Constrained multi-objective aerofoil design using a multi-level optimisation strategy

L. Cameron
lcameron05@qub.ac.uk
J. Early, R. McRoberts and M. Price
School of Mechanical and Aerospace Engineering
Queen's University
Belfast
UK

ABSTRACT

A novel approach for the multi-objective design optimisation of aerofoil profiles is presented. The proposed method aims to exploit the relative strengths of global and local optimisation algorithms, whilst using surrogate models to limit the number of computationally expensive CFD simulations required. The local search stage utilises a re-parameterisation scheme that increases the flexibility of the geometry description by iteratively increasing the number of design variables, enabling superior designs to be generated with minimal user intervention. Capability of the algorithm is demonstrated via the conceptual design of aerofoil sections for use on a lightweight laminar flow business jet. The design case is formulated to account for take-off performance while reducing sensitivity to leading edge contamination. The algorithm successfully manipulates boundary layer transition location to provide a potential set of aerofoils that represent the trade-offs between drag at cruise and climb conditions in the presence of a challenging constraint set. Variations in the underlying flow physics between Pareto-optimal aerofoils are examined to aid understanding of the mechanisms that drive the trade-offs in objective functions.

Paper 3982. Manuscript received 26 March 2013, revised version received 28 November 2013, accepted 19 February 2015.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

NOMENCLATURE

A	weighting coefficient attached to component of Bernstein polynomial
c	aerofoil chord length
C	class function in class-shape transformation
c_d	drag coefficient
c_l	lift coefficient
c_m	moment coefficient
c_p	pressure coefficient
F1	objective function 1
F2	objective function 2
h	kriging basis function
k	number of design variables
M	Mach number
n	order of Bernstein polynomial
N	critical amplification factor
$N1$	class function coefficient 1
$N2$	class function coefficient 2
p	vector of exponents in Kriging basis function
$p1$	penalty function 1
$p2$	penalty function 2
$p3$	penalty function 3
$p4$	penalty function 4
Re	Reynolds number
S	shape function in Class-Shape Transformation
th	aerofoil thickness
$\mathbf{v}$	Vector of design variables
x	aerofoil x-co-ordinate
y	aerofoil y-co-ordinate
Δ_b	parameter controlling smart Pareto filter ($\geq \Delta s$)
Δ_s	parameter controlling smart Pareto filter ($\leq \Delta b$)
θ	vector controlling width of Kriging basis functions
μ	Pareto-optimal solution

Subscripts

1	cruise flight condition
2	climb flight condition
3	take-off flight condition
4	high-speed flight condition
tr	at boundary layer transition onset location
$tr = 0$	boundary layer transition at leading edge
max	maximum value

List of Acronyms

CFD	Computational fluid dynamics
UAV	Unmanned Aerial Vehicle
NLF	Natural Laminar Flow

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

PSO	Particle Swarm Optimiser
MOGA	Multi-Objective Genetic Algorithm
NSGA	Non-Dominated Sorting Genetic Algorithm
CST	Class-Shape Transformation
LHC	Latin Hypercube
AMMA	Adaptive Metamodel Assisted
SPF	Smart Pareto Filter
NACA	National Advisory Committee for Aeronautics
NASA	National Aeronautics and Space Administration
PI	Probability of Improvement
MSE	Mean Squared Error
RAE	Royal Aircraft Establishment

1.0 INTRODUCTION

The application of both gradient-based and global optimisation algorithms to aerodynamic shape optimisation has gained considerable attention in recent years. In particular, studies with the intent to produce designs with desirable performance at a single flight condition are well documented. However, such single-point optimisation exercises tend to lead to poor off-design performance as the quality of any particular geometry is defined only by its efficiency at one flight condition. In aviation, it is of obvious importance that a design performs at least adequately throughout the entire flight envelope; with this in mind, multi-point optimisation in which the performance indicator to be optimised is posed as a weighted-sum of the performance metric at different flight conditions[(1)] has been proven to be a viable approach. This multi-point approach leads to designs with more balanced performance, but ambiguity exists with regards the values that should be used for the weighting coefficients in the objective function definition. Furthermore, at the preliminary design stage it is often desirable for designers to have access to information relating to the possible trade-offs that can be made between performance metrics at different flight conditions, and indeed different points in the aircraft flight envelope. The need for such trade-off information at early stages of the design process has become greater in recent times, with many classes of aircraft being required to perform well over a number of different mission profiles. For example Unmanned Aerial Vehicles (UAVs) and certain business jets may be required to fly missions in which the cruise condition is responsible for the majority of fuel burn, but also perform well in missions for which loiter or climb conditions dominate due to the flexible nature of their use. The development of multi-objective optimisation algorithms based on the Pareto dominance concept has made the generation of the aforementioned trade-off information possible and indeed these multi-objective algorithms have been successfully applied in a number of aerodynamic optimisation exercises[(2,3)].

The emerging interest in laminar flow technologies[(4,5)] raises a number of challenges in how to effectively implement accurate transition solvers within such complex multi-objective environments without the computational effort becoming excessive[(6)]. The use of full stability-based transition solvers alongside a Navier-Stokes or coupled inviscid/viscid solver can induce a substantial increase in computational burden and induce a lack of robustness which can make implementation within an optimisation algorithm non-trivial.

Despite these difficulties, there has been some research into the utilisation of physics-based transition solvers in aerodynamic optimisation studies. Green *et al*[(7)] used inverse design methodologies to match a user-specified target N-factor distribution by altering aerofoil geometry. Gradient-based optimisation of Natural Laminar Flow (NLF) aerofoils in which the adjoints of the

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

full stability Equations were utilised to enable efficient calculation of objective function gradients has been carried out and found to be effective in improving the performance of a baseline design[8,9]. The multi-modal nature of the objective functions in aerofoil design exercises has been verified by Khurana and Winarto[10] and McRoberts *et al*[11] who compared results obtained by local and global optimisers for aerofoil design tasks, and found the global optimisation routines to yield superior results.

However, comparatively little work has been undertaken in the use of multi-point and multi-objective optimisation of laminar flow aerofoils. Khurana and Winarto[10] used a Particle Swarm Optimiser (PSO) and the weighted-sum objective function approach to design multi-point high altitude long endurance (HALE) aerofoils. Gonzalez *et al*[2] used a multi-objective genetic algorithm (MOGA) to design a set of UAV aerofoils that were Pareto-optimal with regards drag minimisation at cruise and loiter flight conditions. In both of these studies the open-source aerofoil analysis tool XFOIL[12] was used as the aerodynamic solver.

In this work, the potential use of a multi-objective optimisation framework for the preliminary design of natural laminar flow aerofoil sections which are less sensitive to leading-edge contamination (minor damage, insect accretion and other) is addressed. The paper presents a novel implementation of a multiobjective genetic algorithm which can be used in order to ascertain the optimal filtered set of leading edge contaminate-tolerant laminar flow profiles based a constrained design case. This is demonstrated through the design of aerofoils for a laminar flow business jet, demonstrating the ability of the method to discover non-intuitive design solutions that manipulate boundary layer transition location to improve performance throughout the flight envelope. The optimisation framework is based on the NSGA-II multi-objective genetic algorithm of Deb *et al*[13] coupled with several Kriging surrogate models to alleviate the computational burden associated with searching the global design space. Once those regions of the design space are identified which are most promising based on the applied constraint set, an intensive local search of these regions is carried out whilst sequentially increasing the order (and therefore flexibility) of the geometry parameterisation. An adaptation of the Class-Shape Transformation (CST)[14] geometry definition technique is used to achieve this.

