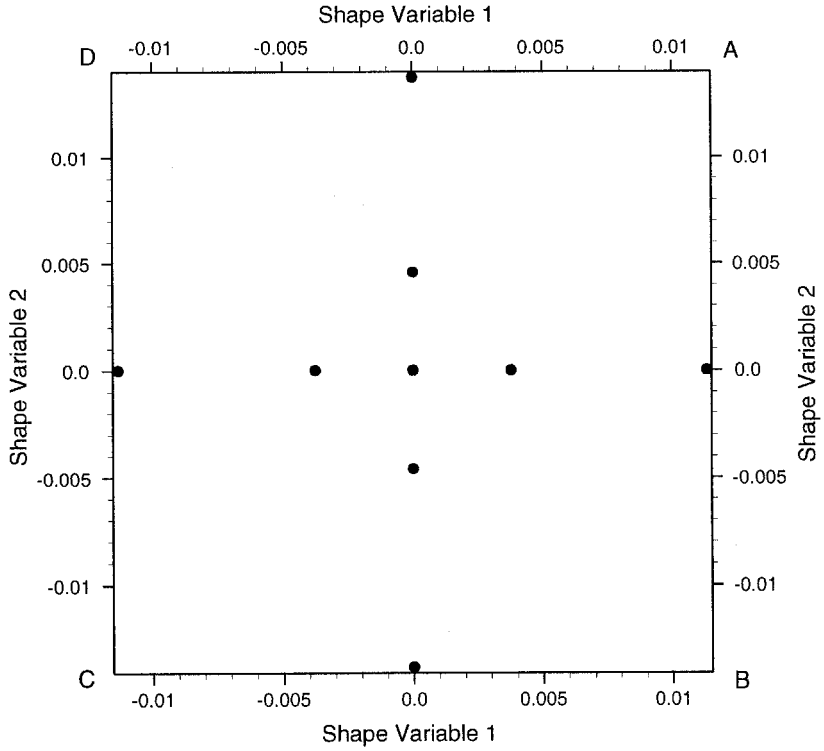

Fig. 2 Positions in the two-dimensional domain where data from the CFD code (center point) and the PIR runs are used to approximate the function.

workstation. The time taken to complete the PIR DSAT stage is approximately 90% of the time of one CFD run.

Representing parameters of interest as continuous functions provides a number of advantages: the information is available at any point in the design space, and there is easy evaluation of the first- and second-order gradients with respect to each design variable for use in the optimization scheme.

The analytical representation used here is advantageous because of its simplicity, and it is suitable for many types of optimization problems. It is, however, ony one of many possible functional representations. For example, when studying buffet, a functional representation with cross terms in the angle of attack and grid vertices is advisable. Then the gradient of the lift coefficient with respect to the angle of attack can be utilized in the modeling of the onset of separation. Once an analytic formulation for the lift coefficient has been formulated, the corresponding expression for the first derivative with respect to angle of attack can be derived and used in the optimization scheme as a constraint on the maximum angle of attack.

Optimization Technique

The optimization program is loosely coupled to the analysis codes. A sequential quadratic programming algorithm from the NAg libraries[9] is used as the optimizer. For the cases considered here, the objective function to be minimized is always the drag coefficient. Equality constraints are set on the lift coefficient, pitching moment, and aerofoil area. The upper and lower bounds imposed on the angle of attack and on the movement of each vertex correspond to the limits of the design space considered in the PIR/DSAT run. In some of the cases considered, further (inequality) constraints are imposed on other parameters, e.g., maximum Mach number.

In the design space there often exists more than one local minima in the functional representation for the objective function. Because the SQP algorithm is a gradient-based scheme, there is no guarantee that the global minimum will be found. The optimization program automatically performs many optimizations, each time starting the search procedure from a different initial point in the design space. The SQP algorithm required between 20 and 30 major iterations (converging to the first-order Kuhn–Tucker point) for the single-point optimizations and between 20 to 60 iterations for the multipoint optimizations.

There are two reasons the results from these optimization studies should be considered as design improvements rather than the optimal solutions: the result may be a local rather than a global minimum, and the result is obtained using the analytic representations of the parameters in the database and there will be an error associated with these approximation techniques.

The true measure of the improvement made in the optimization is by comparing the values of the aerodynamic coefficients (and other chosen parameters) obtained from the CFD results on the original aerofoil with the corresponding values obtained from running the CFD code on the optimized aerofoil (obtained from the CFD/PIR/DSAT data).

In the examples considered in this paper, the aim of the CFD/PIR/DSAT optimizations is to identify one or more points of interest in a design space that show promise in providing a solution to an optimization problem. Once these points have been identified, a second, more accurate, investigation can be performed about each point in turn. The CFD/PIR/DSAT is one such tool that can be used successfully in such an investigation. However, the design space must be reduced from that considered in these problems to a more localized domain to ensure that differences between the PIR and CFD results at any design point considered in a local domain are no more than, say, 1–2%. For example, for the angle of attack, the range considered would reduce from the 12 deg considered here to a range of about 2–3 deg. Using the design point (and the corresponding CFD solution) from the optimization over the large domain as the central point of the localized domain, a new, smaller domain is chosen and the PIR/DSAT procedure can be repeated. Note that the smaller domain can either be chosen by the user or set automatically within the code.

