

AIRFOIL AND WING DESIGN THROUGH HYBRID OPTIMIZATION STRATEGIES

A. Vicini*, D. Quagliarella†

C.I.R.A., Centro Italiano Ricerche Aerospaziali, 81043 Capua, Italy

Real world design problems need robust and effective system-level optimization tools, as they are ruled by several criteria, most often in multidisciplinary environments. In this work a hybrid optimization algorithm has been obtained by adding a gradient based technique among the set of operators of a multiobjective genetic algorithm. In this way it has been possible to increase the computational efficiency of the genetic algorithm, while preserving its favourable features of robustness, problem independence and multiobjective optimization capabilities. The results here illustrated regard aerodynamic shape design problems, including both airfoil and wing design.

1. Introduction

Several techniques are today available for design through numerical optimization;¹ concerning in particular the field of aerodynamic design, beyond methods developed *ad hoc* and characterized by inverse design capabilities, the techniques more properly related to direct optimization include mature gradient based methods, and more recent approaches like automatic differentiation, control theory based methods and genetic algorithms (GAs). Generally speaking, it is not possible to state the superiority of one method over the others, if not with reference to a specific problem that needs to be faced. The characteristics of importance that need to be evaluated are numerous:

- generality of the formulation vs. dependence on the problem;
- robustness, intended as the capability of avoiding local optima, vs. the need of human interaction and expertise;
- capability of multiple objective optimization vs. single-objective one;
- computational efficiency vs. the need of large computational resources.

From this point of view, the choice of one particular optimization technique implies the renunciation of some possible advantages in favour of some others. On the other hand, due to the fact that aerodynamic shape design represents only a part of the

overall design of a flying vehicle, and that the need for an effective multidisciplinary approach to the design task is arising, it is important for an optimization tool to combine as much as possible all the favourable characteristics stated above while avoiding the shortcomings. In this sense, a hybrid approach to optimization, in which techniques of different nature are used at the same time, may result extremely beneficial. Hybrid optimization may in fact exploit the most favourable features of the methods which are combined while masking the corresponding shortcomings.²

In this work a hybrid optimization tool, developed by incorporating a gradient based optimization routine among the operators of a multiobjective genetic algorithm,³ is applied to aerodynamic shape design problems. Genetic algorithms belong to the class of evolutionary optimization procedures, which finds its philosophical basis in Darwin's theory of survival of the fittest.⁴ In an attempt to mimic the process of biological evolution, a set of design alternatives, representing a population in this metaphorical transposition, is let evolve through successive generations so as to promote the individuals which better adapt themselves to the environment, i.e. those which better meet the design requirements. Each element is characterized by the value of its *fitness*, which is the measure of how fit it is for the given environment - in other words, how good the corresponding solution is for the problem at hand. The process of evolution is realized in the reproduction phase using a selection criteria driven by the value of the fitness of the individuals, so that bias is allocated to the best fit members of the population. The individuals selected for reproduction are recombined using genetic operators (*crossover*, *mutation*), so that a combination of their most desirable characteristics

*Research Engineer. E-mail: a.vicini@cira.it

†Research Engineer. E-mail: d.quagliarella@cira.it. Member AIAA.

Copyright © 1998 The American Institute of Aeronautics and Astronautics Inc. All rights reserved.

may be obtained in the offsprings, and hence elements characterized by higher fitness are produced in subsequent generations.

These search methods rely only on the evaluation of the fitness of the elements and do not require the computation of gradients; therefore, they are less susceptible to pitfalls of convergence to a local optimum, and can successfully deal with disjoint or non-convex design spaces. Moreover, they are capable of facing the problem of multiple objectives optimization in a straightforward fashion, using the notion of domination among solutions.⁵ These characteristics make GAs very attractive optimization tools, and explain the considerable growth of interest which has been devoted to them in recent years for applications of engineering interest.⁶

The major weakness of GAs lies in their relatively poor computational efficiency, as they generally require a very high number of evaluations of the objective function. For this reason, the use of GAs may become unpractical when this evaluation is “expensive”, as happens for aerodynamic optimization applications where the solution of complex partial differential equation systems is necessary.

Coupling a genetic algorithm with a different optimization technique can be an effective way to overcome its lack of efficiency while preserving its favourable features. Many different strategies to hybridize the GA can be realized; the simplest one is that of using the best solution found by the GA as starting point for a subsequent optimization with the other method adopted.⁷ However, a closer interaction between the different algorithms, rather than the two-stage optimization described, may more favourably combine the best features of both methods, and provide results better than those obtainable using either of the two techniques.

