

CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation

A. M. Morris, C. B. Allen^{*,†} and T. C. S. Rendall

University of Bristol, Avon BS8 1TR, U.K.

SUMMARY

A novel domain element shape parameterization method is presented for computational fluid dynamics-based shape optimization. The method is to achieve two aims: (1) provide a generic ‘wrap-around’ optimization tool that is independent of both flow solver and grid generation package and (2) provide a method that allows high-fidelity aerodynamic optimization of two- and three-dimensional bodies with a low number of design variables. The parameterization technique uses radial basis functions to transfer domain element movements into deformations of the design surface and corresponding aerodynamic mesh, thus allowing total independence from the grid generation package (structured or unstructured). Independence from the flow solver (either inviscid, viscous, aeroelastic) is achieved by obtaining sensitivity information for an advanced gradient-based optimizer (feasible sequential quadratic programming) by finite-differences.

Results are presented for two-dimensional aerofoil inverse design and drag optimization problems. Inverse design results demonstrate that a large proportion of the design space is feasible with a relatively low number of design variables using the domain element parameterization. Heavily constrained (in lift, volume, and moment) two-dimensional aerofoil drag optimization has shown that significant improvements over existing designs can be achieved using this method, through the use of various objective functions. Copyright © 2008 John Wiley & Sons, Ltd.

Received 21 September 2007; Revised 17 December 2007; Accepted 1 January 2008

KEY WORDS: aerodynamic optimization; domain element; parameterization; mesh deformation; radial basis functions; feasible sequential quadratic programming

*Correspondence to: C. B. Allen, University of Bristol, Avon BS8 1TR, U.K.

†E-mail: c.b.allen@bristol.ac.uk

Contract/grant sponsor: Agusta–Westland

Contract/grant sponsor: EPSRC (The UK Engineering and Physical Sciences Council); contract/grant number: GR/S61294

Contract/grant sponsor: U.K. MoD Joint Grant Scheme (JGS)

Contract/grant sponsor: University of Bristol

1. INTRODUCTION

Within the aerospace industry, individual companies are continually striving for increased competitive strength, and this is frequently achieved by an advancement in design. While some factors are difficult to determine, for example, aesthetics, many of the factors associated with a design's performance can be quantified and often predicted numerically. The injection of numerical optimization into the design process along with advances in computer hardware has allowed engineers and researchers to attempt problems with increased complexity and to maximize a design's predicted performance within pre-specified limits and constraints.

However, the quality and applicability of results gained by numerical optimization techniques are inherently dependent on the fidelity of the analysis tool used. Computational fluid dynamics (CFD) is at the forefront of aerodynamic analysis capabilities, and optimization algorithms can be applied to such codes with the aim of improving the numeric value of an objective function. The objective function in this sense is defined as an indicator of aerodynamic performance and is typically based on drag, lift and moment coefficients. High-fidelity aerodynamic shape optimization has previously been applied to two-dimensional aerofoil sections [1–4] and aspects of three-dimensional aircraft [5–11], 3D aeroelastic aircraft [12], and even proof-of-concept optimization of rotors [13, 14].

Within the UK Rotorcraft DARP (Defence and Aerospace Research Partnership), a collaborative numerical simulation research theme exists between the Universities of Bristol and Liverpool and Agusta–Westland Helicopters. Liverpool have provided the flow solver, an unsteady implicit, Navier–Stokes, parallel multiblock code [15], and Bristol the mesh generation tools [16]. The objective of this research theme is to use CFD in the main industrial design loop at Agusta–Westland, and one aspect of this is the development of aerodynamic optimization tools, to be wrapped around the flow solver and grid generation software already under development.

There are three critical areas involved in the development of the optimization tool: suitable parameterization of the aerodynamic body, including effective surface and mesh deformation; computation of the required sensitivities for each parameter; and effective stable optimization of the design. Each area is reviewed and described in detail in the following sections. The tools required are developed and applied to a two-dimensional aerofoil, including a study of sensitivity dependence on perturbation sizing and grid density. Inverse design results illustrate a large and appropriate region of design space is achievable with the parameterization method, and heavily constrained aerofoil optimization shows significant drag reductions that can be achieved over initial designs. The use of various objective functions to achieve this drag reduction is also considered.

2. PARAMETERIZATION OPTIONS

The choice of parameterization method has implications to both the success and the computational cost of an optimization. With regard to success, the method must allow sufficient free-form design such that any likely optimum design that may exist is achievable, and inverse design methods are usually applied to gauge the range of design space achievable by any method. To keep computational cost to a minimum, it is essential to use both an efficient optimization strategy and have as few design variables as possible: especially if sensitivity information is to be evaluated by finite-difference.

Parameterization methods used for CFD shape optimization either parameterize the design geometry from which a mesh is generated or they parameterize the aerodynamic mesh itself. Geometry parameterization methods uniquely define a surface geometry by the values of the design variables. These methods are inherently linked with the grid generation package and optimization of a design requires automatic grid generation tools. Methods of this nature include partial differential equation (PDE) methods [17, 18], polynomial or spline [19], CAD and recently CST [20, 21] methods.

The use of polynomial or spline methods are common throughout the field of shape optimization [19, 22]. An initial design is approximated by a spline or polynomial function and design variables are either the values of the polynomial coefficients or coordinate location of the spline control points respectively. Castonguay and Nadarajah [23] determined 32 control points were the minimum sufficient for aerofoil design and optimization using spline control points, and even then, Reuther [24] has shown that optimization may still be prematurely halted due to the creation of wavy design surfaces. Another drawback can be the excessive number of design variables required to represent a three-dimensional geometry correctly. The PARSEC [25] method forms another subset of a polynomial method, with two separate polynomials defining the upper and lower surfaces. This, however, can be a disadvantage as there is no constraint on the two polynomials crossing, producing figure-of-eight aerofoils during optimization [26], and only a limited design space is achievable [23].

CAD methods utilize the well-established commercial packages to create design surfaces; however, parameterization of complex geometries can still be a challenging task even for modern CAD systems.

The class function/shape function transformation (CST) [20, 21] method is a recent addition to parameterization methods. The method is fundamental and effectively a limitless design space can be achieved by composition of entirely smooth geometries. The method has been recently investigated by Nadarajah and Castonguay [27] and found to be effective at aerofoil parameterization but, though it can parameterize three-dimensional geometries with only a few design variables, it proved not as successful as either spline or mesh point methods.

Methods that parameterize the grid are generally independent of the grid generation package. The grid topology remains unaltered throughout the optimization, but the mesh requires a deformation algorithm. Grid generation is still a major and time-consuming challenge in industry and these types of methods allow the use of previously generated grids for optimization—a significant advantage. Parameterization methods of this nature include, discrete [24, 28, 29], analytical, basis vector [30], free-form deformation algorithms (FFD) [31], and domain element methods [32].

The discrete method uses mesh points themselves, i.e. the (x, y, z) locations of the design surface mesh points to form the design variables [24, 28, 29, 33]. There is no real limit on the attainable geometry produced using this method. This approach is easy to implement, although the number of design points can become very large, and maintaining smooth geometries can be difficult as the surface is not defined by a continuous function. Sensitivity information can be difficult to obtain if using a finite-difference methodology as the independent displacement of a single mesh point may cause the flow solver to become ill conditioned.

The basis vector method [30] uses several design vectors based on proposed grids to represent the design shape, and free-form design methods are derived from soft object animation (SOA) methods [31], where computer graphics programs are used to morph images and deform designs to achieve the desired geometry.

The analytical parameterization method provides a compact formulation for the representation of two-dimensional aerofoils. The method is commonly used for shape optimization and forms the

basis for some commercial two-dimensional aerodynamic shape optimization programs, such as LINDOP [34] with MSES [35]. The method is centred around adding shape functions (analytical functions) linearly to the initial design. Hicks–Henne ‘bump’ functions [1, 2, 9, 24] also form a subset of this type of method and again Castonguay and Nadarajah [23] have shown difficulty in recreating many aerofoils by inverse design using this parameterization method.

Domain element-type techniques have been seen in commercial shape optimization software packages, such as SCULPTOR [32]; however, little information on this particular parameterization method is available due to its commercial nature. A comprehensive review of available parameterization methods is presented by Samareh [36, 37].

The parameterization method developed in the current work, and presented here, involves a domain element that is wrapped around the exterior of a design. At the centre of this parameterization technique is a multivariate function kernel using radial basis functions (RBFs). This provides an inverse mapping between the domain element, the surface geometry and the locations of grid points in the volume mesh. The result is an efficient parameterization technique that is independent from the grid generation package. The inverse mapping is only required once for the initial design as the values of the parametric coordinates of the grid points with respect to the domain element remain constant throughout the optimization. Updates to the geometry and the corresponding mesh are provided simultaneously by the application of multivariate interpolation; this is extremely fast and efficient and results in very high-quality mesh deformation [38]. The long-term objective of the optimization technique presented is to allow aerodynamic optimization of rotors, with a minimum number of design variables while remaining independent of both flow solver and grid generation package. The parameterization technique developed has several key advantages:

- (1) The number of design variables required to allow free-form design can be very low with two-dimensional inverse design problems demonstrating that 22 is sufficient (see later).
- (2) Design variables can range in fidelity from the coordinate of a single domain element location, providing a detailed surface geometry change, to a gross deformation such as chord or twist design variable. This can allow designers to choose the fidelity of the optimization, to change only large-scale parameters or this can be combined alongside detailed aerofoil profile change.
- (3) The parameterization technique is independent of the initial geometry, with the domain element positioned automatically according to the solid surface. This can allow very complex or multi-element geometries to be parameterized.
- (4) The parameterization method has a low number of control points (in this case the number of domain element points), and this allows the possibility of incorporating a high-quality interpolation to deform the surface and mesh directly from these points. This would be impractical if all the surface mesh points were control points, for example.
- (5) Updates to the surface geometry and the required volume mesh are independent of the grid generation package. Grid generation is still a major and time-consuming challenge in industry and the domain element method allows the use of previously generated grids for optimization. Both structured and unstructured grids are equally applicable.
- (6) The interpolation dependence needs to be computed only once, prior to any simulation. Updates to the geometry and the corresponding volume mesh are then provided simultaneously by matrix–vector multiplication; this is extremely fast and efficient and results in high-quality mesh deformation.

