IMPROVED COMPUTATION OF MIXED μ BOUNDS FOR FLIGHT CONTROL LAW ANALYSIS

T. Mannchen*, D.G. Bates†

* Institute of Flight Mechanics and Control, University of Stuttgart, Pfaffenwaldring 7a, 70550 Stuttgart, Germany
† Department of Engineering, University of Leicester, University Road, Leicester LE1 7RH, UK, dgb3@le.ac.uk

Keywords: Robustness Analysis, Flight Control, Aircraft

Abstract

In this paper we describe an approach for improving the quality of upper and lower bounds on the structured singular value in the case of mixed real and complex uncertainty. The proposed approach uses constrained non-linear optimisation to improve the bounds on μ generated by standard algorithms, and is shown to be simple and computationally efficient to implement. The usefulness of the approach is demonstrated on the problem of checking a stability margin clearance criterion for a V/STOL aircraft flight control law.

1 Introduction

It is generally possible to arrange any *linear time invariant* (LTI) closed-loop system which is subject to some unstructured and/or structured type of norm-bounded uncertainty in the form shown in Fig. 1, where P, K_1 and K_2 denote the plant, prefilter and feedback controller respectively. With respect to this figure, unstructured uncertainty means that the uncertainty matrix Δ is fully populated, while structured uncertainty means that it has some (typically diagonal or block diagonal) structure. In the context of a flight control clearance problem, unstructured uncertainty could correspond, for example, to unmodelled high frequency aircraft dynamics, while structured uncertainty is used to represent particular aircraft parameters such as stability derivatives, inertias, etc, which are subject to change, or known only to within a certain tolerance. Given a model in this form, it is then straightforward to rearrange the system into the form shown in Fig. 2, where M represents the known part of the system (aircraft model and controller) and Δ represents the uncertainty present in the system.

Partitioning M compatibly with the Δ matrix, the relationship between the input and output signals of the closed-loop system shown in Fig. 2 is then given by the upper linear fractional transformation (LFT):

$$y \;=\; \mathcal{F}_u(M,\Delta)\, r \;=\; (M_{22} + M_{21}\Delta(I - M_{11}\Delta)^{-1}M_{12})\, r \tag{1}$$

Now, assuming that the nominal closed-loop system $M(s)$ in Fig. 2 is asymptotically stable and that Δ is a complex *unstructured* uncertainty block, the *small gain theorem* (SGT), [1], gives the following result:

The closed-loop system in Fig. 2 is stable if and only if

$$\overline{\sigma}(\Delta(j\omega)) \;<\; \frac{1}{\overline{\sigma}(M_{11}(j\omega))} \tag{2}$$

This result defines a test for stability (and thus a robustness measure) for a closed-loop system subject to *unstructured uncertainty* in terms of the maximum *singular value* of the matrix M_{11}.

For aerospace systems it is often the case that uncertainty can be related to variations in specific aircraft parameters, such as centre of gravity, inertias, stability derivatives etc. In such cases, it is possible to generate models of uncertainty which have a particular structure, and thus reduce the level of conservatism in the robustness analysis. The generation of a structured LFT-based uncertainty model means that we have been able to place all of the uncertainty affecting the system into an uncertainty matrix Δ which has a diagonal or block diagonal structure, i.e.,

$$\Delta(j\omega) \;=\; diag(\Delta_1(j\omega), ..., \Delta_p(j\omega)),\; \overline{\sigma}(\Delta_i(j\omega)) \le k \tag{3}$$

where k defines an upper bound on the size of the maximum singular value of any uncertainty block Δ_i. Now again assume that the nominal closed-loop system is stable, and consider the question: What is the maximum value of k for which the closed-loop system will remain stable? We can still apply the SGT to the above problem, but the result will be conservative, since the structure of the matrix Δ will not be taken into account. The SGT will in effect assume that all of the elements of the matrix Δ are allowed to be non-zero, when we know that most of the elements are in fact zero. Thus the SGT will consider a larger set of uncertainty than is in fact possible, and the resulting robustness measure will be conservative, i.e. pessimistic.