Some representative results obtained with the hybrid genetic algorithm (HGA) will be demonstrated for aerodynamic shape design problems, including both airfoil and wing optimization test cases.

2. Hybridization of the genetic algorithm

A simple GA may by itself be considered as the combination of two different search techniques, namely crossover and mutation, that are characterized by different behaviours when searching the parameter space. Crossover generates the candidate solutions (offsprings) through a combination of two existing ones (parents); the solutions thus obtained, independently from the way the parents are selected and combined, can be very far from the starting ones. Thus, crossover is a powerful

tool to search the design space and single out the region where the global optima lie, but it lacks the capability of effectively refine the sub-optimal solutions found. On the other hand, mutation has a more local effect, since the modifications it produces are generally small in the coded parameter space. Hence, mutation has two important roles in simple GAs:

- to provide the capability to effectively refine sub-optimal solutions;
- to re-introduce in the population the alleles lost by the repeated application of crossover, maintaining population diversity.

However, the rate of mutation needed for these two tasks may be different; in particular, while mutation is very good for maintaining population diversity, its refining capabilities may not be optimal for every class of problems. There is in fact a broad class of problems, namely the ones where the fitness function is differentiable, for which gradient based techniques are much more efficient to locally improve a given solution. This suggests the introduction of a gradient based routine among the set of operators of the GA; mutation is then prevalently left with the role of keeping the diversity among population elements at an optimum level.

The genetic algorithm developed adopts a bit string codification of the design variables; anyway, this does not prevent the use of operators requiring real number list encoding, such as extended intermediate crossover and word level mutation.⁸ In these cases the binary string is decoded into a real number list, the operator is applied and the set of modified variables is encoded back into a bit string. This scheme allows the use of a free mix of different type of operators; among these, as said before, a routine performing a gradient based optimization (with a conjugate gradients technique) has been included, and called “hill climbing operator” (HcO). The HcO operates as follows: through the application of the selection, crossover and mutation operators, an intermediate generation is created from the current one; afterwards, if the hybrid option is activated, some individuals may be selected and fed into the HcO to be improved, and then introduced into the new generation, as sketched in fig. 1. Regarding the choice of the elements to be fed into the HcO, in the case of single objective optimization three different strategies are possible:

- 1) only the best fit individual of the current generation is chosen;
- 2) a number of elements determined by an assigned probability is picked using the selection operator;

- 3) a number of elements determined by an assigned probability is picked in a purely random fashion.

Of course, these strategies determine different levels of selection pressure, decreasing from strategy #1 to strategy #3; the relative performance will therefore depend on the optimization problem. The above described scheme can be naturally extended for multiobjective optimization. In this case the elements are not ranked on the basis of a scalar fitness function, but are just divided into two classes: the dominated and the not-dominated ones.⁵ The set of not-dominated individuals (Pareto front), updated after each new generation, is composed by all potential solutions of the problem, satisfying the design criteria at different levels of compromise. When multiobjective problems are formulated, strategy 1) becomes the (random) selection of a number of elements determined by an assigned probability from the current set of Pareto optimal solutions, while strategies 2) and 3) remain the same. Of course, the HcO is by its nature capable of dealing only with scalar objective functions; thus, when multiobjective problems are faced, the objective function fed into the HcO is obtained through a weighted linear combination of the n problem objectives, i.e. as $obj = \alpha obj_1 + (1 - \alpha) obj_2$ in the case of $n = 2$. The weighting factor α can be chosen at random or assigned explicitly to favour one of the objectives.

It must be remarked, on the basis of what previously stated, that the aid provided by the HcO simply consists in its capability to improve to some extent the selected individuals; in other words, the role of the HcO is that of introducing improvements which will then be processed and exploited by the GA, which remains the driving engine of the procedure. From this point of view, the use of the HcO has to be limited to the minimum necessary to provide the desired effect, which is that of improving the convergence characteristics of the procedure without causing premature convergence to local minima. On the basis of these considerations, the use of the HcO has been subject to the following rules:

- the hybrid mechanism sketched in fig. 1 is not used at each generation, but only after each assigned block of K generations;
- it is not necessary – and it would be detrimental for the computational performances – for the HcO to carry out each time a converged optimization. Therefore, only one or two gradient iterations are generally prescribed.

