JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 24, No. 6, November–December 2001

New Tools for Computing Tight Bounds on the Real Structured Singular Value

Martin J. Hayes*
University of Limerick, Limerick, Ireland
and
Declan G. Bates† and Ian Postlethwaite‡
University of Leicester, Leicester, England LE1 7RH, United Kingdom

New tools are presented for the computation of tight lower bounds on the structured singular value μ, for high-order plants subject to purely real parametric uncertainty. The first approach uses the μ-sensitivity function to systematically reduce the order of the real uncertainty matrix, so that exponential time lower bound algorithms can be applied. The second approach formulates the search for a worst-case real destabilizing perturbation as a constrained nonlinear optimization problem. Both approaches are applied to the problem of analyzing the stability robustness properties of an integrated flight and propulsion control system for an experimental vertical/short takeoff and landing aircraft configuration. Currently available software tools for calculating lower bounds on real μ fail for this problem, whereas both new approaches deliver tight bounds over the frequency range of interest.

Nomenclature

K	=	controller
P	=	plant transfer function matrix
r	=	reference demands
u	=	control inputs
w	=	ouputs of uncertainty matrix
y	=	controlled variables
z	=	inputs to uncertainty matrix
Δ	=	uncertainty matrix
Δ_i	=	ith element of uncertainty matrix
$\bar{\sigma}$	=	maximum singular value

I. Introduction

IN general, any linear time invariant closed loop-system of the type shown in Fig. 1, which is subject to some unstructured and/or structured type of norm-bounded uncertainty, can be rearranged into the form shown in Fig. 2, where M represents the known part of the system (plant and controller) and Δ represents the uncertainty present in the system. When M is partitioned 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):

$$y = \mathcal{F}_u(M, \Delta)r = \left[M_{22} + M_{21}\Delta(I - M_{11}\Delta)^{-1}M_{12}\right]r \quad (1)$$

We assume that the nominal closed loop in Fig. 2 is asymptotically stable and that the uncertainty matrix Δ has a block diagonal structure, that is,

$$\Delta(j\omega) = \text{diag}[\Delta_1(j\omega), \ldots, \Delta_n(j\omega)\big], \qquad \bar{\sigma}[\Delta_i(j\omega)] \leq k \quad (2)$$

Now consider the following question: What is the maximum value of k for which the closed-loop system will remain stable? We can apply the small gain theorem[1] to the preceding problem, but the result will be conservative because the block diagonal structure of the matrix Δ will not be taken into account. To get a nonconservative solution, Doyle[2] and Packard and Doyle[3] introduced the structured singular value μ defined as

$$\mu_\Delta(M_{11}) = 1/\min[k \text{ subject to } \det(I - M_{11}\Delta) = 0] \quad (3)$$

The preceding result defines a test for stability of a closed-loop system subject to structured uncertainty in terms of the maximum structured singular value of the matrix M_{11}. Because 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 non-polynominal (NP) hard,[4] 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 exploit fully 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, that is, a worst-case combination of uncertain parameters for the problem.[5]

The degree of difficulty involved in computing good bounds on μ depends on (1) the size of the Δ matrix and (2) 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.[6] Both bounds converge to exact μ for low-order problems, and extensive computational experience[7] has shown that the bounds remain quite tight even for complex high-order problems. 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 linear matrix inequality- (LMI-) based optimization,[8,9] whereas the lower bounds are generated via power algorithms.[10,11] The upper bound is generally quite tight, but the quality of the lower bound depends heavily on the amount of complex vs real uncertainty present in Δ.

For many well-motivated stability robustness analysis problems,[12] however, LFT uncertainty modeling gives rise to purely real Δ matrices, and here the situation becomes much more complicated. For purely real μ problems, examples appear in the literature that show that μ can even be a discontinuous function of the problem data.[13,14] For real μ problems with a physical engineering motivation, however, it is shown in Ref. 13 that discontinuity problems do not arise, and convergent upper[15] and lower[16] bound algorithms for μ exist. Unfortunately, both of these algorithms are exponential

Received 13 June 2000; revision received 1 November 2000; accepted for publication 1 November 2000.

