HARRIER AIRCRAFT CONTROL LAW CLEARANCE ANALYSIS USING A BIFURCATION-BASED METHOD

M.H. Lowenberg* **C.D.C. Jones†**

Department of Aerospace Engineering, University of Bristol,
Queens Building, University Walk, Bristol, BS8 1TR, UK.

Keywords: control law clearance; bifurcation analysis; continuation methods.

Abstract

This paper provides a demonstration of control law clearance analysis on an industry-scale aircraft model using a technique based on bifurcation analysis and continuation methods. The approach was developed as part of the GARTEUR Flight Mechanics Action Group 11, set up to find more efficient means of clearing aircraft control laws. Results are presented for criteria based on eigenvalues and Nichols stability margin. It is shown that the worst case combination of uncertainties found by this analysis technique is 'worse' than that chosen by the conventional (baseline) method. Significantly, these uncertainty combinations violated the criteria for three cases in which the conventional method cleared the model. The technique offers substantial time savings over the baseline, identifies the exact conditions at which a criterion is violated and retains a strong link with the physics of the system.

1. Introduction

A lengthy and costly task during development of aircraft with full-authority control laws is clearance of the closed-loop system. This involves assessing the behaviour of the complete system – airframe, propulsion system, physical control devices, control laws, actuators, sensors, etc. – against agreed criteria. It is a time-consuming and complex task because it must consider the aircraft in all its configurations, across the flight envelope, for all controller modes and incorporating each potential failure case. Flight Mechanics Action Group 11, or FM(AG11), of the Group for Aeronautical Research and Technology in Europe (GARTEUR) was formed in 1999 to identify and research new analysis techniques for clearance of flight control laws. The intention was that these might contribute to more effective clearance, either in terms of improved performance (relative to the conventional gridding approach) in determining worst-case scenarios or significant time savings, or both.

The University of Bristol, supported by QinetiQ Bedford, developed a technique based on the numerical continuation methods used in bifurcation analysis. This so-called bifurcation-based analysis technique for control law clearance was demonstrated on a theoretical model – the HIRM+ aircraft with Robust Inverse Dynamics Estimator (RIDE) controller [1] – before being applied to the Harrier Wide Envelope Model (WEM).

The work carried out by all participants in FM(AG11) on the HIRM+/RIDE system is documented in [2].

2. WEM aircraft model

The Harrier WEM is a representation of the Vectored-thrust Aircraft Advanced flight Control (VAAC) Harrier. Its purpose, in the context of FM(AG11), was to provide a challenging task to the analysis technique due to the complexity of the nonlinear model.

Part of the complexity of the model arises from the fact that the Harrier operates over a speed range down to hover conditions, necessitating a comprehensive aerothermodynamic engine model which is itself a closed-loop system.

The control law under investigation is based on VAAC Control Law 002 (CL002). It emulates the standard Harrier's 'three-inceptor' control strategy but with some additional augmentation to reduce workload. CL002 is a classical three-axis (pitch, roll and yaw) manual thrust-vectoring control law; the pilot retains explicit control of thrust magnitude and direction via the throttle and nozzle levers. Of most concern in this partial clearance study was longitudinal control, and the flight conditions covered both pitch attitude (low speed) and pitch rate (higher speed) stick command functions.

A full description of the Harrier WEM and the control law can be found in [3].

The WEM models were supplied to the Action Group in MATLAB/Simulink form and contained both discrete and continuous states (66 in total for the full model with controller). These hybrid models caused difficulties with linearisation – affecting not only the bifurcation-based analysis but the baseline solutions too. Ultimately a continuous version was created (only for the purposes of linearisation) [4]: the three discrete engine states, arising from 1/z blocks, were replaced with fast 1st order lags; nonlinearities such as transport delays and deadbands were removed; and some memory blocks were replaced.

3. Clearance task

A number of clearance criteria were provided by industry members of FM(AG11) for application to the aircraft models. These covered stability, response and handling issues; most

* Senior Lecturer in Flight Dynamics, tel. 0117 9289765, fax. 0117 9272771 m.lowenberg@bristol.ac.uk
† Postgraduate Student, tel. 0117 9289860, fax. 0117 9272771 chris.jones@bristol.ac.uk

were linear and a couple nonlinear. The three linear criteria stipulated in [3] for use in the WEM clearance task are discussed briefly below.