3. OPTIMIZATION APPROACH

Owing to the volume of industrial and academic effort, it is expected that modern aerodynamic designs lie in proximity of an optimum solution (though not necessarily the global optimum). Therefore, to search an entire design space in pursuit of a global optimum design shape may be unnecessarily expensive to gain an improvement over current designs. Additionally, if high-fidelity CFD is used as the analysis tool and the optimization is extended to three-dimensional shapes, the sheer cost and volume of function evaluations required by an evolutionary-type algorithm may render such optimization as unfeasible with current and available computational power.

Constrained gradient-based optimizers used to minimize an objective function are fast and efficient at providing solutions to local optimization problems [39, 40]. However, they can still require the evaluation of many functions and gradients.

Each function evaluation requires the steady-state solution of a different design, but the gradient information can be obtained in a variety of ways. The gradient (or sensitivity) is defined as the derivative of a response (objective function) with respect to a design variable. The finite-difference methodology is the most traditional, where the sensitivity for each design variable is easily obtained by the relative change in the value of steady-state objective function due to a geometric perturbation. The cost of evaluating all the sensitivities with this method is approximately linear with the number of design variables. Hence, computation can prove expensive for a large number of design variables, but due to the nature of the methodology, it is independent of the flow solver used and implementation is straightforward, although perturbation size and sensitivity scaling can be an issue. It also lends itself easily to parallelization; in a data sense, rather than domain decomposition sense. The complex-step methodology pioneered in 1967 by Lyness and Moler [41], which is conceptually similar to the finite-difference approach except the perturbation is made to the i th component of the design variable. Subtractive cancellation is avoided, and the perturbation size is less critical. However, the cost is of the same magnitude as that of a finite-difference methodology, and an optimization becomes inherently linked to the flow solver. The most computationally efficient method lies in the solution of an adjoint equation [28, 42] where, after an initial evaluation of the objective function for a given geometry, the sensitivities of any number of design variables can be obtained from a single adjoint solution. Although an adjoint approach to optimization has been demonstrated by several researchers to be effective, and in some cases, efficient way of evaluating sensitivities, development and implementation of such a methodology is complex. This requires much development of a chosen specific code, which is not practical within the DARP collaboration, and is mainly limited to inviscid analysis due to convergence issues for viscous adjoint solvers [43].

The development of generic optimization tools encompassing a wide range of applicability has been the principle aim of the current research. This has ultimately required the use of a finite-difference technique for evaluating sensitivities to enable independence from the flow solver. This easily allows scope to upgrade the CFD flow solver or grid generation package as they develop.

3.1. Gradient-based Optimizer feasible sequential quadratic programming

When considering practical and relevant optimization of aerodynamic performance of a solid body (airfoil/ wing /rotor blade), there are usually constraints that need to be imposed (minimum thickness, minimum volume, minimum lift, maximum moment, etc). Unconstrained optimizations can incorporate constraints by using a penalty function for constraints that are near or beyond

the constraint boundary. However, these methods are now considered inefficient and have been replaced by methods that focus on the solution of the Kuhn–Tucker equations [44].

The aim of an optimization process is to obtain the values of a set of design variables, $\underline{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_n)^T$, that in some way can be defined as optimal. An objective function(s), determined by the numeric values of the design variables, $J(\underline{\alpha})$, is to be minimized and can be subject to equality constraints, inequality constraints and/or parameter bounds.

A general constrained problem description is stated as

$$\begin{aligned} & \text{minimize } J(\underline{\alpha}) \\ & \text{subject to: Equality constraints } G_i(\underline{\alpha}) = 0, \quad i = 1, \dots, m_e \\ & \quad \quad \quad \text{Inequality constraints } G_i(\underline{\alpha}) \leq 0, \quad i = m_e + 1, \dots, m \\ & \quad \quad \quad \text{Parameter bounds } \underline{\alpha}_l \leq \underline{\alpha} \leq \underline{\alpha}_u \end{aligned} \quad (1)$$

where the vector function $\underline{G}(\underline{\alpha})$ returns a vector of length m containing the values of the equality and inequality constraints evaluated at $\underline{\alpha}$.

The Kuhn–Tucker equations are necessary conditions for optimality for a constrained optimization problem and can be stated as (in addition to the original constraints of Equation (1))

$$\nabla J(\underline{\alpha}^*) + \sum_{i=1}^m \lambda_i^* \cdot \nabla G_i(\underline{\alpha}^*) = 0 \quad (2)$$

$$\lambda_i^* \cdot G_i(\underline{\alpha}^*) = 0, \quad i = 1, \dots, m \quad (3)$$

$$\lambda_i^* \geq 0, \quad i = m_e + 1, \dots, m \quad (4)$$

where $\underline{\alpha}^*$ refers to the optimum values of $\underline{\alpha}$.

The solution of these equations forms the basis of nonlinear programming algorithms. Constrained quasi-Newton methods guarantee superlinear convergence by accumulating second-order information relating to the Kuhn–Tucker equations using a quasi-Newton updating procedure, i.e. at each major iteration, an approximation is made of the Hessian of the Lagrangian function. This is then used to generate a quadratic programming (QP) subproblem where the solution is used to form a search direction for a line search procedure. This forms the basis of the sequential quadratic programming (SQP) algorithm, since a QP subproblem is solved at each major iteration.

SQP methods represent the forefront of nonlinear programming techniques. Schittkowski [45] has implemented and tested an SQP algorithm that outperforms all other tested methods in terms of efficiency, accuracy and percentage of successful solutions, over a large number of test problems. The feasible sequential quadratic programming (FSQP) algorithm used in the current research has been written by Zhou *et al.* [46], and Zhou and Tits [47] and Panier and Tits [39]. The feasibility aspect of the optimizer relates to a generated design satisfying all constraints, i.e. if an initial design does not satisfy the specified constraints, the optimizer first achieves a satisfactory design, and then all subsequent iterates generated also satisfy all constraints simultaneously. This particular algorithm has been implemented across a wide range of optimization problems, most relevant and notable is the work of Qin and Le Moigne [5–8, 43] where the algorithm is used for CFD constrained optimization of a blended wing body using an inviscid adjoint solver to obtain the sensitivities. Prior to implementation, the algorithm has been numerically validated against two test cases proposed in [46].

Owing to the generic nature of the optimization tool developed and the use of a finite-difference technique for gradient evaluation, the cost of the optimization process is ultimately determined by the flow solver cost. Each evolution of the optimization requires one flow solution for calculation of objective and constraint functions, and two flow solutions per design variable to evaluate gradient information. On top of this, there may be several more flow solutions required to evaluate a suitable step length to take in the search direction (based on gradient information) or re-evaluation if constraints are not satisfied. The update to the geometry surface and volume mesh is computationally very cheap and almost insignificant when compared with a flow solution calculation.

4. PARAMETERIZATION FORMULATION

A domain element is placed around the initial body to be optimized. The element currently consists of 16 nodes for a two-dimensional aerofoil section. A method has been developed to locate domain element nodes automatically for aerodynamic geometries (aerofoils, wings and rotors [48]) based on the surface geometry. The automatic location method and the development of design variables based on those node locations were developed by the application of inverse design methods to various aerofoil geometries [48]. The node coordinates can be adjusted or positioned manually if so desired, or to aid parameterization of complex geometries, or when considering a specific design situation. Figure 1(a) shows an example of parameterization for a RAE2822 aerofoil. A multivariate interpolation method using RBFs has also been developed. The global dependence between the domain element nodes and the aerodynamic mesh points is evaluated and translates a deformation of the element due to a design variable change to smoothly alter the aerodynamic shape and its corresponding volume mesh. Using the domain element method, only an initial mesh of the original design is required to allow optimization.