*Lecturer, Department of Electronic and Computer Engineering.
†Lecturer, Department of Engineering, University Road; dgb3@le.ac.uk. Member AIAA.
‡Professor, Department of Engineering, University Road. Member AIAA.

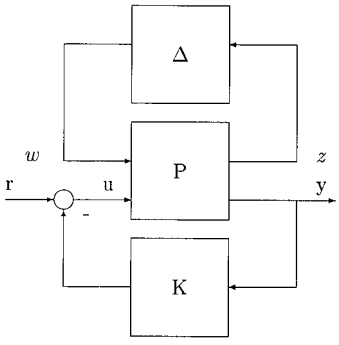

Fig. 1 Interconnection structure of a general uncertain closed-loop system.

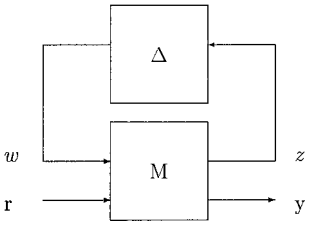

Fig. 2 Upper LFT uncertainty description.

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. A computational algorithm to calculate the exact value of real μ using a branch and bound approach is described in Ref. 17, but this approach is again exponential time. Application of polynomial time mixed μ upper bound algorithms to problems containing purely real Δ generally gives quite good results, even for high-order systems. However, the lower bound on real μ obtained from mixed μ algorithms is generally poor and often fails to converge, particularly for high-dimensional problems. A common fix is to add small amounts of artificial complex uncertainty to each real Δ_i to improve the lower bound derived from the mixed μ algorithm.[6] Whereas this approach can work reasonably well,[12] it has some serious disadvantages. First is that the introduction of enough complex uncertainty to generate a tight lower bound can significantly increase the associated upper bound, thus making the analysis results more conservative and more difficult to interpret. Second, it has been shown that the real part of the worst-case structured Δ determined using existing μ software is particularly sensitive to the addition of uncertainty in this manner.[14] This can result in the computation of a combination of uncertain real parameters that may be very different from the real worst-case combination.

In Refs. 5 and 18, a state-space approach for computing the peak values of the lower bound on real μ is proposed. The approach basically consists of extracting the real part from a destabilizing mixed uncertainty matrix and increasing it until one of the closed loop poles migrates through the imaginary axis. Whereas this approach is polynomial time and can thus be applied to high-dimensional problems, it only returns the peak value of the lower bound over a specified frequency range and cannot be used to generate a tight lower bound at each point of a frequency gridding. Information on the shape of the μ lower bound as a function of frequency can provide insight into the type of uncertainty that is causing the problem, for example, narrow peaks on the μ plot due to aircraft bending modes.

This paper proposes some new tools that may be used to compute tight lower bounds on μ at each point of a frequency gridding, for plants subject to high-dimensional real parametric uncertainty. The first approach uses μ-sensitivity functions to approximate a high-order μ-analysis problem by a smaller one that can be solved exactly using existing exponential time methods. The second approach recasts the computation of the μ lower bound for high-order systems as a constrained nonlinear optimization problem. Both approaches are computationally efficient and, when applied in combination with current mixed-μ upper bound algorithms, allow the powerful structured singular value theory to be applied successfully to realistic high-order robustness analysis problems.

The paper is organized as follows. In the next section we describe the robustness analysis problem that has been used to compare the currently available μ computational software with the new tools proposed in this paper. In Sec. III we describe the μ-sensitivity functions and their use in reducing the order of the Δ uncertainty matrix so that exponential time lower bound algorithms can be applied. In Sec. IV we describe an alternative approach using constrained nonlinear optimization, which is seen to generate tight lower bounds for the full high-order system. Finally, in Sec. V, we present some conclusions.