Furthermore, as the HcO has basically to behave as an improved mutation operator, the beneficial ef-

fects of hybridization can be obtained also by making it operate only on a subset of the active design variables, thus reducing the number of objective function evaluations that the HcO needs to carry out each time; the success in this case will depend on the degree of cross correlation among the design variables. The total number of evaluations of the objective function needed, N_e , can be estimated as follows:

$$N_e = N_{\text{pop}} N_{\text{gen}} \left[1 + \frac{p_{\text{hco}}}{K} N_{\text{it}} (\eta N_{\text{var}} + 4) \right] \quad (1)$$

where N_{pop} and N_{gen} are the population size and the total number of generations, p_{hco} is the HcO probability, K is the frequency in terms of number of generations for the activation of the HcO, N_{it} is the number of gradient iterations, N_{var} is the number of design variables, and $\eta \in [0, 1]$ the factor determining the size of the subset of design variables that is passed to the HcO. In the applications that will be illustrated, the design variables that need to be frozen are chosen at random each time the HcO is used; in this way, a different subset of design variables is passed each time to the HcO.

3. Applications to airfoil design

In the airfoil inverse design problem, a pressure distribution is given corresponding to a design point determined by the values of Mach number and angle of attack, and the geometry of the airfoil producing this target pressure distribution must be found. In this case, the objective function to be minimized is computed by:

$$obj = 10 \int_S \left(c_p - c_p^{(t)} \right)^2 ds \quad (2)$$

where c_p and $c_p^{(t)}$ are the current and target pressure distributions, respectively, and S is the current airfoil contour; the fitness is then obtained as $f = 1/obj^2$. A full potential transonic flow solver, with non-conservative formulation, has been used to calculate the flow field. The airfoil geometry is represented by means of two 5th order B-spline curves, for the upper and lower parts. The coordinates of the control points of the B-spline constitute the design variables;⁸ 7 control points are used both for the upper and lower surfaces of the airfoil, including those fixed at the leading and trailing edges, for a total of 18 design variables (the first control points at the leading edge can move only in direction y). The problem here presented consists in the reconstruction of the CAST-10 airfoil⁹ at $M = 0.765$, $\alpha = 0$. This problem has been solved using a NACA 0012 as initial guess, which can be considered an absolutely generic starting point.

The design variables have been encoded using 8 bit strings (giving a chromosome length of 144 bits), and a 50 individuals population evolved for 100 generations; the hybrid strategies have been activated so as to select on average only one individual every other generation, and carry out 2 gradient iterations ($p_{\text{hco}} = 0.02$, $K = 2$, $N_{\text{it}} = 2$, $\eta = 1$). Hence, to consider the same total number of objective function evaluations, the hybrid strategies must be judged approximately at generation 70. Two different GAs, characterized by the set of operators described in table 1, have been used with and without hybridization. Figure 2 illustrates the convergence histories, each one averaged over 10 successive trials characterized by different starting populations; the convergence history obtained by the application of the gradient based method by itself is also shown in the same figure; besides, it must be noted that a restart procedure had to be used in this case to take the solution out of a local minimum where it got stuck after a few iterations.

As can be seen, for a given GA, hybridisation is always beneficial, meaning that a better result can be found with the same amount of computations, or that the same result can be obtained with a substantial reduction of computation needed (ranging in this case from 30 to 75%). In particular, strategy #1, when the HcO is applied only to the best fit individuals, appear as the less effective, probably due to an excessive selection pressure. At the same time, the behaviour of the gradient based method is considerably improved from the point of view of the robustness.

Another important characteristic that needs to be considered is the statistical dispersion of the results obtained starting from different initial populations; in fact, if it is correct to judge the convergence characteristics of a given GA by averaging the results of a number of runs, from an application-oriented point of view it is more important for the algorithm to guarantee satisfactory convergence performances even on a single run basis. Figure 3 shows all the values of the objective function obtained at the end of each of the 10 different runs, for each one of the algorithms used; it can be observed how the scatter of the results provided by both basic GAs is much higher than that obtained using the corresponding hybrid algorithms. In particular, the best behaviour from this point of view is obtained when the elements to be fed into the HcO are chosen at random, so that the level of selection pressure is not increased too much.

The same runs have then been repeated using $\eta = 0.5$; in this way, the HcO acts on a subset of 9 design variables out of 18, that are chosen at random each time. The convergences obtained are

illustrated in fig. 4, limitedly to the hybrid strategies #2 and #3, i.e. those giving the best performances. In this case the actual number of fitness evaluations has been used for the x axis, in order to better compare the results. We see how freezing some of the design variables has a positive effect when GA #1 is used, whereas for GA #2 convergence is slowed down to some extent. Considering that the design variables for this problem are strongly cross-correlated, as it is not possible to move one control point of the B-splines independently from the others, this result shows that this approach can generally be used with success. In fig. 5 one of the pressure distributions obtained is shown together with the target and initial ones.

