μ-Tools for the Clearance of Flight Control Laws

Declan G. Bates[†], Thomas Mannchen[‡], Ridwan Kureemun[†], Ian Postlethwaite[†]

***Abstract*— A recent industry-led European research project investigated the potential of a number of advanced analysis techniques for use in the clearance of flight control laws for highly augmented aircraft. Among the analysis techniques considered in the project were methods based on μ-analysis and the structured singular value. This paper summarises the approach taken by the authors in attempting to develop advanced robustness analysis tools for highly complex, and highly formalised, industrial flight clearance problems.**

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 controller 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 non-linear 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.

The aircraft models used for clearance purposes describe the actual aircraft dynamics, but only within given uncertainty bounds. One reason for this is the limited accuracy of the aerodynamic data set determined from theoretical calculations and wind tunnel tests. Moreover, the employed sensor, actuator and hydraulic models are usually only approximations, where the nonlinear effects are not fully modelled because they are either not known or it would make the model unacceptably complex. The goal of the clearance process is to demonstrate that a set of selected criteria expressing stability and handling requirements is fulfilled in the presence of all of the above sources of uncertainty. Typically, the clearance criteria can be grouped into four classes: (I) Linear stability criteria, (II) Aircraft handling/Pilot Induced Oscillation (PIO) criteria, (III) Non-linear stability criteria and (IV) Non-linear handling criteria.

To perform the clearance, 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, [1], whereby the various clearance criteria are evaluated for all combinations of the extreme points of the aircraft's uncertain parameters. This process is then repeated over a gridding of the aircraft's flight envelope. The effort involved in the resulting clearance assessment thus increases exponentially with the number of uncertain parameters. Another difficulty is the fact that there is no guarantee that the worst case uncertainty combination has in fact been found, since only a certain number of selected points in the aircraft's flight envelope can be checked.

Motivated by the above issues, the Group for Aeronautical Research and Technology in Europe (GARTEUR) set up Flight Mechanics Action Group FM-AG11 to investigate the usefulness of a number of advanced analysis techniques for flight clearance problems, [1]. The Action Group consisted of participants from Universities, Research Laboratories and Industry, with all of the main European airframe manufacturers being actively involved in the project. The benefits and limitations of the various analysis techniques were evaluated based on their application to a number of realistic flight clearance problems which were defined by the industrial partners.

II. LINEAR STABILITY CLEARANCE CRITERIA

A basic requirement of the flight clearance process is to prove that the aircraft is stable over the entire flight envelope with sufficient margin against instability for all known uncertainties (worst-case combinations), [1]. The process consists of calculating linear stability margins for the open-loop frequency response in pitch, roll and yaw. These frequency responses are obtained by breaking the loop at the input of each actuator or of each sensor and are then plotted in Nichols diagrams where the required phase and gain margins are shown as exclusion regions which must not be violated by the plot.

In single loop analysis, the open-loop frequency response is obtained by breaking the loop at the input of each actuator or sensor, one at a time, while leaving the other loops closed. For the nominal case, these Nichols plots should not violate the outer exclusion region shown in Figure 1, which corresponds to a minimum gain margin of ± 6 dB and a minimum phase margin of $\pm 35^o$. When uncertainties are taken into account, a boundary corresponding to ± 4.5 dB is used, as shown by the inner exclusion region in the same figure.

In multi loop analysis, the closed-loop system is required to withstand the application of simultaneous gain and phase offsets, again defined by trapezoidal shaped regions in the Nichols plane, at the actuators or sensors without becoming unstable. To test for violations of this criterion, a perturbation of the form $K\left(\frac{1-Ts}{1+Ts}\right)$ is inserted at, for example, the input of each actuator. With K set to 1, T is then varied simultaneously in each loop until the eigenvalues of the closed-loop system go unstable. The phase margin is calculated as

$$\phi_{PM} = 2 * tan^{-1}\omega T \tag{1}$$

where ω is the frequency of the generated undamped oscillation. K is then increased and decreased by 1 dB (corresponding to the right corner points of the Nichols exclusion region) and T is again varied for the new fixed gain until the eigenvalues become unstable. By setting $T = 0$ and varying K, the upper and lower gain margins can be obtained (corresponding to the left corner points of the exclusion region). These steps can be repeated for any number of points around the required Nichols exclusion region. Due to the

[†]Department of Engineering, University of Leicester, University Road, Leicester LE1 7RH, UK, dgb3@le.ac.uk

[‡]Institute of Flight Mechanics and Control, University of Stuttgart, Pfaffenwaldring 7a, 70550 Stuttgart, Germany

0-7803-8335-4/04/$17.00 ©2004 AACC

fact that the criterion must be evaluated over all combinations of the aircraft's uncertain parameters, this test is in practice usually restricted to only a few points of the exclusion regions (e.g. the four corners of each exclusion region). In addition, the same gain and phase offsets are usually applied simultaneously in all loops, to avoid testing over too large a number of different combinations.

