Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Hybrid Evolutionary Optimisation Methods for the Clearance of Nonlinear Flight Control Laws

Prathyush P. Menon, Declan G. Bates and Ian Postlethwaite

Abstract— **The application of two evolutionary optimisation methods, namely differential evolution and genetic algorithms, to the clearance of nonlinear flight control laws for highly augmented aircraft is described. The algorithms are applied to the problem of evaluating a nonlinear handling qualities clearance criterion for a simulation model of a high performance aircraft with a delta canard configuration and a full-authority flight control law. Hybrid versions of both algorithms, incorporating local gradient-based optimisation, are also developed and evaluated. Statistical comparisons of computational complexity and global convergence properties reveal the benefits of hybridisation for both algorithms. The differential evolution approach in particular, 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.**

I. 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 dependent on the controller, it must be proven to the clearance authorities that the controller functions correctly throughout the specified flight envelope in all normal and various failure conditions, and in the presence of all possible parameter variations.

This task is a very lengthy and expensive process, particularly for high performance aircraft, where 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.) must be investigated so that guarantees about worst-case stability and performance can be made, [1].

The goal of the clearance process is to demonstrate that a set of selected criteria expressing stability and handling requirements is fulfilled. Typically, criteria covering both linear and nonlinear stability, as well as various handling and performance requirements are used for the purpose of clearance. The clearance criteria can be grouped into four classes, (i) linear stability criteria, (ii) aircraft handling/Pilot Induced Oscillation (PIO) criteria, (iii) nonlinear stability criteria, and (iv) nonlinear handling criteria. This paper focuses on the evaluation of a nonlinear handling criterion, which is described in detail in the next section. Details of the other clearance criteria can be found in [1].

In the clearance process, for each point of the flight envelope, for all possible configurations and for all combinations of parameter variations and uncertainties, violations of the clearance criteria and the worst-case result for each criterion must be found. Based on the clearance results, flight restrictions are imposed where necessary. Faced with limited time and resources, the current flight clearance process employed by the European aerospace industry uses a gridding approach, whereby the various clearance criteria are evaluated for all combinations of the extreme points of the aircraft's uncertain parameters [1]. This process is then repeated over a gridding 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 the 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 clearance problem based on the use of hybrid optimisation techniques, which will be shown to have the capability to significantly improve both the reliability and efficiency of the current flight clearance process. Note that this paper is a summary of the complete results of this study, which are presented in [2].

II. ADMIRE - AIRCRAFT MODEL

The aircraft model used in this study is the ADMIRE (Aero-Data Model In a Research Environment), [3], a nonlinear, six degree of freedom simulation model 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. ADMIRE is augmented with a full-authority flight control system and includes engine dynamics and detailed nonlinear 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. Table I shows the uncertain parameters considered in this study. For more details of the ADMIRE model, control law and flight envelope, the reader is referred to [3] and [4].

A. Nonlinear Clearance Criterion

The clearance criterion considered in this study is the AoA limit exceedence criterion [1, 5, 6], which aims to assess the effectiveness of the AoA limiting scheme in the flight control system. For this criterion it is required to identify the flight cases where, for a defined pull-up manoeuvre, the maximum overshoot occurs in AoA. The corresponding worst-case combination of uncertainties must also be computed. The defined pull up manoeuvre is a rapid pull in longitudinal stick to a defined level (40N) at a 640N/sec stick rate with

This work was supported by EPSRC research grant GR/S61874/01.

The authors are with Department of Engineering, University of Leicester, Leicester, LE1 7RH, UK. (email: ppm6, dgb3, ixp@le.ac.uk)

0-7803-9568-9/05/$20.00 ©2005 IEEE

TABLE I

AIRCRAFT MODEL UNCERTAIN PARAMETERS [5]

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]
$\Delta_{C_{m_{\delta_e}}}$	[-0.05, +0.05]	uncertainty in pitching moment due to elevator deflection [1/rad]
$\Delta_{I_{yy}}$	[-0.2, +0.2]	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]

stick hold for 10 seconds. The analysis aims to estimate the clearance criterion [1]:

$$\alpha_{\max} = \max(\alpha(t)) \quad \text{for} \quad t \leq 10 \text{ [sec]} \tag{1}$$

for all possible combinations of aircraft parametric uncertainty.