An example of multiobjective optimization is then presented, consisting in reducing the wave drag of the airfoil while keeping the corresponding pitching moment under control, for a fixed lift coefficient and maximum thickness. The airfoil chosen as initial geometry is the RAE 2822,¹⁰ at a design point $M = 0.78$, $c_l = 0.75$. The constraint on lift coefficient is satisfied by letting the flow solver find the angle of attack that produces the specified lift, while the thickness of the airfoil is scaled to the desired value after each geometry modification; in this way, every solution is a feasible one. The two objective functions have been evaluated as $obj_1 = c_{dw}/c_l^2$, $obj_2 = c_m^2$. A population of 100 individuals was let evolve for 100 generations; selection was carried out by means of a 3 steps random walk, with one-point crossover ($p_c = 1$) and bit mutation ($p_m = 0.02$). Differently from the inverse design previously described, the geometry of the airfoil has been represented as a linear combination of the initial one, y_o , and a number of modification functions, y_i :

$$y = y_o + \sum_{i=1}^N x_i y_i \quad (3)$$

The coefficients x_i of this combination are the design variables; the functions y_i have been obtained as the difference between the initial geometry and the geometries of other airfoils chosen from a database, so that a particular aerodynamic effect can be associated to each design variable. The allowable range assigned was $x_i \in [-0.2, 1.2]$, $i = 1, N$, and 12 design variables were used. The same run was then repeated using the HcO, on average, on one element per generation, chosen at random from the current Pareto front, and carrying out only one gradient iteration. The run in this case was stopped at generation #86, in order to establish the comparison for the same total number of evaluations of the objective functions.

The Pareto fronts thus obtained are illustrated in fig. 6, together with the starting point. It can be

seen how the result provided by the hybrid algorithm is a Pareto front characterized by solutions of higher quality, and more uniformly distributed; only 12 of the solutions found by the GA are not dominated by those obtained with the HGA.

4. Applications to wing design

Wing design is a highly multidisciplinary task; the use of designer expertise is therefore necessary to obtain realistic results, unless the various design criteria and off-design considerations can be included in the formulation of the optimization problem. The multiobjective optimization approach offers great advantages for these kind of problems, avoiding the need of arbitrarily interrelating the different design criteria into a single scalar objective function.

Genetic algorithms have already been applied to the problem of planform wing design, taking into account aerodynamic and structural requirements. In Ref. [11] structural rigidity considerations are included in the optimization, but a single-objective GA is used, with a selection of design variables which is not representative for a complete definition of the planform shape. In Ref. [12] non-planar wing shapes are allowed to maximize the L/D ratio with the condition that the wing doesn't break under the applied loads. In both cases, the aerodynamic models are limited to subsonic flow, with structural models based on simple beam theory. In Ref. [13] a parallel Pareto GA is used for the planform optimization of a transonic wing, minimizing aerodynamic drag and structural weight, and maximizing tank volume; however, the dimensions of the design space are limited by the use of only 3 design variables.

In this work the HGA has been applied to the optimization of the shape of a wing for transonic flow conditions, modifying both the planform and the wing section. The results here presented have been obtained by coupling the HGA with a finite-difference full potential flow solver.¹⁴ First, the wing planform design has been accomplished by minimizing aerodynamic drag, which is both induced and wave drag, and structural weight, at a given Mach number $M = 0.85$ and lift coefficient $c_L = 0.5$. The starting point chosen is a straight, untwisted and untapered wing of aspect ratio $AR=7$, with a RAE 2822 airfoil; for simplicity, the wing planform is maintained trapezoidal, so that all geometric characteristics vary linearly from the root section to the tip. A total of 5 design variables have been used: 4 of these act directly on the wing planform, namely the taper ratio λ , the sweep angle at 25% of the chord Λ , the aspect ratio AR and the twist angle θ ; moreover, the thickness

at the wing root has also been included among the design parameters, while the thickness at the wing tip has been fixed at $t/c|_t = 10\%$. The wing surface is kept constant, so that the average wing loading is not changed during optimization. In table 2 the initial values of the design parameters are reported together with the prescribed allowable ranges. The wing twist is distributed symmetrically between the root and the tip; in other words, a twist angle θ corresponds to an increase of local incidence of $\theta/2$ at the tip, and a decrease of $\theta/2$ at the root. The wing weight is computed using the algebraic equation of Ref. [15]; this equation combines analytical and empirical (statistical) methods, and shows design sensitivity and prediction accuracy that make it possible to use it with success for preliminary design. As can be seen from table 2, two separate runs have been carried out exploring separately the positive or negative sweep design spaces; in fact, the choice of a positive or negative swept wing is based on considerations of different nature, including stability and handling characteristics.