Although RBFs have been used for mesh deformation techniques in shape optimization [49] and structural–aerodynamic coupling [38, 50], their use in the current research is not to strictly or uniquely parameterize an aerodynamic shape, but to allow that initial shape and its corresponding aerodynamic mesh to deform in a relevant way to enable non-code and mesh-type specific aerodynamic shape optimization. The interpolation method developed here requires no connectivity information and can therefore be applied equally well to either structured or unstructured grid

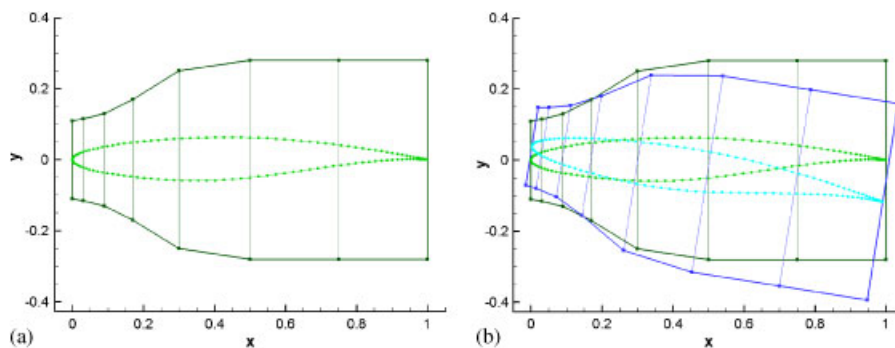


Figure 1. Domain element around (a) RAE2822 and (b) perturbed aerofoil.

topologies. Domain element points and volume mesh points are simply treated as independent point clouds with the dependence matrix computed only once.

The general theory of RBFs is presented by Buhmann [51] and Wendland [52], and the method used is comprehensively detailed in Rendall and Allen [38].

4.1. Formulation

The solution of an interpolation problem using RBFs begins with the form of the required interpolation

$$s(\mathbf{x}) = \sum_{i=1}^{i=N} \beta_i \phi(\|\mathbf{x} - \mathbf{x}_i\|) + p(\mathbf{x}) \quad (5)$$

where $s(\mathbf{x})$ is the function approximated, the index i identifies the centres for the RBFs (the domain element nodes in this case), while \mathbf{x}_i is the location of that centre. For $p(x)$, a linear polynomial is used so that translation and rotation are recovered, i.e. if the domain element translates or rotates then the corresponding design shape and its mesh deform retaining exact translation and rotation. The coefficients β_i are found by requiring exact recovery of the original function; this means if the domain element returns to its initial configuration then the design shape returns to its original form. When the polynomial term is included, the system is completed by the additional requirement (sometimes referred to as a 'side condition')

$$\sum_{i=1}^{i=N} \beta_i q(\mathbf{x}) = 0 \quad (6)$$

This applies for all polynomials $q(\mathbf{x})$ with degree less than or equal to that of $p(\mathbf{x})$.

For easy computational treatment that maximizes the use of matrix multiplication, the problem is written in the following fashion. Exact recovery of the centres gives (using only linear polynomial terms)

$$\mathbf{X}_{\text{DE}} = \mathbf{C} \mathbf{a}_x \quad (7)$$

$$\mathbf{Y}_{\text{DE}} = \mathbf{C} \mathbf{a}_y \quad (8)$$

$$\mathbf{Z}_{\text{DE}} = \mathbf{C} \mathbf{a}_z \quad (9)$$

where

$$\mathbf{X}_{\text{DE}} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \mathbf{x}_{\text{DE}} \end{pmatrix}, \quad \mathbf{x}_{\text{DE}} = \begin{pmatrix} x_{\text{DE}_1} \\ \vdots \\ x_{\text{DE}_N} \end{pmatrix}, \quad \mathbf{a}_x = \begin{pmatrix} \gamma_0^x \\ \gamma_x^x \\ \gamma_y^x \\ \gamma_z^x \\ \beta_{\text{DE}_1}^x \\ \vdots \\ \beta_{\text{DE}_N}^x \end{pmatrix} \quad (10)$$

(analogous definitions hold for \mathbf{Y}_{DE} and \mathbf{Z}_{DE} and their \mathbf{a} vectors)

$$\mathbf{C} = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 & \cdots & 1 \\ 0 & 0 & 0 & 0 & x_{DE_1} & x_{DE_2} & \cdots & x_{DE_N} \\ 0 & 0 & 0 & 0 & y_{DE_1} & y_{DE_2} & \cdots & y_{DE_N} \\ 0 & 0 & 0 & 0 & z_{DE_1} & z_{DE_2} & \cdots & z_{DE_N} \\ 1 & x_{DE_1} & y_{DE_1} & z_{DE_1} & \phi_{DE_1DE_1} & \phi_{DE_1DE_2} & \cdots & \phi_{DE_1DE_N} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{DE_N} & y_{DE_N} & z_{DE_N} & \phi_{DE_NDE_1} & \phi_{DE_NDE_2} & \cdots & \phi_{DE_NDE_N} \end{pmatrix} \quad (11)$$

with

$$\phi_{DE_1DE_2} = \phi(\|\mathbf{x}_{DE_1} - \mathbf{x}_{DE_2}\|) \quad (12)$$

indicating the basis function evaluated on the distance between DE_1 and DE_2 , and subscript DE representing a domain element control point. To locate the deformed aerodynamic mesh points, the following matrix must be formed, where subscript a indicates an aerodynamic node:

$$\mathbf{A} = \begin{pmatrix} 1 & x_{a_1} & y_{a_1} & z_{a_1} & \phi_{a_1DE_1} & \phi_{a_1DE_2} & \cdots & \phi_{a_1DE_N} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{a_N} & y_{a_N} & z_{a_N} & \phi_{a_NDE_1} & \phi_{a_NDE_2} & \cdots & \phi_{a_NDE_N} \end{pmatrix} \quad (13)$$

The positions of the aerodynamic mesh points, given by the vectors \mathbf{X}_a , \mathbf{Y}_a and \mathbf{Z}_a , are

$$\mathbf{X}_a = \mathbf{A}\mathbf{a}_x = \mathbf{A}\mathbf{C}^{-1}\mathbf{X}_{DE} = \mathbf{H}\mathbf{X}_{DE} \quad (14)$$

$$\mathbf{Y}_a = \mathbf{A}\mathbf{a}_y = \mathbf{A}\mathbf{C}^{-1}\mathbf{Y}_{DE} = \mathbf{H}\mathbf{Y}_{DE} \quad (15)$$

$$\mathbf{Z}_a = \mathbf{A}\mathbf{a}_z = \mathbf{A}\mathbf{C}^{-1}\mathbf{Z}_{DE} = \mathbf{H}\mathbf{Z}_{DE} \quad (16)$$

Since matrix \mathbf{C} has a zero block, it is normal to extract this and solve separately, so that \mathbf{C} is subdivided as

$$\mathbf{C} = \begin{pmatrix} \mathbf{0} & \mathbf{P} \\ \mathbf{P}^T & \mathbf{M} \end{pmatrix} \quad (17)$$

Hence the unknowns for the x displacements are (analogous expressions hold for the y and z displacements)

$$\mathbf{a}_x = \mathbf{C}^{-1}\mathbf{X}_{DE} \quad (18)$$

The solution may be expressed in terms of the component matrices of \mathbf{C} , \mathbf{P} and \mathbf{M} . The matrix products now pre-multiply \mathbf{x}_{DE} rather than \mathbf{X}_{DE} and that the vector $\mathbf{a}_{x/y/z}$ has been split into the parts that multiply the RBF and polynomial terms, so that

$$\mathbf{a}_x^{\text{PLY}} = \mathbf{M}_p\mathbf{P}\mathbf{M}^{-1}\mathbf{x}_{DE} \quad (19)$$

$$\mathbf{a}_x^{\text{RBF}} = \mathbf{M}^{-1} - \mathbf{M}^{-1}\mathbf{P}^T\mathbf{M}_p\mathbf{P}\mathbf{M}^{-1}\mathbf{x}_{DE} \quad (20)$$

where

$$\mathbf{M}_p = (\mathbf{P}\mathbf{M}^{-1}\mathbf{P}^T)^{-1} \quad (21)$$

Forming the coupling matrix is now straightforward

$$\mathbf{H} = \mathbf{A} \begin{pmatrix} \mathbf{M}_p \mathbf{P} \mathbf{M}^{-1} \\ \mathbf{M}^{-1} - \mathbf{M}^{-1} \mathbf{P}^T \mathbf{M}_p \mathbf{P} \mathbf{M}^{-1} \end{pmatrix} \quad (22)$$

In this fashion as many zeros as possible are removed from matrices that need to be inverted. Interesting features of RBF interpolation worth noting are as follows:

- The interpolation is unique. This implies no additional constraints may be applied.
- The interpolation in all three coordinate directions is independent. This means that it is only necessary to find and store one matrix, \mathbf{H} , for all coordinate directions.
- The interpolation is time invariant and so only needs to be computed once, prior to any simulation. The surface geometry and mesh are then deformed very efficiently by matrix–vector multiplication.

Owing to the nature of forming an inverse matrix, the number of points used to link with the aerodynamic mesh should not be excessive. Domain element methods work well with RBF mesh deformation due to the low number of domain element node points. However, if one were to use spline control points or surface mesh points in three dimensions, then computational memory issues must be considered.

5. IMPLEMENTATION OF PARAMETERIZATION

The parameterization method relies upon a domain element being placed around the original design to be optimized. This is illustrated in Figure 1 for a two-dimensional aerofoil. Up to 22 design variables are used conforming to various deformations of the element that can be quantized; these include twist, chord, effects of chamber and various local thicknesses. It is important to note that not all design variables need be active, but 22 is considered to be optimum to achieve a large and applicable range of the design space [48]. The use of more design variables was considered, however, little benefit was gained. These design variables have been developed by experimentation with inverse design methods. Comparison is drawn here to a common method of parameterization, when using spline control points, typically a two-dimensional aerofoil section requires 30–50 parameters [23] to define the shape adequately, and for three-dimensional bodies this can be several hundred. The computational implications of this number of design variables can be significant if sensitivity information is obtained via finite-difference. Other parameterization methods can too suffer this limitation.