B. Optimisation Based Flight Clearance

In this paper the flight clearance problem defined above is formulated as an optimisation problem and solved using a number of different approaches. The optimisation problem itself is to find the combination of real parametric uncertainties that gives the worst value of the criterion defined in (1). Since this and many other clearance criteria must be checked over a huge number of envelope points and aircraft configurations, it is imperative to find the most computationally efficient approach to the problem. Previous efforts to apply optimisation methods to this problem, [1] 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 [1] Chapter 21, local optimisation methods such as SQP (Sequential Quadratic Programming), and L-BFGS-B (Limited memory Broyden-Fletcher-Goldfarb-Shanno method with Bounded constraints) were used to evaluate a range of linear clearance criteria for the HIRM+ (High Incidence Research Model) aircraft model. In [1] Chapter 22, global optimisation schemes such as Genetic Algorithms (GA), Adaptive Simulated Annealing (ASA) and Multi Coordinate Search (MCS) were also applied to evaluate nonlinear clearance criteria for the same aircraft model. In [5, 6] global optimisation methods such as GA and ASA were applied to the ADMIRE model with a different flight clearance criterion. To demonstrate the limitations of using local optimisation methods for the type of problem considered in this study results using Sequential Quadratic Programming (as implemented in the function *"fmincon"* from [7]) are shown in Table II - the last column shows the total number of simulations, i.e., the number of cost function evaluations. Later, it will be shown, via exhaustive global optimisation trials, that the parameter combination in the second row is (as far as can be established) the global solution. As shown in the table, however, for each different initial guess for the values of the uncertain parameters, the local optimisation algorithm converges to a different point in the uncertain parameter space. These results show that using local optimisation methods in isolation allows very little confidence to be established that the true worst-case violation of the clearance criterion has been found.

III. GLOBAL OPTIMISATION

A. Genetic Algorithms

The first global optimisation method we consider in this study is Genetic Algorithms (GA), which are general purpose stochastic search and optimisation procedures, based on genetic and evolutionary principles [8]. In a genetic search technique, 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. The candidates are encoded as artificial chromosomes, and a fitness function is defined to assign a performance index to each candidate - this function is specific to the problem and is formed from the knowledge domain. GA have become a popular, robust search and optimisation technique for problems with large as well as small parameter search spaces. The recent survey paper [12] reports that GA have also become a very popular search and optimisation technique for problems in control engineering. Due to their stochastic nature, global optimisation schemes such as GA can be expected to have a much better chance of converging to a global optimal. The price to be paid for this improved performance is a dramatic increase in computation time when compared to local methods. In the sequel, the genetic operators employed to generate and handle the population in the GA for the clearance problem are described. The reader is referred to [8] for more details of different operators, binary coding schemes and the theory of genetic search in general. For details of the variable representation, selection, crossover and mutation strategies used in the present study, the reader is referred to [2].

Figure 1 shows the number of simulations versus the best fitness for 100 GA trials. The statistics of the results, from the 100 independent trials, are given in Table III. The upper histogram of Fig.2 shows the percentage distribution of the maximum value of AoA achieved over the 100 trials. The lower histogram of Fig.2 shows the percentage distribution of the total number of simulations required to obtain the solution over the 100 independent trials. A large number of simulations, an average of 4485 simulations in this case, is required to obtain the global, or near global solution. The probability of success in attaining the true global solution is

TABLE II

RESULTS FOR LOCAL OPTIMISATION ALGORITHM

Starting Point $\left[\Delta^0_{mass}, \Delta^0_{xcg}, \Delta^0_{C_{m_{\delta_e}}}, \Delta^0_{I_{yy}}, \Delta^0_{C_{m\alpha}}\right]$	Convergent Point $\left[\Delta^*_{mass}, \Delta^*_{xcg}, \Delta^*_{C_{m_{\delta_e}}}, \Delta^*_{I_{yy}}, \Delta^*_{C_{m\alpha}}\right]$	Number of Simulations
[0, 0, 0, 0, 0]	[0.100, 0.0750, 0.050, 0.06084, 0.050]	375
[0.100, 0.0750, 0.050, 0.200, 0.050]	[0.100, 0.0750, 0.050, 0.18309, 0.050]	366
[-0.100, -0.0750, -0.050, -0.200, -0.050]	[0.100, 0.0750, 0.050, -0.12634, 0.050]	322

