HYBRID OPTIMISATION SCHEMES FOR THE CLEARANCE OF FLIGHT CONTROL LAWS

Prathyush P. Menon * **Declan G. Bates** *
Ian Postlethwaite *

* *Control and Instrumentation Research Group,*
Department of Engineering,
University of Leicester, Leicester, U.K., LE1 7RH
Email: ppm6, dgb3, ixp @le.ac.uk

Abstract: The application of two evolutionary optimisation schemes to the problem of flight clearance of nonlinear control laws for highly augmented aircraft is described. Hybrid versions of both schemes incorporating local gradient-based optimisation are also developed and evaluated. Comparisons of computational complexity and global convergence are conducted by checking for violations of a nonlinear clearance criterion for a detailed simulation model of a high performance aircraft with a full authority flight control law. The recently introduced differential evolution approach, when appropriately augmented with local optimisation methods, is shown to have significant potential for improving the reliability and efficiency of the current industrial flight clearance process.

Keywords: Flight control, Optimisation, Robustness Analysis, Simulation

1. INTRODUCTION

Modern high performance aircraft are often designed to be naturally unstable due to performance reasons and, therefore, can only be flown by means of a flight control system which provides artificial stability. As the safety of the aircraft is thus dependent on the controller, a formal certification or clearance process must be completed using a detailed simulation model of the aircraft and its flight control system before any flight tests can be conducted. The goal of the clearance process is to demonstrate that a set of selected criteria expressing stability and handling requirements is fulfilled under all foreseeable operating conditions. In this process, for each point of the flight envelope, for all possible configurations and for all combinations of parameter variations and uncertainties, violations of all clearance criteria and the worst-case result for each criterion must be found. Based on the clearance results, flight restrictions are imposed where necessary.

Flight clearance of control laws for high performance aircraft is a very lengthy and expensive process, due to the many different combinations of flight parameters (e.g. large variations in mass, inertia, centre of gravity positions, highly nonlinear aerodynamics, aerodynamic tolerances, air data system tolerances, structural modes, failure cases, etc.) that must be investigated so that guarantees about worst-case stability and performance can be made. Faced with limited time and resources, the current flight clearance process employed by the European aerospace industry uses a griding approach, whereby the various clearance criteria are evaluated for all combinations of the extreme points of the aircraft's uncertain parameters (Fielding *et al.*, 2002). This process is then repeated over a griding of the aircraft's flight envelope. Clearly, the effort involved in the resulting clearance assessment increases exponentially with the number of uncertain parameters. Another difficulty with this approach is the fact that there is no guarantee that the worst case uncertainty combination has in fact been found, since (a) it is possible that the worst-case combination of uncertain parameters does not lie on their extreme points, and (b) only a few selected points in the aircraft's flight envelope can be checked.

This paper presents a new approach to the flight clearance problem based on the use of two different hybrid optimisation techniques, which has the capability to significantly improve both the reliability and efficiency of the current flight clearance process. This study focuses on the evaluation of

Table 1. Aircraft Model Uncertain Parameters

Parameter	Bound	Description
Δ_{mass}	[-0.1 +0.1]	variation in aircraft mass from nominal one (9100 kg) [%]
Δ_{xcg}	[-0.075 +0.075]	variation in position of center of mass [m]
$\Delta_{C_{m_{\delta_e}}}$	[-0.005 +0.005]	uncertainty in pitching moment due to elevator deflection [1/rad]
$\Delta_{I_{yy}}$	[-0.025 +0.025]	uncertainty in aircraft inertia around y-axis from nominal one (81000 kg·m^2) [%]
$\Delta_{C_{m_\alpha}}$	[-0.05 +0.05]	uncertainty in pitching moment due to AoA [1/rad]

a nonlinear handling criterion, which is described in detail in the next section.

2. FLIGHT CLEARANCE APPLICATION

2.1 ADMIRE Aircraft Model

The aircraft model used in the present study is the ADMIRE (Aero-Data Model In a Research Environment), a non-linear, six degree of freedom simulation model (Forssell *et al.*, 2001), developed by the Swedish Aeronautical Research Institute (FOI) using aero data obtained from a generic single seated, single engine fighter aircraft with a delta-canard configuration. The ADMIRE simulation model is augmented with a full-authority flight control system and also includes engine dynamics and actuator models. The model includes a large number of uncertain aerodynamic, actuator, sensor and inertia parameters, whose values, within specified ranges, can be set by the user.