Domain element nodes located around the leading edge are nearer and more clustered to the aerofoil surface; this allows for finer control of the leading edge. This is in comparison with nodes over the main body of the aerofoil where smooth, gross changes are generally more applicable.

The RBF interpolation is used to link the domain element nodes with the entire aerodynamic mesh (and spline control points) shown in Figure 2. Figure 1(b) depicts a slightly exaggerated example of the deformation caused by one of the design variables, in this case one associated with

angle of attack (about the $\frac{1}{4}$ chord) of the aerofoil. Figure 3 shows the resulting RBF deformation in the aerodynamic mesh.

In order to incorporate the different design parameters in optimization in a meaningful and relevant way, so that there is no dominance of one design parameter over any other, requires scaling. All design variables are divided by their initial value for the original design, giving each the value of unity. Each design variable is scaled such that an increment of one to each design variable, taken in isolation, produces a displacement of equivalent to 1% of the aerofoil chord. This gives approximately the same weight to all the design variables so they can then be combined into the optimization problem in a consistent and well-conditioned manner. Parameter bounds are placed on each design variable such that each design variable in isolation could produce a maximum deformation to the original aerofoil equivalent to a length of $\pm 25\%$ of the chord length.

To evaluate sensitivity information via finite-difference requires the evaluation of the steady-state objective function due to a geometric perturbation. The flow solver HMB [15] is used to




Figure 2. Domain element, spline control points and mesh points: (a) entire mesh view RAE2822 and (b) zoomed view RAE2822.




Figure 3. RBF deformation: RAE2822 (mesh—purple) and perturbed aerofoil (mesh—dark blue).

provide a fully converged flow field for all two-dimensional solutions. First-order finite-difference are computationally the simplest to evaluate; however, to ensure no biasing towards one direction and to increase accuracy, a second-order accurate stencil is used. Design variables, α_n , give rise to objective function $\underline{J}(\alpha_n, U(\alpha_n))$, where $U(\alpha_n)$ represents the evaluation of the flow field. Perturbations, $\delta\alpha$, are made to the design variables in order to calculate the sensitivity, defined as $d\underline{J}/d\alpha_n$.

$$\frac{d\underline{J}}{d\alpha_n} = \frac{\underline{J}(\alpha_n + \delta\alpha_n, U(\alpha_n + \delta\alpha_n)) - \underline{J}(\alpha_n - \delta\alpha_n, U(\alpha_n - \delta\alpha_n))}{2\delta\alpha_n} + O(\delta\alpha_n^2) \quad (23)$$

When considering aerodynamic shape optimization, angle of attack is an important variable. However, within the majority of parameterization techniques it is often excluded as a design variable. When lift and/or moment are aerodynamically constrained, the angle of attack is frequently adjusted externally to the optimizer to meet the constraint. This can compromise the efficiency of the optimizer and often leads it in misleading directions. Using the domain element parameterization technique adopted here, angle of attack is included in the optimization process as a design variable. Consequently gradient information with respect to this design variable (with regard to both the objective function and the constraint function) are included within the FSQP algorithm—a significant advantage.

5.1. Optimization parameters

The perturbation size is of critical importance to ensure the correct value for sensitivity information and has to be investigated. Figure 4 shows the viscous drag sensitivity of each design variable across the perturbation size scale. Below a perturbation size of 1×10^{-6} , it is apparent that the perturbation is too small; the change in shape of the design is insufficient to disturb the flow such as to record an accurate change. In the region of 1×10^{-5} – 1×10^{-3} , the sensitivity of each design variable is relatively constant, and the correct value of the gradient is obtained, Figure 4(b). Depending on the design variable, solutions for perturbation sizes of 1×10^{-2} or greater may not be shown in Figure 4(b), this is due to the aerofoil becoming so highly deformed in one particular aspect that the viscous flow solver is unable to converge. A central-difference stencil was used




Figure 4. Viscous drag sensitivity of each design variable for differing perturbation sizing.

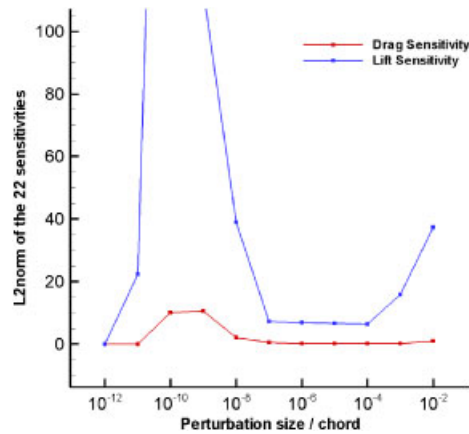


Figure 5. L2norm of the 22 sensitivities across perturbation sizing scale.




Figure 6. L2norm of the 22 sensitivities for differing grid densities: (a) lift sensitivity and (b) drag sensitivity.

about a NACA0012 aerofoil at 0° incidence, Mach number = 0.4, $Re = 2 \times 10^6$, grid size was 50 000 cells. The L2norm of the 22 sensitivities for these results is shown in Figure 5.

The effect of grid density is also an important issue to study with regard to obtaining the correct sensitivities. Results from a grid dependence study consisting of five levels of grid density are shown in Figure 6. The five grids consisted of 72×12 , 144×24 , 288×48 , 576×96 and 1152×192 mesh cells, respectively. Results indicate that a medium grid density ($\geq 50\,000$ points) is sufficient to obtain correct gradient information. Again the study revolved around a second-order central-difference stencil used about a NACA0012 aerofoil at 0° incidence, Mach number = 0.4, $Re = 2 \times 10^6$.

6. INVERSE DESIGN

The inverse design method, pioneered by Lighthill [53], is rarely used for optimization purposes any longer due to requirement of *a priori* knowledge of an improved or desirable pressure distribution.

However, inverse design techniques are applied here to serve to evaluate the applicability of both the parameterization method and the choice of the design variables. Although only minor deformation of an initial design may be necessary to achieve significant improvements with regard to drag optimization, the parameterization method must be free-form to such an extent that any likely optimum design must be achievable. The FSQP algorithm along with an inviscid Euler solution provided by HMB, at Mach number=0.4, on a grid consisting of 288×48 cells was used. The grids were produced by GHMB [16]; using a three block structured C-mesh with the far field set to 15 chords, a leading and trailing edge spacing of 1×10^{-4} and 3×10^{-3} , respectively. The normal spacing is set to give aspect ratio one cell at the leading edge.

Two cases are shown here:

Inverse (1): initial aerofoil NACA2411 at 0° ; target aerofoil NACA0012 at 2° , Figure 7.

Inverse (2): initial aerofoil NACA0012 at 0° ; target aerofoil NACA0006 at 2° , Figure 8.




Figure 7. Inverse (1) (initial aerofoil NACA2411 at 0° , target aerofoil NACA0012 at 2°): (a) C_p distribution: initial—green, target—black, optimized—blue and (b) domain element and aerofoil surface: initial—green and light green, optimized—blue and light blue.




Figure 8. Inverse (2) (initial aerofoil NACA0012 at 0° , target aerofoil NACA0006 at 2°): (a) C_p distribution: initial—green, target—black, optimized—blue and (b) domain element and aerofoil surface: initial—green and light green, optimized—blue and light blue.

The inverse design problem is of the form

$$\text{minimize } \underline{J} = \sum_{N=1}^{N_{\text{surface}}} (Cp_{\text{target}} - Cp)_N^2 \quad (24)$$

Results from Inverse (1) are shown in Figure 7. Here an asymmetric aerofoil at 0° incidence is optimized until the Cp distribution matches that of a symmetric aerofoil at incidence. The geometry profile change required to achieve this can be seen as quite considerable. Results from Inverse (2) are shown in Figure 8. In this case, although both the initial and target aerofoil are symmetric, the incidence and thickness change required are again considerable.

The closeness of the matches considering the low number of design variables is extremely encouraging and demonstrates that a large and appropriate range of the applicable design space is achievable. It must be noted that the results are presented after no more than 100 evolutions, and although solutions lie near to an optimum, the objective function in many cases was still reducing. More inverse design cases are presented in Morris *et al.* [48].

7. TWO-DIMENSIONAL AEROFOIL OPTIMIZATION

Two-dimensional aerofoil sections have been optimized using the FSQP algorithm under various flight conditions, with several different objective functions and constraints. For many fixed wing applications, the objective is to reduce the drag experienced by aircraft without detriment to total lift. The pitching moment is another key aerodynamic quantity that is desired to remain low and frequently ignored by many optimization efforts. The internal volume of the aerodynamic body is again a necessary constraint that needs imposing. Internal volume needs preserving so that sufficient structure or fuel can be contained within the body. As important, aerodynamic drag can be easily reduced by a reduction in volume; therefore, an 'optimum' design with lower internal volume generally may not constitute a true improvement in design.