As better-quality information is acquired about the design space, the database can be updated and extended, and is available for further investigations.

Test Case Considered

The test case chosen is a generic aerofoil at Mach 0.8 and 2-deg angle of attack. This transonic case has a strong shock on the upper surface of the aerofoil and can be considered toward the extreme end of the typical test cases used to validate the computational methods.

Single-Point Optimization Results

Case 1

Equality constraints are set on the aerofoil area, lift coefficient, and pitching moment such that these values are the same on the optimized aerofoil as on the original. Hence, the center point of the design domain (with the original aerofoil shape and the angle of attack of 2 deg) is a solution to the optimization problem. The objective function is taken as the drag coefficient.

In case 1, the starting point for the optimization is at $M0.8$ and 2-deg angle of attack. The angle of attack and shape parameters are allowed to vary, but the freestream Mach number is held constant. The aim is to minimize the objective function. The optimization process gives a 37% reduction in the value of the drag coefficient for a 4% loss in the value of the lift coefficient (comparison made between the CFD results on the optimized and original aerofoils, results are presented in Table 1). This reduction in the drag is achieved by flattening the upper surface of the aerofoil compared with the original that creates a weaker shock near the trailing edge. The original and optimized aerofoils are shown in Fig. 3. The main picture en-

Table 1 Comparison of CFD and PIR data for the original and optimized aerofoils[a]

Aerofoil shape	Method	Freestream Mach number	Angle of attack, deg	Aerodynamic coefficients		
				C_L	C_D	C_M
Original	CFD	0.80	2.0	0.9772	0.0735	−0.2936
Optimized	PIR/DSAT	0.80	1.41	0.9772	0.0517	−0.2936
Optimized	CFD	0.80	1.41	0.9358	0.0462	−0.2944
Original	CFD	0.75	2.0	0.8950	0.0245	−0.1706
Optimized	PIR/DSAT	0.75	1.48	0.8950	0.0203	−0.1706
Optimized	CFD	0.75	1.48	0.8153	0.0197	−0.2001
Original	CFD	0.85	2.0	0.7959	0.1026	−0.3251
Optimized	PIR/DSAT	0.85	1.65	0.7959	0.0836	−0.3251
Optimized	CFD	0.85	1.65	0.7950	0.0878	−0.3274

[a]Single-point optimization, case 1.

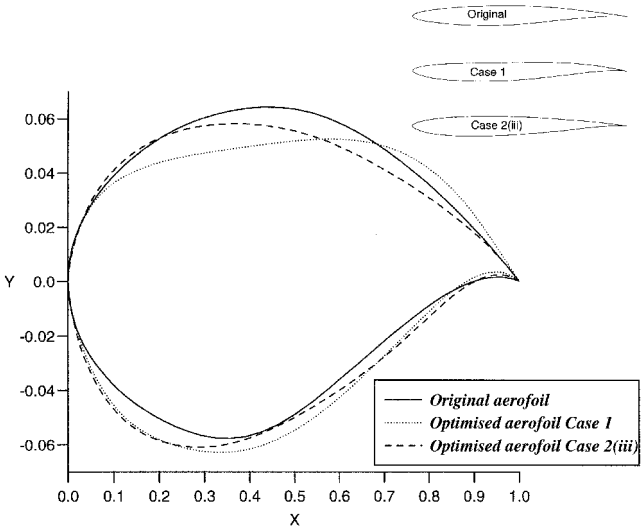

Fig. 3 Comparison of aerofoil shape between the original and the results of the optimizations at *M*0.8, with and without constraints on the maximum Mach number on both surfaces.

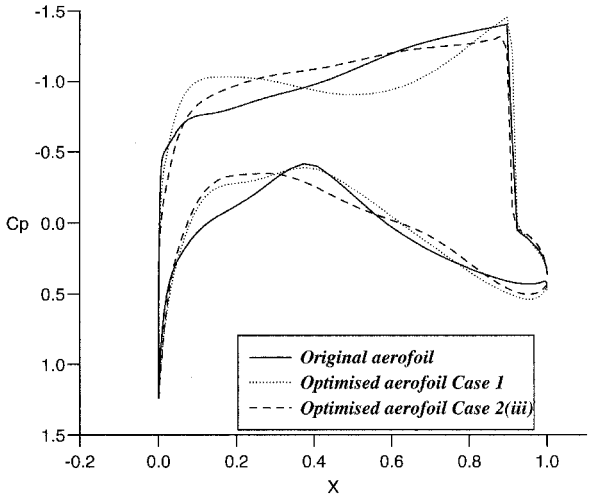

Fig. 4 Surface C_p distributions around the original aerofoil at 2 deg and optimized aerofoils at 1.41 deg (case 1) and at 1.71 deg [case 2(iii)]. Freestream Mach number is 0.8.

hances the small changes in the shape, whereas the aerofoils at the top right of the picture correspond to the actual shape. The amount of supersonic flow over the aerofoil is reduced. Because of the area constraint on the aerofoil, the lower surface is displaced further below the chord line to compensate for the reduction in area created by the modifications to the upper surface. The angle of attack in the optimized result is 1.41 deg.