In the ADMIRE, the aircraft dynamics are modelled as a set of twelve first order coupled nonlinear differential equations:

$$\dot{x}(t) = f(x(t), u(t), \Delta);\ \ y(t) = h(x(t), u(t)) \quad (1)$$

where $x(t)$ is the state vector with twelve components, i.e., velocity, angle of attack (AoA), sideslip angle, angular rate , attitude angle, and position vectors. Δ represents the uncertain parameters and Table 1 shows the uncertain parameters considered in this paper. $y(t)$ is the output vector, and $u(t)$ is the control input vector, whose components are left and right canard deflection angle, left and right inboard/outboard elevon deflection angle, leading edge flap deflection angle, rudder deflection angle, landing gear status (extract/retract), and vertical and horizontal thrust vectoring. The control input is determined by

$$u(t) = g(x(t), y_{\mathrm{REF}}(t)) \quad (2)$$

where $g(\cdot,\cdot)$ is an industry standard flight control law, which is provided with the ADMIRE model, and $y_{\mathrm{REF}}(t)$ is the reference demand that consists of the pilot inputs such as pitch stick, roll stick, rudder pedal, and thrust demands. Equations (1) and (2) together represent the closed loop dynamics of the aircraft with the flight control law in the loop.

The augmented ADMIRE operational flight envelope is defined up to Mach 1.2 and altitude 6000 meters (Forssell *et al.*, 2001). The longitudinal control law is gain scheduled over the whole flight envelope with respect to Mach and altitude variations and is designed to ensure robust stability and handling qualities over the entire flight envelope. The model also contains rate limiting and saturation blocks as well as nonlinear stick shaping elements in its forward path.

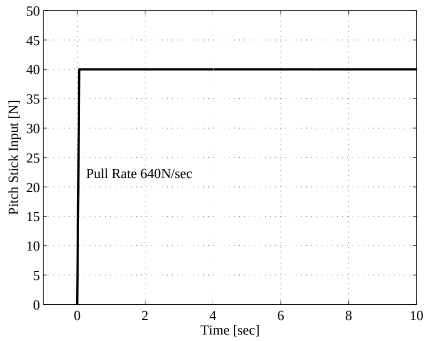

Fig. 1. Pitch Stick Pull Command

2.2 Nonlinear clearance criterion

The clearance criterion considered in this study is the AoA-Limit exceedence criterion (Fielding *et al.*, 2002; Forssell, 2003). For this criterion, it is required to identify the flight cases where, for the pull-up manoeuvre shown in Figure 1, the maximum overshoot occurs in AoA. In particular, the combination of uncertainties that yields the largest exceedence of the defined limits must be identified. The test aims to assess the effectiveness of the incidence limiting scheme in the control system, in terms of the peak overshoot in AoA that occurs in response to the specified manoeuvre. Figure 1 shows the specified pitch stick command, a rapid pull in longitudinal stick to a defined level (40N) at a 640N/sec stick rate with stick hold for 10 seconds. The present analysis aims to estimate the clearance criterion (Fielding *et al.*, 2002):

$$\alpha_{\max} = \max(\alpha(t)) \quad \text{for} \quad t \leq 10 \text{ [sec]} \quad (3)$$

for all possible combinations of aircraft parametric uncertainty.

2.3 Optimisation-Based Flight Clearance

In this paper, the flight clearance problem defined above is formulated as an optimisation problem and solved using two different global and hybrid local-global optimisation schemes. Since this and many other clearance criteria must be checked over a huge number of aircraft envelope points and configurations, it is imperative to find the most computationally efficient approach to the problem. Previous efforts to apply optimisation methods to this problem, ((Fielding *et al.*, 2002); Chapter 7) have revealed that the nonlinear optimisation problems arising in flight clearance, while having relatively low order, often have multiple local optima and expensive function evaluations. Therefore, the issue of whether to use local or global optimisation, and the associated impact on computation times, is a key consideration for this problem.