In addition to the stability margin criterion, the eigenvalues of the closed-loop system must be calculated in order to identify possible unstable (i.e. those with positive real part) eigenvalues which do not appear in the Nichols plots. It is required to identify the flight cases where unstable eigenvalues occur and for what tolerance combination these eigenvalues have the largest real part, since this can be directly linked to existing handling qualities requirements on the minimum time to double amplitude of unstable modes. This test aims to determine the most severe cases of divergent modes in the closed-loop system in order to allow an assessment of their acceptability in terms of their influence on aircraft handling. A typical boundary on the real part of the eigenvalues is shown in Figure 2.

III. HWEM AIRCRAFT MODEL AND CONTROL LAW

The Harrier Wide-Envelope Model (HWEM) used in this study is a full non-linear model of the Vectored-thrust Aircraft Advanced flight Control (VAAC) Harrier, developed by QinetiQ Ltd. for research on various aspects of flight control that are relevant to Short Take-Off and Vertical Landing (STOVL) operations. Recent research milestones with the VAAC Harrier include the first reported flight of an H^{∞} control law, [2], and the first reported flight of an LPV control law, [3]. More details of this aircraft model can be found in [4]. The flight control law for the HWEM analysed in this study is based on VAAC Control Law 002 (CL002). It is a full three-axis (pitch, roll and yaw) control law designed using classical methods. The study reported in this paper considers the pitch axis only - for more details of this control law and of the lateral/directional control laws the reader is referred to [4]. Points in the flight envelope of the HWEM from 200 knots down to hover are to be analysed in the clearance task. All flight conditions are defined for $1g$ straight and level flight at an altitude of 200 ft AMSL. The angle of attack range for all flight conditions is $[-4^o,+16^o]$. Five Category 1 (most significant) uncertain parameters are specified for the longitudinal axis analysis and are shown in Table I. Further information about the uncertain aircraft parameters can be found in [4], [5].

IV. LFT-BASED UNCERTAINTY MODELLING

In order to apply μ-analysis tools to the HWEM model, the uncertainties in the original non-linear aircraft model must be represented in the form of linear fractional transformations (LFT's), [6], [7]. The approach to generating LFT-based uncertainty models used for this study is based on physical modelling principles. In this approach, the uncertainties are directly introduced in the non-linear SIMULINK model of the aircraft in the form of multiplicative (or additive) uncertainties. For example, the uncertainty in the aircraft dynamics due to $C_{M_{TAIL}}$ can be represented physically in block diagram form as shown in Figure 3. In Figure 3, an extra 'fictitious' input and output w_1 and z_1 have been added, respectively, at the point in the system where the uncertainty Δ_1 occurs. This step is then repeated for the other uncertain parameters. Using standard block diagram manipulation software the resulting non-linear model can then be linearised to calculate the transfer matrix of the system M with inputs u = [w_1, ..., w_n, u_c] and outputs y = [z_1, ..., z_n, y_m], where u_c are the control inputs and y_m are the measured outputs. The LFT-based uncertainty model for the system, shown in Figure 4, is then given by the relation

$$y_m = F_u(M(s),\Delta)u_c \tag{2}$$

where $\Delta = diag(\Delta_1, ..., \Delta_n)$

The above approach is simple and intuitive and allows an exact description of joint parametric dependencies in the model. If each uncertainty appears in only one location in the SIMULINK block diagram, the resulting LFT-based uncertainty model will also be of minimal order. In addition, this physical modelling approach allows additional uncertainties in the physical parameters (such as products of the uncertainties) to be easily implemented in the model. The main limitation of the approach is that detailed information about the model and the uncertainties is required. Hence, its application is restricted to those models that can be implemented in a SIMULINK block diagram type representation. Another drawback is that the dependence of the linearisations on the uncertain parameters is ignored. To determine the accuracy of the derived LFT-based uncertainty model, time-domain simulations were used to validate it against the original linearised and nonlinear model. Figure 5, for example, shows good matches between the various models for a small magnitude step demand on pitch rate.