Note that if the lift coefficient was required to be within 1% of the original value, then a localized PIR/DSAT calculation would be advised about the optimized point found earlier, and a further optimization performed to find a new design point at which the lift coefficient was at the original value. As more PIR data are produced than used in the DSAT continuous representation, the design space can therefore be reduced when required. This means that a second optimization run, using a smaller design space and the results of the first run as the starting point, takes a matter of seconds. The design space is reduced automatically to that containing the initial point, the DSAT code evaluates the new coefficient values by Gauss–Jordan elimination, and the optimization procedes directly.

The surface pressure distributions of the original and optimized aerofoils are presented in Fig. 4. Flowfield contour diagrams of the pressure coefficient over the original aerofoil at 2 deg and over the optimized aerofoil at 1.41 deg are shown in Figs. 5 and 6.

The pressure coefficient field over the original aerofoil at 2-deg angle of attack and freestream Mach number *M*0.75 is given in Fig. 7. When the optimization is performed at *M*0.75, the result shows a double shock over the new aerofoil shape (Fig. 8). This is related to the tendency in these inviscid cases for the optimized result to show a saddle-like C_p distribution on the upper surface (as in Fig. 4).

The results of optimizations at *M*0.75, *M*0.8, and *M*0.85 are listed in Table 1. In each case, the values obtained from the CFD code for the coefficients of lift, drag, and pitching moment are given on the original aerofoil at 2-deg angle of attack. The optimizations using the PIR/PSAT data produce a new shape and a new value for the angle of attack. Two further sets of corresponding data are provided in Table 1 at these optimized conditions: the results from the PIR/DSAT data and the result from an individual CFD run of the new design configuration.

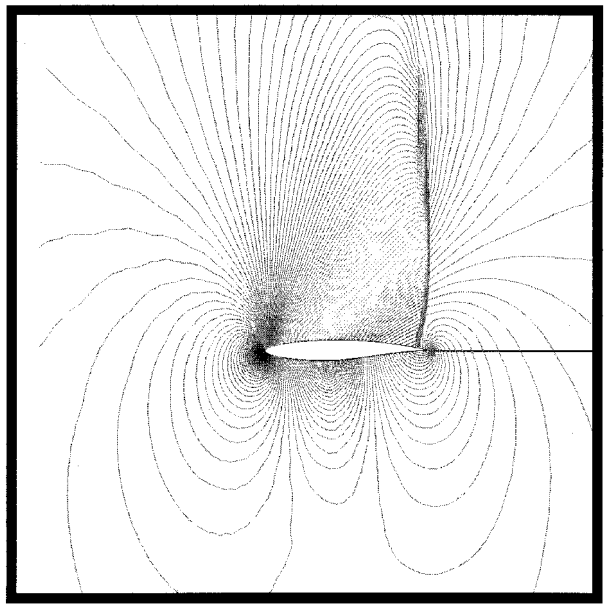

Fig. 5 Pressure coefficient contours around the original aerofoil at *M*0.8 and 2 deg. The C_p range is from −1.4034 to 1.2416.

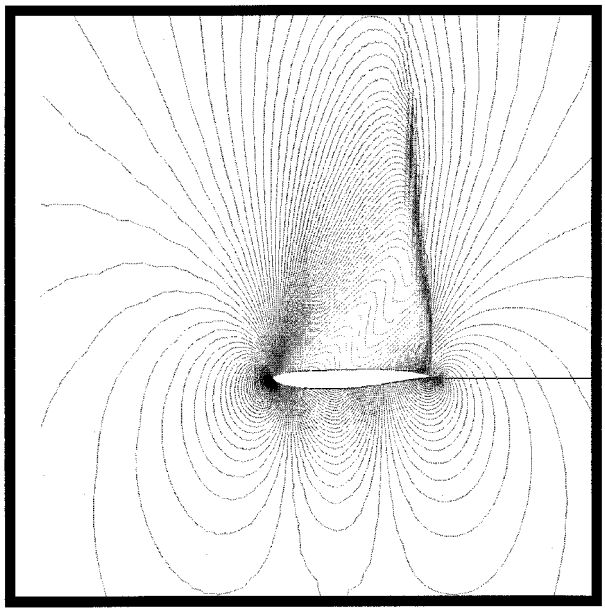

Fig. 6 Pressure coefficient contours around the optimized aerofoil at *M*0.8 and 1.41 deg. The C_p range is from −1.4277 to 1.2588.