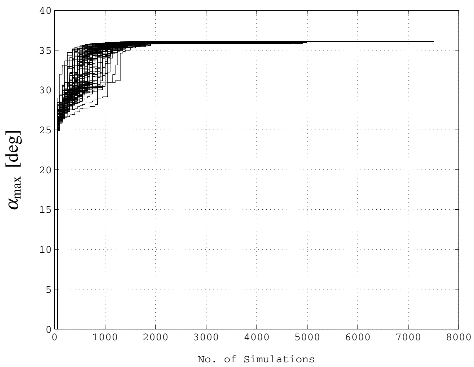

Fig. 1. GA - No. of simulations vs. best fitness

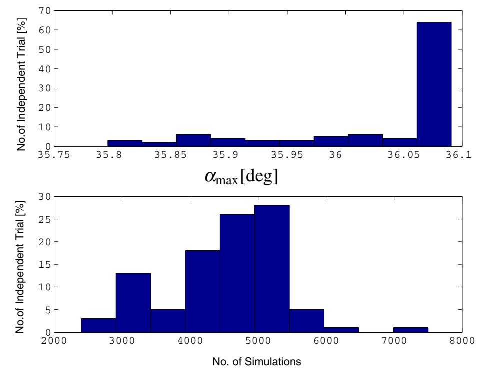

Fig. 2. GA results histogram

also rather low, at only 65%. The global solution found in this example is the following:

$$\left[\Delta^*_{mass}, \Delta^*_{xcg}, \Delta^*_{C_{m_{\delta_e}}}, \Delta^*_{I_{yy}}, \Delta^*_{C_{m\alpha}}\right] = [0.1000,\ 0.0750,\ 0.0500,\ 0.18309,\ 0.0500] \tag{2}$$

and $\alpha_{\max}$ is 36.0908°. Note that four of the uncertain parameters in this case are on their upper bounds and $\Delta^*_{I_{yy}}$ is inside its bound. A sensitivity analysis is performed about the solution and is shown in Figure 3, where the x-axis is normalized. As shown in the figure, the uncertain parameter $\Delta_{I_{yy}}$ has many local maxima.

Tuning of the GA optimisation parameters, such as the different GA-operator probabilities may, of course, improve the above results to a certain extent. However, there are few available guidelines as to how to do this tuning. Another possible approach would be to use alternate selection schemes and scaling and ranking procedures, such as those described in Ref. [8] Chapter 4. However, for the present problem the advantage to be gained from these techniques is not expected to be significant. Finally, we note that in the context of the current flight clearance process, the computational cost of the number of simulations required by the above approach would be likely to prove prohibitive to its widespread adoption by industry, [1] Chapter 1.

B. Differential Evolution

The second global optimisation method considered in this study is Differential Evolution (DE), a relatively new global optimisation method, introduced by Storn and Price in [14]. It belongs to the same class of evolutionary global optimisation techniques as GA, but unlike GA it does not require either a selection operator or a particular encoding scheme. Despite its apparent simplicity, the quality of the solutions computed using this approach has been claimed to be often better than that achieved using other evolutionary algorithms, both in terms of accuracy and computational overhead [14].

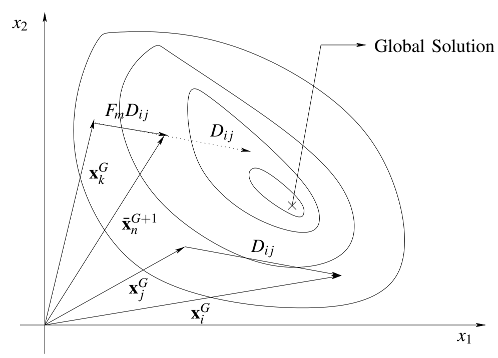

Fig. 4. DE mutation strategy

The DE method has recently been applied to several problems in different fields of engineering design, with promising results. In [13], for example, it was applied to find the optimal solution for a mechanical design example formulated as a mixed integer discrete continuous optimization problem. In [15], the DE method has been applied and compared with other local and global optimization schemes in an aerodynamic shape optimization problem for an aerofoil. For more details of the DE methodology, and the particular initialization, mutation, crossover and selection strategies employed in this study, the reader is referred to [2].

Figure 5 shows the number of simulations versus the best fitness for 100 DE trials. 90 trials converged to the true global solution given in Eq. 2, giving the maximum AoA overshoot. Seven trials converged to solutions very close to the global solution, and 3 trials gave different solutions.