High-quality structured grids were generated by GHMB [16]. For viscous simulations, grids contain 576×96 mesh cells with appropriate viscous wall spacing (leading and trailing edge spacing of 2×10^{-4} and 6×10^{-4} , respectively, and a $y+$ spacing of 5×10^{-6}). $k-\omega$ [15] is used for turbulence modelling and the flow is assumed wholly turbulent. The choice of the turbulence model was swayed by two arguments, first by industrial experience from Agusta–Westland, with whom this research is performed, and second by evidence presented in the following section. For inviscid simulations grids consist of 288×48 mesh cells with appropriate wall spacing as used in the inviscid inverse design cases. The pitching moment is taken about the initial aerofoil's $\frac{1}{4}$ chord position. Percentage changes to pitching moment are not given in the results' tables as the constraint is in the form of an absolute magnitude, and a percentage change value can appear misleading when the sign of the moment has changed (Table II–XII).

Eleven optimization cases were examined, ranging from subsonic to transonic, viscous and inviscid, and also various objective functions, details of each optimization are given in Table I.

The use of three different objective functions has been considered, namely, drag coefficient, volume entropy and surface entropy. Both volume and surface entropy objective functions are based on a measure of normalized local entropy. This is a measure of energy loss in the flow due

Table I. Summary of optimization results, where subscript I indicates values of the initial aerofoil.

Case	Aerofoil	Flow	Objective	Constraint1	Constraint2	Constraint3
1	NACA0012	Subsonic viscous	Cd	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
2	NACA0012	Subsonic viscous	Cd	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq \frac{2}{3} Vol_I$
3	NACA0012	Transonic inviscid	Cd	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
4	NACA0012	Transonic viscous	Cd	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
5	RAE2822	Transonic inviscid	Cd	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
6	RAE2822	Transonic inviscid	e_{vol}	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
7	RAE2822	Transonic inviscid	e_{surf}	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
8	RAE2822	Transonic viscous	Cd	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
9	RAE2822	Transonic viscous	e_{vol}	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
10	RAE2822	Transonic viscous	e_{surf}	$Cl \geq Cl_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$
11	NACA23012	Subsonic viscous	$1/Cd$	$Cd \leq Cd_I$	$ Cm \leq Cm_I $	$Vol \geq Vol_I$

Table II. Optimization (1) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.5937	0.5943	+0.1
Cm	0.0297	0.0293	—
Volume	0.0805	0.0805	+0.02
Cd	0.0135	0.0121	-10.37

Table III. Optimization (2) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.5937	0.5938	+0.01
Cm	0.0297	0.0296	—
Volume	0.0805	0.0594	-26.31
Cd	0.0135	0.0110	-18.67

to the aerodynamic body and is defined as

$$e = \frac{P_{local}/\rho_{local}^\gamma}{P_\infty/\rho_\infty^\gamma} - 1 \quad (25)$$

Values of zero for normalized entropy indicate no loss in the flow, conversely, where shocks or retardation of flow due to boundary layers or wakes are present, the value of entropy is higher due to these non-isentropic processes. The volume entropy values are then computed by

$$e_{vol} = \frac{\sum_{N=1}^{N_{cells}} e_N * V_N}{\sum_{N=1}^{N_{cells}} V_N} \quad \text{where } V_N \text{ is cell volume} \quad (26)$$

Surface entropy is a measure of the entropy experienced at the surface of the design geometry and is evaluated by

$$e_{\text{surf}} = \frac{\sum_{N=1}^{N_{\text{surface}}} e_N * \Delta S_N}{\sum_{N=1}^{N_{\text{surface}}} \Delta S_N} \quad \text{where } \Delta S_N \text{ is the cell surface length} \quad (27)$$

7.1. Subsonic drag optimization (optimizations (1) and (2))

A NACA0012 aerofoil is optimized for two different volume constraints under same flow conditions ($\alpha_I = 5^\circ$, $\text{Mach} = 0.4$, $Re = 2 \times 10^6$) with the objective of reducing the viscous drag coefficient. Both aerofoils are subject to the same lift and moment constraints; however, to highlight the effect that reduction of a body's volume can have on the drag coefficient; optimization (1) is subject to a rigid constraint of $\text{Vol} \geq \text{Vol}_I$, whereas optimization (2) is subject to a less stringent requirement of $\text{Vol} \geq \frac{2}{3} \text{Vol}_I$.

Figure 9 compares the initial and optimized domain elements and the corresponding aerofoils for the two subsonic optimizations. Figures 10 and 11 depict flooded C_p and surface C_p plots for initial and final aerofoils, respectively. Optimized aerofoils are seen to reduce the 'in-flight' angle of attack and incorporate effects of chamber to maintain lift and reduce drag.

The achievement of over 10% reduction in viscous drag for a heavily constrained aerofoil in subsonic conditions is an extremely encouraging result. When the volume of the aerofoil is allowed to decrease, over an 18% reduction is observed, this extra reduction in drag is expected due to an allowable decrease in volume and hence maximum thickness. These results are not only encouraging but they also highlight the need for a volume constraint to be imposed so that optimized designs can be compared with initial aerofoils on a truly like-for-like basis.

Observation of the history of the percentage reduction in objective function (and internal volume of the aerofoil) through the optimizations, Figure 12, shows some insight, and possibly advantages, to the way the FSQP algorithm works. Seen for both optimizations (1) and (2), the objective function decreases in a step-like fashion. Over the first few evolutions, the optimizer is guided by first-order gradient information only, and a plateau is soon reached (it is likely that a first-order gradient optimizer may halt here). Around this new design point, a few evolutions are required for the optimizer to gain accurate approximations of the second-order gradients. Once the Hessian is sufficiently accurate, a new search direction can be found and another rapid decrease in

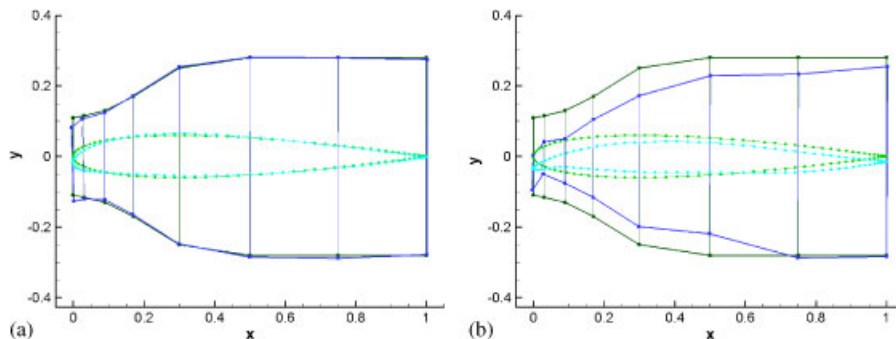


Figure 9. Domain element and aerofoil points: NACA0012 aerofoil—green and light green, optimized aerofoil—blue and light blue: (a) optimization (1) and (b) optimization (2).




Figure 10. Flooded C_p plot of initial and optimized aerofoils: (a) NACA0012 aerofoil; (b) optimization (1); and (c) optimization (2).

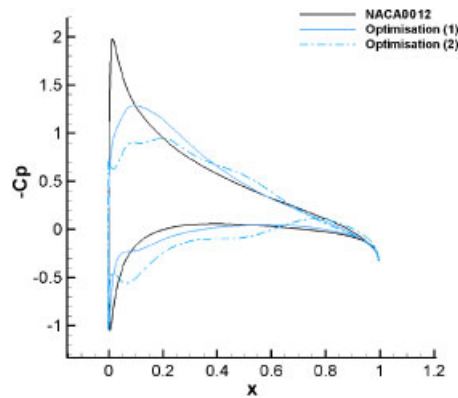


Figure 11. C_p plot of initial and optimized aerofoils.




Figure 12. History of objective function through optimization: (a) optimization (1) and (b) optimization (2).

objective function is achieved, and this pattern continues. Figure 12(a) for optimization (1) clearly shows the volume constraint, and an optimum design is achieved after 67 evolutions. However, for optimization (2), Figure 12(b), where the volume is allowed to decrease significantly, this increases the size of the design space considerably and 133 evolutions are required to achieve an optimum design. Interestingly the design does not lie on the volume constraint boundary, only having achieved a 27% reduction.

7.2. Transonic drag optimization (optimizations (3) and (4))

In a transonic flow regime aerofoil designs frequently exhibit shockwaves, this in-turn can lead the wave drag to be a significant proportion of the total drag. Fixed-wing aircraft typically cruise at transonic Mach numbers, and low total drag values in this flow regime are crucial to economic success, and thus a significant proportion of aerodynamic design effort is placed around this design point. Inviscid flow calculations can calculate wave drag efficiently and they are typically much faster than viscous calculations not only due to simpler equations to solve but also due to a lower grid density required. To introduce optimization methods into industrial design loops, results need to be obtained quickly; hence, viscous calculations are typically prohibitive. Therefore in this subsection a comparison is made between optimized aerofoils obtained under viscous and inviscid flow calculations.

Both optimizations centre around on NACA0012 aerofoil at $\alpha_I = 5^\circ$, Mach = 0.65, however, optimization (3) is achieved under inviscid conditions, whereas optimization (4) is obtained using viscous calculations ($Re = 1 \times 10^7$). Results are given in Tables IV and V, respectively.

Figure 13 compares the initial and optimized domain elements and the corresponding aerofoils for the transonic optimizations. Figure 14(a) and (b) shows the flooded C_p plots for the initial and final aerofoils of optimization (4) and Figure 14(c) and (d) for optimization (3). Figure 15 shows initial and optimum aerofoil surface C_p plots.