In (Fielding *et al.*, 2002) Chapter 21, local optimisation methods such as SQP (Sequential Quadratic Programming), and L-BFGS-B were used to evaluate a range of linear clearance criteria for the HIRM+ (High Incidence Research Model) aircraft model. In (Fielding *et al.*, 2002) Chapter 22, global optimisation schemes such as genetic algorithms (GA's), adaptive simulated annealing (ASA) and multi coordinate search (MCS) were also applied to evaluate nonlinear clearance criteria for the same aircraft model. In addition, in (Forssell, 2003) global optimisation methods such as GA's and ASA were applied to the ADMIRE model with a different flight clearance criterion. In this paper, we demonstrate conclusively, for a realistic, industry-standard aircraft simulation model, that incorporation of local optimisation methods into global optimisation algorithms, can significantly reduce computation times *and* improve convergence to the true global solution. In addition, we show the effectiveness of the relatively new differential evolution (DE) global optimisation technique over more standard GA's when applied to the flight clearance problem.

For the rest of this paper, the ADMIRE model is trimmed at Mach 0.4 and altitude 3000 meters in straight and level flight. Once the trim is achieved, the pull up manoeuvre shown in Figure 1 is applied to the model and the optimisation cost function is given by Equation (3), i.e., maximum AoA.

3. GLOBAL OPTIMIZATION

3.1 Genetic Algorithms

GA's are general purpose stochastic search and optimisation procedures that use genetic and evolutionary principles. They are based on the assumption that the evolutionary process observed in nature can be simulated on a computer to generate a population of fittest candidates (Goldberg, 1989). In genetic search techniques, a randomly sourced population of candidates undergoes a repetitive evolutionary process of reproduction through selection for mating according to a fitness function, and recombination via crossover with mutation. A complete repetitive sequence of these genetic operations is called a generation. To use this evolutionary method, it is necessary to have a means of encoding the candidate as an artificial chromosome, as well as a means of discriminating between the fitness of different candidates. A fitness function is thus defined to assign a performance index to each candidate - this function is specific to the problem and is formed from the knowledge domain.

For the problem considered in this paper, the optimisation variables are the uncertain parameters given in Table 1. Each optimisation variable, or gene, is binary coded depending on a desired accuracy level, presently fixed at 1e-6, and combined sequentially to form the chromosome, which represents a potential candidate solution. A typical chromosome length used in the present study is 105 bits, with 5 genes each of 21 bits. The search starts from an initial random number of candidates of size N_{size}, which in this case was fixed at 50. The candidates from the current generation are qualified to produce the successive generations depending on the selection scheme. A roulette wheel selection scheme with a selection probability of 0.6 is applied in this study. During crossover, the recombination operator ensures mixing up of the information content between two different binary coded chromosomes. A single point crossover with a probability of 0.9 is used here. The point of crossover is determined randomly over the length of bits. Mutation introduces random variations in the population in the search space, by randomly flipping a bit value in the binary GA's. The probability of mutation is usually kept low and in this study was fixed at 0.05. The number of maximum generations is used as the termination criterion and is fixed at 100 generations, giving a total number of 5000 simulations (100 times 50 simulations).

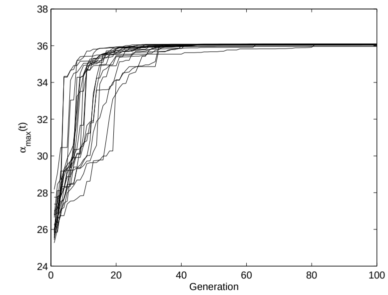

Fig. 2. Generation vs. Best fitness for the GA

The recent survey paper (Fleming and Purshouse, 2002) reports that GA's have become a very popular search and optimisation technique for problems in control engineering with large as well as small parameter search spaces. Due to their stochastic nature, GA's can be expected to have a much better chance of converging to a global optimum. The reader is referred to (Goldberg, 1989) for more details of different GA operators, binary coding schemes and the theory of genetic search.

3.1.1. Results 25 different trials, each involving 5000 simulations of the closed loop aircraft model, were conducted in this part of the study. The GA's were seen to converge to the true global solution $\left[\Delta^*_{mass}\,\Delta^*_{xcg}\,\Delta^*_{C_{m_{\delta_e}}}\,\Delta^*_{I_{yy}}\,\Delta^*_{C_{m\alpha}}\right]$ = [0.100 0.0750 0.050 0.18309 0.050] in 8 of the 25 trials. Note that, at the global solution, one of the uncertain parameters I_{yy} lies inside, and not at the vertices of the hyperbox (pentacube) defined by the uncertain parameters. In the other 17 trials the optimisation converged to points near the global solution. Figure 2 shows the best fitness over the number of generations for the multiple trials. The rate of convergence to the true global solution in 25 different trials over the 100 generations, can be assessed from figure 2.