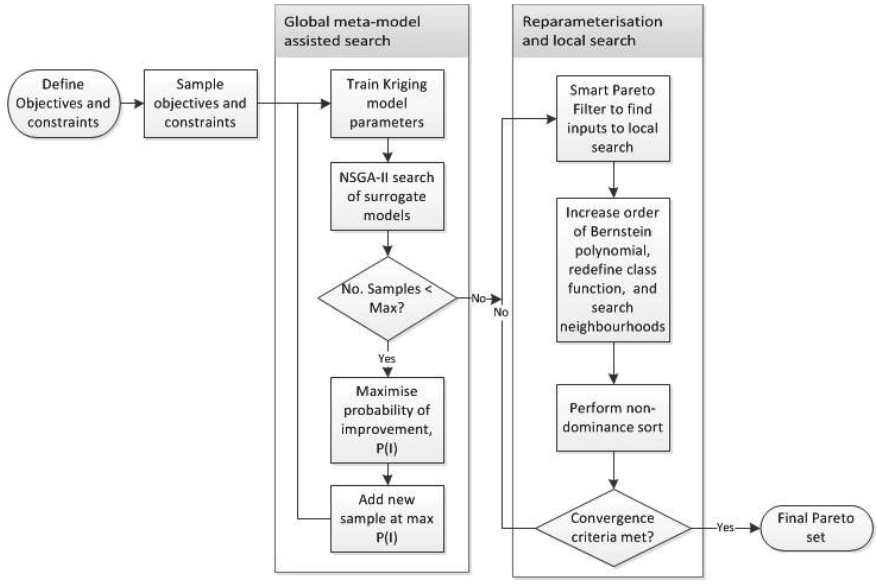


Figure 1. Adaptive metamodel assisted multi-objective optimisation approach.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

2.0 THE OPTIMISATION FRAMEWORK

Multi-objective evolutionary algorithms that converge a population of candidate designs towards a Pareto optimal set are well-suited for providing information on possible performance trade-offs between flight conditions. However, these MOGAs typically require a large number of function evaluations to converge on an optimal set. This proves problematic when utilising MOGAs for aerodynamic design tasks as in many cases the solvers used for appraisal of the performance metrics and constraints are computationally expensive to run. Implementing cheap-to-evaluate surrogate models of the objective and constraint functions provides a viable means to alleviate this computational burden. Figure 1 presents an overview of the process and the interaction between the key solvers; each step in this process is described below.

The Kriging interpolation technique is a well validated surrogate modelling technique that has been applied extensively in engineering design tasks[15]. In this work, Kriging metamodels are built to emulate the objective functions and computationally expensive constraints. The true functions are sampled at a pre-defined number of points within the design space according to an optimal Latin Hypercube (LHC). The LHC is optimised according to the maxi-min criterion. The Kriging models are then constructed by selecting the hyper-parameters that maximise the maximum likelihood measure (discussed in Section 3.4).

A global multi-objective optimisation based on evaluation of the Kriging models is then carried out using the fast and elitist non-dominated sorting genetic algorithm (NSGA-II) of Deb *et al*[13]. The constraints are accounted for by means of a penalty function formulation. After the NSGA-II search has completed a set number of generations, refinement of the surrogate models in promising regions is performed using an adaptive sampling technique[16] which maximises the probability of improvement over the current Pareto front. This refinement of the surrogate models ensures that the available computational resources are not wasted by using the aerodynamic solvers to evaluate individuals that are unlikely to lead to optimal performance. In order to ensure that the Pareto front predicted by the NSGA-II is accurate, a new sample is also taken from the Pareto front at each iteration. This cycle between evolutionary optimisation of the Kriging models and adaptive sampling of the objective functions is iterated until a set number of samples have been reached, or the probability of improvement over the current Pareto front is found to be less than a user defined threshold. It is also worth noting that the probability of constraint satisfaction is also taken into account using the single objective probability of improvement metric described by Jones[17].

The class-shape transformation (CST)[14] geometry definition technique is used to parameterise the aerofoil geometry. A third order Bernstein polynomial is used for the shape function in the initial design space (a total of eight variables define aerofoil geometry), but the geometry is successively refined with increasing order of Bernstein polynomial as the optimisation progresses (Sections 3.2 and 3.7). Performance characteristics of the aerofoils over a range of flight conditions are assessed using the coupled Euler/boundary layer solver MSES[18]. Transition prediction is obtained through the envelope method coupled with the e^N criterion[19,20]. For the test case presented here, maximum lift coefficient at low speed forms one of the constraints, and XFOIL[12] is used for this purpose as it provides accurate predictions of lift coefficient to just beyond the stall at the Mach and Reynolds number considered and requires minimal time for computation. The outputs from these solvers form the sample databases used to build the Kriging surrogate models.

On completion of the global metamodel assisted optimisation, a local multi-objective search with a higher order geometry parameterisation is performed. This local search takes advantage

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

of the mathematical formulation of the CST representation to define a neighbourhood centred about the current Pareto-optimal section under consideration. A smart Pareto filter (SPF) based on the technique presented by Mattson *et al*[21] is used to select the individuals for input to the local search (Section 3.6). As this new local design space is defined using a higher order Bernstein polynomial, a greater degree of flexibility can be achieved in the aerofoil geometry. This is especially important when considering laminar flow aerofoils as transition onset location is extremely sensitive to small changes in the distribution of pressure gradient imposed on the boundary layer. This local search can be iterated until no improvement in the Pareto front is observed, effectively optimising the order of the geometry parameterisation. The re-parameterisation and local search routine is described in detail in Section 3.7.

In industrial scale applications, coupled re-parameterisation and optimisation methods of the type described herein can help bridge the gap between preliminary design, when relatively few design variables are used to identify promising designs, and final detailed design in which a full description of the geometry, including small localised details, is required. It is also worth noting that although the accuracy of the numerical CFD simulations is not increased as the optimisation algorithm progresses, the fidelity of the objective function evaluation is increased from utilisation of Kriging surrogates at the global search stage, to direct calls to the CFD programs during the local search.

To benchmark the performance of the proposed Adaptive Metamodel Assisted NSGA-II (AMMA-NSGA-II) a simple test case was run comparing AMMA-NSGA-II to a traditional MOGA search of the true objective function using the baseline genetic algorithm. Furthermore, a single iteration of the local search was executed to evaluate the potential to improve an existing set of Pareto-optimal aerofoils using this method.

Simultaneous reduction of drag coefficient at two flight conditions (cruise and loiter) was specified with constraints on maximum lift coefficient at a take-off condition, moment coefficient and maximum thickness. The flight conditions at which the drag was to be minimised are defined in Table 1 below and are representative of a simplified HALE UAV mission (similar to the test case presented in Gonzalez[2]). It should be noted that the main purpose of this study is not to develop a fully defined final aerofoil profile, but to demonstrate the capability of the framework to elicit sets of candidate profiles which are less sensitive to variations in the leading edge induced by damage.

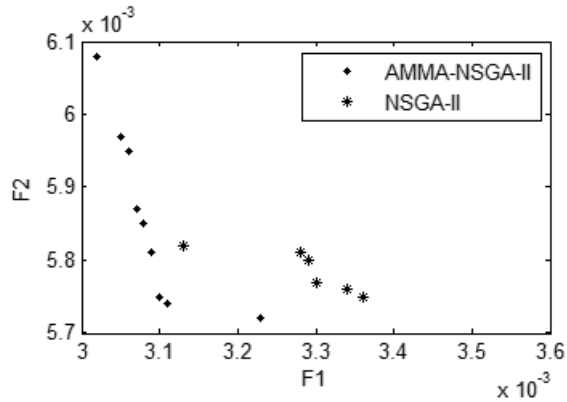


Figure 2. AMMA-NSGA-II versus NSGA-II ($F1 = c_{d\ Cruise}$ + penalty, $F2 = c_{d\ Loiter}$ + penalty functions) Pareto Fronts.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