V. μ-TOOLS FOR STABILITY CLEARANCE CRITERIA

In the approach first proposed in [8], [9], [10], the original Nichols exclusion regions shown in Figure 1 are replaced with elliptical regions of the form shown in Figure 6. Thus, any feedback system whose open-loop frequency response avoids the regions A and B in Figure 6 provides gain and phase margins of $\pm 6dB/\pm 36.87^{\circ}$ and $\pm 4.5dB/\pm 28.44^{\circ}$ respectively. Also, for these particular choices of gain and phase margins the corresponding exclusion regions in the Nyquist plane are *circles* with (centre,radius) given by (-1.25,0.75) for region *A*, and (-1.14,0.54) for region B, [8], [11]. Now, as shown in [8], another way to interpret the requirement for avoidance of, for example, the circle corresponding to region B in the Nyquist plane by the open-loop frequency response $L(j\omega)$, is to consider a plant subject to disc uncertainty of (centre,radius) given by (+1.14,0.54) at each frequency. Avoidance of the (-1,0) critical point in the Nyquist plane by $L(j\omega)$ for all possible plants in this set is exactly equivalent to avoidance of the exclusion region B by $L(j\omega)$ for the original plant. The set of possible plants can be represented as

$$P(s) = P_1(s)(1.14+\Delta_N) \tag{3}$$

where P_1 is the original plant, Δ_N is complex and $\| \Delta_N \|_{\infty} \leq 0.54$. This is of course the same as writing

$$P(s) = 1.14P_1(s)(1+W_N\Delta_N) \tag{4}$$

with $W_N = 0.47$ and $\| \Delta_N \|_{\infty} \leq 1$. In this way the Nichols exclusion region is represented as a 'fictitious' multiplicative input uncertainty for the scaled nominal plant which can be pulled out of the closed loop system along with all the other uncertainties to form an LFT-based representation of the uncertain system in the usual way.

A second approach to casting Nichols plane exclusion region specifications as a μ problem was developed in [12]. This method models the Nichols exclusion regions of Figure 1 using a Padé approximation. The variations in the phase and gain are represented

by equations (5) and (6) respectively. The phase offset is given by

$$\phi = \left(\frac{\phi_{max} - \phi_{min}}{2}\right)\delta_2 + \left(\frac{\phi_{max} + \phi_{min}}{2}\right) \quad (5)$$

The gain offset a (in dB) is represented as

$$a = \delta_1(t - m\delta_2) \quad (6)$$

where δ_1 and δ_2 are normalised real uncertainties, and t and m characterise the top limit line of the exclusion region. For example, the inner exclusion region in Figure 1 for the single loop analysis requires that t = 3 and m = 1.5. To cast this problem into a μ framework, it is necessary to convert these equations to the polar form $ae^{-j\phi}$, where the negative sign denotes phase lag. This gives

$$\begin{aligned} ae^{-j\phi} &= e^{c\delta_1(t-m\delta_2)-j(\gamma_1\delta_2+\gamma_2)} \\ &= e^{-j\gamma_2}e^{c\delta_1(t-m\delta_2)-j\gamma_1\delta_2} \end{aligned} \quad (7)$$

where c = (ln10)/20, $\gamma_1 = \frac{\phi_{max}-\phi_{min}}{2}$ and $\gamma_2 = \frac{\phi_{max}+\phi_{min}}{2}$
To generate the LFT-based uncertainty description, a first order Padé approximation is used:

$$e^{-Ts} = 1 - \frac{Ts}{1 + \frac{Ts}{2}} \quad (8)$$

where $-Ts$ is given by

$$-Ts = c\delta_1(t - m\delta_2) - j\gamma_1\delta_2 \quad (9)$$

This first order approximation is adequate for phase margins of up to 90^o. The resulting LFT-based uncertainty model for this first order Padé approximation is shown in Figure 7. The uncertainty block, $\Delta_{margins}$, is made up of two real scalars, δ_1 and δ_2, the latter being repeated twice.

$$\Delta_{margins} = \begin{bmatrix} \delta_1 & 0 \\ 0 & \delta_2 I_2 \end{bmatrix} \quad (10)$$

For multi-loop analysis, the criterion is checked by scaling the exclusion region by applying a scaling factor to m, t, γ_1 and γ_2 until μ = 1. The multi-loop gain and phase margins can then be computed by back-substituting these values in equations (5) and (6).

μ-analysis methods can also be used to address the worst-case eigenvalue clearance criterion defined in Section 2 by shifting the imaginary axis into the left and right half planes until an uncertainty combination is found which places a closed loop pole on the axis. Other tests are also possible, for example, by sweeping s_0 along a line of constant damping one may find the smallest perturbation which reduces damping below this level, [13].

VI. HWEM ANALYSIS RESULTS

For the considered ranges of uncertainties, the worst-case eigenvalue criterion was satisfied for all seven flight conditions, and almost identical results were obtained using the μ-analysis and classical techniques. For FC1, for example, the nominal, μ worst-case and classical worst-case eigenvalue positions are shown in Table II. Also shown in the table are the worst-case values of the uncertain parameters found using both approaches. Although the results are very similar, the μ worst-case uncertainty combination places all the eigenvalues slightly nearer the boundaries.