The selection has been carried out through a 3 steps random walk, with one-point crossover ($p_c = 1$) and bit mutation ($p_m = 0.1$); a population of 64 individuals was let evolve for 50 generations. The HcO has been used on one element for each generation, selected from the current Pareto front, with $N_{it} = 1$. The Pareto fronts obtained are illustrated in fig. 7 (where W_o is a reference weight), together with the planform of the wings corresponding to the extremities and to the center of the fronts. It can be seen that, for a given value of aerodynamic drag, the negative swept wings are heavier than the corresponding positive swept ones; therefore, almost all the solutions with a negative sweep angle would be dominated if the two Pareto fronts were merged. The fronts are populated by 116 and 100 individuals, in the case of positive and negative swept wings, respectively. The use of the HcO doesn't prevent the development of the complete Pareto front; on the contrary, the solutions are uniformly distributed along the fronts without the need of niching techniques. In order to evaluate how well these solutions are representative of the real Pareto fronts, the same test case has been solved using the gradient based method by itself, with the problem formulated through the weighted linear combination approach (i.e. the same used by the HcO): $obj = \alpha obj_1 + (1 - \alpha) obj_2$; 5 different values for α have been used: 1, 0.7, 0.5, 0.3, 0. The solutions thus obtained are compared in fig. 8 with the Pareto fronts provided by the HGA. As can be observed, these solutions lie at most on the Pareto Fronts, and in some cases they fall in the dominated solutions region. It is also interesting to observe that in neither cases the gradient method

is capable of finding the solution of minimum drag, i.e. that corresponding to $\alpha = 1$.

In fig. 9 the values of the design parameters of the solutions belonging to the Pareto fronts are shown as a function of aerodynamic drag. As can be expected, the sweep angle varies in an almost linear fashion from the maximum allowable values (positive or negative) to zero. Similarly, the aspect ratio is at its maximum at the low-drag end of the front, and rapidly diminishes to the minimum as drag increases. It can be seen how changing the aspect ratio from 8 to 6 implies an increase of aerodynamic drag of about 80%; this increase is composed for 60% by induced drag, and by wave drag for the remaining 40%. The taper ratio remains approximately at the minimum allowable value, $\lambda = 0.1$, for most part of the front, assuming higher values only for the solutions corresponding to minimum drag; in the case of positive sweep minimum drag is obtained for a taper $\lambda = 0.5$, whereas for negative sweep a higher value is necessary, $\lambda = 0.78$. The role of twist is essentially that of redistributing the spanwise loading so as to better approach the elliptic distribution; this explains the opposite sign of the twist angle that is obtained when positive or negative swept wings are considered. It can also be observed how higher values of twist are necessary in the latter case. Finally, the behaviour of the thickness at the root section appears less intuitive; only at the low-weight end of the front a clear trend can be observed, with an almost linear increase of drag with the thickness.

As anticipated, after optimization of the planform a further improvement of the aerodynamic characteristics has been obtained by modifying the shape of the wing section. One of the solutions belonging to the Pareto front has been selected as starting point; attention has been focused on the positive sweep angles, and the geometry chosen lies approximately at the center of the front, being characterized by $c_D/c_L^2 = 0.776$ and $W/W_o = 0.65$. The wing section has been modified using the same shape functions technique described in §3; 12 design variables have been used also in this case, and for simplicity the wing profile has been maintained constant in the spanwise direction. As modifying the wing section may have a strong impact on the aerodynamic characteristics but not on the structural weight, which is going to remain (almost) constant, the optimization problem in this case has been formulated so as to reduce wave drag with control on pitching moment; the latter in fact determines the level of trim drag. The design objectives have then been formulated as $obj_1 = c_{Dw}/c_L^2$, $obj_2 = (c_M - 0.5)^2$; like in the previous case, the lift coefficient has been fixed to $c_L = 0.5$, and the maximum thickness has been maintained at the value