3.2 Differential Evolution

Differential evolution (DE) is a relatively new global optimisation method, introduced by Storn and Price in (Storn and Price, 1997). It belongs to the same class of evolutionary global optimisation techniques as GA's, but unlike GA's it does not require either a selection operator or a

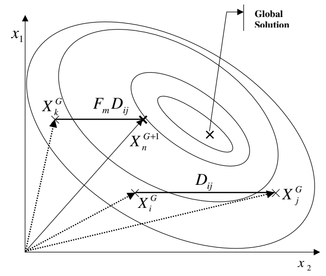

Fig. 3. DE mutation strategy

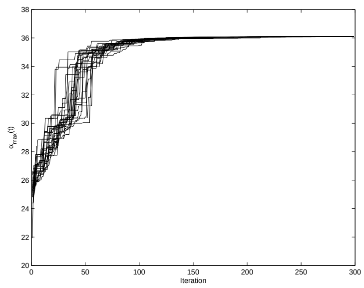

Fig. 4. DE robustness Iterations vs. Best cost

particular encoding scheme. Despite its apparent simplicity, the quality of the solutions computed using this approach is claimed to be generally better than those achieved using other evolutionary algorithms, both in terms of accuracy and computational overhead (Storn and Price, 1997).

The methodology has recently been applied to several problems in different fields of engineering design, with promising results. For example, the DE methodology has been applied to find the optimal solution for a mechanical design example formulated as a mixed integer discrete continuous optimization problem(Lampinen and Zelinka, 1999). In (Rogalsky *et al.*, 2000), the DE method has been applied and compared with other local and global optimization schemes in an aerodynamic shape optimization problem for an aerofoil. The DE methodology consists of the following four main steps 1) Random initialization, 2) Mutation 3) Crossover 4) Evaluation and Selection. These steps will be described in detail in the sequel:

3.2.1. Random initialization Like other evolutionary algorithms, DE works with a fixed number, N_p, of potential solution vectors, initially generated at random according to

$$\mathbf{x}_i = \mathbf{x}^L + \rho_i(\mathbf{x}^U - \mathbf{x}^L), \; i = 1, 2, ..., N_p \quad (4)$$

where $\mathbf{x}^U$ and $\mathbf{x}^L$ are the upper and lower bounds of the parameters of the solution vector and ρ_i is a vector of random numbers in the range [0 1]. N_p is fixed at 12 in the current study. Each $\mathbf{x}_i$ consists of elements $(x_{1i}, x_{2i}, ..., x_{di})$, which are the uncertain parameters defined in Table 1. The dimension d of the optimization problem considered is 5. The fitness of each of these N_p solution vectors is evaluated using the cost function given in Eq. 3.

3.2.2. Mutation The scaled difference vector $F_m D_{ij}$ between two random solution vectors $\mathbf{x}_i$ and $\mathbf{x}_j$ is added to another randomly selected solution vector $\mathbf{x}_k$ to generate the new mutated solution vector $\bar{\mathbf{x}}_n^{G+1}$ as given in Eq. 5. F_m is the mutation scale factor, a real valued number in the range [0, 1], (fixed at 0.8 in this study). The subscript G represents the iteration number.

$$\bar{\mathbf{x}}_n^{G+1} = \mathbf{x}_k{}^G + F_m D_{ij}, \quad D_{ij} = \mathbf{x}_i{}^G - \mathbf{x}_j{}^G \quad (5)$$

Figure 3 shows a simple two dimensional example of the mutation operation used in the DE scheme. The difference vector D_{ij} determines the search direction and F_m determines the step size in that direction from the point $\mathbf{x}_k$.