Compared to the GA results, DE can be seen to offer

TABLE III
GLOBAL OPTIMISATION COMPARISON STATISTICS : NO. OF SIMULATIONS

Optimisation Method	Trails	Average	Minimum	Maximum	Standard Deviation	Probability of Success
GA	100	4485	2400	7500	828.364	65%
DE	100	3086	1152	4176	567.57	90%

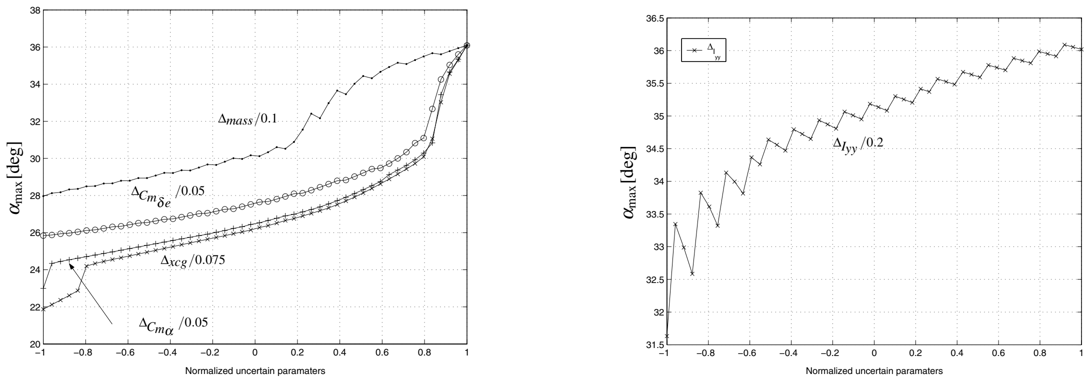

(a) (b)

Fig. 3. Sensitivity Plots about the Global Solution

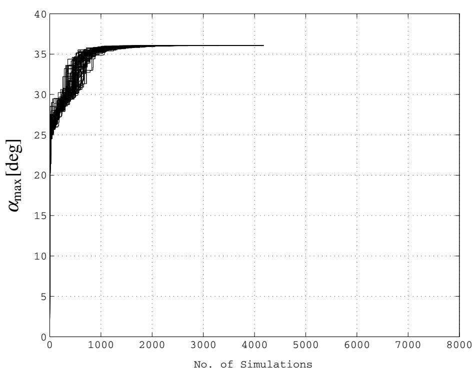

Fig. 5. DE - No. of simulations vs. best fitness

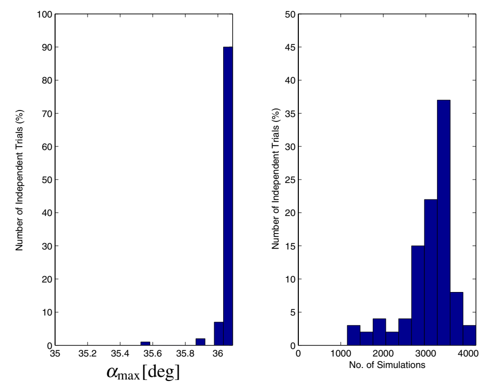

Fig. 6. DE results histogram

significantly improved convergence properties, while the reduced number of initial random starting points (only 12 initial random points against 50 random initial points for the GA) means that the total number of simulations required in each trial was also significantly reduced. Table III provides the statistics of the results obtained from the 100 trials of the DE algorithm, and also compares them to those from the GA. The average number of simulations required for DE, 3086 in this case, is 31% less than required by the GA. The probability of success of the DE algorithm is also much higher, at 90%. The left subplot in Fig.6 shows the distribution of the maximum value of AoA achieved. The right subplot shows the distribution of the number of simulations over 100 independent trials of the DE algorithm. Note that, in addition to the improved results, another advantage of this method compared to GA is the reduced number of optimisation parameters that must be adjusted by the user.

IV. 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. As shown by the preceeding results, however, the rate of convergence can be very slow, 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 [9]–[11]. In such hybrid schemes there is the possibility of incorporating domain knowledge, which gives them an advantage over a pure blind search based on evolutionary principles.

A. 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 at each iteration. The reward associated to each optimisation scheme will determine the probability for that optimisation scheme to be chosen at the next iteration. The reward for each optimisation scheme thus keeps varying depending on how well it is performing. 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 [11]:

$$W^{k+1}_{\text{GA/Local}} = W^{k}_{\text{GA/Local}}(1-c) + cR^{k}_{\text{GA/Local}} \qquad (3)$$

where W^k and R^k are the weighted reward and the improvement in the solution at the iteration k, respectively, and c is a constant in $[0, 1]$. R^k is computed based on the improvement in the best solution attained over each iteration/generation. In case no improvement occurs, the value of R^k is set equal to zero. 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 is taken based on the previous reward and by calculating the associated probability. A pseudocode listing for the algorithm used in the hybrid switching scheme is given in [2].