II. Stability Robustness Analysis of an Integrated Flight and Propulsion Control Law

We describe a robustness analysis problem typical of those encountered in the flight control law clearance process for modern civil and military aircraft. The aircraft simulation model used has been developed by the U.K. Defence and Evaluation Research Agency (DERA) to explore the challenges and possibilities associated with the design of integrated flight and propulsion control systems for future vertical/short takeoff and landing (V/STOL) aircraft. The airframe model is based on the nonlinear DERA Bedford Harrier T.Mk4 wide envelope model (WEM). This model has been used extensively in the vectored thrust aircraft advanced flight control Harrier research program[19] and has been established through flight trials as being representative of the real aircraft. To allow the development of more advanced engine control concepts, the original Pegasus engine previously included in the WEM has been replaced with a high-fidelity thermodynamic model of the Rolls–Royce Spey engine. The Spey is a two-spool reheated turbofan engine with the same basic architecture, for the purposes of control, as the EJ200, which is used to power the Eurofighter. The work described in this paper focuses on the analysis of the longitudinal axis controller only. A centralized integrated flight and propulsion controller[20,21] was designed for a linearized representation of the Spey-WEM model (omitting lateral/directional airframe states) at the 80-kn trim point of the flight envelope, using the method of $\mathcal{H}_\infty$ loop shaping.[22] The control law follows a two-inceptor strategy. In this scheme fore/aft displacement of the center stick produces a change in flight path angle rate $\dot{\gamma}$, and displacement of the left-hand inceptor demands aircraft velocity VT parallel to this flight path. The state-space model of the integrated airframe and engine system has 35 states: 14 engine states, 16 actuator states, and 5 airframe states. The eight control inputs are elevator, pitch reaction control system, front nozzles, rear nozzles, thrust split, main fuel flow, exit nozzle area, and inlet guide-vane angle. The seven controlled variables were chosen as VT, $\dot{\gamma}$, and angle of attack α (which is to be limited within specified minimum and maximum values), and four internal engine variables, which are to kept within their maximum safety limits. For more detail about the Spey-WEM model and the design of the centralized integrated flight and propulsion control (IFPC) system, refer to Ref. 21.

For the purposes of analyzing the stability robustness properties of the IFPC system, uncertainties in the nominal plant model arising from two distinct sources are considered: variations in the values of the aircraft mass and center of gravity and variations in the airframe/engine dynamics over different regions of the V/STOL flight envelope. To use μ to analyze the robustness of the IFPC system, it is necessary to generate LFT- based parametric uncertainty descriptions for the model and its associated uncertainty. The approach used for generating such descriptions will necessarily depend on the nature (and size) of the plant under consideration. For a general nonlinear plant of the form

$$\dot{\boldsymbol{x}}(t) = F[\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{p}], \qquad \boldsymbol{y}(t) = G[\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{p}] \quad (4)$$

where $\boldsymbol{x}$, $\boldsymbol{u}$, and $\boldsymbol{y}$ are the state, input, and output vectors, respectively, and $\boldsymbol{p}$ is a vector of uncertain parameters, symbolic linearization methods[23] can be used to generate linear, rationally parameter-dependent representations about some equilibrium point of the form

$$\delta\dot{\boldsymbol{x}} = A(\boldsymbol{p})\delta\boldsymbol{x} + B(\boldsymbol{p})\delta\boldsymbol{u}, \qquad \delta\boldsymbol{y} = C(\boldsymbol{p})\delta\boldsymbol{x} + D(\boldsymbol{p})\delta\boldsymbol{u} \quad (5)$$

The model (5) is nonconservative in that the entries of the state-space matrices (which are rational functions of the uncertain parameters $\boldsymbol{p}$) preserve an exact description of joint parametric dependencies in the original nonlinear model. Transformation of the model (5) into

an LFT form for μ analysis is then relatively straightforward. For plants of the type considered in this paper, however, several problems arise with the preceding approach. First, because of the high order of the plant, the computational burden associated with the required symbolic manipulations (even for relatively few parameters in p) can easily become prohibitive. Second, precise models relating the effect of the parameters p on the nonlinear system, as in Eq. (4), may not be readily available. Third, models of the form (5) are only valid in a neighborhood around a particular linearization point and so are not generally appropriate for investigating robustness over large regions of the operating envelope. An alternative approach that avoids these problems, at the expense of a certain amount of conservatism, is therefore used in this case.