To compare the worst-case stability margin criterion results, Nichols curves were plotted for (i) the worst-case obtained using μ and (ii) every combination of the extreme points of the Five Category 1 uncertain parameters. For three flight conditions (FC4, FC5 and FC6), the worst-case uncertainties did not lie on the extreme points of the parameters. Sample results for FC4 are shown in Figure 8 - zoomed in Figure 9. In all three cases, the classical approach produces optimistic results, i.e. the worst-case Nichols plots found by μ are closer to the exclusion regions. For example, the tailplane loop cut analysis at FC4 using μ (Figures 8 and 9) generated a worst-case uncertainty of $\delta_{C_{m_q}}$ = 0.9845, $\delta_{C_{m_\alpha}}$ = 0.9976, $\delta_{C_{m_{tail}}}$ = 1, $\delta_{I_{yy}}$ = 0.0931 and $\delta_{X_{Dxcg}}$ = 0.1438. When the classical approach was used, the worst-case was found to be $\delta_{C_{m_q}}$ = 1, $\delta_{C_{m_\alpha}}$ = 1, $\delta_{C_{m_{tail}}}$ = 1, $\delta_{I_{yy}}$ = -1 and $\delta_{X_{cg}}$ = 1. Using μ-sensitivities, [14], it can be show that $\delta_{X_{cg}}$ and $\delta_{I_{yy}}$ are in fact the second and third most important elements in the set of uncertain parameters, [11] - the fact that their worst-case values do not correspond to thir maximum or minimum values clearly calls into question the implicit assumptions made in the classical approach.

The computation times for finding the worst cases was recorded for (i) the classical approach using only minimum and maximum values of each parameter, and (ii) μ analysis with 100 frequency points. The results are plotted in Figure 10. As expected, the computation time for the classical approach increases exponentially with the number of uncertainties, so that for a Δ size > 8, computation of mixed μ bounds is seen to be faster than the classical technique. This fact becomes significant when we seek to also include Category 2 uncertainties in the analysis, or when we seek to analyse the effect of longitudinal and lateral uncertainties together.

For the worst-case stability margin multi-loop analysis using the μ-analysis approach, the uncertainties associated with the elliptical Nichols exclusion regions were increased in each loop simultaneously until μ = 1. At flight condition FC5, for example, the gain and phase margins were found to be 13.97 dB and 37.86^o respectively. Using the classical approach, the corner points of the trapezoidal Nichols exclusion region were checked and the gain and phase margins were computed as 15.5 dB and 41.2^o respectively. The results obtained using the classical approach are more optimistic than those computed from μ since (a) every possible combination of the phase/gain offsets is considered in the μ-analysis and (b) the elliptical exclusion regions used by μ are slightly bigger than the diamond-shaped exclusion regions used by the classical approach. Complete results from the analysis of the longitudinal dynamics of the CL002 control law over the HWEM flight envelope can be found in [11].

VII. IMPROVING MIXED μ LOWER BOUNDS

Formulating the the stability margin clearance criterion in a μ-analysis framework using elliptical Nichols exclusion regions results in a mixed μ-analysis problem. The complex part of the uncertainty matrix Δ represents the "fictitious" uncertainty associated with the elliptical exclusion regions, while the real part represents the uncertain aircraft parameters. A potential problem with this approach is that the standard lower bound algorithm in the MATLAB μ-Analysis and Synthesis Toolbox, [15], sometimes produces a bound which is far from the assiciated upper bound, for cases where the real uncertainty in the Δ matrix is dominant. An example of such a result is shown in Figure 11(a) - clearly, in this case, the "worst-case" uncertainty combination returned by the lower bound software will be far from the true worst-case.

In order to improve the lower bound returned by the standard algorithm, the problem of computing a lower bound for μ can be formulated as a search for the worst case (i.e. smallest) destabilising uncertainty matrix Δ. Denote the real Δ_i entries of

Δ by the vector p, and the complex Δ_i entries by q. Thus, for an $n \times n$ Δ matrix, define the vector x as

$$[x] \;=\; [p,q]^T \;\; p \in \mathscr{R}^l,\; q \in \mathscr{C}^m,\; l+m=n \tag{11}$$