In order to get a non-conservative solution, Doyle [2], introduced the structured singular value μ:

$$\mu_\Delta(M_{11}) \;=\; \frac{1}{\min(k \text{ s.t. } \det(I - M_{11}\Delta) = 0)} \tag{4}$$

The above result defines a test for stability (robustness measure) of a closed-loop system subject to *structured uncertainty* in terms of the maximum *structured singular value* of the matrix M_{11}. Since it is always possible to introduce scalings to make k equal to 1, the test for robust stability reduces to checking that $\mu_\Delta(M_{11})$ is less than one at all frequencies of interest. A problematic issue in applying the structured singular value

theory is that its exact computation is NP hard [3], so that the computational burden of the algorithms, which compute the exact value of μ, is necessarily an exponential function of the size of the problem. It is consequently impossible to compute the exact value of μ for large dimensional problems associated with complex industrial systems. A usual solution in this case is to compute upper and lower bounds on μ - if these are sufficiently tight, then little information is lost. Note that to fully exploit the power of the structured singular value theory, tight upper *and* lower bounds on μ are required. The upper bound provides only a sufficient condition for stability in the presence of a specified level of structured uncertainty. The lower bound provides a sufficient condition for *instability*, and also returns a worst-case Δ, i.e. a worst-case combination of uncertain parameters for the problem, [4].

The degree of difficulty involved in computing good bounds on μ depends on (a) the size of the Δ matrix, and (b) whether Δ is complex, real or mixed. For systems whose uncertain dynamics give rise to purely complex Δ matrices, polynomial time algorithms are available to compute upper and lower bounds, [5]. Both bounds converge to exact μ for low order problems and extensive computational experience, [6], has shown that the bounds remain quite tight even for complex high order problems.

For purely real μ problems, examples appear in the literature which show that μ can even be a discontinuous function of the problem data, [8]. For real μ problems with a physical engineering motivation, however, it is shown in [7] that discontinuity problems do not arise, and convergent upper, [9], and lower, [10], bound algorithms for μ exist. Unfortunately both of these algorithms are exponential time, and thus in practice this limits the size of the Δ matrix to about 11, which is much too small for many practical problems. Various approaches have been proposed to address this problem - see for example, [11, 4, 12],

For mixed real and complex uncertainty, polynomial time algorithms are available for calculating both upper and lower bounds on μ. The upper bound algorithms use LMI based optimisation, [13], while the lower bounds are generated via power algorithms, [15, 16]. The upper bound is generally quite tight, but the quality of the lower bound depends heavily on the amount of complex versus real uncertainty present in Δ. If the real uncertainty dominates, the lower bound will often be quite poor, and thus the true worst-case uncertainty combination can not be calculated. Another problem with standard algorithms for computing μ is that they employ a frequency gridding, so that for problems which give rise to sharp peaks in the μ plot, the true worst case can easily be missed. In this paper we propose an approach based on constrained non-linear optimisation to address both of the above problems.

2 Worst-Case Stability Margin Criterion

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). 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 this paper, we consider single loop analysis, where 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 3, which corresponds to a minimum gain margin of ± 6 dB and a minimum phase margin of $\pm 35^{\circ}$. When uncertainties are taken into account, a boundary corresponding to ± 4.5 dB is used, as shown by the inner exclusion region in the figure. For details of the corresponding multi loop analysis, see [18].

In order to cast the above clearance criterion as a μ problem, the original Nichols exclusion regions shown in Figure 3 are replaced with elliptical regions of the form shown in Figure 4, [20, 21]. These elliptical regions are centered around the critical point (-180,0) and satisfy the equation

$$\frac{\mid L(j\omega) \mid^2_{dB}}{G_m^{\,2}} + \frac{(\angle L(j\omega) \; + \; 180)^2}{P_m^{\,2}} = 1 \qquad (5)$$

where $L(j\omega)$ is the open-loop frequency response, G_m is the desired gain margin and P_m is the desired phase margin.

Thus for example, any feedback system whose open-loop frequency response avoids the regions A and B in Figure 4 provides gain and phase margins of $\pm 6dB / \pm 36.87^{\circ}$ and $\pm 4.5dB / \pm 28.44^{\circ}$ respectively (note that these values are very close to those required under the classical exclusion regions defined in Section 2). A key point is that 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 - see Figure 5. Thus, as shown in [20], we can represent the Nichols exclusion region as a 'fictitious' multiplicative input uncertainty for the scaled nominal plant. This uncertainty can then 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 - see [21] for a simple example. For single-loop analysis, the 'fictitious' uncertainty representing the Nichols exclusion region is inserted in one loop at a time, while in the multi-loop case it is applied to all loops simultaneously. Note that this approach allows simultaneous variations of the uncertainty in all the loops of the system and thus every possible combination of the phase/gain offset is considered. In contrast, the classical approach assumes the same phase and gain margin variation in each loop and also checks for only a few points (usually the corners) in the exclusion region. For a complete discussion of the flight control law clearance problem the reader is referred to [19].