In this approach numerical linearizations are repeatedly performed over several points in the aircraft operating envelope and/or over all combinations of the extreme points of the uncertain parameters. For an uncertain parameter vector $\boldsymbol{p}$ of size n, with each parameter lying between some minimum and maximum value, we thus generate a set of 2^n linear models. These models form a so-called multimodel state description

$$\delta\dot{\boldsymbol{x}}(t) = A_i\delta\boldsymbol{x}(t) + B_i\delta\boldsymbol{u}(t), \qquad \delta\boldsymbol{y}(t) = C_i\delta\boldsymbol{x}(t) + D_i\delta\boldsymbol{u}(t) \quad (6)$$

For each varying element of each state-space matrix, we can now calculate its minimum, for example, $a_{ij}^{\min}$; maximum, for example, $a_{ij}^{\max}$; and nominal, for example, $(a_{ij}^{\max} + a_{ij}^{\min})/2$, values. Thus, we can replace the multimodel system (6) by an affine parameter-dependent representation of the form

$$\delta\dot{\boldsymbol{x}}(t) = A_{p_A}\delta\boldsymbol{x}(t) + B_{p_B}\delta\boldsymbol{u}(t)$$

$$\delta\boldsymbol{y}(t) = C_{p_C}\delta\boldsymbol{x}(t) + D_{p_D}\delta\boldsymbol{u}(t) \quad (7)$$

where the state matrices are in the form

$$\begin{bmatrix} A_{p_A} & B_{p_B} \\ C_{p_C} & D_{p_D} \end{bmatrix} = \begin{bmatrix} A_{p0} & B_{p0} \\ C_{p0} & D_{p0} \end{bmatrix} + \sum_{i=1}^{n_A} \Delta_i \begin{bmatrix} A_{pi} & 0 \\ 0 & 0 \end{bmatrix} + \sum_{i=n_A+1}^{n_B} \Delta_i \begin{bmatrix} 0 & B_{ip} \\ 0 & 0 \end{bmatrix} + \sum_{i=n_B+1}^{n_C} \Delta_i \begin{bmatrix} 0 & 0 \\ C_{ip} & 0 \end{bmatrix} + \sum_{i=n_C+1}^{n} \Delta_i \begin{bmatrix} 0 & 0 \\ 0 & D_{ip} \end{bmatrix} \quad (8)$$

Thus, for example, for each varying entry in the matrices A_i, we have a Δ_i and an A_{pi} in the preceding expression, where Δ_i is an uncertain real scalar parameter that varies between 1 and -1 and A_{pi} is equal to $[(a_{ij}^{\max} - a_{ij}^{\min})/2]$. Each of the matrices associated with each Δ_i has rank one and can be factored using the singular value decomposition (SVD) into row and column vectors:

$$\begin{bmatrix} A_{pi} & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} E_i \\ F_i \end{bmatrix} \begin{bmatrix} G_i & H_i \end{bmatrix} \quad (9)$$

Now proceeding according to the method of Ref. 24, we define the linear system P with extra inputs and outputs via the equations

$$\begin{bmatrix} \dot{x} \\ y \\ z_1 \\ \vdots \\ \vdots \\ z_n \end{bmatrix} = \begin{bmatrix} A_0 & B_0 & E_1 & .... & E_n \\ C_0 & D_0 & F_1 & .... & F_n \\ G_1 & H_1 & 0 & .... & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ G_n & H_n & 0 & .... & 0 \end{bmatrix} \begin{bmatrix} x \\ u \\ w_1 \\ \vdots \\ \vdots \\ w_n \end{bmatrix} \quad (10)$$

It is then easy to form the standard closed-loop interconnection structure for μ analysis shown in Fig. 3. In this study, this method was used to generate an uncertainty description for the Spey-WEM model, to assess the robustness of its IFPC system to variations in the aircraft's mass m and center of gravity x c.g., that is, $p = [m, x \text{ c.g.}]$. The final uncertainty description for variations in p produced a Δ matrix of dimension 76, consisting of nonrepeated purely real scalars.