3.1 Single loop stability margin (Nichols)

This criterion involves the identification of all flight cases where the Nichols plot stability margin boundary – a trap-ezoidal exclusion zone defining minimum allowable phase and gain margins – is violated. A *normalised stability margin*, ρ, is defined by scaling the boundary until it just touches the frequency response whilst preserving its aspect ratio; ρ is defined as the ratio by which it is scaled. ρ=1 corresponds to the nearest point on the response just touching the boundary; $\rho > 1$ means that there is no violation, and $\rho < 1$ means that the criterion is violated.

The single loop frequency response is obtained by breaking the loop at the input of each actuator, one at a time while leaving the other loops closed. In order to be cleared, the exclusion zone must be avoided under all combinations of uncertainty parameter. The task in each flight case is to identify the combination of uncertainty parameters that give the worst case violation (smallest ρ).

The purpose of this criterion is to assess sensitivity to changes in the dynamics of each actuation system and ensure that it maintains adequate stability margins; it gives an indication of sensitivity to changes in control power [2].

In applying the criterion to the WEM, the loop between demanded tailplane (from the CL002 control law) and the tailplane actuator was broken. There are no other actuators incorporated within CL002 for longitudinal control. A worst case combination of the uncertainty parameters was identified for each flight case (i.e. each angle of attack, AoA, in each Flight Condition).

3.2 Worst case eigenvalue criterion

This criterion requires identification of flight cases where unstable closed loop eigenvalues occur, and the combination of uncertainty parameters for which these eigenvalues have the largest value of their real part [2]. The purpose of this test is to determine the most severe cases of divergent closed loop modes, allowing an assessment of aircraft handling sensitivity. The criterion is violated if a real or complex eigenvalue crosses the specified boundary:

Let $\lambda = \sigma + j\omega$ be an eigenvalue of the linearised model; to comply with the criterion, the real part of λ, σ, must satisfy:

$$\sigma < \begin{cases} \sigma_1 = 0 & ,\text{if } \omega \in \Omega_1 = \{\omega : |\omega| \geq 0.15\,\text{rad/s}\} \\ \sigma_2 = (\ln 2)/20 & ,\text{if } \omega \in \Omega_2 = \{\omega : 0 < |\omega| < 0.15\,\text{rad/s}\} \\ \sigma_3 = (\ln 2)/7 & ,\text{if } \omega \in \Omega_3 = \{0\} \end{cases}$$

3.3 Average phase rate criterion

This test, designed to detect pilot-induced oscillation in pitch and roll, is defined in terms of pitch and bank attitude-to-stick force frequency responses [2]. For the WEM it is only to be implemented in fully wing-borne flight, which includes just one of the Flight Conditions defined for the task (at 200 kts). When applied, it was found to give level 1 behaviour which was insensitive to incidence angle and uncertainty parameter values. Therefore this criterion is not considered further here.

The clearance task defined for the WEM specified seven Flight Conditions (FCs) at which the criteria were to be applied – all at a height of 200 feet above mean sea level. Their respective flight velocities are listed in Table 1.

Table 1: Flight condition velocities for WEM clearance task.

Flt condition	FC1	FC2	FC3	FC4	FC5	FC6	FC7
Airspeed (kts)	200	150	130	110	90	60	0[†]

[†] May use 0.01 knots to avoid singularities in hover.

The two linear clearance criteria were to be applied at trim points at 2° increments in angle of attack for all FCs, over a maximum AoA range of [–4°, +16°]. The fact that a range of AoA in true 1g straight and level flight trim is achievable is due to the existence of two independent longi-tudinal control effectors: horizontal tailplane and nozzle angle. At low speeds the full range is obtainable but not at higher flight velocities.

Only longitudinal uncertainty parameters were considered in this study. The five that were specified for the FM(AG11) task are given in Table 2; see [4] for further details.

Table 2: WEM pitch uncertainty parameters.

Parameter	**Variable name**
Longitudinal position of centre of gravity	U_dxcg
Pitch moment of inertia	U_Iyy
Uncert. on tailplane effectiveness in pitch	U_CMTAIL
Uncertainty on pitching damping derivative	U_CMQ
Uncertainty on pitching stiffness derivative	U_CMALFA

For convenience, we refer to both the aerodynamic uncertain-ty parameters (U_CMTAIL, U_CMQ, U_CMALFA) and the two model variabilities (U_dxcg, U_Iyy) as 'uncertainties'.