3 HWEM Aircraft and CL002 Control Law

The Harrier Wide-Envelope Model (HWEM) is a full nonlinear 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. 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 the lateral/directional control laws the reader is referred to [17]. Points in the flight envelope of the HWEM from 200 knots to hover were 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 1. Further information about the uncertain aircraft parameters can be found in [17]. A "physical modelling" approach was used to generate an LFT-based uncertainty model to represent the uncertain aircraft dynamics at a particular flight condition, see [18] for details. Further details of the aircraft model, control law and flight clearance process can be found in [17, 18, 19].

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 1: Most relevant longitudinal uncertainties

4 Improving Mixed μ Lower Bounds

Using the approach described in Section 2, the stability margin clearance criterion was formulated as a mixed μ-analysis problem for the HWEM aircraft model. The complex part of the uncertainty matrix Δ represents the "fictitious" uncertainty associated with the elliptical Nichols plane exclusion regions, while the real part represents the uncertain aircraft parameters. Using the standard algorithms in the MATLAB μ-Analysis and Synthesis Toolbox, [5], upper and lower bounds on μ were computed for this problem, and are shown in Figure 6. Clearly, the standard lower bound algorithm is performing very poorly, due to the fact that the real uncertainty in the Δ matrix is dominant.

In order to improve the lower bound returned by the standard algorithm, we formulate the problem of computing a lower bound for μ 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

$$[x] \;=\; [p,q]^T \; p \in \mathcal{R}^l, \; q \in \mathcal{C}^m, \; l+m=n \qquad (6)$$

For real scalar uncertainty this search can be formulated as an equivalent constrained minimisation problem, $f(x)$, over a frequency range Ω:

$$\min f(x) \;=\; \min_{\Delta_{i=1..l} \in \mathcal{R}, \; \Delta_{i=l+1..n} \in \mathcal{C}, \; \omega \in \Omega} \bar{\sigma}(\Delta) \qquad (7)$$

$$\text{subject to } \underline{\sigma}(I - M_{11}\Delta) \leq \epsilon \qquad (8)$$

ϵ in the above constraint is a user defined parameter (typical values are in the range 10^{-6} to 10^{-9}) which allows the set of admissable Δ's to be expanded or contracted as required. To locate the minimising x, it is common for optimisation algorithms to consider the first two terms of the Taylor approximation of f at a candidate x. This recasts the minimisation as the well known quadratic programming problem:

$$f(x) \;\approx\; \frac{1}{2}x^T H x \;+\; x^T g \qquad (9)$$

where H is the symmetric matrix of second derivatives of f and g is the direction of the gradient of f. There is a comprehensive literature relating to the solution of this problem [22, 23]. In this paper, commercially available optimisation software, [24, 25], has been used to solve eqn. (9). Here, f(x) is minimised on a two dimensional subspace $S \;=< s_1, s_2 >$. s_1 is a vector in the steepest descent direction g so that the algorithm demonstrates fast convergence, while s_2 considers the approximate Newton direction, i.e., $H.s_2 \;=\; -g$, in an attempt to locate a global minimum. f is minimised using a line search on S:

$$\min_{\alpha \in [0,\ldots,1]} \; f : < \alpha s_1, (1-\alpha) s_2 > \qquad (10)$$

where a golden section approach is used on α to force fast convergence of the line search. A feature of this approach is that a destabilising Δ of appropriate structure is computed after each iteration of the line search algorithm. Exit criteria can easily be chosen for a particular problem so that at each frequency a good estimate of the worst case destabilising Δ will be computed. 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. If the Δ computed by the standard algorithms was used as the starting point for the optimisation, however, dramatic improvements in the quality of the lower bound can be achieved. As shown in Figure 7, 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 $\epsilon \;=\; 10^{-6}$) - an improvement of almost 60%.

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.