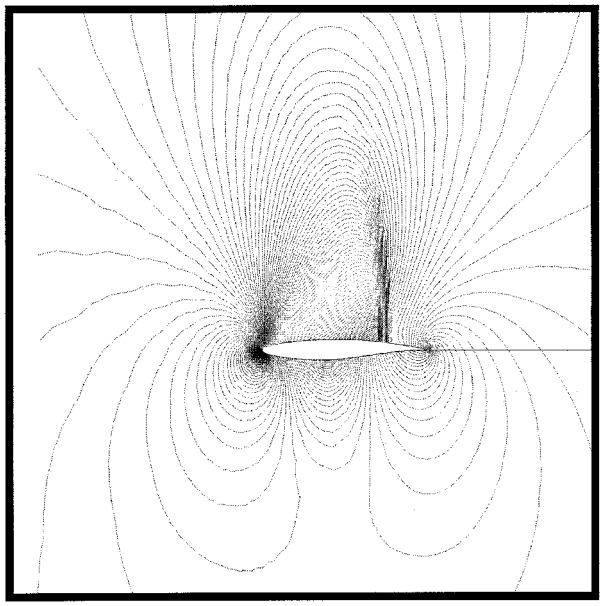

Fig. 7 Pressure coefficient contours around the original aerofoil at $M0.75$ and 2 deg. The C_p range is from −1.3832 to 1.2135.

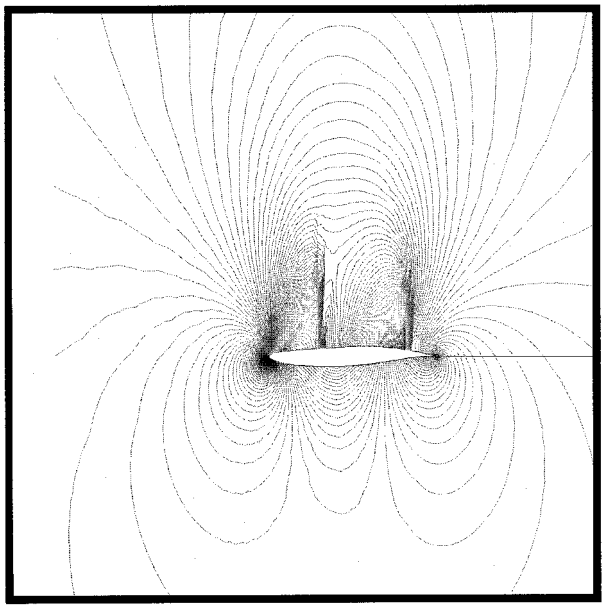

Fig. 8 Pressure coefficient contours around the optimized aerofoil at $M0.75$ and 1.48 deg. The C_p range is from −1.4176 to 1.221.

The percentage error between the CFD and the PIR/DSAT results on the optimized aerofoil are up to 11.7%, which occurs in the $M0.8$ case for the coefficient of drag. The error range can be considerably reduced when design domains of smaller extent are used.

Case 2

It is important in optimization exercises to build constraints on all of the parameters of interest into the model. For example, in the previous test case at $M0.8$, the optimization algorithm successfully found a configuration that significantly reduced the objective function. However, when the flowfields around the aerofoil are studied, it is found that the maximum Mach number (at the attachment point of the shock to the aerofoil close to the trailing edge) had increased from 1.5684 on the original aerofoil (at 2-deg angle of attack) to 1.6118 on the final configuration. One advantage of the PIR/DSAT approach is the ease with which further constraints on flow parameters can be added into the optimization problem. These constraints use the analytic representations for the parameters that are calculated at the same time as those for the coefficients of lift and drag using the PIR/DSAT method.

To illustrate how the optimization model can be enhanced, the maximum surface Mach number is constrained in the chosen region of the chord line $X = [0.4, 1.0]$ (the chord length is always one unit). Cases are run where (i) only the lower surface, (ii) only the upper surface, and (iii) both surfaces are considered. On the original aerofoil, in the required surface section, the maximum Mach number on upper and lower surfaces are 1.57 and 0.98, respectively. Hence, when constraints are applied to the Mach number, the upper bound is first set to that of the original aerofoil. The optimization code produces a solution that satisfies the constraints of case 1 and the additional constraint(s) on the surface Mach number. Further runs are then performed, reducing the upper bound on the Mach number. The code contains the option of allowing the user to define the upper bound, or of permitting the code to successively reduce the Mach number bound until no solutions are found. Hence, for each case, there are a set of results corresponding to different bounds on the Mach number(s).

A constraint is also applied to the maximum surface pressure gradient on the lower surface [case (iv)], which had a value of 10 on the original aerofoil. A constraint could not be set on the surface pressure gradient on the upper surface because of the presence of the shock discontinuity.

One result has been chosen from each of the sets of the results to cases (i), (ii), and (iv), and are presented in Table 2. One result from the case 2(iii) optimization problem is plotted in Figs. 3 and 4.