In the transonic test cases, the initial NACA0012 aerofoil exhibits a strong shock at $\frac{1}{4}$ chord, this shock is slightly stronger and sharper for the inviscid case, as to be expected. The inviscid aerofoil has a drag coefficient of 0.0188, approximately 20% lower than the viscous drag coefficient for the same geometry under the same flow conditions. The shockwave appears only marginally weaker

Table IV. Optimization (3) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.7956	0.7957	+0.01 (-2.06)
Cm	0.0132	-0.0117	—
Volume	0.0805	0.0817	+1.49
Cd	0.0188	0.0039	-79.25 (-42.67)

Table V. Optimization (4) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.7164	0.7182	+0.25
Cm	0.0214	-0.0211	—
Volume	0.0805	0.0813	+1.00
Cd	0.0239	0.0131	-45.18

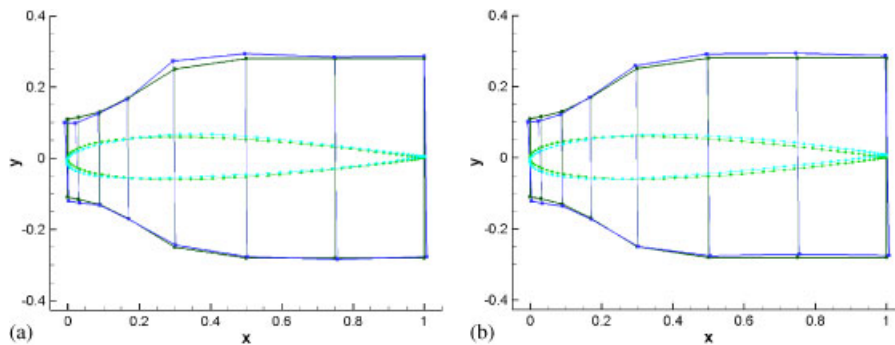


Figure 13. Domain element and aerofoil points: NACA0012 aerofoil—green and light green, optimized aerofoil—blue and light blue: (a) optimization (3) and (b) optimization (4).

for the viscous case, and with a significant amount skin friction drag with a total drag coefficient of 0.0239.

Optimization (4) produced a final aerofoil with more than a 45% reduction in viscous drag, and optimization (3) produced an aerofoil with nearly an 80% reduction in inviscid drag. To verify the result from inviscid optimization (optimization (3)), this final aerofoil geometry was also solved in viscous flow (on a new viscous mesh containing 576×96 cells), and a flooded C_p plot of this solution can be seen in Figure 14(e). The difference between the results of this viscous solution of the aerofoil geometry produced by inviscid optimization and the initial viscous NACA0012 aerofoil are given in brackets in Table IV.

It is noted that the design resulting from an inviscid optimization still produces notable reductions in viscous drag—comparable to the aerofoil from viscous optimization. This result is significant, and optimizations can therefore be legitimately performed with inviscid analysis, which is vital




Figure 14. Flooded C_p plots of initial and optimized aerofoils: (a) NACA0012 inviscid; (b) optimization (3); (c) NACA0012 viscous; (d) optimization (4); and (e) optimization (3) aerofoil when solved under viscous conditions.

when considering cost implications if this method is extended to three-dimensional bodies, or when attempting to implement optimization in industrial design loops. It can be clearly seen from the C_p plots that for both viscous and inviscid optimizations shock strength has been reduced dramatically.

It is noted here that the initial NACA0012 aerofoil is operating well outside its intended design range; hence such large improvements in drag can be achieved. However, it serves here to demonstrate and validate the concepts.

7.3. RAE2822 Transonic drag optimization (optimizations (5)–(10))

Test case 9 from the AGARD experimental database [54] for the RAE2822 aerofoil is commonly used as a benchmark for optimization. The test conditions are $\alpha_I = 2.67^\circ$, $\text{Mach} = 0.73$ and $Re = 6.5 \times 10^6$. The experimental data is compared to inviscid and various viscous CFD calculations, Figure 16. The $k-\omega$ turbulence model provides superior results with regard to matching the experimental C_p distribution and thus is the turbulence model employed throughout all viscous optimizations.

Optimizations (5)–(10) concern the optimization of the AGARD test case aerofoil with rigid constraints on moment, volume and lift. However, the effectiveness of different objective functions

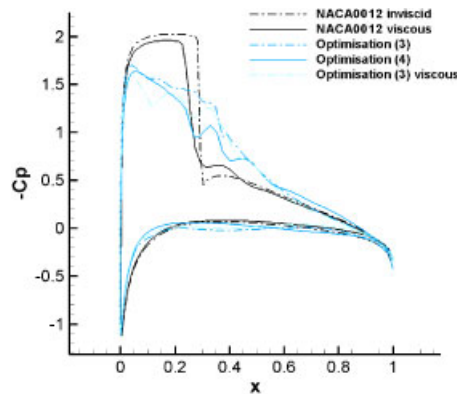


Figure 15. C_p plots of initial and optimized aerofoils.




Figure 16. C_p distributions for RAE2822.

is assessed in both an inviscid and a viscous environment. Optimizations (5)–(7) are performed with inviscid calculations and are analogous to optimizations (8)–(10) performed under viscous conditions. Refer to Tables VI–XI for details of optimization results.

Table VI. Optimization (5) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.9571	0.9571	0.00 (−2.63)
Cm	−0.1371	−0.1371	—
Volume	0.0773	0.0781	+0.01
Cd	0.0182	0.0073	−59.85 (−15.85)
e_{vol}	0.003758	0.000617	−83.58 (−22.56)
e_{surf}	0.009009	0.005332	−41.11 (−0.65)

Table VII. Optimization (8) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.8017	0.8029	+0.15
Cm	−0.0987	−0.0986	—
Volume	0.0773	0.0776	+0.40
Cd	0.0183	0.0136	−25.68
e_{vol}	0.005054	0.003151	−37.65
e_{surf}	0.304841	0.302112	−0.88

Table VIII. Optimization (6) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.9571	0.9600	+0.3 (−3.03)
Cm	−0.1371	−0.1370	—
Volume	0.0773	0.0781	+0.01
Cd	0.0182	0.0078	−57.14 (−14.75)
e_{vol}	0.003758	0.000530	−85.89 (−21.20)
e_{surf}	0.009009	0.004553	−50.04 (−0.62)

Table IX. Optimization (9) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.8017	0.8033	+0.20
Cm	−0.0987	−0.0987	—
Volume	0.0773	0.0774	+0.13
Cd	0.0183	0.0137	−25.13
e_{vol}	0.005054	0.003130	−38.06
e_{surf}	0.304841	0.301802	−0.98

Table X. Optimization (7) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.9571	0.9577	+0.06 (−2.11)
Cm	−0.1371	−0.1365	—
Volume	0.0773	0.0789	+0.02
Cd	0.0182	0.0097	−46.70 (−14.21)
e_{vol}	0.003758	0.000947	−74.80 (−20.33)
e_{surf}	0.009009	0.003509	−59.20 (−0.59)

Table XI. Optimization (10) aerofoil results.

	Initial	Optimized	%Diff
Cl	0.8017	0.8031	+0.17
Cm	−0.0987	−0.0986	—
Volume	0.0773	0.0774	+0.13
Cd	0.0183	0.0150	−18.03
e_{vol}	0.005054	0.003766	−25.48
e_{surf}	0.304841	0.302007	−0.91

Optimizations (5) and (8) concern drag coefficient minimization, optimizations (6) and (9) concern the minimization of volume entropy, defined by Equation (26), and optimizations (7) and (10) are also concerned with entropy minimization, but through surface entropy defined by Equation (27).

The resulting geometries produced under inviscid conditions (optimizations (5)–(8)) are subsequently solved under viscous conditions (on new viscous meshes containing 576×96 cells) and compared with the forces resulting from the initial viscous RAE2822 aerofoil. These results are then highlighted in brackets in the relevant tables.

Figure 17 shows initial and optimized domain element and aerofoil geometries. The optimized designs in all six cases required only minor deformation from the initial RAE2882 aerofoil, and interestingly, all six designs are very similar. Figures 18 and 21 depict C_p plots for initial and optimized geometries. Figures 20 and 23 are analogous to those figures just mentioned; however, they depict normalized entropy.

Comparing the initial RAE2822 aerofoil under inviscid and viscous conditions, Figure 18(a) and Figure 21(a), respectively, in both circumstances the upper surface of the aerofoil is dominated by transonic flow, but the inviscid solution results in a stronger shockwave further downstream, as expected. This is also observed from the surface C_p plots, Figure 19. Hence, this is a harsh test case for an inviscid-viscous optimization comparison, as the two initial flow fields are substantially different due to the effect of viscosity. It is important to note that the long-term aim of the optimization methods developed are to perform rotor optimizations, where it is expected that any transonic regions will be small, and so similar to those of the NACA0012 aerofoil used in the previous two sections.

For optimizations (5)–(10) all six objective functions result in a design that achieves a reduction in the viscous drag coefficient. The entropy plots, Figures 20 and 23, clearly highlight that each optimized aerofoil has significantly reduced losses, and in nearly all cases removes most of the losses associated with the initial strong shock wave of the RAE2822 aerofoil. For inviscid cases,




Figure 17. Domain element and aerofoil points: RAE2822 aerofoil—green and light green, optimized aerofoil—blue and light blue: (a) optimization (5); (b) optimization (8); (c) optimization (6); (d) optimization (9); (e) optimization (7); and (f) optimization (10).

all high levels of entropy are removed from the flow, whereas for viscous optimizations almost all losses associated with the shock are removed, but some high losses contained within the boundary layer remain, Figure 22.