obtained by the previous run at each spanwise station. The same GA parameters used for the wing planform optimization have been adopted, except for the mutation rate which has been reduced to $p_m = 0.04$, and for the population size which has been increased to 100. The Pareto front obtained is illustrated in fig. 10, where some of the corresponding wing section shapes are also shown. Depending on the actual design requirements, it is now possible to extract from this front the solution with the desired characteristics. In particular, in fig. 11 the drag rise curve (at $c_L = 0.5$) of the wing before optimization of the section is compared with those of three wings extracted from the front: the low-drag end of the front, corresponding to an unconstrained optimization, and the two solutions characterized by the same pitching moment and wave drag coefficients, respectively, of the initial wing. As can be seen the first two of these solutions provide an overall improvement of the drag rise curve; at the design point $M = 0.85$ the reduction of wave drag is 32 drag counts for the unconstrained solution, and 22 for the fixed c_M one. On the other hand, when the drag coefficient is kept constant so that a reduction of c_M can be achieved, lower wave drag values are obtained for Mach numbers lower than the design one, but a steeper increase at higher Mach numbers. Finally, in fig. 12 the pressure distributions on three wing sections, for the initial wing and the low-drag end of the front, are shown.

5. Conclusions

In most practical applications, design problems are governed by several criteria, most often deriving from different disciplines; to approach such design tasks, robust and effective system-level optimization tools are needed. Genetic algorithms are characterized by a number of favourable features that make them attracting for this class of problems; besides, multiobjective optimization, which is a peculiar feature of GAs, appears particularly suited for multidisciplinary environments, as it allows to determine sets of Pareto solutions in the design space where tradeoffs can be conveniently examined a-posteriori. In fact the generated solutions (Pareto front) represent different levels of compromise among the design goals or constraints. Therefore, the designer can make his/her choice introducing an a-posteriori selection criteria. The flexibility of the design process can thus be increased, as the need of interrelating criteria of different natures is avoided, and the effect of changing constraints can be evaluated off-line.

In this work an effective algorithm for multiobjective applications has been developed through a hybrid approach, by coupling a multiobjective ge-

netic algorithm to a gradient based operator. Applications to multiobjective airfoil and wing design have been presented. The basic idea to use a gradient based routine in the fashion of a genetic operator derives from the observation that mutation by itself is not very effective as a refinement operator, leading to generally poor convergency speed; on the other hand, it has been demonstrated how the beneficial effects of the HcO can be exploited even when its use is considerably limited, in terms of number of elements processed and computations carried out for each element. For the class of problems that has been investigated, significant improvements have been obtained both with respect to simple genetic algorithms, in terms of computational efficiency, and with respect to gradient based approaches in terms of robustness. In particular, it has been possible to use the HGA with success even in the case of multiobjective problems, when a weighting approach must be used to compose a scalar objective function each time resort is made to the HcO.

References

- ¹Optimum Design Methods for Aerodynamics, AGARD-R-803, Nov. 1994.
- ²Davis, L., *Handbook of genetic Algorithms*, Van Nostrand Reinhold, New York, 1991.
- ³Quagliarella, D., and Vicini, A., "Coupling Genetic Algorithms and Gradient Based Optimization Techniques", in Quagliarella D. et al., editors, *Genetic Algorithms and Evolution Strategies in Engineering and Computer Science*, John Wiley & Sons Ltd., England, Nov. 1997, pp. 289–309.
- ⁴Holland, J. H., *Adaptation in Natural and Artificial Systems*, The University of Michigan Press, Ann Arbor, Michigan, 1975.
- ⁵Goldberg, D. E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, Massachusetts, Jan. 1989.
- ⁶Ahmed, Q., Krishnakumar, K., and Neidhoefer, J., "Applications of Evolutionary Algorithms to Aerospace Problems - A Survey", in *Computational Methods in Applied Sciences '96*, John Wiley & Sons Ltd., England, 1996, pp. 236–242.
- ⁷Poloni, C., "Hybrid GA for Multi Objective Aerodynamic Shape Optimisation", in Winter G. et al., editors, *Genetic Algorithms in Engineering and Computer Science*, John Wiley & Sons Ltd., England, Dec. 1995, pp. 397–416.
- ⁸Vicini, A., and Quagliarella, D., "Inverse and Direct Airfoil Design Using a Multiobjective Genetic Algorithm", *AIAA Journal*, Vol. 35, No. 9, Sep. 1997, pp. 1499–1505.
- ⁹CAST-10-2/DOA 2, *Airfoil Studies Workshop Results*, NASA CP 3052, 1989.
- ¹⁰Cook, P. H., McDonald, M. A., and Firmin, M. C. P., "Aerofoil RAE 2822 — Pressure Distributions, Boundary Layer and Wake Measurements", in *AGARD AR-138*, Paper A6, AGARD, May 1979.
- ¹¹Doorly, D. J., Peiró, J., and Oesterle, J.-P., "Optimisation of Aerodynamic and Coupled Aerodynamic-Structural Design Using Parallel Genetic Algorithms", in *6th AIAA/NASA/USAF Multidisciplinary Analysis & Optimization Symposium*, AIAA Paper 96-4027, Bellevue, Seattle, WA, U.S.A., Sep. 1996.
- ¹²Anderson, M. B., and Gebert, G. A., "Using Pareto Genetic Algorithms for Preliminary Subsonic Wing Design", AIAA Paper 96-4023, Sep. 1996.
- ¹³Obayashi, S., Yamaguchi, Y., and Nakamura, T., "Multiobjective Genetic Algorithm for Multidisciplinary Design of Transonic Wing Planform", *Journal of Aircraft*, Vol. 34, No. 5, 1997, pp. 690–693.
- ¹⁴Jameson, A., and Caughey, D. A., "Numerical Calculation of the Transonic Flow Past a Swept Wing", Technical Report NASA-CR-153297, NASA, June 1988.
- ¹⁵Torenbeek, E., "Development and Application of a Comprehensive, Design-Sensitive Weight Prediction Method for Wing Structures of Transport Category Aircraft", Technical Report LR-693, TU Delft, Sep. 1992.