3.2.3. Crossover During crossover, each element of the n^{th} solution vector of the new iteration, $\mathbf{x}_n{}^{G+1}$, is reproduced from the mutant vector $\bar{\mathbf{x}}_n^{G+1}$, and a chosen parent individual $\mathbf{x}_n{}^G$ as given in Eq. 6, where $j = 1, ..., dn$ and $i = 1, ..., N_p$. Note that $\bar{\mathbf{x}}_n^{G+1}$ has elements $(\bar{x}_{1n}^{G+1}, \bar{x}_{2n}^{G+1}, ..., \bar{x}_{dn}^{G+1})$ and $\mathbf{x}_n{}^G$ has elements $(x_{1n}{}^G, x_{2n}{}^G, ..., x_{dn}{}^G)$.

$$x_{ji}{}^{G+1} = \begin{cases} x_{ji}{}^G, \text{if random number} > \rho_c \\ \bar{x}_{ji}^{G+1}, \text{otherwise;} \end{cases} \quad (6)$$

$\rho_c \in [0, 1]$ is the crossover factor, which is fixed at 0.8 in the present study.

3.2.4. Evaluation and Selection After crossover, the fitness of the new candidate $\mathbf{x}_n{}^{G+1}$ is evaluated. The cost function is as given in Eq. 3. If the new candidate $\mathbf{x}_n{}^{G+1}$ has a better fitness than the parent candidate $\mathbf{x}_n{}^G$, then $\mathbf{x}_n{}^{G+1}$ is selected to become part of the next iteration. Otherwise $\mathbf{x}_n{}^G$ is selected and subsequently identified as $\mathbf{x}_n{}^{G+1}$.

3.2.5. Termination criteria The maximum number of iterations is considered as the termination criterion. Here the maximum number of generations is fixed at 50.

3.2.6. Results Figure 4 shows the best fitness found over iterations, for 25 trials. In contrast to the trials using the GA's, all 25 trials with the DE algorithm converged to the true global solution giving the maximum AoA overshoot. Compared to the GA's, DE was seen to offer significantly improved convergence properties, while the reduced number of initial random starting points meant that the total number of simulations required in each trial was also significantly reduced (3600 against 5000). Another advantage of this method is the reduced number of optimisation parameters that must be adjusted by the user.

4. HYBRID OPTIMISATION

Global optimisation methods based on evolutionary principles are generally accepted as having a high probability of converging to the global or near global solution, if allowed to run for a long enough time with sufficient initial candidates and reasonably correct probabilities for the evolutionary optimisation parameters. The rate of convergence is often very slow, however, and moreover, there is still no guarantee of convergence to the true global solution. Local optimisation methods, on the other hand, can very rapidly find optimal solutions, but the quality of those solutions entirely depends on the starting point chosen for the optimisation routine. In order to try to extract the best from both schemes, several researchers have proposed combining the two approaches ((Davis, 1991; Yen *et al.*, 1995; Lobo and Goldberg, 1996). In such hybrid schemes there is the possibility of incorporating domain knowledge, which gives them an advantage over a pure blind searcher based on evolutionary principles such as GA's. In (Menon *et al.*, 2004), a hybrid GA (HGA) scheme was developed using a switching strategy which was originally proposed in (Lobo and Goldberg, 1996), and applied successfully to a nonlinear flight clearance problem. In the next section, we compare the performance of this HGA scheme with a novel hybrid DE (HDE) scheme developed for this study.

4.1 Hybrid GA

The HGA scheme is based on the idea of associating with both the global and local methods, a reward, or gain. The reward associated with a method is a measure of how well the method helped in providing a solution which is better than the one previously found. The reward associated to each optimisation scheme will determine the probability for that optimisation scheme to be chosen. The reward for each optimisation scheme keeps varying depending on how well it has performed. A simple way to assign a reward, is with a weighted geometric average. The following equation is used to update the weighted reward for each optimisation scheme (Lobo and Goldberg, 1996):

$$W^{k+1}_{\mathrm{GA/Local}} = W^{k}_{\mathrm{GA/Local}}(1-c) + cR^{k}_{\mathrm{GA/Local}} \quad (7)$$

where W^k and R^k are the weighted reward and the reward at the iteration k, respectively, and c is a constant in [0, 1]. $R^{k}_{\mathrm{GA/Local}}$ is decided based on the improvement in the best solution attained over each iteration. If one knows at each time step which optimisation method is going to give most improvement towards the global solution, that particular method can be chosen to accelerate the convergence. When it is not known beforehand, a decision has to be taken based on the previous reward and by calculating the associated probability. The algorithm is summarized in Table 2.