Comparing the three inviscid optimizations (5)–(7), it is optimization (5) using inviscid drag coefficient as objective function that results in the superior aerofoil. For this aerofoil, inviscid drag is reduced by 60%, and when solved under viscous conditions approximately a 16% reduction in drag is observed when compared with the initial viscous RAE2822 aerofoil. Constrained in lift, inviscid optimizations aim at reducing wave drag, and all three optimizations here have managed




Figure 18. Flooded C_p plot of initial and optimized aerofoils (inviscid): (a) RAE2822 inviscid; (b) optimization (5); (c) optimization (6); and (d) optimization (7).

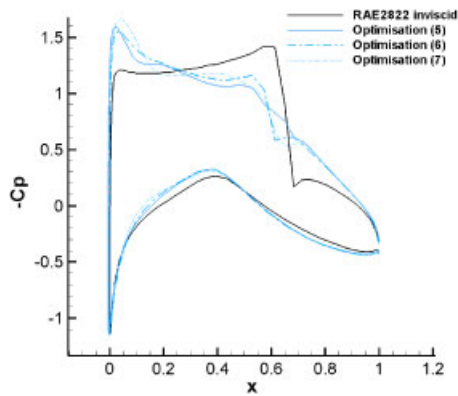


Figure 19. C_p plots of initial and optimized aerofoils.




Figure 20. Flooded entropy plot of initial and optimized aerofoils (inviscid): (a) RAE2822 inviscid; (b) optimization (5); (c) optimization (6); and (d) optimization (7).

to reduce shock strength significantly with only minor shape adjustment. However, when these designs are subsequently solved in a viscous environment skin friction drag, unaccounted for in the optimizations and a major contributor to total drag, reduces the impact of shock strength reduction on total drag coefficient. Furthermore, the optimum designs may have affected boundary layer—shock interaction in an adverse way. Although the lift coefficient is maintained throughout the inviscid optimization, the optimized design when placed in a viscous environment produces an aerofoil that has a lower lift coefficient than the initial viscous RAE2822 aerofoil. Lift can simply be augmented by an increase in incidence but produces an associated detriment with drag.

Interestingly, optimizations (6) and (7), using entropy objective functions, produce similar aerofoil geometries to optimization (5); therefore, similar performance characteristics to the design achieved through an inviscid drag coefficient optimization.

Out of the three viscous optimizations (8)–(10), again it is the design resulting from the drag coefficient as the objective function that produces the superior design (optimization (8) with over 25% reduction in viscous drag), although the aerofoil from optimization (9) using volume entropy as objective function is only marginally worse.




Figure 21. Flooded C_p plot of initial and optimized aerofoils (viscous): (a) RAE2822 viscous; (b) optimization (8); (c) optimization (9); and (d) optimization (10).




Figure 22. C_p plots of initial and optimized aerofoils.




Figure 23. Flooded entropy plot of initial and optimized aerofoils (viscous): (a) RAE2822 viscous; (b) optimization (8); (c) optimization (9); and (d) optimization (10).

It is noted from all optimizations that the value of surface entropy decreases only slightly ($\leq 1\%$) even for optimizations where over a 25% reduction in drag coefficient is observed. Of the three different objective functions investigated, it is surface entropy that produces the least improvement in design, in both an inviscid and a viscous situations. More importantly, in a viscous situation, using volume entropy as the objective function actually results in lower surface entropy than when surface entropy is used as the objective function itself. Surface entropy values are relatively insensitive to drag changes, so all design variables have a relatively similar and small gradient, therefore making it difficult to establish a good search direction to drive the shape deformation for ultimately reduced drag coefficient.

Again it is worth noting that strict constraints on volume, moment and lift were imposed on all optimizations here, and reductions in viscous drag between 14 and 25% are significant results.

7.4. Subsonic lift optimization (optimization (11))

A rotor-type aerofoil case is now considered. This is optimization of an aerofoil section to provide a rotor with improved aerodynamic performance in forward flight. On the retreating side of a rotor

Table XII. Optimization (11) aerofoil results.

	Initial	Optimized	%Diff
Cl	1.305	1.386	+5.98
Cm	0.0063	-0.0062	—
Volume	0.0797	0.0799	+2.5
Cd	0.0415	0.0415	+0.0




Figure 24. Domain element and aerofoil points: NACA23012 aerofoil—green and light green, optimized aerofoil—blue and light blue.

blade at inboard locations, it is often desirable to have an aerofoil that provides a large amount of lift at a significant incidence angle in low Mach numbers but must also exhibit low drag at low incidence when on the forward side of the rotor disc. The NACA23012 aerofoil is a typical example of an aerofoil section used for this type of application. In the following optimization, the NACA23012 aerofoil is optimized for lift at a high incidence at a low subsonic Mach number ($M=0.3$, $Re=3 \times 10^6$, $\alpha_I=12^\circ$). A rigid constraint is placed on moment, volume and importantly drag; this final constraint should ensure good working performance while the aerofoil is on the forward side of the rotor disc. Results of the optimization are given in Table XII, and initial and final domain elements and aerofoil geometries are shown in Figure 24.

It is clear from the C_p distribution plots in Figure 25 that there is an increased region of low pressure when compared with the initial aerofoil; however, the suction peak has diminished somewhat. Therefore, the optimized aerofoil has been able to improve upon the already well-suited NACA23012 by reducing viscous stresses associated with the strong suction peak allowing an increased lift by nearly 6% without any volume, moment or drag penalty.

8. CONCLUSIONS

A completely generic ‘wrap-around’ aerodynamic optimization tool has been developed and presented. This comprises a domain element method with a new geometric parameterization technique for application to CFD-based aerodynamic optimization. The parameterization uses radial basis functions (RBFs) to interpolate positions of the domain element and the grid coordinates to




Figure 25. C_p plots of initial and optimized aerofoils: (a) NACA23012 aerofoil; (b) optimization (11); and (c) C_p distribution: NACA23012-black, optimization (11)-dark blue.

provide simultaneous deformation of the design surface and its corresponding aerodynamic mesh. The inverse mapping is only required once for the initial design as the values of the parametric coordinates of the grid points with respect to the domain element remain constant throughout the optimization. Updates to the geometry and the corresponding mesh are provided simultaneously by matrix–vector multiplication; this is extremely fast and efficient and results in very high-quality mesh deformation. This domain element technique allows for optimization to be performed totally independently from the mesh generation package and topology, requiring only an initial mesh. Two-dimensional aerofoil inverse design and heavily constrained aerofoil optimization is centred around a feasible sequential quadratic programming (FSQP) gradient-based optimizer, where sensitivity information is provided by finite-difference; this allows for independence from the flow solver used.

The domain element is automatically located around the exterior of a design. This is a key advantage, allowing complex or even multi-element geometries to be parameterized. Design variables can range in fidelity, from the gross deformation of a design to local detailed surface changes, and inverse design problems have demonstrated that 22 design variables are sufficient to achieve a large and applicable range of design space.

Constrained drag optimization of aerofoil sections has shown that significant reductions in drag are achievable. Viscous drag reductions in the order of 10% are observed for a wholly subsonic aerofoils with rigid constraints on lift, moment and volume. If the volume constraint is relaxed slightly, a further 8% reduction in drag is achievable. For aerofoils in transonic Mach numbers (with stringent constraints on lift, moment and volume), viscous drag reductions range from 14 to 45% depending on the aerofoil and flow conditions. Using measures of entropy in the flow as an objective function has also shown to produce aerofoils with similar performance characteristics to those optimized with drag coefficient as the objective function. More notable, this is also true for inviscid optimizations. Although no design resulting from an inviscid optimization is superior to the design resulting from a viscous optimization, the sometimes small difference in performance is nonetheless impressive considering the extra computational overhead of a viscous optimization. This last finding has significant computational cost implications once the method is extended to three-dimensional geometries.

Although two-dimensional results are presented here, the parameterization method presented can be equally applied to three-dimensional bodies, where the reduction in the required number of design variables will be even more significant. The parameterization technique cannot only combine variables of different scales and types with only a few parameterization nodes, but it can lead to a significantly reduced number of design variables for a three-dimensional application when compared with many other types of parameterization method.

ACKNOWLEDGEMENTS

The authors would like to thank Agusta–Westland, EPSRC (the UK Engineering and Physical Sciences Council), and the U.K. MoD Joint Grant Scheme (JGS) funding, under contract GR/S61294, for their financial support of Asa Morris through the Rotorcraft DARP; the University of Bristol with whom Thomas Rendall has been awarded a postgraduate research scholarship; former postgraduate student James Rodgers for his contribution to the surface grid generation codes.