When more than one *aerodynamic* uncertainty parameter is applied simultaneously, a 'reduction factor' is used to avoid unduly pessimistic conditions; this is based on a probability argument [2]. Each aerodynamic uncertainty value is multiplied by a factor that decreases in value as more such uncertainties are included.

The GARTEUR task requires only the vertices of uncertainty parameter space to be considered, i.e. combinations of minimum and maximum uncertainty values. We adopt the notation '–' for the minimum uncertainty value of a parameter and '+' for its maximum value; a '0' refers to its nominal value (i.e. when no uncertainty is applied).

4. Bifurcation-based analysis technique

The methodology is founded upon the use of 'continuation methods', which are a fundamental tool in numerical bifur-cation analysis. Bifurcation analysis is a process used to study the behaviour of nonlinear dynamical systems in terms of the geometry of their underlying structure, as characterised by the evolution of steady state solutions as parameters vary. Steady states include in general stationary point equilibria and periodic orbits (and other attractors) and nonlinear systems can have multiple steady states for the same values of input parameter.

One means of visualising the numerical output is the 'one-parameter bifurcation diagram': projections of the steady state solution paths as a parameter varies, plotted as one state component at a time versus the parameter. The algorithms used to generate this information are known as 'continuation methods' – and it is principally this that is adapted to form the bifurcation-based analysis technique.

Given a nonlinear dynamical system, $\dot{x} = f(x, \delta)$, where x is the state vector, δ is a vector of parameters and f is a smooth vector function, we choose one of the δ as the parameter to vary ('continuation parameter', μ) and fix the remaining members of δ. For equilibria steady states (the only type considered in this paper), we solve for $\dot{x} = f(x, \mu) = 0$ as μ varies; the idea is to find *all* solutions within the required range of μ. The continuation method is thus a path-following algorithm which, given a starting guess, attempts to continue along the solution branch. Bifurcation points are identified along the path and often it is required to solve for the new solution branches that arise from them. Local stability along the branches is indicated by use of different line types on the bifurcation diagrams; bifurcation points are also indicated where necessary.

When applied to aircraft flight dynamics models the state includes the vehicle translational and rotational velocity and orientation, while the parameters are usually the inputs to the system (control surfaces or pilot demands). However, for the purposes of control law clearance analysis, the parameters include uncertainty parameters. The process of applying continuation methods to clearance analysis requires first generating the steady state solution branch, as in standard bifurcation analysis. The model used will be set up to represent whatever form of 'trim' is specified for the clearance task. For the WEM, this refers to true trims but in some cases a form of quasi-trim is required (e.g. where a large AoA range is to be analysed, accelerated trims with non-zero pitch rate may be used).

Once each solution point is found, one or more clearance criterion is evaluated at that point. The criteria may use a different form of the model, such as true trim with controller command path omitted, to match the clearance requirement. Thus the versatility of continuation methods is exploited in the process: using one form of the model for finding the steady state solution and one or more others for application of criteria at each solution. Note that the criteria are implemented as in a conventional baseline clearance process, so there is no conservatism involved. The 'bifurcation diagrams' generated during the analysis may adopt line-type definitions corresponding to the outcome of a clearance criterion: e.g. solid line for cleared, dashed for uncleared.

A detailed description of the analysis cycle is given in [2]. In principle, the process is as follows: first, for each FC, evaluate each clearance criterion along the required trim points across the specified AoA range *for the nominal model* (no uncertainties applied). This involves a continuation run, with an appropriate pilot input as continuation parameter (e.g. stick position); it shows AoAs where the nominal system violates the criteria, or values where it comes closest to doing so. These points may be referred to as *nominal critical points* and suggest where the system should be studied further (it is this logic that provides the majority of time saving relative to the conventional gridded approach*).

The next step is to evaluate each criterion in the neighbourhood of each nominal critical point, with uncertainties applied. The continuation method is now run at each such point, with AoA fixed, and the uncertainty parameters used as continuation parameter, one at a time. In the first iteration, the remaining uncertainties are fixed at their nominal value. Each of these *nonlinear sensitivity* bifurcation diagrams indicates the change in clearance criterion as the variable uncertainty ranges from its minimum to maximum value; it reveals the value of this uncertainty that gives the worst case (biggest degradation in criterion measure) while the others are fixed at their nominal value. We repeat this step of varying one uncertainty at a time but now the others take on their worst-case value from the first iteration. Although this approach allows the worst-case value of each uncertainty to lie anywhere between its minimum and maximum values, we follow the conventional clearance process and choose either the minimum or the maximum value. Iterations continue until there is no change relative to the previous iteration.