	GA #1	GA #2
selection operator	random-walk, 5 steps	random-walk, 3 steps
crossover operator	extended intermediate	one-point
crossover probability	1	1
mutation operator	bit	word
mutation probability	0.02	0.02

Table 1 - Set of operators of the two GAs used

design variable	initial value	allowable range	selected wing
λ	0.0	[0.1, 1.0]	0.10
Λ	1.0	[0.0, ± 50]	39.6
θ	0.0	[-10, 10]	-5.2
AR	7.0	[6.0, 8.0]	6.41
$t/c \mid_r \%$	12	[12, 15]	12.1

Table 2 - Design parameters for the wing planform optimization

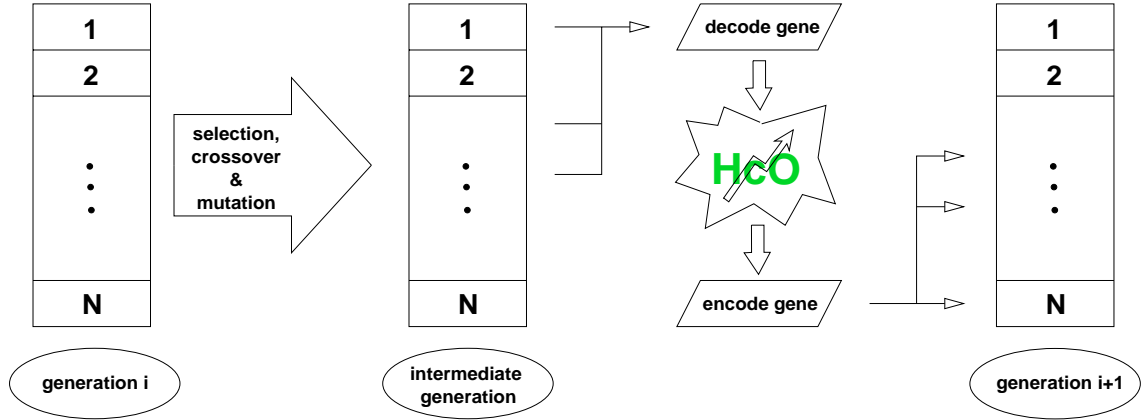


Fig.1 - Sketch of the hybrid genetic algorithm

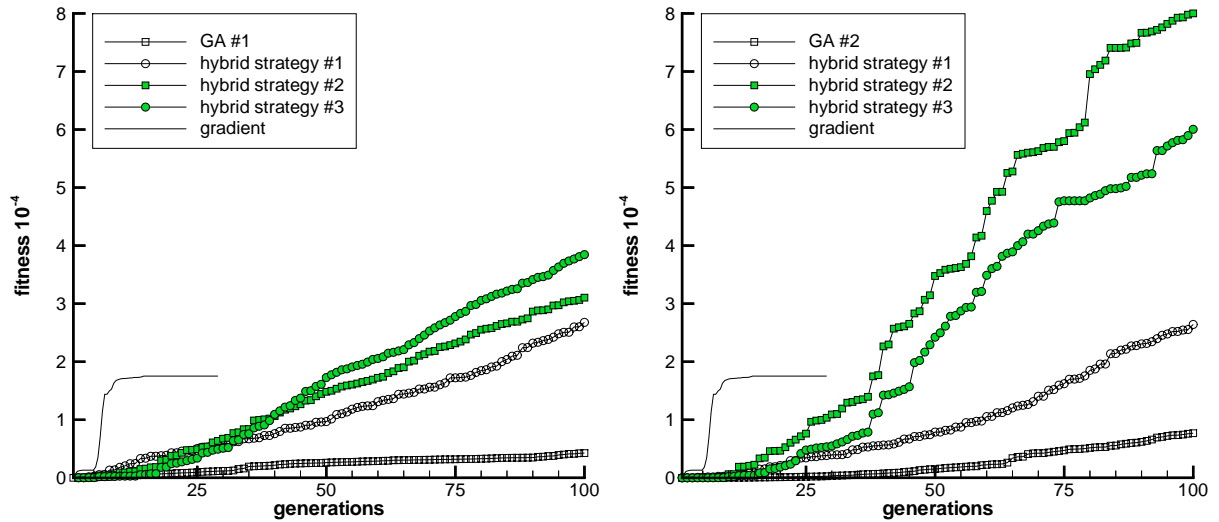


Fig.2 - Convergence histories for the CAST 10 inverse design problem




Fig. 3 - Scatter of the results obtained in 10 different runs for the CAST 10 inverse design problem




Fig. 4 - Comparison of the convergence histories obtained by letting the HcO operate on all design variables or on a 50% subset