Due to the frequent occurrence of local maxima in flight clearance problems, it is desirable that, initially, the GA should have a higher probability of being chosen than the local algorithm. Hence, initially the weights for GA and the local algorithm are given as 0.9 and 0.1, respectively. The local algorithm used in the present study is the implementation of the SQP method [7], described in Section III.

Due to the improved convergence properties of the HGA algorithm (see below), it was possible to reduce the size of the initial population to 40 candidates. 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 at any iteration can be calculated from the following equation [11]:

$$P^{k}_{\text{GA}} = W^{k}_{\text{GA}} / \left(W^{k}_{\text{GA}} + W^{k}_{\text{Local}} \right) \qquad (4)$$

A random number generator simulates a coin toss and depending on the result 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 runs until it either converges or reaches the defined maximum number of cost function evaluations. At the end of a run of either of the optimisation schemes, the improvement achieved above

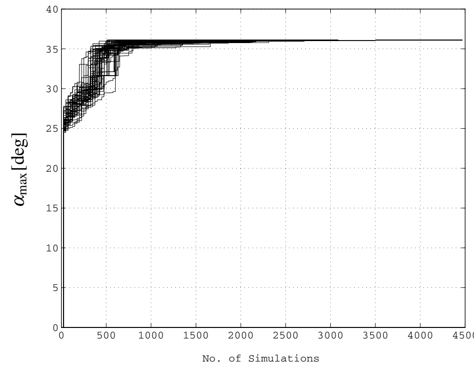

Fig. 7. HGA - No. of simulations vs. best fitness

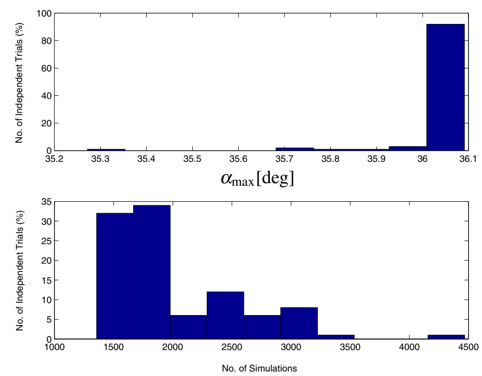

Fig. 8. HGA histogram

the value of the best solution prior to the optimisation run is checked. The reward for a particular, local or global, optimisation is assigned, the probabilities are updated and the sequence is repeated until no improvement occurs from either of the two methods.

Figure 7 shows the number of simulations versus the best fitness for 100 trials of the HGA. Table IV provides the statistical results. The average number of cost function evaluations required was 2011, an improvement of 55% when compared with the standard GA. The success rate in finding the true global solution is also dramatically improved, from 65% to 92%. The upper and lower subplots of Fig. 8 show the histogram distributions of maximum AoA obtained and the number of simulations taken respectively, over the 100 independent trials.

B. Hybrid DE

In [16], the conventional DE methodology was augmented by combining it with a downhill simplex local optimisation scheme. This hybrid scheme was applied to an aerofoil shape optimization problem and was found to significantly improve the convergence properties of the method. At each iteration, local optimisation was applied to the best individual in a current random set. 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 our problem, this was seen to give better results than using the *best* solution vector, as proposed in [16]. When the local scheme is chosen, the optimisation

TABLE IV

HYBRID OPTIMISATION COMPARISON STATISTICS : NO. OF SIMULATIONS

Optimisation Method	Trails	Average	Minimum	Maximum	Standard Deviation	Probability of Success
HGA	100	2011	1357	4468	547.42	92%
HDE	100	1106	477	1434	192.42	98%