This search can then be formulated as an equivalent constrained minimisation problem, $f(x)$, over a frequency range Ω:

$$\min f(x) \;=\; \min_{\Delta_{i=1..l} \in \mathscr{R},\; \Delta_{i=l+1..n} \in \mathscr{C},\; \omega \in \Omega} \bar{\sigma}(\Delta)$$
$$\text{subject to} \;\; \underline{\sigma}(I - M_{11}\Delta) \leq \varepsilon \tag{12}$$

ε in the above constraint is a user defined parameter which can be used to trade-off computation time versus tightness of the resulting lower bound. The above optimisation problem can then be cast as a well known quadratic programming problem:

$$f(x) \;\approx\; \frac{1}{2}x^T H x \;+\; x^T g \tag{13}$$

where H is the symmetric matrix of second derivatives of f and g is the direction of the gradient of f. Commercially available optimisation software, [16], can then be used to solve eqn. (13). As the search for a worst case destabilising Δ is non-convex, local minima can occur. A key issue with this approach is therefore the selection of a good initial guess for the worst-case Δ at each frequency. In fact, for the example considered in this paper, application of the above approach with a random initial guess for the worst-case Δ at each frequency produced results that were generally not much better than those produced by the standard algorithms. By using the Δ computed by the standard algorithms as the starting point for the optimisation, however, dramatic improvements in the quality of the lower bound were achieved. As shown in Figure 11(b), for example, the gap between the peak values of the upper and lower bounds has been reduced from 0.2147 (with the standard algorithms) to 0.0901 (using the proposed approach with $\varepsilon = 1e-6$) - an improvement of almost 60%. Note that a similar approach to improving the quality of μ lower bounds can be adopted in the case of purely real uncertainty, as for example in the case of using trapezoidal exclusion regions in a μ framework, [17].

Computing times for the numerical optimisations involved in the above approach are a function of the problem size, ε, the number of optimisation restarts, and also depend on internal algorithm settings. For the example considered in this paper, computing times to generate the improved lower bounds were comparable to those required by the standard algorithms.

A second problem in calculating μ-bounds encountered by the authors was that of using a frequency grid to compute bounds on μ in the case where the μ plot has narrow and high peaks, and indeed this issue has been well documented in the literature, [6], [18]. This problem is of particular concern for many aeronautical applications of μ-analysis, where, for example, aircraft structural modes can cause just such fine peaks in the μ plot. Two standard solutions to the problem are available. The first is to simply increase the resolution of the frequency grid. The second is to transform the original μ-analysis problem into a so-called "skewed μ" problem, where the frequency is introduced as an uncertainty into the Δ matrix of the LFT-based uncertainty model. The maximum value of μ over frequency can then be computed directly, as shown in [18]. Although useful, both of the above approaches have some serious drawbacks. Increasing the number of points in the frequency grid is computationally expensive, and (as we shall demonstrate) provides no guarantee of improving the accuracy of the μ bounds. Incorporation of frequency into the LFT-based uncertainty model requires a repeated real scalar uncertain parameter to be included in the Δ matrix. The number of times this parameter is repeated is equal to the number of states in the plant, therefore for high order systems the result is often a huge increase in the size of the Δ matrix. More importantly, the existence of repeated real uncertainties is known to produce conservatism (and sometimes convergence problems) for the standard mixed μ upper bound algorithm, [15]. To get around these difficulties we propose a new approach, based on constrained non-linear optimisation, which allows "safer" computation of both upper and lower μ bounds without any extra conservatism being introduced.

We illustrate our approach using the stability margin clearance criterion for the HWEM aircraft model. For a frequency grid of 50 points over the frequency range 0.01 rads/sec to 1 rad/sec, upper and lower bounds were computed as shown in Figure 12(a). Increasing the number of points in the frequency grid to 100 actually results in decreased (i.e. less accurate) bounds, as shown in Figure 12(b). This surprising result can be explained by noting that, due to the logarithmic spacing between 10^{-2} and 10^{0}, these 100 points will not include the original 50 points as a subset. This therefore allows the theoretical possibility that some point in the original 50 points could be nearer to the μ peak than any point in the subsequent 100. This is in fact what occurred for our example in Figure 12(a) and (b). To address this problem, we formulate an optimisation problem

$$\max_{\omega \in \mathscr{R}} \mu(\omega) \;, \tag{14}$$

where the cost function is a μ upper bound calculation using standard μ algorithms for a single frequency. Maximising this cost function with respect to frequency is then expected to converge to the frequency of the maximum μ value. The same approach is applied to the lower bound optimisation problem simply by including ω in the search vector x for the optimisation problem 12. Solving the above optimisation problems using standard software and an initial grid of 50 frequency points led to the results shown in Figure 12(c). Extremely tight upper and lower bounds were computed (0.8602698 and 0.8602691 respectively). In addition, the peak value of μ has now been correctly identified, in contrast to the results obtained using a frequency gridding approach.