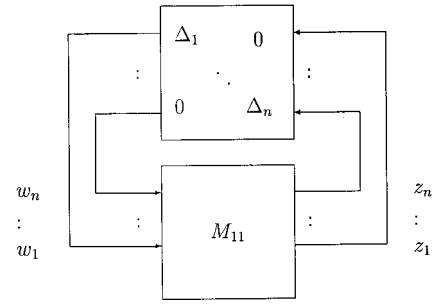

Fig. 3 Standard block diagram for μ analysis.

The approach adopted to evaluate the robustness of the IFPC system to variations in the airframe and engine dynamics over the aircraft's flight envelope is slightly different to that described earlier. In this analysis, differences between the linear model at the 80-kn design point and the linear models at various other points in the envelope are to be represented to analyze the extent of the expected degradation in stability and performance robustness as the aircraft moves away from its nominal design point. In standard μ analysis, the uncertainty matrix Δ is required to be normalized so that each Δ_i can assume any value in the interval $[-1, +1]$. For the particular problem under consideration, this is not appropriate because we seek to represent changes in the plant state-space matrices from 80 to 120 kn (for example) and not from 80 ± 40 kn, that is, a one-sided uncertainty. Hence, the following modifications are made to the standard configuration to prevent μ considering an uncertainty set that is too large:

1) Offset the appropriate elements of M_{11} with an additive Δ perturbation where $\Delta_i = 0.5\,\forall i$.

2) Scale the outputs from Δ to M_{11} by 0.5. This has the effect of constraining $\Delta_i \in [-0.5, 0.5]$.

M_{11} is now in a suitable form so that μ can be determined with a correct uncertainty set. By use of this method, the uncertainty description representing variations in the airframe and engine dynamics between 80 and 120 kn produced a Δ matrix of dimension 30, consisting of nonrepeated purely real scalars. The Δ matrix corresponding to the variations between 80 and 50 kn was of the same dimension.

III. Real μ Lower Bounds by Δ-Matrix Order Reduction Using μ Sensitivities

LFT uncertainty descriptions of the type described in the preceding section present significant difficulties for current μ computational software. The high dimensions of the Δ matrices involved mean that the use of exponential time lower bound algorithms for real μ (Ref. 16) is ruled out a priori. On the other hand, the polynomial time mixed μ lower bound algorithms available in Ref. 6 do not produce useful results and often fail to converge, because of the size of the uncertainty matrix and because it is purely real. In this section we propose a way of overcoming these problems by selectively reducing the size of the Δ matrix until the use of exponential time lower bound algorithms becomes feasible. The lower bound generated for this reduced-order problem will then of course still be a valid lower bound for the original problem. To ensure that the reduced-order lower bound is a tight lower bound for the original problem, we seek to identify and eliminate those elements Δ_i of the Δ matrix that make little or no contribution to the maximum value of μ. The tool that we use for identifying these insignificant Δ_i is the μ-sensitivity function[25] defined as

$$\delta\mu =: \lim_{\Delta\alpha \to 0^+} \frac{\mu[M_{11}(\alpha)] - \mu[M_{11}(\alpha - \Delta\alpha)]}{\Delta\alpha} \Bigg| \alpha = 1 \quad (11)$$

where α is a scalar introduced into the standard block diagram for μ analysis as shown in Fig. 4. For the purposes of μ computation, α is then absorbed into the M_{11} matrix in the usual fashion. Note that $\mu[M_{11}(\alpha)]$ is a nondecreasing function of α, so that $\delta\mu$ is necessarily nonnegative. Moreover, it is shown in Ref. 25 that the μ sensitivities are well defined and equal to the corresponding full derivatives,

almost everywhere on any interval of α. Thus, $\delta\mu$ can be loosely interpreted as the derivative of μ with respect to α. Two alternatives exist for computing $\delta\mu$. The first computes $\delta\mu$ exactly, using linear algebra manipulations to derive equations for the derivatives.[26] A much simpler approach is to calculate $\delta\mu$ approximately using finite differences, as proposed in Ref. 25. Because we are actually computing $\delta\mu$ with respect to the upper bound for μ (which we can calculate to a high precision) and because we are essentially interested in the relative (as opposed to the absolute) values of $\delta\mu$ for each uncertainty, this approach is adopted for the current problem.