Table 1
Mission profile for benchmarking test case

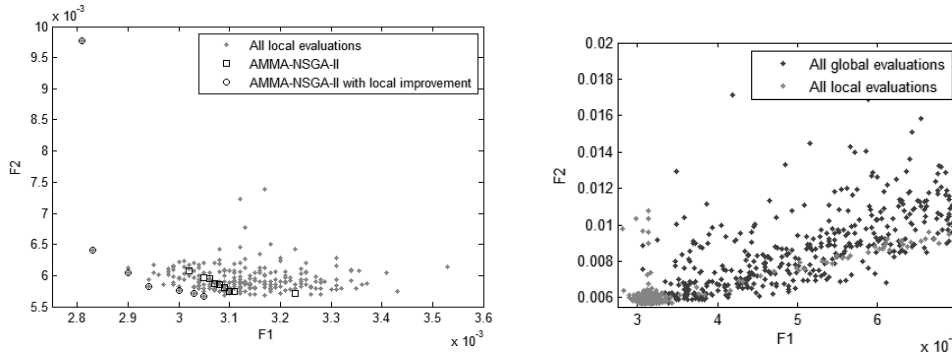
	Mach Number	Lift Coefficient	Reynolds Number
Cruise Condition	0·48	0·17	$6{\cdot}92 \times 10^6$
Loiter Condition	0·15	0·5	$6{\cdot}44 \times 10^5$

For this benchmarking case, the initial global design space was defined by six variables (2nd order Bernstein polynomials used as shape function) each of which was allowed to vary between 0·1 and 0·4 for the upper surface and –0·1 and –0·4 for the lower surface coefficients.

The results obtained by the AMMA-NSGA-II are compared to those from a traditional NSGA-II search of the true objective function. The NSGA-II algorithm is run for 50 generations with a population size of 50 (2,500 evaluations of each objective). The AMMA-NSGA-II was allowed to run until the objective functions were sampled 530 times each. Figure 2 compares the results obtained by NSGA-II and AMMA-NSGA-II:

It can be seen from Fig. 2 that the AMMA-NSGA-II outperforms the traditional search even with significantly less calls of the true function made. This owes largely to the extra information that the adaptive algorithm incorporates in the form of the Kriging models and adaptive sampling routine. It is worth noting that the Pareto front displayed here for the AMMA-NSGA-II is calculated by performing a non-dominance sort on the sample database rather than the results of evaluating the surrogate models; this decision was taken to ensure the results displayed did not correspond to inaccurate predictions from the meta-model which would lead to invalid conclusions being drawn. This also demonstrates the capability of the adaptive sampling algorithm to effectively target promising regions of the design space.

To gain some indication of the potential effectiveness of the local search routine, one iteration was performed with the order of the Bernstein polynomials increased to four (ten variables defining an aerofoil). The neighbourhood around all nine Pareto-optimal solutions was populated with two-hundred individuals meaning that the AMMA-NSGA-II with local search required a total of 2,330 evaluations of each objective function. Figure 3 compares the Pareto front obtained using the metamodel assisted global search to the front after the local search. The clustering of points in the region of the Pareto front by the local search is also illustrated.



(a) AMMA-NSGA-II with local search (b) Point clustering effect of local search

Figure 3. Illustration of local search results in objective function space.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

Figure 3 shows that the local search results in a marked improvement of the Pareto front, with all members of the original front dominated by members of the newly obtained front. The extent of the Pareto front is also improved in terms of both objectives.

Having outlined the optimisation framework and demonstrated its potential using a simple test case, the following sections will be concerned with describing the components of the framework in greater detail with particular attention afforded to the design space adaptation and local search.

3.0 COMPONENTS OF THE OPTIMISATION FRAMEWORK

3.1 The baseline multi-objective genetic algorithm

Utilisation of the concept of Pareto-dominance in multi-objective evolutionary algorithms allows for the gradual convergence of a population of individuals towards a set of non-dominated Pareto-optimal individuals. In this context, one individual is said to dominate another if it is superior to that individual in terms of at least one objective, and not worse in terms of any of the objectives. Pareto-optimal solutions are those which are not dominated by any other individual. Evolutionary algorithms are well-suited to the task of Pareto optimisation as the genetic operators act on a population of solutions meaning that the architecture of the algorithm does not need to be greatly altered; only a re-casting of the fitness function is required.

Srinivas and Deb[22] used non-dominance sorting and crowding distance to compare the fitness of individuals within a population and applied standard bio-inspired operators to converge the population towards a Pareto-optimal set of solutions. Ranking individuals according to non-dominance count ensures convergence towards the Pareto front is achieved, whilst taking into account the crowding distance measure helps to maintain diversity in the population and encourage a spread of solutions across the optimal front. Deb *et al*[13] further improved the algorithm to address an issue with regards the computational expense of the non-dominance sorting routine and also to include elitism; this new improved MOGA was the NSGA-II. A custom real-coded implementation of the NSGA-II is used in this work; polynomial mutation[23] and simulated binary crossover[24] are the genetic operators used. Selection is by a standard binary tournament procedure with crowding distance used to decide the winner when the individuals are in the same non-dominance front.

3.2 Aerofoil parameterisation and design space definition

To enable aerofoil shape optimisation, the geometry of the wing section needs to be parameterised. In this work the class-shape transformation (CST) method proposed by Kulfan[14] is used. The CST method describes the aerofoil geometry as the product of a class function and a shape function. The class function is defined as:

$$C_{N2}^{N1}\left(\frac{x}{c}\right)=\left(\frac{x}{c}\right)^{N1}\left(1-\frac{x}{c}\right)^{N2} \quad \ldots (1)$$

In this definition, $N1$ and $N2$ are constant coefficients that take on different values depending on the type of geometry that is to be represented. For a NACA type aerofoil $N1 = 0{\cdot}5$ and $N2 = 1$. Kulfan[14] recommends decomposing the shape function into component shape functions using the Bernstein polynomial such that:

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

$$S_{r,n}\left(\frac{x}{c}\right) = K_{r,n}\left(\frac{x}{c}\right)^{r}\left(1-\left(\frac{x}{c}\right)\right)^{n-r} \quad \ldots (2)$$

where $r = 0 \rightarrow n$ and $K_{r,n} = \dfrac{n!}{r!(n-r)!}$

A Bernstein polynomial of n[th] order is therefore defined by $n + 1$ terms, each representing a component shape function. Unique aerofoil geometries can therefore be defined simply by assigning a weighting coefficient to each of the component shape functions in accordance with Equation (3):

$$S\left(\frac{x}{c}\right) = \sum_{i=1}^{n} A_i S_i\left(\frac{x}{c}\right) \quad \ldots (3)$$

The y-coordinates of the aerofoil are then given by the product of the class function and the shape function:

$$\left(\frac{y}{c}\right) = C_{N2}^{N1}\left(\frac{x}{c}\right) S\left(\frac{x}{c}\right) \quad \ldots (4)$$

As the framework presented here is intended for use at the preliminary design stage, the design space definition for the initial global search must include a wide range of aerofoil types that are likely to be effective during the different flight conditions that are to be catered for in the objective function definitions. With this in mind, the weighting coefficients corresponding to various well-known sections (NACA 4-series, NACA 6-series, NACA 7-series, RAE 2822, NASA NLF, NASA SC) were computed by performing a least-squares fit to the geometries. Third order Bernstein polynomials were used as shape functions for both upper and lower surfaces during this exercise as this will be the level of parameterisation used in the initial design space for the design case presented in Section 4. The distribution of these coefficients was then used to define the bounds on the 8 design variables as illustrated in Fig. 4. The coefficients defining the upper surface are numbered 1 to 4 and the lower surface coefficients are numbered 5 to 8.