VIII. CONCLUSIONS

This paper has described new μ-analysis tools for the clearance of flight control laws for highly augmented aircraft. These tools have been used to analyse the stability robustness properties of a flight control law for a vertical/short take-off and landing aircraft. Comparisons between the μ-analysis techniques and the classical industrial approach show that the new analysis tools can provide more rigorous and efficient analysis of worst-case aircraft stability characteristics in the presence of multiple sources of parametric uncertainty. An interesting result is the fact that worst-cases were shown by the μ-tools to occur sometimes in the interior of the uncertain parameter space, i.e. *not* at some combination of the extreme values of the uncertain parameters. This result contradicts the basic assumption of the classical gridding approach, and further motivates the use of the proposed new analysis tools. Remaining obstacles to the widespread adoption of μ-analysis tools by industry are centred on difficulties associated with efficiently generating accurate LFT-based uncertainty models for non-linear systems, and problems with routinely obtaining tight bounds on μ for all types

of analysis problems - further research in these areas is certainly needed.

IX. ACKNOWLEDGEMENTS

The authors are indebted to all the members of FM-AG11 for many helpful discussion over the course of the project.

REFERENCES

[1] Korte, U., *Tasks and Needs of the Industrial Clearance Process*, Chapter 2 of *Advanced Techniques for Clearance of Flight Control Laws*, Fielding, C. et al (Eds.), Springer Verlag, 2002, pp. 13-33.

[2] Hyde, R.A., *H^∞ Aerospace Control Design - A VSTOL Flight Application*, Springer-Verlag, 1995, pp. 123-202.

[3] Papageorgiou, G., Glover, K., D'Mello, G. and Patel, Y., "Taking Robust LPV Control into Flight on the VAAC Harrier", *in proceedings of the IEEE Conference on Decision and Control*, Sydney, 2000.

[4] D'Mello, G., *The Harrier Wide Envelope Model*, GARTEUR Technical-Publication, TP-119-08, 2000.

[5] D'Mello, G., *Selected Criteria For Clearance of the HWEM Flight Control Laws*, GARTEUR Technical-Publication, TP-119-08-A, 2002.

[6] Ferreres, G., *A Practical Approach To Robustness Analysis*, Kluwer Academic/Plenum Publishers, 1999, pp. 43-52.

[7] Mannchen, T., Bates, D.G. and Postlethwaite, I., "Modelling and Computing Worst-Case Uncertainty Combinations for Flight Control Systems Analysis", *AIAA Journal of Guidance, Control and Dynamics*, Vol. 26, No. 6, pp. 1029-1039, 2002.

[8] Deodhare, G. and Patel, V.V., 1998, "A 'Modern' Look at Gain and Phase Margins: an $\mathscr{H}^\infty$ / μ Approach", *Proceedings of the AIAA Conference on Guidance, Navigation and Control, Boston, USA*, 1998.

[9] Kureemun, R., Bates, D. G. and Postlethwaite, I., "Quantifying the Robustness of Flight Control Systems using Nichols Exclusion Regions and the Structured Singular Value", *IMechE Journal of Systems and Control Engineering*, 215(16), 2001.

[10] Bates, D.G. and Postlethwaite, I., *Robust Multivariable Control of Aerospace Systems*, Delft University Press, 2002, pp. 78-88.

[11] Kureemun, R., *μ-Analysis Tools for the Flight Clearance of Highly Augmented Aircraft*, Ph.D. Thesis, Department of Engineering, University of Leicester, November 2002.

[12] Mannchen, T., Petermann, C., Weinert, B. and Zobelein, T., *Flight Control Law Clearance of the HIRM+ Fighter Aircraft Model using μ-Analysis*, GARTEUR Technical Publication, TP-119-12, 2001.

[13] Dailey, R. and Gangsaas, D., "Worst Case Analysis of Flight Control Systems using the Structured Singular Value ", AIAA Paper No. A89-49406, 1989.

[14] Braatz, R. D. and Morari, M., "μ-Sensitivities as an Aid for Robust Identification", *Proceedings of the American Control Conference*, IEEE, Vol. 1, pp. 231-236, 1991.

[15] Balas, G., J., Doyle, J., C., Glover, K., Packard, A. and Smith, R., *μ-analysis and synthesis toolbox*, The Math Works, 1995.

[16] Branch, M., A. and Grace, A., *MATLAB Optimization toolbox user's guide*, The MathWorks, 1996.

[17] Hayes, M. J., Bates, D. G. and Postlethwaite, I., "New Tools for Computing Tight Bounds on the Real Structured Singular Value", *AIAA Journal of Guidance, Control and Dynamics*, 24(6), 2001.