Extensive computational experience suggests that existing mixed μ upper bound software, [5], generally yields tight upper bounds. It is therefore reasonable to conclude that where a lower bound on μ does not track the upper bound quite closely then the optimisation algorithm is locating a local minimum. A number of additions to the basic algorithm can be tried in an attempt to improve the lower bound on μ at the expense of some increase in computing times: (a) Reduce ϵ in eqn. (7). A smaller ϵ reduces the gap between the upper and lower bounds on μ as the set of local minima satisfying eqn. (7) will now be smaller. (b) Consider frequencies that are 'close' to a local minimum. It is a straightforward extra step to consider a region very close to a candidate frequency that exhibits a poor local minimum. The maximum lower bound in this region can be significantly better than the initial candidate lower bound. (c) Edit internal optimisation settings.

5 Elimination of Frequency Gridding

The problem of using a frequency grid to compute bounds on μ in the case where the μ plot can have narrow and high peaks has been well documented in the literature, [4, 26]. This issue 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 this 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 [26]. 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, [5]. In this section 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 same stability margin clearance criterion for the HWEM aircraft model, formulated as a mixed μ-analysis problem, that was described in the previous section. 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 8. Note that increasing the number of points in the frequency grid to 100 actually results in decreased (i.e. less accurate) bounds, as shown in Figure 9. This somewhat 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 theoretically possibility that some point in the original 50 points could be nearer to the μ peak than any point in the subsequent 100. This seemingly unlikely result is exactly what is observed to occur for our example in Figures 8 and 9. To address this issue, we formulate an optimisation problem

$$\max_{\omega \in \mathcal{R}} \mu(\omega) \ , \tag{11}$$

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 defined in Section 3. Solving the above optimisation problems using standard software and an initial grid of 50 frequency points led to the results shown in Figure 10. Note that extremely tight upper and lower bounds have been computed (0.86026981 and 0.86026913 respectively). In addition, the peak value of μ has now been correctly identified, in stark contrast to the results obtained using a frequency gridding approach.

6 Conclusions

This paper has described a new approach to computing bounds on the value of the structured singular value μ, in the case of mixed real and complex uncertainties. Improved lower bounds on μ are obtained by using the worst-case uncertainties found by standard algorithms as initial guesses for a constrained non-linear optimisation problem. This optimisation problem can be solved using standard software and has been found to be simple and computationally efficient to implement. A new approach to computing upper and lower bounds on μ without the need for a frequency grid was also described, which again makes use of constrained non-linear optimisation. When applied to a realistic aircraft control law analysis problem, the proposed approach was shown to give much more reliable results, without introducing any additional conservatism into the μ bound computation.

References

[1] Skogestad, S. and Postlethwaite, I., *Multivariable Feedback Control*, Wiley, 1996.

[2] Doyle, J. C., "Analysis of feedback systems with structured uncertainty", *IEE Proceedings on Control Theory and Applications, Part D*, 129(6), pp. 242-250, 1982.

[3] Braatz, R., Young, P., Doyle, J. and Morari, M., "Computational complexity of μ calculation", *IEEE Transactions on Automatic Control*, AC-39(5), pp. 1000-1002, 1994.

[4] Ferreres, G., *A Practical approach to robustness analysis with aeronautical applications*, Kluwer Academic, 1999.

[5] Balas, G., Doyle, J.C., Glover, K., Packard, A. and Smith, R., *μ-analysis and synthesis toolbox*, MathWorks, 1995.

[6] Packard, A., Zhou, K., Pandey, P., Leonhardsen, J. and Balas, G., "Optimal I/O similarity scaling for full information and state-feedback problems", *Systems and Control Letters*, 19, pp. 271-280, 1991.

[7] Packard, A. and Pandey P., "Continuity properties of the real/complex structured singular value", *IEEE Transactions on Automatic Control*,(38)3, pp. 415-428, 1993.

[8] Barmish, B. R., P. P. Khargonekar, P. P., and Shi, Z. "Robustness margin need not be a continuous function of the problem data", *Systems and Control Letters*, Vol. 15, pp. 91-98, 1990.

[9] Jones, R., "Structured singular value analysis for real parameter variations", *Proceedings of the AIAA Conference on Guidance, Navigation and Control* , AIAA, Vol. 2, pp. 1424-1432, 1987.

[10] Dailey, R. "A new algorithm for the real structured singular value", *Proceedings of the American Control Conference*, IEEE, Vol. 3, pp. 3036-3040, 1990.