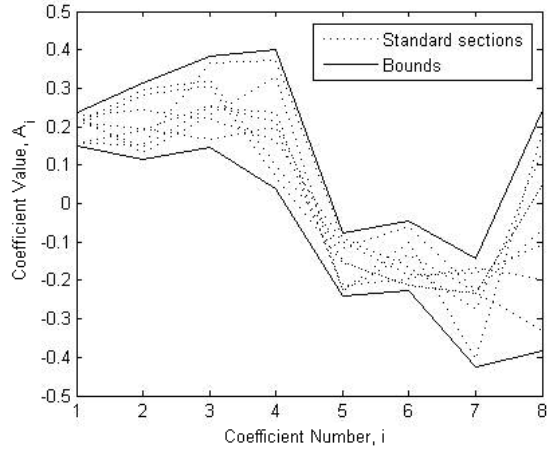


Figure 4. Illustration of bounds on design variables along with a number of commonly used sections.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

The bounds were defined by adding and subtracting 10% of the range of each variable (from the set of coefficients obtained from the least-squares fits) to the maximum and minimum values respectively. These standard sections occur in the first generation of solutions and are seeded as a starting position, but have not (within the studies undertaken) occurred in the final Pareto optimal sets but are substituted with 'superior' solutions during the optimization process.

3.3 Aerodynamic solvers

The coupled compressible Euler and boundary layer solver MSES[18] is used for evaluation of the objective functions in this work. A streamline based Euler solver accounts for external inviscid flow whilst a two-equation integral boundary layer method accounts for viscous effects. A final flow solution is converged upon by implementation of a global Newton method. MSES is also used to calculate the moment constraint at transonic conditions in this work, implemented in order to more closely replicate the design conditions as used by Fujino *et al*[25].

The open-source aerofoil analysis code XFOIL[12] is used to compute the maximum lift coefficient at take-off condition which is implemented as a constraint in the design case. XFOIL is used as opposed to MSES for this constraint as a sweep of angle of attack is required for calculation of maximum lift coefficient and using XFOIL as opposed to MSES will result in considerable computational savings.

In both solvers, transition prediction is performed using an envelope method based on solutions to the full linear stability equations. Amplification rates are calculated based on the boundary layer profile at each discretised station downstream of the neutral point (being the point at which boundary layer disturbances first become unstable) and integrated along the streamwise direction. Transition onset is defined as the location at which boundary layer disturbances are amplified by a factor of e^{N}[24,25]. A value of $N = 9$ is used in this work. Additionally, any solutions that have not achieved convergence are rejected by the optimisation algorithm.·

3.4 Kriging

As has been mentioned, the direct use of the aerodynamic solvers for evaluation of objectives and constraints would require too great a computational effort, particularly in situations where computational intensive solvers are being implemented. For this reason, the Kriging interpolation technique is used to dynamically build surrogates approximating the true functions.

Given a set of observed samples of the true function, $f(v)$, at locations v_i in the design space, Kriging is used to make an estimation, $\hat{f}(v)$, of the true function at an unsampled location v_0 using a weighted linear combination of basis functions. The basis functions are expressed as:

$$h_i = \exp\left(-\sum_{j=1}^{k} \theta_j \left| v_{i,j} - v_{0,j} \right|^{Pj}\right) \quad \ldots (5)$$

where k represents the number of design variables.

In order to give reliable predictions, appropriate values of the parameters in θ and p need to be adopted. In this work, model training is achieved by maximising the ln-likelihood function as detailed by Forrester and Keane[15]. In order to perform this maximisation a single-objective genetic algorithm is used. It is also worth noting that as the Kriging method assumes a Gaussian process, a measure of the mean-square-error (MSE) is given as part of the prediction.

3.5 Adaptive sampling algorithm

The Kriging models are adaptively refined in regions of interest using the multi-objective probability of improvement metric introduced by Keane[16]. The probability of improvement criterion takes

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

advantage of the Gaussian nature of the predictor to calculate the probability that adding a new sample to the database at any given point within the design space will lead to an improvement over the current Pareto front. This is achieved by integrating the joint probability distribution of the objective function predictions over the regions of the objective space that result in augmentation of or dominance over the current optimal front. Keane recommends implementing the probability of improvement metric by maximising the average of probability of augmenting the front (PI_{aug}) and that of dominating current optimal solutions (PI_{dom}); this approach is taken in this work.

The probability of constraint satisfaction is also accounted for using the single objective probability of improvement metric described by Jones[14] by integrating the part of the Gaussian distribution associated with each constraint prediction is across the constraint threshold. The metric used to identify infill points is then given by Equation (6):

$$PI = \left(\frac{PI_{aug} + PI_{dom}}{2} \right) \left(\Pi_{i=1}^{n_{con}} PI_{con}^{(i)} \right) \quad \ldots (6)$$

By finding the set of design variables that maximise the probability of improvement using a single objective optimiser (a genetic algorithm is used in this work), re-sampling the true function at this point and retraining the Kriging models, the surrogates are adaptively refined in promising regions during the optimisation routine.

3.6 Smart Pareto filter

The multi-objective local search implemented in this work consists of searching the neighbourhood around certain Pareto-optimal sections generated using the AMMA-NSGA-II. The optimal solutions from this global search are determined by performing a non-dominance sort on the sample database, and may result in the identification of aerofoils that are extremely similar in terms of geometry

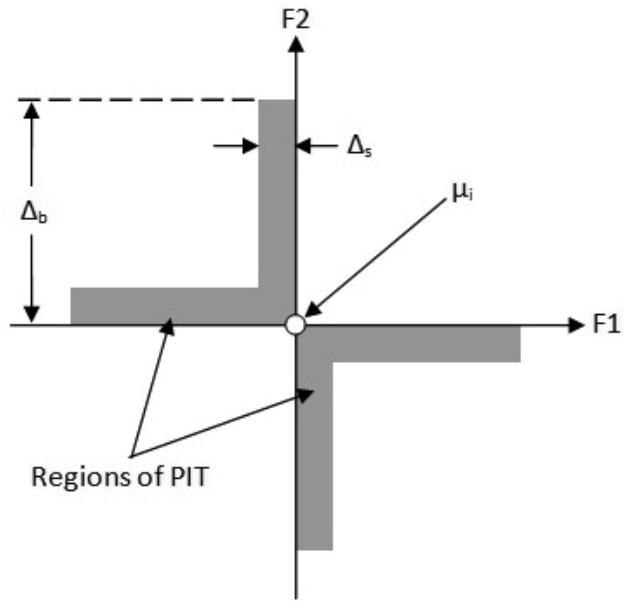


Figure 5. Region of PIT around a Pareto-optimal solution.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

and objective function values. It is also not known a priori how many optimal designs will be identified. For this reason it is not feasible or beneficial to search the neighbourhood around all the optimal individuals from the global search as many of the resulting neighbourhoods would overlap. For this reason the smart Pareto filter[21] is utilised to reduce the Pareto set to a small subset that adequately represent the trade-offs in objective functions across the optimal front; the resulting smart Pareto set then forms the input into the local search algorithm.

The SPF requires the definition of two parameters, Δ_s and Δ_b, that are used to define a region of practically insignificant trade-off (PIT) around any given Pareto optimal solution, μ_i, as shown in Fig. 5.

The SPF is applied to reduce a Pareto-optimal front to a set of smart Pareto solutions using the algorithm described by the following pseudo-code:

```
smart_set=[]
rejected_set=[]
for i = 1 to length(Pareto_set)
        if Pareto_set(i) is in rejected_set
                // do nothing
        else
                add Pareto_set(i) to smart_set
                for j = 1 to length(Pareto_set)
                        if i=j
                                //do nothing (no need to compare individual to itself)
                        elseif Pareto_set(j) is in region of PIT of Pareto_set(i)
                                add Pareto_set(j) to rejected_set
                        end if
                end for
        end if
end for

Output smart_set and rejected_set
```

The smart Pareto filter is applicable to problems with any number of objective functions.

3.7 Re-parameterisation and local search

Having identified the region of the design space in which the Pareto front is situated via implementation of the previously described AMMA-NSGA-II algorithm, a local search of this region using a higher fidelity parameterisation and the 'true' aerodynamic solvers is carried out to further improve the objective function values and ensure all optimum results are based on calls of the true functions.