This yields the worst-case combination of uncertainties for that specific solution point for the criterion under consideration (although it does not guarantee that this is the global worst case). Furthermore, since it gives a quantitative change in criterion measure for each uncertainty, it is possible to invoke the reduction factors for aerodynamic uncertainties. This allows the choice of *all* the uncertainties to be compared with a selection of a *subset* of the uncertainties – something that the conventional baseline method does not do.

Finally, a continuation run with the pilot input as continuation parameter is conducted again but this time using the worst-case combination of uncertainties. This identifies the AoA at which the system violates the criterion under worst-case conditions. It is only strictly applicable local to the nominal critical point because the worst-case combination was determined at that specific AoA. This is repeated in the region of each nominal critical point for each criterion at each FC, giving the desired clearance results (cleared and uncleared AoA regions).

The analysis cycle, whilst seemingly complicated when described in words, is actually rather simple. It is illustrated by a set of sample results in the next section.

---

* Violation of a criterion with uncertainties applied at an AoA far from the nominal critical points is not likely unless there is a discontinuity in the system – e.g. a non-smooth mode change – that occurs when uncertainties are applied but not in the nominal case. Such situations can be missed also in the gridding method, and prior knowledge of the existence of discontinuities should be obtained and acted on, whatever clearance technique is used.

5. Results

The bifurcation-based clearance process is illustrated using sample results for FC1, followed by a summary for all FCs.

5.1 Nominal Results

Figure 1 is a set of bifurcation diagrams for the nominal case. The continuation parameter, longitudinal stick movement (ALONG), is plotted on the x-axis. Six state components (y-axes) have been selected to provide some information on the physics of the solutions: angle of attack (α), angle of sideslip (β), pitch rate (q), roll rate (p), throttle lever movement (ATHROT) and nozzle lever movement (ANOZZ). Each point on the solution line is an equilibrium (trimmed) condition. The line type in the plots represents the maximum unstable eigenvalue criterion, the solid line indicating that all eigenvalues are acceptable in terms of the criterion. Limits on the continuation were imposed by the trimmable α range (smallest for FC1, at [–2.2°, 3.6°]); in some cases the limits are due to maximum throttle or nozzle angle range [0°, 98.5°].

We note from Figure 1 that the trim conditions are all symmetric, i.e. zero values for lateral-directional variables, as expected for this model. We can observe the changes in throttle and nozzle values through the trim range. More generally, violations of criteria can be indicated on the diagrams, saturation of control surfaces can be detected and certain nonlinearities identified.

Figure 2 shows the (Nichols) stability margin for the tailplane loop. The stability margin criterion for the tailplane loop does not reach the exclusion zone boundary (ρ=1) but is closest to it at $\alpha = -2.2^{o}$.

Figure 3 shows the worst case real eigenvalue for the nominal case, as the continuation parameter is varied. A plot of the variation of eigenvalues in the s-plane is omitted here, in the interests of brevity, but it reveals that no complex roots come close to violating the criterion. The real eigenvalues, whilst positive, do not reach the criterion boundary (+0.099). The unstable eigenvalues criterion is closest to violating the boundary at $\alpha = -1.3^{o}$.

Thus for the nominal case, FC1 is cleared over the entire α range. The results are summarised in Table 3. Note that for this, and several other cases, numerical conditions result in a non-smooth variation of criterion measure with α.

Table 3 – FC1 nominal clearance results.

Nominal critical α (eigenvalues) Maximum real eigenvalue	–1.3° 0.04058
Nominal critical α (stability margin) Minimum stability margin	–2.2° 1.257
Cleared α range	[–2.2°, 3.6°]

5.2 Worst case uncertainties applied

Figure 4 shows the nonlinear sensitivity of the stability margin criterion to one longitudinal uncertainty at a time, at the nominal critical α. Each uncertainty is varied from its minimum to its maximum normalised value of –1 to +1 whilst the others are fixed at 0. Gaps in the middle of the sensitivity analysis plots are due to an initially large choice of the continuation parameter step size (can be reduced if necessary).