Here μ sensitivities were computed for each of the real scalar Δ_i for each of the three Δ matrices in the LFT-based uncertainty descriptions associated with the three robustness analysis problems derived in the preceding section. The extent of the spread in the values of the μ sensitivities is shown in Table 1, where the five largest and smallest μ sensitivities for the 80–50 kn uncertainty Δ are shown. Each Δ matrix was, thus, reduced to a size for which the exponential time lower-bound algorithm of Ref. 16 can be applied by eliminating those elements Δ_i with smallest μ sensitivities. The effect of eliminating different numbers of elements from Δ can be seen by the change (if any) in the associated upper bounds: If no significant Δ_i have been eliminated, the upper bounds for the full and reduced systems should remain close.

Table 1 Comparison of five largest and five smallest μ sensitivities for plant dynamics variations between 80 and 50 kn

Largest	μ sensitivities	Smallest	μ sensitivities
Δ_3	0.1164	Δ_{30}	7.7716×10^{-13}
Δ_2	0.1123	Δ_{22}	8.8818×10^{-13}
Δ_{17}	3.8751×10^{-1}	Δ_{29}	1.3323×10^{-12}
Δ_{15}	1.0759×10^{-1}	Δ_{23}	1.5543×10^{-12}
Δ_{18}	4.1405×10^{-2}	Δ_{27}	1.1779×10^{-9}

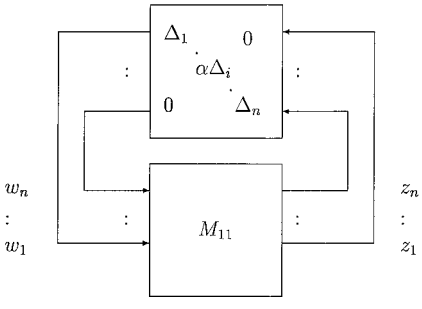

Fig. 4 Introduction of a scalar weight in the interconnection structure.

Figures 5–10 show the robustness of the IFPC system to variations in the plant dynamics, mass, and xc.g. Figures 5, 7, and 9 show the relative effect of eliminating increasing numbers of elements from the Δ matrix for each of the LFT-based uncertainty descriptions derived in the preceding section. In all cases, the uncertainty matrix could be reduced sufficiently to allow reasonably tight lower bounds to be computed using the exponential time algorithm of Ref. 16. Figures 6, 8, and 10 show the resulting exponential time lower bounds for each Δ matrix reduced to dimensions 9 and 10. Computing times for the lower bound on μ using the algorithm in Ref. 16, for 100 frequency points on a 133-MHz personal computer were approximately 30 min for Δ reduced to dimension 9 and 60 min for Δ of dimension 10. Note that the lower bounds generated by the polynomial time mixed μ algorithm currently available in Ref. 6 are everywhere equal to zero for all three problems, that is, the algorithm failed to converge. This is particularly a problem for the robustness analysis with respect to plant dynamics variations between 80 and 120 kn (Fig. 6) because the upper bound result alone does not allow any positive conclusions to be made about the robust stability of the IFPC system. Clearly, the proposed approach offers significant improvements over current tools, generating tight bounds on μ for each of the three problems addressed.

Although the results described are encouraging, it is clearly not possible to guarantee that the Δ order reduction procedure will work for all problems: There may simply be no way of reducing Δ to a size for which exponential time algorithms can be applied without eliminating significant Δ_i. In the case of such problems, an alternative approach may be employed as described in the following section.