Due to the frequent occurrence of local maxima in this type of problem, initially, the GA's should have a higher probability to be chosen than the local algorithm. Hence, initially the weights for GA's and the local algorithm are given as 0.9 and 0.1, respectively. The local algorithm used in the present study is the function *"fmincon"* which finds the constrained minimum of a scalar function of several variables starting at an initial estimate. In the present analysis, constraints are due only to the upper and lower bounds of the uncertainty in the variables. A medium scale optimisation scheme is chosen where the gradients are estimated by the function itself using the finite difference method. The function uses the sequential quadratic programming (SQP) method - for further details of the *"fmincon"* optimisation strategy, the reader is referred to (*Optimization Toolbox User's Guide, Version 2*, 2000).

Table 2. Hybrid Genetic Algorithm

(1) Initialize $W^0_{\mathrm{GA}} = 0.9$, $W^0_{\mathrm{Local}} = 0.1$, $c = 0.3$, $k = 1$, set the calculation mode "Search", the number of confirmation zero, and generate initial population for GA
(2) While the confirmation number is less than a certain number (e.g. 20)
  (a) Calculate P^k_{GA}, (8)
  (b) (Flip Coin) = a random number between zero and one
  (c) If (Flip Coin) $< P^k_{\mathrm{GA}}$ then run GA and update W^k_{GA}, (7)
  (d) else choose the local algorithm with the following initial guess
    (i) If the calculation mode is "Search", choose one randomly from two best in the population,
    (ii) else choose one randomly from the subset of population where the distance of each element from the current best is out of 1σ (standard deviation of the population from the current best)
    (iii) Update W^k_{Local}, (7)
  (e) If the cost does not improve,
    (i) Initialize the following every five confirmation: population, $W^0_{\mathrm{GA}} = 0.5$, $W^0_{\mathrm{Local}} = 0.5$, $c = 0.6$ and set calculation mode equal to "Confirm"
    (ii) Increase the number of confirmation
  (f) else set the number of confirmation equal to zero
(3) end of While

The hybrid scheme starts from a randomly generated population of 40, which is fewer than that used in the GA case alone which was 50. The initial guess for the local algorithm is taken from the population depending on the calculation mode. There are two modes in the algorithm, search and confirm. In search mode the initial guess is chosen from the two best in the population. In confirm mode the initial guess is chosen from a subset of the population, chosen to be far away from the current best. From here onwards the decision-making is done based on probability matching depending on the rewards associated with each of the optimisation schemes. The probability of selecting the GA can be calculated from the following equation (Lobo and Goldberg, 1996):

$$P^k_{\mathrm{GA}} = W^k_{\mathrm{GA}} / \left(W^k_{\mathrm{GA}} + W^k_{\mathrm{Local}}\right) \quad (8)$$

A random number generator simulates a coin toss and depending on this flip one of the optimisation schemes is chosen and proceeded with. If the scheme chosen is global optimisation, it proceeds with only one generation. If the local scheme is chosen, then the optimisation starts from the initial condition until it either converges or reaches the maximum number of cost function evaluations. At the end of a run of either of the optimisation schemes, the improvement achieved above the value of the best solution prior to the optimisation run is checked. The reward for a particular, local or global, optimisation and the probabilities are updated and the sequence is repeated until no improvement occurs from either of the two methods.

4.1.1. Results 25 different trials of the HGA were performed - the results are summarised in Table 3. Only 2 out of 25 trials failed to find the true global solution. The number of cost function

Table 3. Hybrid optimisation statistics

	Trial	Avg.	Max.	Min.	Std.	Success
HGA	25	3080	4266	2109	517	92%
HDE	25	1867	2053	1644	91.43	100%

evaluations can be reduced by decreasing the population size and/or the number of confirmations, however, this will also reduce the resulting success rate. When compared to the standard GA method, there is a 38.4% reduction in computational effort when the average number of simulations are considered.

4.2 Hybrid DE

As shown previously, for the flight clearance problem considered in this paper, the global DE scheme converged to the global solution in all 25 trials conducted. In this section, therefore, we focus on evaluating the impact on the DE convergence rate of incorporating local optimisation methods via a hybrid scheme.