A re-parameterisation of the true function is performed based on an adaptation of the CST technique; the re-parameterisation results in an aerofoil that is identical in its mathematical description but defined using a higher order Bernstein polynomial. The neighbourhood around this newly parameterised aerofoil is then searched according to an optimal LHC. This is effectively searching a new 'local design space' around the original aerofoil, with a more flexible geometry parameterisation than was used in the initial global search.

Members of the Pareto front generated by the AMMA-NSGA-II are selected for the local search routine using the SPF of Mattson[21]. This ensures distinct individuals are assessed and that the candidate pool is minimised. For each of the individuals selected for input to the local search the steps performed are as follows:

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

1. Individuals are selected from the current Pareto front using the SPF.
2. For each of the selected foils set the class function for the lower and upper surface equal to the coordinates of the lower and upper surface of the selected optimal aerofoil. This differs from the standard CST method in which a fixed analytical function (Equation (1)) defines the class function for both upper and lower surfaces.
3. The order of the Bernstein polynomial used to represent the component shape functions is then increased and the weighting coefficients attached to each component function are set to equal one; this results in the definition of an aerofoil that is mathematically identical to that selected from the smart Pareto set, but with a higher order parameterisation.
4. By setting tight bounds on the weighting coefficients, a small neighbourhood around the section selected in step 1 is defined. As the parameterisation now utilises a higher order shape function, the flexibility that can be achieved when altering the geometry is increased.
5. Each of these constricted design spaces is then sampled according to a space filling technique (optimal LHC) and a non-dominance sort is performed to update the Pareto front.
6. If there is no improvement in the Pareto front over a pre-defined number of iterations the local search is terminated. A maximum number of total local search evaluations is also prescribed. If the Pareto front has been improved by the local search (members of old front dominated by new individuals from local search) return to step 1 and repeat the process until convergence is achieved or a pre-defined number of iterations have been performed.

Having described all the components that make up the optimisation framework, the next section will deal with its application to a realistic design case.

4.0 APPLICATION TO NLF AEROFOIL DESIGN FOR LIGHTWEIGHT BUSINESS JET

Presented in this section are the results garnered from application of the proposed method to a realistic aerofoil design case. The case has been developed based on the HondaJet laminar aerofoil section design requirements[25]. Full details of the design case are presented in Ref. 25 for reference.

- "*Low section profile drag coefficients are desired at* $C_l = 0{\cdot}26$ *for* Re $= 11{\cdot}7 \times 10^6$ *and* M $= 0{\cdot}69$ *(a cruise condition) and at* $C_l = 0{\cdot}35$ *for* Re $= 13{\cdot}6 \times 10^6$ *and* M $= 0{\cdot}31$ *(a climb condition).*"[25]
- "*The section maximum lift coefficient with no flap deflection should be at least 1·6 for* Re $= 4{\cdot}8 \times 10^6$ *and* M $= 0{\cdot}134$*. The loss in maximum lift coefficient due to leading-edge contamination should be less than 7%.*"[25]
- "*The section pitching moment coefficient should be no more negative than –0·04 at* $C_l = 0{\cdot}38$ *for* Re $= 7{\cdot}93 \times 10^6$ *and* M $= 0{\cdot}70$*.*"[25]
- "*The airfoil thickness must be 15% chord to ensure sufficient fuel volume to satisfy the range requirement.*"[25]

These requirements were used to formulate the objective and constraint functions that were to be searched by the optimisation algorithm presented in Section 2; the optimisation problem is posed formally below:

$$\text{Minimise } F1 = C_{d1} + p1 + p2 + p3 + p4$$

$$\text{Minimise } F2 = C_{d2} + p1 + p2 + p3 + p4$$

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

where:

$$if\ C_{l\,max\,3} < 1{\cdot}6$$
$$p1 = 0{\cdot}25\ (1{\cdot}6 - C_{l\,max\,3})$$
$$else$$
$$p1 = 0$$

$$if\ (C_{l\,max\,3} - C_{l\,max\,3\,tr=0}) > (0{\cdot}07 \times C_{l\,max\,3})$$
$$p2 = 0{\cdot}25\ (C_{l\,max\,3} - C_{l\,max\,3\,tr=0})$$
$$else$$
$$p2 = 0$$

$$if\ C_{m\,4} \leq -0{\cdot}04$$
$$p3 = 0{\cdot}25\ (-0{\cdot}04 - C_{m\,4})$$
$$else$$
$$p3 = 0$$

$$if\ th_{max} < 0{\cdot}15$$
$$p4 = 0{\cdot}4\ [0{\cdot}15 - th_{max}]$$
$$else$$
$$p4 = 0$$

In the above definitions the subscripts 1-4 indicate the four flight conditions summarised in Table 2. The '*tr*' subscript in $C_{l\,max\,3\,tr\,=\,0}$ implies transition at the leading edge due to contamination.

Leading edge contamination in this case refers to physical contamination of the aerofoil surface at the leading edge. Some commonly encountered sources of such physical contamination are damage incurred by debris and insect contamination. In the case of laminar flow aerofoil design, this can prove particularly problematic as it can result in a loss of laminar flow, particularly in take-off configurations where large quantities of debris are displaced from the ground. This immediate transition of the boundary layer to turbulence can lead to significant loss of lift during take-off and must be accounted for during NLF design exercises. Indeed, the problem of loss of lift due to leading edge contamination in NLF aerofoils has been acknowledged since early attempts at NLF aerofoil design and continues to be an important design consideration today[26].

Table 2
Summary of flight conditions for NLF business jet design case

Flight Condition	Mach Number	Lift Coefficient	Reynolds Number
1 Cruise	0·69	0·26	$11{\cdot}7 \times 10^6$
2 Climb	0·31	0·35	$13{\cdot}6 \times 10^6$
3 Take-off	0·134	$C_{l\,max}$	$4{\cdot}8 \times 10^6$
4 High trim drag cruise condition	0·7	0·38	$7{\cdot}93 \times 10^6$

The AMMA-NSGA-II algorithm was run until the objective functions had been sampled 500 times. 250 of these samples were generated in the initial sampling plan according to an optimal

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

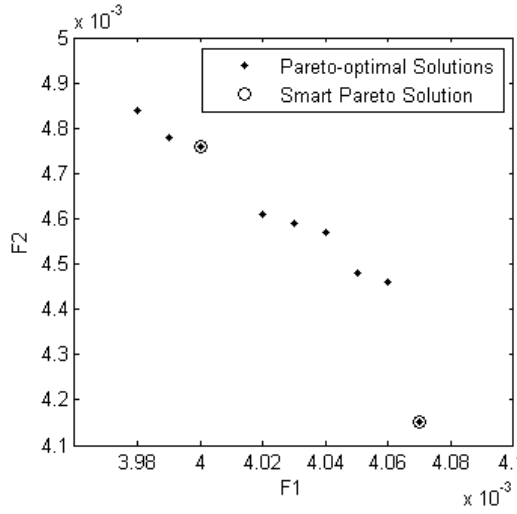


Figure 6. Pareto front from AMMA-NSGA-II (with smart Pareto solutions marked).

LHC with a further 250 samples added using the adaptive refinement routine described in Section 3.5. Surrogate models were also built for $C_{l\ max\ 3}$, $C_{l\ max\ 3\ tr\ =\ 0}$ and $C_{m\ 4}$; each of these constraint functions were also adaptively refined up to a maximum of 500 samples. During the adaptive Kriging-assisted search a population size of 200 was used to search the surrogates using NSGA-II, with additional objective function evaluations added to the sample database every ten generations via maximisation of the probability of improvement metric and by sampling one of the individuals from the current Pareto front using the aerodynamic solvers. At the end of the metamodel assisted search the Pareto front based on a non-dominance sort of the sample database is shown in Fig. 6:

In Fig. 6 the smart Pareto solutions to be used as input into the local search are also indicated; these two individuals have been selected by implementing values of $\Delta_s = \Delta_b = 0{\cdot}0005$ in the SPF and form the input to the local search. The ability of the Pareto filter to effectively reduce the number of unique solutions to reduce the computational expense associated with assessing samples from the Pareto front is clearly demonstrated in this case. The trade-offs achieved in $F1$ are significantly smaller than those achieved in $F2$. This is due to the difficulty of satisfying the challenging constraint set, particularly in finding sections that satisfy the constraint related to loss in maximum lift coefficient due to leading-edge contamination. It is thought that the local search, with the increased flexibility induced in the geometry description by the re-parameterisation scheme, will be able to identify designs that entail greater trade-off in the objectives while still satisfying the constraint set.