The worst-case uncertainty in each case is the value for which the change in Nichols margin is negative, i.e. margin, ρ, closer to 1. In the first plot (U_dxcg) nonlinearity in the system is evident; the worst case value of U_dxcg is approximately 0.7. This demonstrates the ability for this technique to pick out worst case uncertainties at values other than the vertices. However, to compare with the baseline results, a value of +1 is chosen here. The other uncertainties yield more linear variations in the stability margin.

A 2nd iteration of the nonlinear sensitivity analysis was performed (not shown); this time, the non-varying uncertainties are fixed at the worst-case values picked from Figure 4. The predicted worst cases are unchanged from the 1st iteration – with the exception of U_CMALFA which is now +1 rather than –1 at its worst case value. A 3rd iteration yields the same outcome as the 2nd and the selected worst-case combination of uncertainties is shown in Table 4.

Table 4 – FC1 worst case uncertainties (stab. marg. criterion).

U_dxcg	U_Iyy	U_CMTAIL	U_CMQ	U_CMALFA
+1	–1	+1	–1	+1

The numerical values of the predicted stability margin under worst case conditions are then compared for each uncertainty and the reduction factors are applied to the aerodynamic uncertainties: this allows a decision to be made as to whether one, two or all three gives the smallest stability margin. The shaded region in Table 4 indicates the selected uncertainties in this case. Note that this is a fundamental difference relative to the baseline method: in the latter, *all* the aerodynamic uncertainties are used simultaneously. The bifurcation-based approach almost invariably suggests a worse case with a *subset* of all the aerodynamic uncertainties.

Figure 5 shows the bifurcation run with the worst case combination of uncertainties applied. The equivalent result from the baseline clearance is also indicated (corresponding to those in [4]), as is the nominal situation. We note that, whilst the worst case detected by the baseline method shows no violation of the criterion, the bifurcation-based result shows that the criterion is violated for α less than –1.4°. The results with uncertainties applied are summarised in Table 5.

Table 5 – FC1 worst case stability margin results summary.

Min. ρ (at nominal critical α)	Cleared α range
0.9341	[–1.4°, 3.6°]

When a similar process was applied for the unstable eigenvalues criterion, the worst-case combination was as shown in Table 6. The results of the bifurcation run over the entire α range with this worst case combination of uncertainties appears in Figure 6; the baseline and nominal results are also shown. The case with longitudinal uncertainties applied clearly does not violate the criterion.

Table 6 – FC1 worst case uncert's (unstable eig. criterion).

U_dxcg	U_Iyy	U_CMTAIL	U_CMQ	U_CMALFA
–1	–1	–1	–1	–1

Note from Figure 6 (in the α=3° region) that the bifurcation-based worst-case uncertainty combination does not necessarily hold throughout the trimmed range. It is only strictly valid local to the nominal critical α as it is at this point that the worst case combination of uncertainties was determined.

5.3 Results summary

A summary of the nominal results and those with worst-case uncertainties applied is given in Table 7. The baseline analysis, using the same model and clearance criteria implementations, revealed no violations for any of the FCs. The worst-case ranges shown in bold (FCs 1, 3 and 4) are therefore those for which the bifurcation-based method has given a materially different outcome.

These results demonstrate the effectiveness of selecting worst-case uncertainty combinations via nonlinear sensitivities, accounting for reduction factors. In principle, the bifurcation-based technique has the potential to find even worse cases by accepting normalised uncertainty values between –1 and +1 (although this may not always be worthwhile, as the small improvement in predicting worst cases may require additional iterations to be carried out).

The time taken to set up and conduct the bifurcation-based clearance task summarised here was approximately 2.5 times less than the equivalent time taken for the baseline analysis. The majority of this saving derives from the insight given by the bifurcation diagrams: this allows the search for worst cases to be confined to a limited number of angles of attack.

The general nature of the technique presented in this paper endows it with further potential in helping the industry to reduce the time and cost of control law clearance. For example, a different means of selecting the worst-case uncertainties at nominal worst case points could be incorporated within the method: an optimisation approach could be used, possibly with the nonlinear sensitivities providing the starting guess. If a 'guaranteed' worst case were required, the baseline method could be implemented at the points indicated from the nominal runs. It is in principle also possible to extend the use of the continuation methods to trace out loci of violation points through the flight envelope.