In case 2(i), the added constraint on the lower surface maximum Mach number produced a 5–6% reduction, reducing it to 0.926 compared with the values on the original and case 1 aerofoils. The lift coefficient was within 2% of the value on the original aerofoil. There is also an increase in the drag coefficient compared to the solution found in case 1, but it is still 24% lower than the original value.

Imposing a constraint on the surface Mach number on the upper surface [case 2(ii)] produced a reduction in the Mach number of 7.6% (compared with the optimized solution found in case 1), and an increase in the values of the lift and drag coefficients (the solutions considered are the CFD results on the optimized aerofoils, not the PIR/DSAT values that show the required lift coefficient). Compared with the original solution, the case 2(ii) result had an upper surface maximum Mach number 5.1% lower, a C_L value within 0.4%, and a reduction in the C_D value of 10.8%.

For case 2(iii), applying constraints to the upper and lower surfaces tends to have the same effect as seen in cases 2(i) and 2(ii). Reducing the maximum Mach number on each surface increases the value of C_D compared with the optimized solution of case 1. For example, the drag coefficient rises to 0.0661 when the Mach numbers on the lower and upper surfaces are reduced to 0.959 and 1.501, respectively (all CFD values). Comparing the CFD solutions on the optimized aerofoil with those on the original, shows that the C_L values are within 0.9%, the C_D value was reduced by 10.8%, and the maximum Mach numbers were reduced by 2.4% (lower surface) and 4.3% (upper surface). The optimized aerofoil produced in this case is shown in Fig. 3, whereas the surface pressure distribution is presented in Fig. 4. The inclusion of constraints on the Mach number has removed the saddle-shaped feature on the upper surface and produced a smoother C_p distribution over the whole aerofoil. If the pitching moment is allowed to vary, i.e., changing to an inequality constraint, solutions are possible with more than a 10% reduction in the surface Mach numbers while maintaining low C_D values.

A further constraint is then added to reduce the surface pressure gradient on the lower surface [case 2(iv)]. A number of cases were run, reducing the upper bound on the pressure gradient. As the value of the gradient reduced, the drag coefficient

Table 2 Comparison of CFD and PIR data for the optimized aerofoils at the freestream Mach number of 0.8[a]

		Constraints applied to aerofoil surface			Results				
					Aerodynamic coefficients			Surface values	
No.	Method	M(max)	$\frac{\lvert dp(\max)\rvert}{\lvert ds\rvert}$	Angle of attack, deg	C_L	C_D	C_M	M(max)	$\frac{\lvert dp(\max)\rvert}{\lvert ds\rvert}$
i	PIR/DSAT	Lower	——	1.807	0.9772	0.0596	−0.2936	0.8155	——
	CFD	Lower	——	1.807	0.9579	0.0560	−0.2957	0.9260	——
ii	PIR/DSAT	Upper	——	1.682	0.9772	0.0715	−0.2936	1.4586	——
	CFD	Upper	——	1.682	0.9731	0.0656	−0.3158	1.4891	——
iv	PIR/DSAT	Upper		1.74	0.9772	0.0852	−0.2936	0.894	——
	PIR/DSAT	Lower	Lower	——	——	——	——	1.4272	1.771
	CFD	Upper	——	1.74	0.9721	0.0775	−0.3112	0.9730	——
	CFD	Lower	Lower	——	——	——	——	1.5146	——

[a]Case 2. Constraints are imposed on the surface Mach number and in case 2(iv) on the surface pressure gradient on the lower surface. On the original aerofoil at 2 deg, the maximum Mach number on upper and lower surfaces is 1.5684 and 0.98253, respectively. The surface pressure gradient is 10.043.

increased. For example, when the PIR/DSAT data produced a decrease in the surface pressure gradient of 32%, the CFD run on the optimized aerofoil showed an increase in the drag coefficient of 23% compared with the case 1 result. However, when compared with the result on the original aerofoil, there is a 23% reduction in the drag coefficient for a 1.7% loss in lift coefficient and less than 1% change in the maximum Mach number on the upper and lower surfaces.

The result shown in Table 2 is for the case where the lower surface pressure gradient is reduced to its lowest value: 1.771 (an 82% reduction). The lift coefficient value is within 1% of the required value. The effect on the drag coefficient is to increase it above that on the original aerofoil while producing reductions in the Mach number range of 0.4 and 6% on lower and upper surfaces, respectively.

Hence, the PIR/DSAT functional representations enable a designer to use constraints on many aerodynamic parameters, and quickly evaluate the effect on the objective function and on the other parameters. This promotes an understanding of the relations between the parameters in the design space. The time taken to produce these results with the PIR/DSAT data takes no more than a few seconds on a SUN workstation. However, the real cost is the time taken to produce the CFD solution at the new design point, i.e., at the new angle of attack and modified aerofoil shape.