Each neighbourhood investigated using the local search was populated with 25 individuals and the search was allowed to run until 1,000 evaluations had been performed. As aerofoil sections that result in violation of the constraints are not of interest at this point in the process, the constraints are evaluated before the objectives with the cheapest constraint functions evaluated first. If any of the constraints are not satisfied the individual is immediately rejected and no further analysis is carried out; this prevents computational resources from being wasted on designs that do not meet the requirements. The values of $\Delta_s = \Delta_b = 0{\cdot}0005$ in the SPF were maintained throughout the local search routine. A total of 12 iterations of the local search were executed before the maximum number of evaluations had been performed. The Pareto fronts produced by successive iterations of the local search are summarised in Fig. 7.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

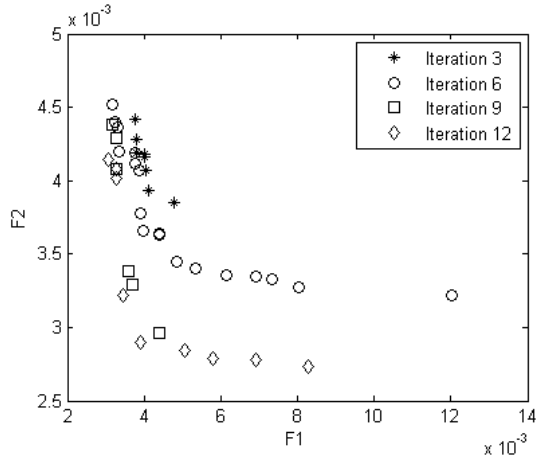


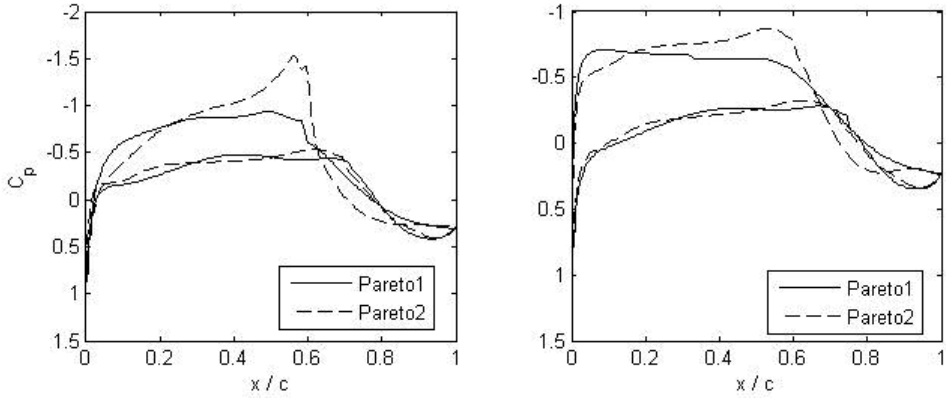
Figure 7. Pareto fronts obtained via local search at successive iterations.

Figure 7 provides a clear illustration of the consistent improvements that are made in the Pareto front throughout the local search routine, both in terms of dominance over previous Pareto-optimal sections and in the extent of the Pareto front in objective function space. Comparing the results of AMMA-NSGA-II (Fig. 6) to those displayed in Fig. 7 shows the ability of local search and re-parameterisation scheme to utilise the additional degrees of freedom in the geometry description to realise large improvements in performance. It is also worth noting that the concentrated neighbourhood search has led to the identification of optimal individuals that not only satisfy the demanding constraint set, but define significant trade-offs in both objectives, bearing in mind that this was not achieved by AMMA-NSGA-II. This owes largely to the increased flexibility that is inherent in the geometry description. For this design case, the order of the shape function is increased by two per iteration of the local search.

It can also be seen that the final Pareto front is populated with seven members. Across the Pareto front, a trade-off of 52·2 drag counts can be achieved in c_{d1}, with a trade-off of 14·1 drag counts possible in c_{d2}. The proposed optimisation framework has proved capable of designing a set of aerofoil sections that efficiently produce low drag in cruise and climb even in the presence of the difficult-to-satisfy constraint set. Inspection of all the results achieved by the local search shows that the objective functions were fully evaluated for 597 individuals, with the constraints on take-off maximum lift coefficient and the allowable degradation in maximum lift due to leading edge contamination proving particularly difficult to satisfy.

The aerodynamic performance of the two individuals at the extremes of the Pareto front will now be examined in order to gain some indication of the mechanisms by which the performance trade-offs across the Pareto front are realised. Again, it should be noted that the purpose of the optimisation framework is to provide insight and guidance to designers to promising design configurations at early design stages for further design refinement when limited design information is available, rather than developing final design solutions. For this purpose the optimal aerofoil achieving the best performance for cruise (objective 1) will be notated as 'Pareto1' and the individual giving the best performance for climb (objective 2) referred to as 'Pareto2', and distinguishing geometric characteristics will be discussed. Figure 8 compares the distribution of pressure coefficient over the two aerofoils at the cruise and climb flight conditions. The aerofoil geometries are illustrated

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press



(a) Cruise condition cp distribution (b) Climb condition cp distribution

Figure 8. c_p distributions for cruise- and climb-optimal aerofoils.

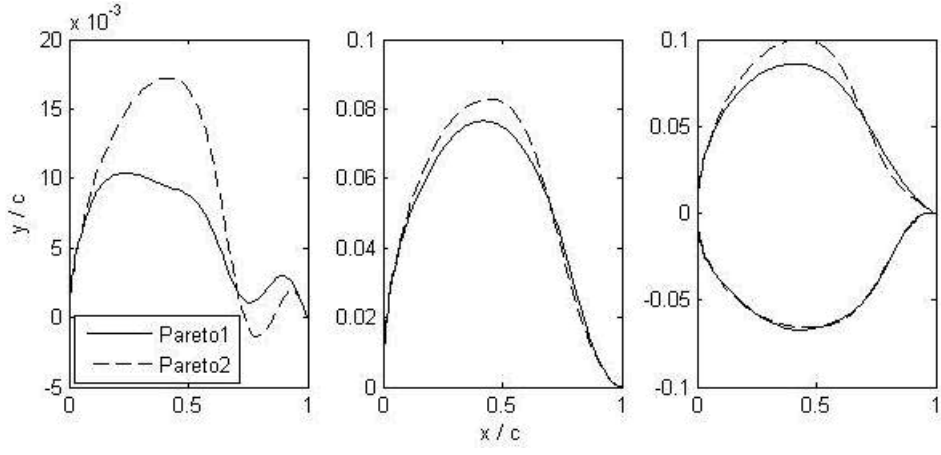


Figure 9. Aerofoil camber distributions, thickness distributions, and geometries for cruise- and climb-optimal aerofoils.

in Fig. 9, with the geometries of all Pareto-optimal sections given to scale in Fig. 10. What is instantly notable is the significant difference in the cruise condition characteristics relative to the differences observed in climb.

Figure 8 shows that the section favouring low drag at the cruise condition produces no shock wave on the upper surface. However, the section favouring the climb condition has a strong shock at approximately 60% of the chord length which will contribute wave drag. Both Pareto1 and Pareto2 exhibit extended regions of favourable pressure gradient on the upper and lower surfaces at the cruise condition.