[18] Sideris, A., "Elimination of Frequency Search from Robustness Tests", *IEEE Transactions on Automatic Control*, 37(10), pp. 1635-1640, 1992.

Variable name	[Min,Max] value	Units
U_Xcg	[-1.72,-11.7]	%MAC
U_I_{yy}	[56887,69529]	kgm^2
U_$C_{m_{tail}}$	[-20,+20]	%
U_C_{m_q}	[-20,+20]	%
U_C_{m_α}	[-20,+20]	%

TABLE I

CATEGORY I LONGITUDINAL UNCERTAINTIES

classical w.c. eig.	μ w.c. eig.	classical w.c. uncertainties	μ w.c. uncertainties
-0.1985	-0.1980	1	0.9682
-0.1669	-0.1660	-1	-0.9985
-0.0631	-0.0630	1	0.9771
-8.5e-5+0.0775j	7.8e-5+0.0775j	-1	-0.9906
-8.5e-5-0.0775j	7.8e-5-0.0775j	1	0.9886

TABLE II

SAMPLE RESULTS FOR WORST-CASE EIGENVALUE CRITERION AT FC1

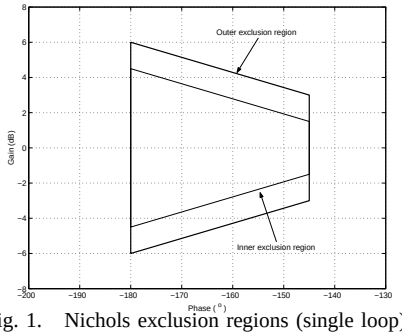

Fig. 1. Nichols exclusion regions (single loop)