Multipoint Optimization Results

In this section, a three-point design problem is considered. The flight conditions on the aerofoil in cruise are such that it flies at $M0.8$, but spends part of its cruise phase in the range ($M0.75$, $M0.85$). The aim is to ensure that the combined drag coefficient is reduced to a minimum for this mission profile. Constraints are thus applied at three Mach numbers: $M0.8$, $M0.75$, and $M0.85$. The objective function used is the sum of the weighted drag coefficients at each Mach number, i.e.,

$$\text{Objective function} = \sum_{i=1}^{3} w_i C_{D_i} \tag{1}$$

$$\sum_{i=1}^{3} w_i = 1 \tag{2}$$

where w is $0 \le w \le 1$. The index $i = 1$–3 corresponds to values in the $M0.75$, $M0.80$, and $M0.85$ cases, respectively. The optimizer requires information about the design space at each of these three Mach numbers. Hence, for each Mach number, the initial design point is taken as the original aerofoil at 2-deg angle of incidence. A converged flowfield around this configuration is obtained using the CFD code and then the PIR/DSAT algorithms are used to find the values for the coefficients of the analytical functions (describing the aerodynamic coefficients, surface pressure gradients, etc.) as a function of the shape variables and angle of attack. The variation in the coefficient of lift as a function of angle of attack for the range 0–4 deg on the original aerofoil is plotted in Fig. 9 for each of the three Mach numbers considered.

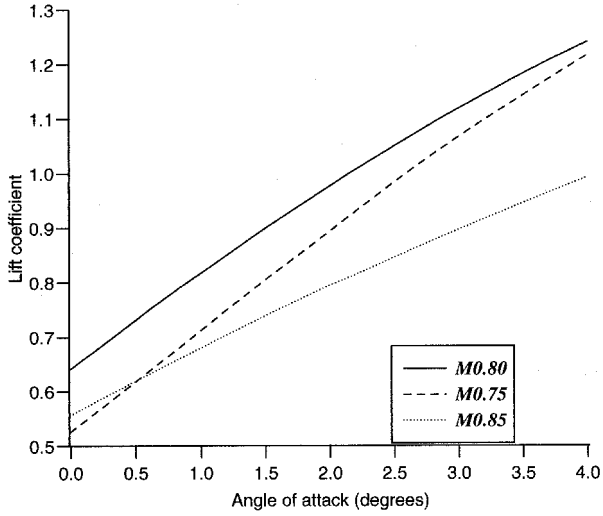

Fig. 9 Coefficient of lift against angle of attack on the original aerofoil for the three Mach numbers considered in the optimization.

Case 3

Constraints are set on the aerofoil area, lift coefficient, and pitching moment. The aerofoil area is taken as the same as the original aerofoil, and is the same at each Mach number. The pitching moment C_M on the optimized aerofoil is required to be within 2% of the value on the original aerofoil; i.e., for each Mach number

$$C_{M_i}^{\text{opt}} = k \cdot C_{M_i}^{\text{orig}}$$

where constant $k = (0.98, 1.02)$ and $i = 1,3$. Finally, the coefficients of lift are set such that the lift at each Mach number design point is the same as at the original design point of $M0.8$. So,

$$C_{L_i}^{\text{opt}} M_i^2 = C_L^{\text{orig}} \text{ (at } M0.8) \cdot (0.8)^2$$

where $i = 1, 3$.

Hence, at $M0.8$, the original aerofoil satisfied the constraints. However, the design points at $M0.75$ and $M0.85$ did not satisfy the constraints on the coefficient of lift. As the coefficient of

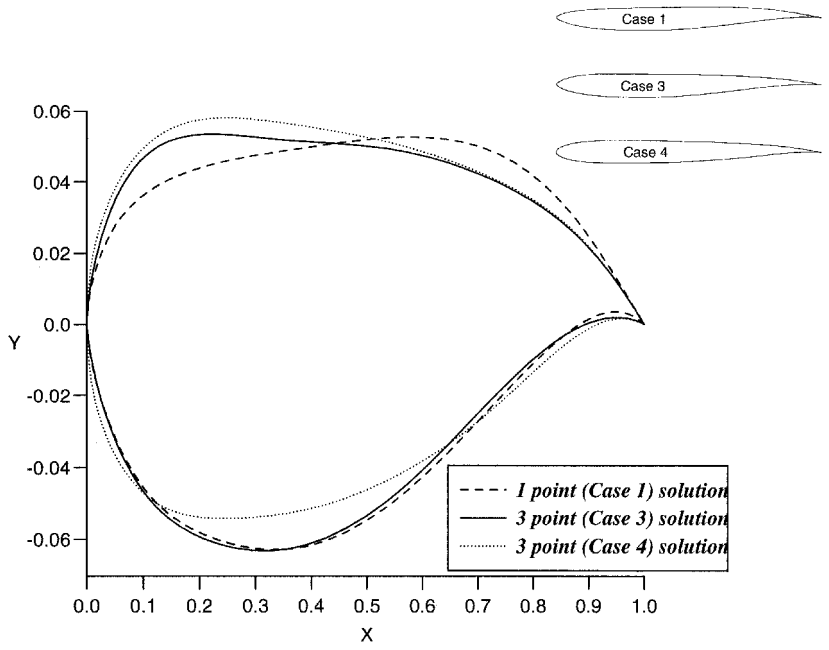

Fig. 10 Comparison of aerofoil shape between the single-point optimization result and the results of the three-point optimizations with and without constraints on the Mach number range on the lower aerofoil surface.

lift on the original aerofoil at $M0.8$ is larger than at $M0.75$ and $M0.85$ at any angle of attack, it is expected that the angle of attack on the optimized aerofoil at $M0.8$ is smaller than the angles at $M0.75$ and $M0.85$ to preserve the same value of the lift.