A convex pressure recovery region is present for both aerofoils resulting in fully attached flow along the along the entire upper surface for Pareto1, although a small region (0·03*c*) of shock-induced separation occurs for Pareto2. The dependence of moment coefficient on cruise drag for all Pareto optimal configurations is shown in Fig. 11, demonstrating that although all aerofoil profiles are similar, as shown in Fig. 10, there is no clear correlation evident.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

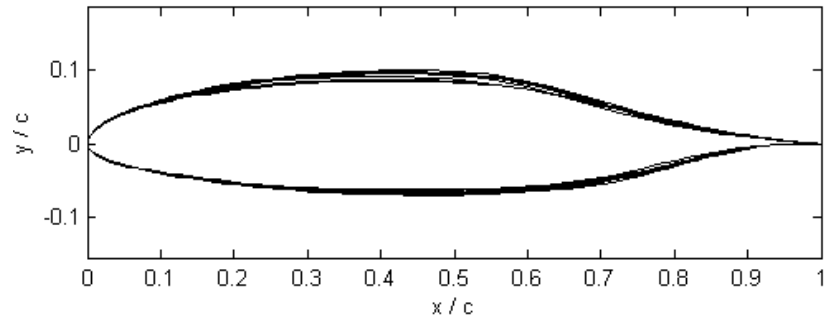


Figure 10. All Pareto-optimal aerofoil geometries.

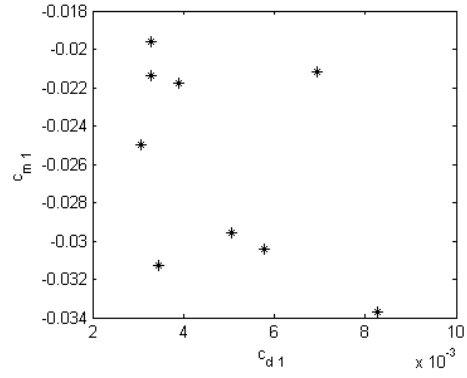


Figure 11. c_m vs c_d 1 for all Pareto-optimal aerofoils.

At the climb condition, both aerofoils show a sharp decrease in c_p at the leading edge, this being slightly more pronounced for Pareto1. For the section Pareto1 adverse pressure gradient dominates downstream of the leading edge. The upper surface pressure distribution for Pareto2 exhibits favourable pressure gradient up to 55% of the chord length. As for the cruise condition, a convex pressure recovery is exhibited by both Pareto1 and Pareto2, resulting in fully attached boundary layer flow over the entire upper surface for both sections.

In terms of the aerofoil geometries, Fig. 9 indicates that the variation in geometry is mostly comprised of differences in both camber and thickness distributions. The camber distribution for both sections is similar in form, with a reflexive distribution near the trailing edge having been realised to achieve the convex pressure recovery. Overall, camber is more pronounced for Pareto2 as it is optimised for the higher lift climb condition. Also, the position of maximum camber is situated further aft for Pareto2 in order to maintain favourable pressure gradient downstream of the leading edge during the climb condition. The thickness distributions are also of similar shape, with the variability being a result of the increased maximum thickness of Pareto2 relative to Pareto1; this enabling it to more easily realise the higher lift coefficient during climb.

In order to further understand and quantify the trade-off information displayed on the Pareto front (Fig. 7), the contributions to the overall drag (c_d) in terms of skin friction drag (c_{df}), pressure drag (c_{dp}) and wave drag (c_{dw}) are illustrated for both flight conditions in Fig. 12. Table 3 shows the locations of boundary layer transition in terms of aerofoil chord for both aerofoils in cruise and climb conditions.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

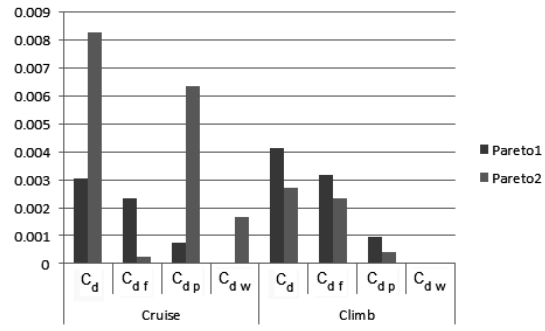


Figure 12. Break down of contributions to drag coefficient.

Figure 12 shows that for the cruise condition the presence of the shock wave on the upper surface of Pareto2 contributes a large amount to the overall section drag coefficient. However, the majority of the increase of drag for Pareto2 relative to Pareto1 at cruise is due to pressure drag increase as a result of the shock induced separation and thickening of the viscous layer. At the cruise condition, Pareto1 actually exhibits slightly increased levels of friction drag relative to Pareto2. This is reflected in Table 3, from which it can be seen that transition onset location on both upper and lower surfaces occur at approximately the same locations for Pareto1 and Pareto2, with both sections achieving extended runs of laminar flow on both surfaces. The high pressure drag contribution observed for Pareto2 at cruise, which occurs as a result of shock induced separation of the boundary layer, signifies the importance of boundary-layer/shock-wave interactions in determining the effectiveness of laminar aerofoils at transonic flight conditions.

As can be seen from Fig. 12, Pareto2 achieves reduced drag relative to Pareto1 during the climb condition as a result of decreased pressure and frictional drag contributions. Table 3 confirms that this is realised by a significantly extended run of laminar boundary layer flow on the upper surface.

Also of interest in this design case is the requirement to satisfy the constraint on loss of $c_{l_{max\ 3}}$ due to leading edge contamination. Fujino *et al*[25] commented that a specially designed leading edge was required to meet this constraint, encouraging upper surface transition near the leading edge at high angles of attack; observing the location of transition location as lift coefficient is increased at the take-off condition for Pareto1 and Pareto2 (Fig. 13) shows that the optimisation algorithm presented in this work has also designed such a leading edge to satisfy this constraint without user intervention. It is this leading edge geometry that causes the c_p suction peak during climb condition that was referred to earlier.

Table 3
Locations of boundary layer transition on upper and lower surfaces

	Cruise		Climb	
	Pareto1	**Pareto2**	**Pareto1**	**Pareto2**
Transition upper	0·5846	0·5970	0·3159	0·6014
Transition lower	0·7059	0·6947	0·7449	0·7178

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

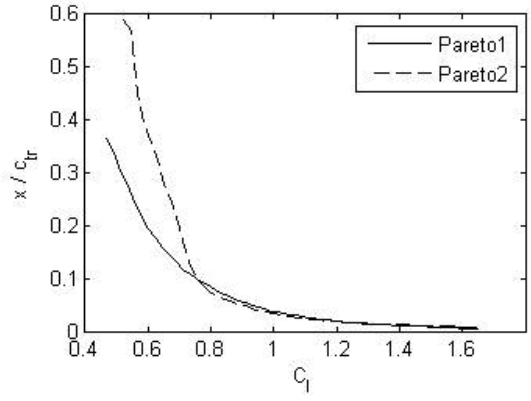


Figure 13. Relationship between lift coefficient and upper surface transition location at take-off.

The fact that the algorithm was able to design an unintuitive reflexive camber distribution that induces significant laminar flow during the climb condition, and a convex pressure recovery, despite the necessary presence of the suction peak at the leading edge shows its ability to manipulate geometry locally to realise changes in transition onset, and therefore pressure and friction drag reductions. It is also worth referring back to Fig. 7 at this point; it can be seen that the sections with lowest drag during climb are not discovered until the latter iterations of the local search routine implying that the increased flexibility in the geometry representation that is achieved using the re-parameterisation scheme is necessary to exploit these small changes in geometry fully.