6. Conclusions

It has been demonstrated in this paper that bifurcation analysis is an extremely powerful tool for control law clearance. The clearest indicator of this in the WEM application is that the technique reproduced the baseline results in significantly less time and, furthermore, found that the stability margin criterion was violated for three flight conditions whereas the baseline method found none. This effectiveness is due to the ability to home in on specific operating points at which violation is most likely to occur using the nominal system, and to the novel way in which the worst case uncertainty combinations are determined. The latter permits full use of the probabilistic concept of reduction factors on aerodynamic uncertainties; it also allows uncertainties between their minimum and maximum values to be selected if desired.

The bifurcation-based method keeps the user in touch with physical aspects of the clearance problem, particularly through the variation of states versus parameter (bifurcation diagrams), and gives knowledge of the relative influence of the various uncertainty parameters. Although implemented here on linear criteria, the non-linear sensitivity approach can also select worst-case uncertainties for nonlinear criteria.

The flexibility of the method, based as it is on continuation methods, means that there is a variety of ways in which it could be integrated with one or more other analysis techniques to facilitate rapid and effective control law clearance.

References

1. Lowenberg, M.H. & Richardson, T.S. "HIRMPRIDE Clearance Using Bifurcation Analysis", GARTEUR/TP-119-13-v3, Oct. February 2001.
2. Fielding, C., Varga, A., Bennani, S. and Selier, M. (Eds.) *Advanced Techniques for Clearance of Flight Control Laws*, Lecture Notes in Control and Information Sciences no. 283, Springer, 2002.
3. D'Mello, G., Dickman, G. & Green, J. "The Harrier Wide Envelope Model (HWEM)", GARTEUR/TP-119-8, July 2001.
4. Lowenberg, M.H., Swain, N. & Bedford, R. "Baseline Clearance Of Harrier Wide Envelope Model", GARTEUR/TP-119-17v2, July 2002.

Table 7: Nominal and worst-case results FC1–FC7.

	Maximum real eigenvalue		Minimum stability margin		Cleared α range (deg)	
	NOMINAL	WORST-CASE	NOMINAL	WORST-CASE	NOMINAL	WORST-CASE
FC1	0.04058	0.0573	1.257	0.9341	[–2.2, 3.6]	**[–1.4 , 3.6]**
FC2	0.06067	0.0808	1.312	1.0076	[–0.4, 11.2]	[–0.4 , 11.2]
FC3	–0.0001	0.00005	1.213	0.9096	[–0.4, 16]	**[4 , 16]**
FC4	–0.00179	–0.00152	1.213	0.9343	[–1.0, 16]	**[4.2 , 16]**
FC5	–0.00386	–0.00357	1.298	1.0077	[–4, 16]	[–4 , 16]
FC6	–0.00446	–0.0043	1.535	1.282	[–4, 16]	[–4 , 16]
FC7	0.02299	0.027	1.287	1.0069	[–4, 16]	[–4 , 16]

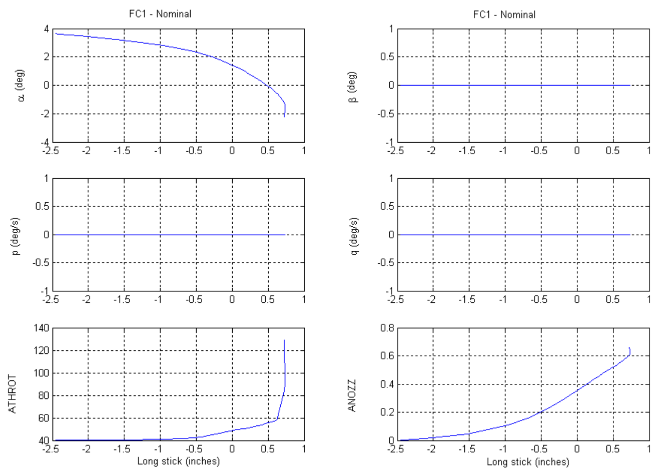

Figure 1: FC1 bifurcation diagram – nominal case.