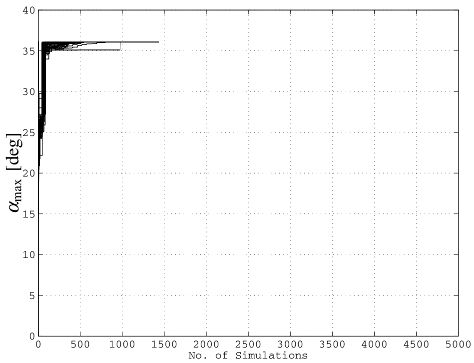

Fig. 9. HDE - No. of simulations vs. best fitness

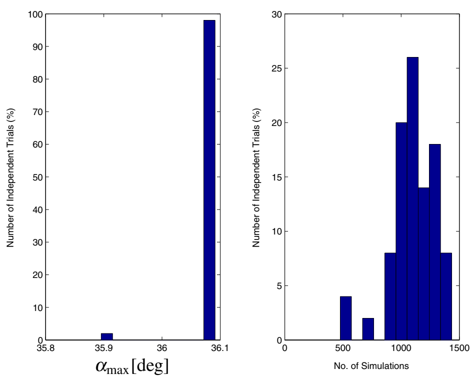

Fig. 10. HDE results histogram

starts from the given 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. A psuedo-code for the hybrid DE algorithm is given in [2].

Figure 9 shows the number of simulations versus the best fitness for 100 trials of the HDE algorithm. Table IV provides the statistical results and compares them with the results of the HGA. The average number of cost function evaluations required was 1106, an improvement of 64% when compared with the standard DE algorithm, and 45% when compared with the HGA. The success rate in finding the true global solution is also extremely high, at 98%. The upper and lower subplots of Fig. 10 show the histogram distributions of maximum AoA obtained and the number of simulations taken, respectively, over the 100 independent trials.

V. CONCLUSIONS

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

REFERENCES

[1] Christopher Fielding, Andras Varga, Samir Bennani, and Michiel Selier (Eds.) *Advanced Techniques for Clearance of Flight Control Laws* Springer, 2002.

[2] P. P. Menon, D. G. Bates and I. Postlethwaite, "Clearance of Nonlinear Flight Control Laws using Hybrid Evolutionary Optimisation", submitted to the IEEE Transactions on Evolutionary Computation, 2005.

[3] Lars S Forssell, Gunnar Hovmark, Åke Hyden, and Fredrik Johansson, *The Aero-Data Model In a Research Environment (ADMIRE) for Flight Control Robustness Evaluation* GARTUER/TP-119-7, August 1^{st} 2001.

[4] Rundqwist, L., Stahl-Gunnarsson, K., and Enhagen, J., "Rate Limiters With Phase Compensation in JAS39 GRIPEN", *Proc. of the European Control Conference*, 1997.

[5] Lars S. Forssell and Åke Hyden, "Flight control system validation using global nonlinear optimisation algorithms", *Proc. of the European Control Conference*, Cambridge, U.K., 2003.

[6] Lars S. Forssell, "Flight clearance analysis using global nonlinear optimisation based search algorithms", *Proc. of the AIAA Guidance, Navigation, and Control Conference*, Austin, Texas, 2003.

[7] *Optimization Toolbox User's Guide, Version 2*, September 2000 The MathWorks.

[8] David E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning* Addison-Wesley, 1989.

[9] L. Davis, Ed., *Handbook of Genetic Algorithms*, Van Nostrand Reinhold, New York, 1991.

[10] J. Yen, J. C. Liao, D. Randolph, and B. Lee, "A hybrid approach to modeling metabolic systems using genetic algorithm and simplex method", *Proc. of the 11th IEEE Conference on Artificial Intelligence for Applications*, pp. 277–283, Los Angeles, CA, Feburary 1995.

[11] Fernando G. Lobo and David E. Goldberg, Decision making in a hybrid genetic algorithm, IlliGAL Report No. 96009, September 1996.

[12] P. J. Fleming and R. C. Purshouse, "Evolutionary algorithms in control systems engineering: a survey", *Control Engineering Practice*, Vol. 10, pp-1223-1241, 2002.

[13] Jouni Lampinen and Ivan Zelinka, "Mechanical Engineering Design by Differential Evolution", David Corne, Marco Dorigo and Fred Glover (eds.) *New Ideas in Optimisation*, McGraw-Hill, London (UK), 1999, pp-127–146.

[14] R. Storn and K. Price, "Differential evolution: a simple and efficient heuristic for global optimization over continuous space", *Journal of Global Optimization*, Vol.11, pp-341–369, 1997.

[15] T. Rogalsky, R. W. Derksen and S. Kocabiyik, "Differential Evolution in Aerodynamic Optimization", *Canadian Aeronautics and Space Institute Journal*, Vol.46, pp-183–190, 2000.

[16] T. Rogalsky and R. W. Derksen, "Hybridization of Differential Evolution for Aerodynamic Design", *Proc. of the 8th Annual Conference of the Computational Fluid Dynamics Society of Canada*, pp-729–736, 2000.