5.0 CONCLUSIONS

A framework for the optimisation of sets of flight envelope balanced NLF aerofoil sections is presented, benchmarked against traditional approaches and demonstrated on a heavily constrained design problem. An adaptive meta-model assisted multi-objective genetic algorithm is used to identify promising regions of the design space with a small number of function evaluations by implementing Kriging models of objectives and constraints alongside a probabilistic adaptive sampling algorithm. Members of the Pareto front produced by this global search are then selected for input to a local multi-objective optimisation routine using a smart Pareto filter. The local search using re-parameterisation of the aerofoil geometry to allow a greater degree of flexibility in the geometry representation is performed. A neighbourhood is then defined around the individual and explored using the true aerodynamic solvers.

A benchmarking exercise based on a simplified mission profile compares the performance of the algorithm to a traditional multi-objective evolutionary search of the true objectives. The AMMA-NSGA-II algorithm is found to identify a superior Pareto front both in terms of extent in objective function space and dominance over the front produced by NSGA-II whilst using considerably less function evaluations (530 compared to 2,500 in this case). The local search routine also leads to further improvement of the Pareto front in terms of both extent and dominance ranking. The total number of function evaluations carried out by the coupled AMMA-NSGA-II and local search (2,330 evaluations) was less than the NSGA-II search utilised and the results produced by the new algorithm were greatly superior.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

When applied to a heavily constrained test problem based on the requirements of the HondaJet natural laminar flow aerofoil section[25] the algorithm produces some very interesting results. The AMMA-NSGA-II search produced a Pareto front that comprised significant trade-offs in climb performance, but extremely minimal trade-offs in terms of cruise drag due to the difficulty in satisfying the constraint defining maximum allowable loss in $c_{l\,max}$ at take-off due to leading edge contamination. The local search routine was able to identify solutions that represented significant trade-offs in terms of both objective functions. Significant improvements are made in the Pareto front with each successive iteration of the re-parameterisation and local search algorithm.

Analysis of the final set of Pareto-optimal aerofoils show that the algorithm is capable of exploiting changes in the boundary layer transition location to achieve significant reductions in skin friction and pressure drag at certain points within the flight envelope. These trade-offs amounted to 52·2 drag counts at cruise and 14·1 drag counts at the climb condition. A reflexive camber distribution that encouraged a convex pressure recovery to encourage fully attached boundary layer flow over the upper surface was also realised. This was achieved whilst still maintaining the geometrical features necessary for satisfaction of the constraints. For example, the leading edge geometry was able to promote transition close to the leading edge at high angles of attack which was necessary to control the sensitivity of take-off $c_{l\,max}$ to leading edge contamination. Overall it has been shown that the proposed algorithm outperforms traditional approaches in terms of efficiency, and is capable of performing extremely well on a realistic natural laminar flow design problem.

REFERENCES

1. NEMEC, M., ZINGG, D.W. and PULLIAM, T.H. Multipoint and multi-objective aerodynamic shape optimization, *AIAA J*, June 2004, **42**, (6), pp 1057-1065.
2. GONZALEZ, L.F., LEE, D.S., SRIVINAS, K. and WONG, K.C. Single and multi-objective UAV aerofoil optimisation via hierarchical asynchronous parallel evolutionary algorithm, *Aeronaut J*, 2006, **110**, pp 659 -672.
3. CAMERON, L., EARLY, J.M. and MCROBERTS, R. Metamodel Assisted Multi-Objective Global Optimisation of Natural Laminar Flow Aerofoils, AIAA 2011-3001, 29th Applied Aerodynamics Conference, Honolulu, Hawaii, USA, June 2011.
4. ARGUELLES, P. *et al* European Aeronautics: A Vision for 2020, European Commission, Belgium, 2001.
5. BRAUN, M., WICKE, K. and KOCH, A. Analysis of Natural Laminar Flow Aircraft Based on Airline Network Design and Fleet Assignment, AIAA 2011-6807, 11th AIAA Aviation Technology, Integration and Operations (ATIO) Conference, Virginia Beach, Virginia, USA, September 2011.
6. ATKIN, C. Laminar Flow Control: Leap or Creep, AIAA 2008-3986, 38th Fluid Dynamics Conference and Exhibit, Seattle, Washington, USA, June 2008.
7. GREEN, B.E., WHITESIDE, J.L., CAMPBELL, R.L. and MINECK, R.E. Method for the constrained design of natural laminar flow airfoils, *J Aircr*, 1997, **34**, (6), pp 706-712.
8. AMOIGNON, O., PRALITS, J., HANIFI, A., BERGGREN, M. and HENNINGSON, D. Shape optimization for delay of laminar-turbulent transition, *AIAA J*, 2006, **44**, (5), pp 1009-1023.
9. DRIVER, J. and ZINGG, D.W. Optimized Natural-Laminar Flow Airfoils, AIAA 2006-247, 44th AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, USA, January 2006.
10. KHURANA, M. and WINARTO, H. Development and validation of an efficient direct numerical optimisation approach for aerofoil shape design, *Aeronaut J*, 2010, **114**, pp 611-627.
11. MCROBERTS, R., EARLY, J.M., SPENCE, S. and MEDINA, H. Investigation of Utilizing a Single Surface Depression in the Optimisation of NLF Aerofoil Design, AIAA 2010-4679, 28th AIAA Applied Aerodynamics Conference, Chicago, Illinois, USA, June 2010.
12. DRELA, M. XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils, MIT Department of Aeronautics and Astronautics, Cambridge, Massachusetts, USA, 1989.
13. DEB, K., PRATAP, A., AGARWAL, S. and MEYARIVIAN, T. A Fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computing*, **6**, (2), 2002, pp 182-197.
14. KULFAN, B.M. Universal parametric geometry representation method, *J Aircraft*, 2008, **45**, (1), pp 142 - 158.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press

15. KEANE, A.J. Statistical improvement criteria for use in multiobjective design optimisation, *AIAA J*. 2006, **44**, (4), pp 879 - 891.
16. FORRESTER, A.I.J. and KEANE, A.J. Recent advances in surrogate-based optimization, *Progress In Aerospace Sciences*, 2009, 45, pp 50-79.
17. JONES, D.R.A Taxonomy of global optimization methods based on response surfaces, *J Global optimization*, **21**, pp 345-383, 2001.
18. DRELA, M. A User's Guide to MSES 3.00, MIT Department of Aeronautics and Astronautics, Cambridge, Massachusetts, 2004.
19. SMITH, A. Transition, Pressure Gradient, and Stability Theory, In the Proceedings of the 9th International Congress of Applied Mechanics, **4**, pp 234-244, 1956.
20. VAN INGEN, V.A. Suggested Semi-Empirical Method for the Calculation of the Boundary-Layer Region, Technical Report, University of Delft, The Netherlands, 1956. Report VTH71, VTH74.
21. MATTSON, C.A., MULLUR, A.A. and MESSAC, A. Minimal Representation of Multiobjective Design Space Using A Smart Pareto Filter, 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Georgia, USA, September 2002.
22. SRINIVAS, N. and DEB, K. Multiobjective optimization using nondominated sorting genetic algorithms, *Evolutionary Computing*, 1994, **2-3**, pp 221-248.
23. DEB, K. and GOYAL, M. A Combined Genetic Adaptive Search (GeneAS) for engineering design, *Computer Science and Informatics*, 1996, **26**, pp 30-45.
24. DEB, K. and AGRAWAL, R.B. Simulated binary crossover for continuous search space, *Complex Systems*, 1995, **9**, pp 115-148.
25. FUJINO, M., YOSHIZAKI, Y. and KAWAMURA, Y. Natural-laminar-flow airfoil development for a lightweight business jet, *J Aircr*, 2003, **40**, (4), pp 609-615.
26. COLLIER, F. NASA Sponsored Activities in Laminar Flow Technologies for Advanced Transport Aircraft 2nd UTIAS-MITAS International Workshop on Aviation and Climate Change, 27-28 May, 2010, Toronto, Canada.

https://doi.org/10.1017/S0001924000010940 Published online by Cambridge University Press