REFERENCES

1. Hicks RM, Henne PA. Wing design by numerical optimization. *Journal of Aircraft* 1978; **15**:407–412.
2. Hicks RM, Murman EM, Vanderplaats GN. An assessment of aerofoil design by numerical optimization. NASA TMX-3092, Ames Research Center, Moffett Field, CA, July 1974.
3. Li W, Huyse L, Padula S. Robust aerofoil optimization to achieve consistent drag reduction over a mach range. NASA/CR-2001-211042, *Icase Report No. 2001-22*.
4. Li W, Huyse L, Padula S. Profile optimization method for robust aerofoil shape optimization in viscous flows. NASA/TM-2003-212408, NASA/CR-2001-211042, *Icase Report No. 2001-22*.
5. Wong WS, Le Moigne A, Qin N. Parallel adjoint-based optimisation of a blended wing body aircraft with shock control bumps. *The Aeronautical Journal* 2007; **111**(1117):165–174.
6. Qin N, Vavalle A, Le Moigne A. Spanwise lift distribution for blended wing body aircraft. *Journal of Aircraft* 2005; **42**(2):356–365.
7. Qin N, Wong WS, LeMoigne A, Sellars N. Validation and optimisation of 3D bumps for transonic wing drag reduction. *CEAS/Ketnet Conference on Key Aerodynamic Technologies*, Bremen, Germany, June 2005.
8. Qin N, Vavalle A, Le Moigne A, Laban M, Huckett K, Weinerfelt P. Aerodynamic studies of blended wing body aircraft. *The 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization Conference*, AIAA-2002-5448, Atlanta, GA, September 2002.
9. Reuther J, Jameson A, Farmer J, Martinelli L, Saunders D. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. *RIACS Technical Report 96.02 January 1996 Presented at the AIAA 34th Aerospace Sciences Meeting and Exhibit*, AIAA Paper 96-0094, January 1996.

10. Chung HS, Alonso J. Multiobjective optimization using approximation model-based genetic algorithms. *Tenth AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Albany, New York, AIAA 2004-4325, 30 August–1 September 2004.
11. Gumbert CR, Hou G, Newman PA. Simultaneous aerodynamic analysis and design optimization (SAADO) for a 3-D rigid wing. *Fourteenth AIAA Computational Fluid Dynamics Conference*, Norfolk, VA, AIAA 99-3296, 28 June–1 July 1999.
12. Gumbert CR, Hou G, Newman PA. Simultaneous aerodynamic analysis and design optimization (SAADO) for a 3-D flexible wing. *AIAA 2001-1107*, 2001.
13. Le Pape A, Beaumier P. Numerical optimisation of helicopter rotor aerodynamic performance in hover. *Aerospace Science and Technology* 2005; **9**:191–201.
14. Le Pape A. Numerical aerodynamic optimization of helicopter rotors. *ERF*, Florence, September 2005.
15. Richards BE, Badcock KJ, Woodgate MA. Elements of computational fluid dynamics on block structured grids using implicit solvers. *Progress in Aerospace Sciences* 2000; **36**:351–392.
16. Allen CB. Towards automatic structured multiblock mesh generation using improved transfinite interpolation. *International Journal for Numerical Methods in Engineering* 2008; DOI: 10.1002/nme.2170.
17. Bloor MIG, Wilson MJ. Efficient parameterization of generic aircraft geometry. *Journal of Aircraft* 1995; **32**(6):1269–1275.
18. Smith RE, Bloor MIG, Wilson MJ, Thomas AT. Rapid airplane parametric input design (RAPID). *Proceedings of 12th AIAA Computational Fluid Dynamics Conference*. AIAA: Washington, DC, 1995; 452–462.
19. Braibant V, Fleury C. Shape optimal design using B-splines. *Computer Methods in Applied Mechanics and Engineering* 1984; **44**(3):247–267.
20. Kulfan BM, Bussoletti JE. Fundamental parametric geometry representations for aircraft component shapes. *Eleventh AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Renaissance Portsmouth, VA, AIAA-2006-6948, 6–8 September 2006.
21. Kulfan BM. A universal parametric geometry representation method—CST. *Forty-fifth AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, 8–11 January 2007.
22. DYoung DP, Huffman WP, Melvin RG, Bieterman MB, Hilmes CL, Johnson FT. Inexactness and global convergence in design optimization. *Fifth AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Panama City, FL, AIAA Paper 94-4286, September 1994.
23. Castonguay P, Nadarajah SK. Effect of shape parameterization on aerodynamic shape optimisation. *Forty-fifth AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, 8–11 January 2007.
24. Reuther J. Aerodynamic shape optimization using control theory. *NASA Technical Report NASA-CR-201064*, 1996.
25. Sobieczky H. *Parametric Airfoils and Wings*. Notes on Numerical Fluid Mechanics, vol. 68. Vieweg Verlag, 1998; 71–88.
26. Wu H-Y, Yang S, Liu F. Comparison of three geometric representations of aerofoils for aerodynamic design for optimization. *Sixteenth Computational Fluid Dynamics Conference*, Orlando FL, AIAA 2003-4095, 23–26 June 2003.
27. Nadarajah S, Castonguay P, Mousavi A. Survey of shape parameterization techniques and its effect on three-dimensional aerodynamic shape optimization. *Eighteenth AIAA Computational Fluid Dynamics Conference*, Miami, FL, AIAA-2007-3837, 25–28 June 2007.
28. Jameson A. Aerodynamic design via control theory. *Journal of Scientific Computing* 1988; **3**:233–260.
29. Jameson A. Automatic design of transonic aerofoils to reduce the shock induced pressure drag. *Proceedings of the 31st Israel Annual Conference on Aviation and Aeronautics*, Tel Aviv, February 1990; 5–17.
30. Pickett RM, Rubinstein MF, Nelson RB. Automated structural synthesis using a reduced number of design coordinates. *AIAA Journal* 1973; **11**(4):494–498.
31. Watt A, Watt M. *Advanced Animation and Rendering Techniques*, Chapter 17. Addison-Wesley: New York, 1992.
32. <http://www.optimalsolutions.us/>.
33. Jameson A. Optimum aerodynamic design using control theory. In *Computational Fluid Dynamics Review*, Hafez M, Oshima K (eds). Wiley: U.K., 1995; 495–528.
34. Drela M. *A Users Guide to LINDOP V2.5*. MIT Computational Aerospace Sciences Laboratory, 1996.
35. Drela M. *A Users Guide to MSES 2.95*. MIT Computational Aerospace Sciences Laboratory, 1996.
36. Samareh JA. Status and future of geometry modeling and grid generation for design and optimisation. *Journal of Aircraft* 1999; **36**(1):97–104.
37. Samareh JA. Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimisation. *AIAA Journal* 2001; **39**(5):877–884.

38. Rendall TCS, Allen CB. Unified fluid–structure interpolation and mesh motion using radial basis functions. *International Journal for Numerical Methods in Engineering* 2007. DOI: 10.1002/nme.2219.
39. Panier E, Tits AL. On combining feasibility, descent and superlinear convergence in inequality constrained optimization. *Mathematical Programming* 1993; **59**:261–276.
40. Vanderplaats GN. *Numerical Optimization Techniques for Engineering Design: With Applications*. McGraw-Hill: New York, 1984.
41. Lyness JN, Moler CB. Numerical differentiation of analytic functions. *SIAM Journal on Numerical Analysis* 1967; **4**:202–210.
42. Pironneau O. *Optimal Shape Design for Elliptic Systems*. Springer: New York, 1984.
43. Le Moigne A. A discrete Navier–Stokes adjoint method for aerodynamic optimisation of blended-wing-body configurations. *Ph.D. Thesis*, Cranfield University College of Aeronautics, 2002.
44. Polak E, Tits AL. *System Modeling and Optimization*. Lecture Notes in Control and Information Sciences, vol. 38/1982. Springer: Berlin, Heidelberg, 1982.
45. Schittkowski K. Test examples for nonlinear programming codes: supplement. *Report*, Institute for Information Technology, University of Stuttgart, 1984.
46. Zhou JL, Tits AL, Lawrence CT. User’s guide for FFSQP version 3.7: a Fortran code for solving optimization programs, possibly minimax, with general inequality constraints and linear equality constraints, generating feasible iterates. *Technical Report SRC-TR-92-107r5*, Institute for Systems Research, University of Maryland, MD, 1997.
47. Zhou JL, Tits AL. Nonmonotone line search for minimax problems. *Journal of Optimization Theory and Applications* 1993; **76**(3):455–476.
48. Morris A, Allen CB, Rendall TCS. Development of Generic CFD-based aerodynamic optimisation tools for helicopter rotor blades. *Twenty-fifth AIAA Applied Aerodynamics Conference*, Miami, AIAA 2007-3809, June 2007.
49. Jakobsson S, Amoignon O. Mesh deformation using radial basis functions for gradient-based aerodynamic shape optimisation. *FOI Report FOI-R-1784-SE*, December 2005.
50. Beckert A, Wendland H. Multivariate interpolation for fluid–structure-interpolation problems using radial basis functions. *Aerospace Science and Technology* 2001; **5**:125–134.
51. Buhmann H. *Radial Basis Functions* (1st edn). Cambridge University Press: Cambridge, 2005.
52. Wendland H. *Scattered Data Approximation* (1st edn). Cambridge University Press: Cambridge, 2005.
53. Lighthill MJ. *A new method for two dimensional aerodynamics design*. *R and M*, 2112, Aeronautical Research Council, London, 1945.
54. Cook PH, McDonald MA, Firmin MCP. Aerofoil RAE 2822—pressure distributions, and boundary layer and wake measurements. *AGARD Report AR 138, Experimental Data Base for Computer Program Assessment*, 1979.