The conventional DE methodology was augmented by combining it with a downhill simplex local optimisation scheme in (Rogalsky and Derksen, 2000). At each iteration, local optimisation was applied to the best individual in a current random set. This hybrid scheme was applied to an aerofoil shape optimization problem and was found to significantly improve the convergence properties of the method. The hybrid DE scheme employed in this study applies gradient-based local optimisation, again using *"fmincon"*, to a solution vector randomly selected from the current set, for iterations when the DE optimisation does not return an improved solution. When the local scheme is chosen, the optimisation starts from the chosen initial condition and continues until it either converges or reaches a defined maximum number of cost function evaluations. The algorithm is simple, and tries to search for the global optimum in a "greedy" way, demanding improvement in the achieved optimum value in every iteration.

4.2.1. Results Again, 25 different trials were conducted: the results are summarized in Table 3 and compared with those obtained from the HGA algorithm. In all 25 trials, the HDE scheme found the true global solution. From Table 3, it can be seen that the HDE scheme is by far the most successful of the various schemes evaluated in this study. The proposed HDE scheme clearly outperforms the HGA scheme, with on average a 39.38% reduction in computational overheads. The standard deviation and success indicator are also better for the HDE scheme. When the HDE scheme is compared to the standard DE scheme, there is on average a 48% reduction in computational complexity.

5. CONCLUSION

This paper has compared the performance of two different evolutionary optimisation schemes, namely genetic algorithms and differential evolution, on a nonlinear flight control law clearance problem. Hybrid versions of both schemes incorporating local gradient-based optimisation were shown to offer significant advantages in terms of both computational complexity and global convergence properties. In particular, the recently introduced differential evolution approach, when appropriately augmented with local optimisation methods, is shown to have significant potential for improving both the reliability and efficiency of the current industrial flight clearance process.

6. ACKNOWLEDGEMENTS

This work was carried out under EPSRC research grant GR/S61874/01.

REFERENCES

Davis, L. (1991). *Handbook of Genetic Algorithms.* Van Nostrand Reinhold.

Fielding, C., Varga A., Bennani S. and Selier M. (Eds.) (2002). *Advanced Techniques for Clearance of Flight Control Laws.* Springer.

Fleming, P. J. and R. C. Purshouse (2002). Evolutionary algorithms in control systems engineering: a survey. *Control Engineering Practice.* **10**, 1223–1241.

Forssell, L. S. (2003). Flight clearance analysis using global nonlinear optimisation based search algorithms. In: *Proc. AIAA Guidance, Navigation and Control Conference.*

Forssell, L. S., Hovmark G., Hyden Å. and Johansson F. (2001). The aero-data model in a research environment (ADMIRE) for flight control robustness evaluation. Technical Report GARTUER/TP-119-7.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley.

Lampinen, Jouni and Ivan Zelinka (1999). Mechanical engineering design by differential evolution. In: *David Corne, Marco Dorigo and Fred Glover (eds.) New Ideas in Optimisation, McGraw-Hill, London (UK).* pp. 127–146.

Lobo, F. G. and D. E. Goldberg (1996). Decision making in a hybrid genetic algorithm. Technical Report IlliGAL Report No. 96009.

Menon, P.P., J. Kim, D.G. Bates and I. Postlethwaite (2004). Improved clearance of flight control laws using hybrid optimisation. In: *Proceedings of IEEE Conference on Cybernetics and Intelligent Systems.*

Optimization Toolbox User's Guide, Version 2 (2000). The MathWorks.

Rogalsky, T. and R. W. Derksen (2000). Hybridization of differential evolution for aerodynamic design. In: *Proc. of the 8th Annual Conference of the Computational Fluid Dynamics Society of Canada.*

Rogalsky, T., R. W. Derksen and S. Kocabiyik (2000). Differential evolution in aerodynamic optimization. *Canadian Aeronautics and Space Institute Journal* **46**, 183–190.

Storn, R. and K. Price (1997). Differential evolution: a simple and efficient heuristic for global optimization over continuous space. *Journal of Global Optimization.* **11**, 341–369.

Yen, J., J. C. Liao, D. Randolph and B. Lee (1995). A hybrid approach to modeling metabolic systems using genetic algorithm and simplex method. In: *In Proceedings of the 11th IEEE Conference on Artificial Intelligence for Applications.*