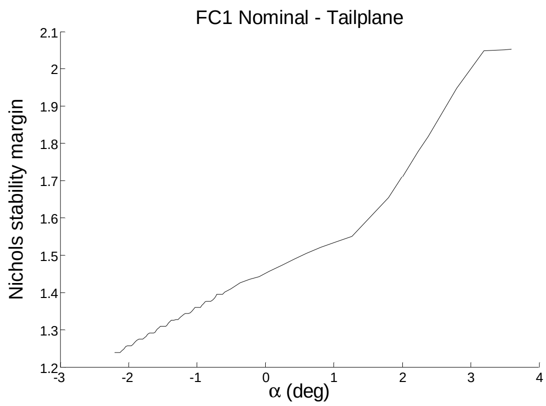

Figure 2: FC1 worst case stability margin (tailplane loop) – nominal case.

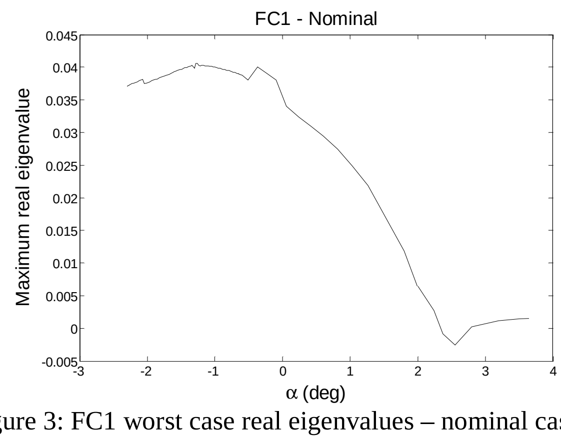

Figure 3: FC1 worst case real eigenvalues – nominal case.

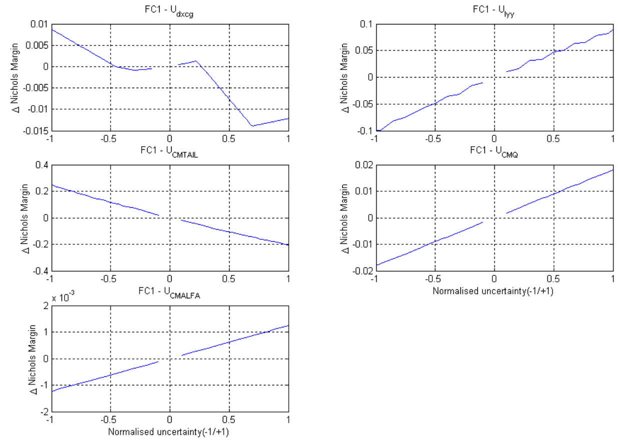

Figure 4: FC1 nonlinear sensitivity analysis – stability margin criterion.

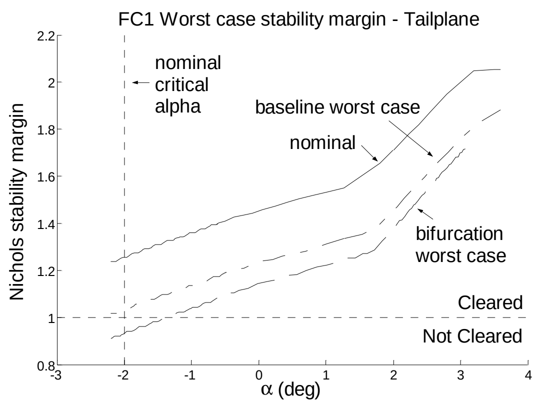

Figure 5: FC1 worst case stability margin – worst case combination of uncertainties applied.

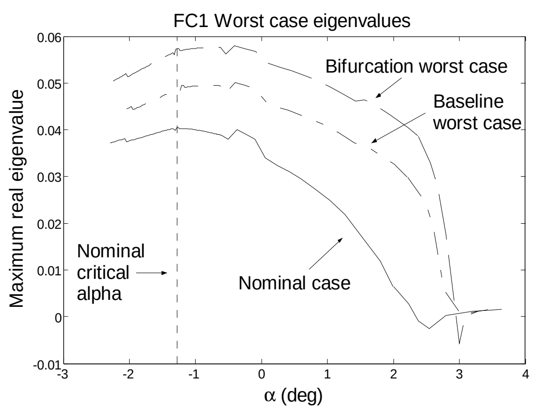

Figure 6: FC1 worst case eigenvalues – worst case combination of uncertainties applied.