Fig. 5 - Target, initial and computed pressure distributions




Fig. 6 - Pareto fronts obtained for the RAE 2822 optimization problem

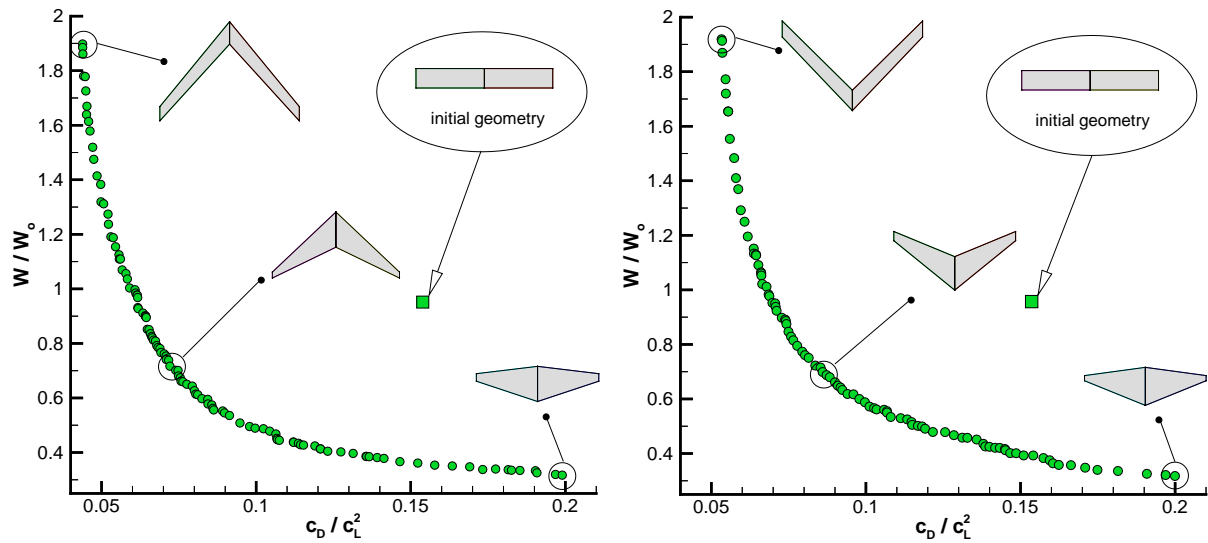


Fig. 7 - Pareto fronts obtained for the wing planform optimization




Fig. 8 - Comparison between the Pareto fronts and the results obtained through the gradient based method




Fig. 9 - Design parameters for the solutions on the final Pareto fronts (continued)




Fig. 9 - Design parameters for the solutions on the final Pareto fronts (concluded)




Fig. 10 - Pareto front obtained for the wing section optimization




Fig. 11 - Drag rise curve of the original wing, and of wings selected from the Pareto front




Fig. 12 - Pressure coefficient on the original and optimized (unconstrained) wings.
(a) - $2y/b=0.16$; (b) - $2y/b=0.44$; (c) - $2y/b=0.76$.