The aerofoil shape and angle of attack are allowed to move within the same domain as used for the PIR/DSAT run (this design space is the same size as used in the single point optimizations). As in the case of single-point optimization, many points throughout the domain are used to start the optimization procedure.

Because of the requirement that the specified values of the lift coefficient must be achieved at the three Mach numbers, the shape of the optimized aerofoil is quite different from the single point case, even when the weighting parameter on the $M0.8$ values is at 100%. This result implies that for cases where aerodynamic coefficient values greatly vary, even for a minute part of the mission, a single-point optimization result may give misleading information on the optimum shape of the aerofoil. Figure 10 compares the aerofoil shape from the single-point and triple-point optimization cases with constraints on C_L, C_M, and area (and a triple-point case where constraints have been applied to the Mach number range on the aerofoil's lower surface). These optimized aerofoils have flatter upper surfaces than the original aerofoil (Fig. 3). The surface C_p distributions at $M0.8$ on the aerofoils in Fig. 10 are presented in Fig. 11. In the multipoint cases, the weighting parameters used at $M0.8$, $M0.75$, and $M0.85$ are 0.9, 0.05, and 0.05, respectively. Only minor variations in these results are obtained for weighting parameters at $M0.8$ between 0.6 and 1. From Fig. 11 it is seen that the optimization procedure preserved the position of the shock in both single-point and multipoint optimization examples. This is not the case for other examples considered by the authors.

The values for the angle of attack and the aerodynamic coefficients at each Mach number are listed in Table 3. The values on the original aerofoil are provided, as are two sets of results on the optimized shape (which is the same as the three Mach numbers). One set of results comes from the optimization run using the PIR/DSAT data. The other set of results comes from running the CFD code at the optimized conditions for the design variables. The differences between the CFD and PIR/DSAT results are 3.5–5% for the coefficient of lift, and

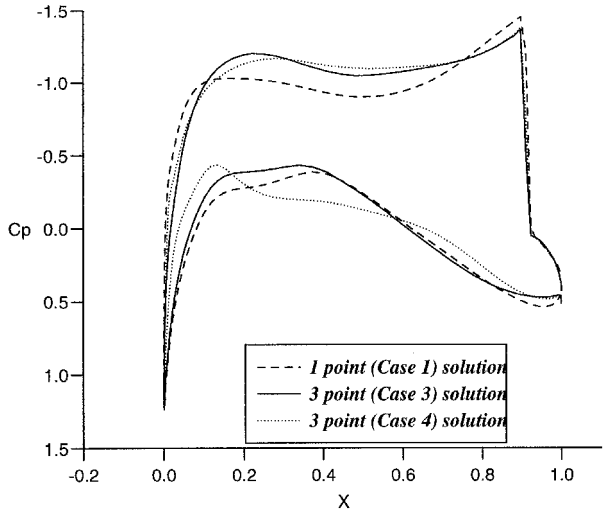

Fig. 11 Surface pressure coefficient distributions around the single-point optimized aerofoil (case 1) at 1.41 deg and the three-point optimized aerofoils from case 3 at 1.17 deg and case 4 at 1.13 deg. In the latter case, constraints are applied to the Mach number range on the lower aerofoil surface. Freestream Mach number is 0.8.

between 0.29 and 21.5% for the coefficient of drag. The latter discrepancy for C_D at $M0.75$ suggests that the results would be improved by considering a smaller domain, particularly in terms of reducing the range for the angle of attack.

The time taken to produce this optimization result (in 24-parameter space) is equivalent to less than six CFD runs (three CFD runs on the original aerofoil at $M0.75$, $M0.8$, and $M0.85$, and three PIR/DSAT runs about these initial design points). Three additional CFD runs are needed to check the validity of the solution. With a total cost of around one day's usage on a DEC 3000/600 workstation, the turnaround is rapid enough to permit another investigation of a more localized domain around the optimized solution.

Table 4 lists the values of the objective function [Eq. (1)] for the case considered in Table 3. The value of the objective function, initially at 0.07252 (and not a solution to the opti-

Table 3 Comparison of CFD and PIR data for the original and optimized aerofoils[a]

Aerofoil shape	Method	Freestream Mach number	Angle of attack, deg	Aerodynamic coefficients		
				C_L	C_D	C_M
Original	CFD	0.80	2.0	0.9772	0.0735	−0.2936
Optimized	PIR/DSAT	0.80	1.17	0.9772	0.0617	−0.2878
Optimized	CFD	0.80	1.17	0.9422	0.0570	−0.3045
Original	CFD	0.75	2.0	0.8950	0.0245	−0.1706
Optimized	PIR/DSAT	0.75	2.93	1.1119	0.0525	−0.1740
Optimized	CFD	0.75	2.93	1.1524	0.0669	−0.2397
Original	CFD	0.85	2.0	0.7959	0.1026	−0.3251
Optimized	PIR/DSAT	0.85	2.04	0.8657	0.1055	−0.3186
Optimized	CFD	0.85	2.04	0.8245	0.1052	−0.3333

[a]Case 3, multipoint optimization.