[11] Magni, J. and Doll, C., "A new simple lower bound of the mixed structured singular value", *Proceedings of the 1997 Asian Control Conference*, Vol. 1, pp. 847-850, 1997.

[12] 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.

[13] Fan, M., Tits, A. and Doyle, J., "Robustness in the presence of mixed parametric uncertainty and unmodelled dynamics",*IEEE Transactions on Automatic Control*, AC-36(1), pp. 25-38, 1991.

[14] Safonov, M. and Lee, P., "A multiplier method for computing real multivariable stability margins", *Proceedings of the IFAC World Congress*, IFAC, pp. 275-278, 1993.

[15] Packard, A. and Doyle, J., "A power method for the structured singular value", *Proceedings of the American Control Conference*, IEEE, Vol. 2, pp. 1213-1218, 1988.

[16] Young, P., M., Newlin, M., P., and Doyle, J., C., "Computing bounds for the mixed μ problem", *International Journal of Robust and Nonlinear Control*, Vol. 5, pp. 573-590, 1995.

[17] D'Mello, G., *The Harrier Wide Envelope Model*, GARTEUR Technical-Publication, TP-119-8, version 1, 2000.

[18] Bates, D.G., Kureemun, R. and Mannchen, T., "Improved clearance of a flight control law using μ-analysis techniques", *to appear in the Journal of Guidance, Control and Dynamics*, 2003.

[19] Korte, U., *Tasks and Needs of the Industrial Clearance Process*, Chapter 2 of *Advanced Techniques for Clearance of Flight Control Laws*, Fielding, C., Varga, A., Bennani, S., and Selier, S., (Eds.), Springer Verlag, Lecture Notes in Control and Information Sciences No. 283, 2002.

[20] Deodhare, G. and Patel, V.V., 1998, "A 'Modern' Look at Gain and Phase Margins: an $\mathcal{H}^\infty$ / μ Approach", *Proc. of the AIAA Conference on GNC, Boston, USA*, 1998.

[21] Bates, D.G. and Postlethwaite, I., *Robust Multivariable Control of Aerospace Systems*, Delft University Press, 2002.

[22] Fletcher, R. *Practical methods of optimisation*, Wiley (2nd ed.), 1987.

[23] Gill, P. E., Murray W. and Wright, M. H. *Practical optimisation*, Academic Press, 1981.

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

[25] Coleman, T. F., and Y. Li, "A reflective newton method for minimising a quadratic function subject to bounds on some of the variables", *SIAM Journal on Optimisation*, Vol. 6, pp. 1040-1058, 1996.

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

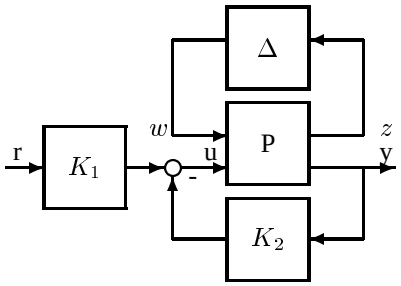

Figure 1: General uncertain closed-loop system

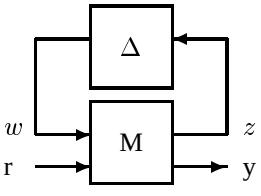

Figure 2: Upper LFT uncertainty description

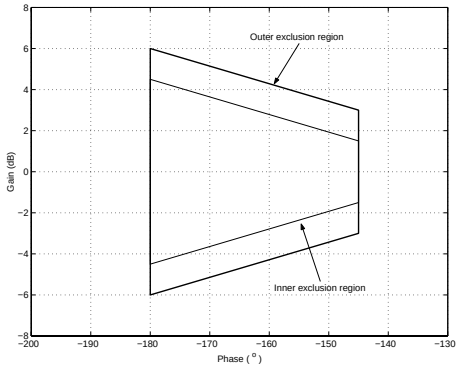

Figure 3: Nichols stability margin boundaries (single loop analysis)