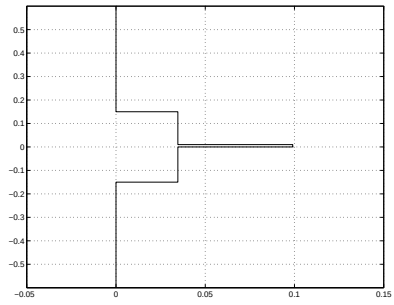

Fig. 2. Boundaries for the unstable eigenvalue requirement

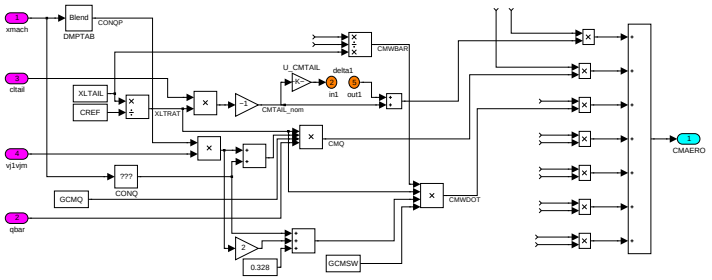

Fig. 3. Part of the HWEM model with uncertainty included

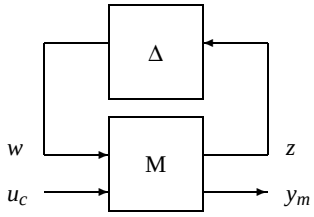

Fig. 4. Upper LFT uncertainty description

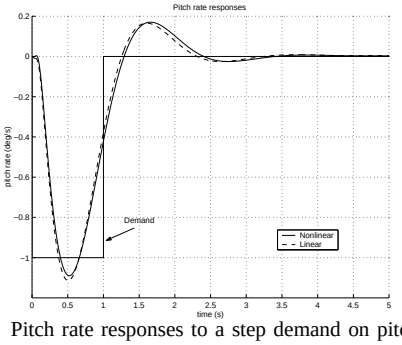

Fig. 5. Pitch rate responses to a step demand on pitch rate

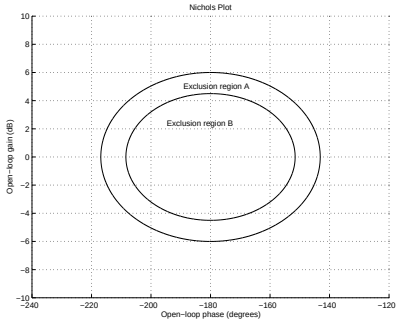

Fig. 6. Elliptical Nichols plane exclusion regions

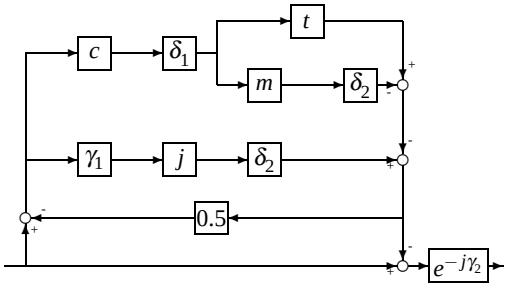

Fig. 7. LFT representation of trapezoidal Nichols exclusion region

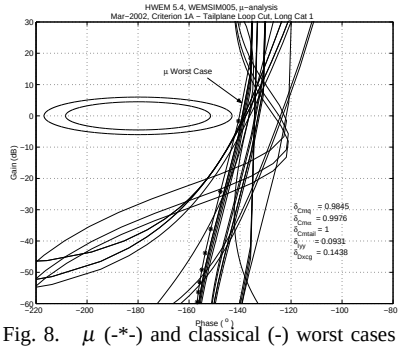

Fig. 8. μ (-*-) and classical (-) worst cases

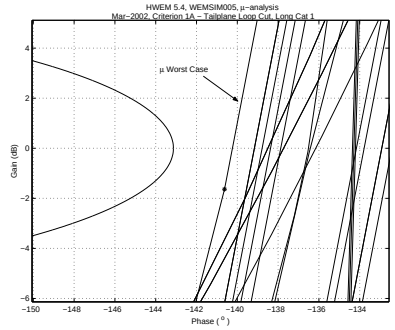

Fig. 9. Close-up of Figure 8

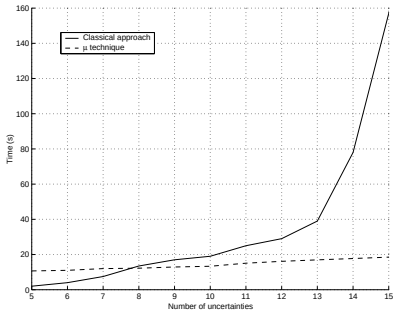

Fig. 10. Computation times for μ and classical techniques

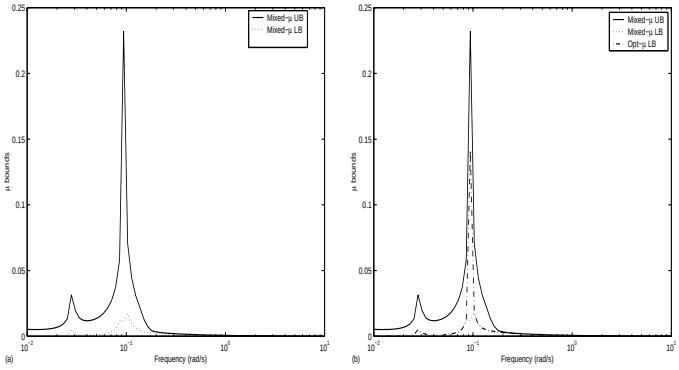

Fig. 11. (a) μ upper and lower bounds (standard mixed μ algorithms), (b) Improved mixed μ lower bound

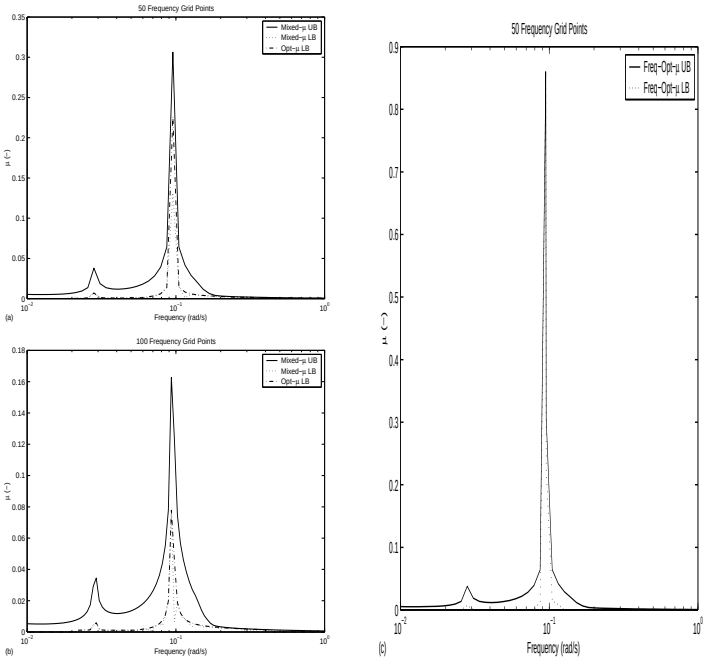

Fig. 12. (a) Standard and improved μ bounds for 50 frequency points, (b) Standard and improved μ bounds for 100 frequency points, (c) Optimisation based μ bounds for an initial 50 point frequency grid