Table 4 Comparison of CFD and PIR data for the original and optimized aerofoils[a]

Aerofoil	Method	Objective function
Original	CFD	0.07252
Optimized	PIR/DSAT	0.06339
Optimized	CFD	0.05989

[a]Case 3, multipoint optimization case.

mization problem) dropped to 0.05989 at the optimized design point, a reduction of 17.4%.

Case 4

As in the single-point optimization case 2, additional constraints are easily added into the multipoint problem because the data are already available. Constraints are set on the maximum Mach number on the lower surface of the aerofoil in the region X = [0.4, 1.0] at M0.75, M0.8, and M0.85. The optimization runs take only a few seconds to produce a result, and then additional CFD runs must be made at the new design point. Although the aim is to minimize the objective function, a number of cases are run in which different upper bounds are set for the maximum Mach number on the lower surfaces.

From the CFD solutions on the optimized result of case 3, the maximum Mach numbers are 0.837, 0.999, and 1.299, respectively. (On the original aerofoil, the values are 0.879, 0.982, and 1.315.) The greater the reduction in these maximum Mach number values, the greater the increase in the value of the objective function. For example, reducing the maximum Mach numbers to 0.793, 0.996, and 1.169 (a percentage decrease with respect to the case 3 solution of 5.26, 0.3, and 10.01%), increases the objective function value to 0.06285, a rise of 4.94%. The resulting aerofoil shape and C_p distribution are shown in Figs. 10 and 11.

Conclusions

A three-stage strategy has been adopted to produce a database of aerodynamic and geometric variables describing a multiparameter design space for aerofoil design. There are numerous advantages in using this three-program approach over the conventional approach of CFD runs. There is a considerable decrease in the time required to investigate a design space, and the user has, if wished, control over the directions taken in this environment. There is also an option within the codes to automatically sweep through the chosen design space. The scheme does not require any a priori knowledge of dynamic, e.g., a pressure distribution, or geometric features. There exists the opportunity to set constraints on aerodynamic quantities on the aerofoil surface and on the geometry of the aerofoil, and the codes have the ability to investigate the sensitivity of a chosen design within user-defined limits, e.g., the effect of reducing the maximum surface Mach number. Finally, the system incorporates more design variables and constraints with ease.

Optimization examples have shown how modifications to a design can be successively built up as knowledge is acquired on the characteristics of each design and, thus, additional constraints can be added to the problem.

Acknowledgments

This work was carried out with Sowerby Research Centre applied research funding. We wish to thank C. Warsop, R. A. Hughes, and A. R. B. Gould for their helpful comments.

References

[1]Rogers, D. F., and Adams, J. A., *Mathematical Elements for Computer Graphics*, McGraw–Hill, New York, 1990.

[2]Topliss, M. E., Toomer, C. A., and Hills, D. P., "Rapid Design Space Approximation for Two-Dimensional Transonic Aerofoil Design," *Journal of Aircraft*, Vol. 33, No. 6, 1996, pp. 1101–1108.

[3]Wada, Y., and Liou, M.-S., "A Flux Splitting Scheme with High-Resolution and Robustness for Discontinuities," AIAA Paper 94-0083, Jan. 1994.

[4]Taylor, III, A. C., Hou, G. J.-W., and Korivi, V. M., "A Methodology for Determining Aerodynamic Sensitivity Derivatives with Respect to Variation of Geometric Shape," *AIAA Journal*, Vol. 30, 1992, No. 10, pp. 2411–2419.

[5]Korivi, V. M., Taylor, III, A. C., Newman, P. A., Hou, G. J.-W., and Jones, H. E., "An Approximately Factored Incremental Strategy Calculating Consistent Discrete Aerodynamic Sensitivity Derivatives," *Journal of Computational Physics*, Vol. 113, No. 2, 1994, pp. 336–346.

[6]Baysal, O., Eleshaky, M. E., and Burgreen, G. W., "Aerodynamic Shape Optimisation Using Sensitivity Analysis on Third-Order Euler Equations," *Journal of Aircraft*, Vol. 30, No. 6, 1993, pp. 953–961.

[7]Pandya, M. J., and Baysal, O., "Gradient-Based Aerodynamic Shape Optimization Using ADI Method for Large-Scale Problems," AIAA Paper 96-0091, Jan. 1996.

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

[9]"NAg E04UCF," *NAg Mark 16*, Numerical Algorithms Group, Ltd., Oxford, England, UK, 1993.