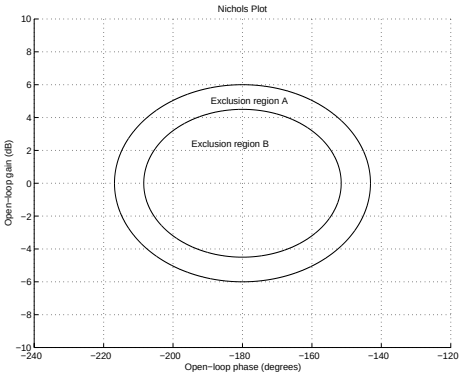

Figure 4: Elliptical Nichols plane exclusion regions

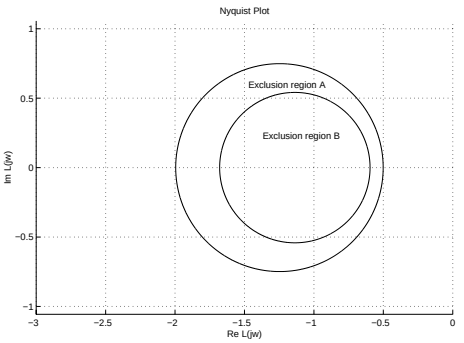

Figure 5: Corresponding circular Nyquist plane exclusion regions

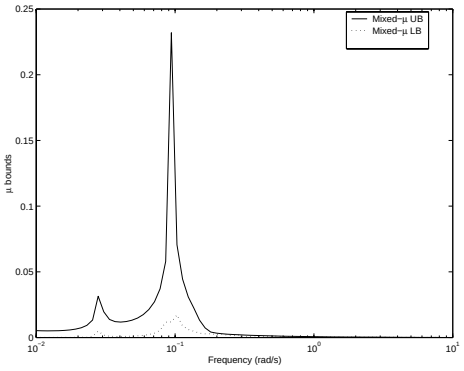

Figure 6: μ upper and lower bounds (standard mixed μ algorithms)

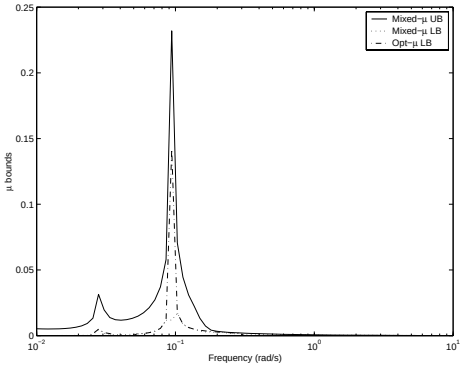

Figure 7: Improved mixed μ lower bound

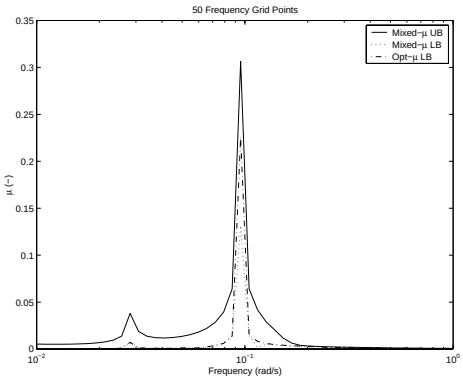

Figure 8: Standard and improved μ bounds for 50 frequency points

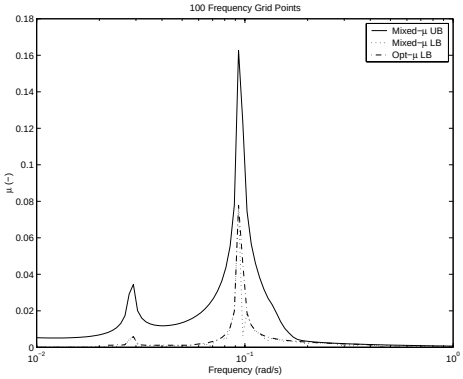

Figure 9: Standard and improved μ bounds for 100 frequency points

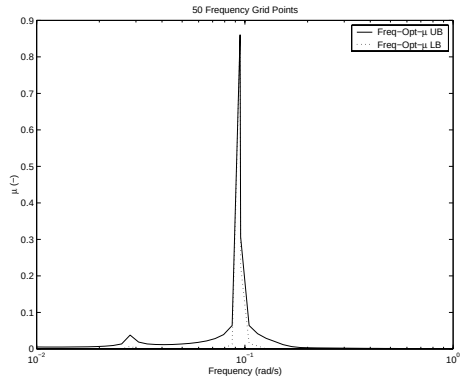

Figure 10: Optimisation based μ bounds for an initial 50 point frequency grid