IV. Real μ Lower Bounds via Constrained Nonlinear Optimization

The problem of computing a lower bound for μ is equivalent to a search for the worst case, that is, smallest, real diagonal destabilizing uncertainty matrix Δ. Denote the vector of Δ_i diagonal entries of Δ by $\boldsymbol{x}$. Thus, if $\Delta \in \mathcal{R}^{m \times m}$

$$[\boldsymbol{x}] = [\Delta_1, \ldots, \Delta_m]^T \in \mathcal{R}^m \tag{12}$$

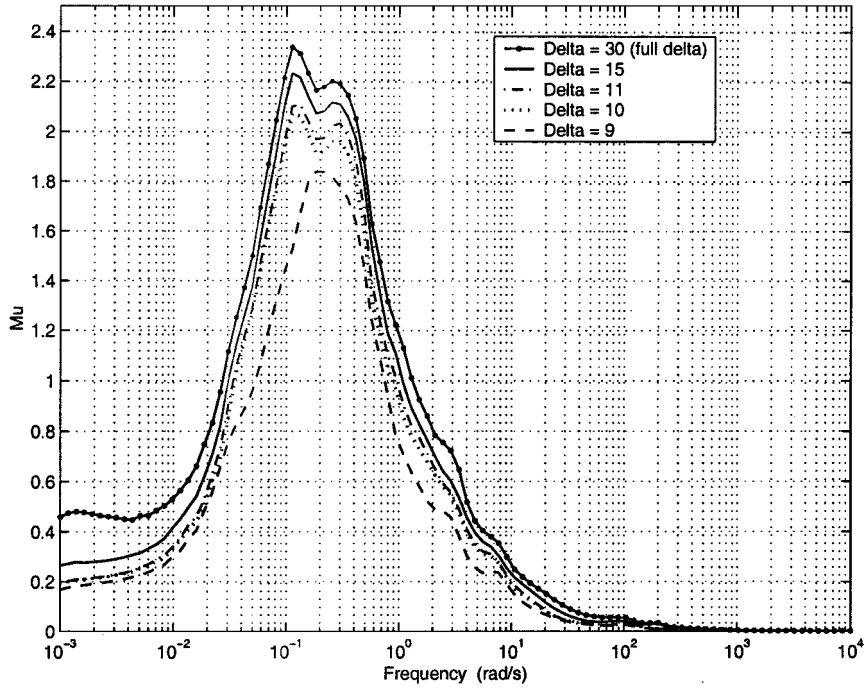

Fig. 5 Robustness of IFPC system to plant dynamics variations between 80 and 120 kn: μ upper bounds for different size Δ.

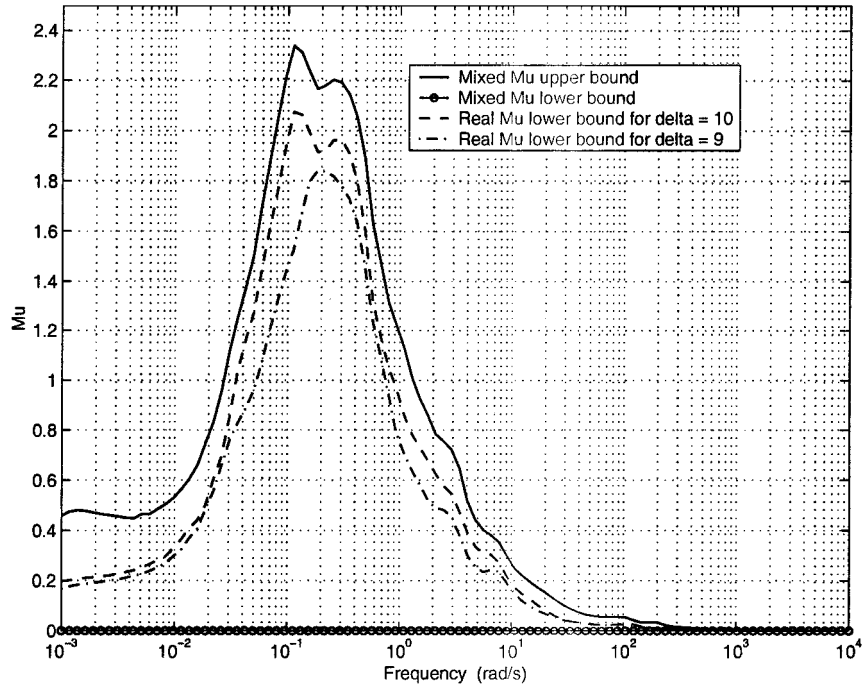

Fig. 6 Robustness of IFPC system to plant dynamics variations between 80 and 120 kn: comparison of μ lower bounds.

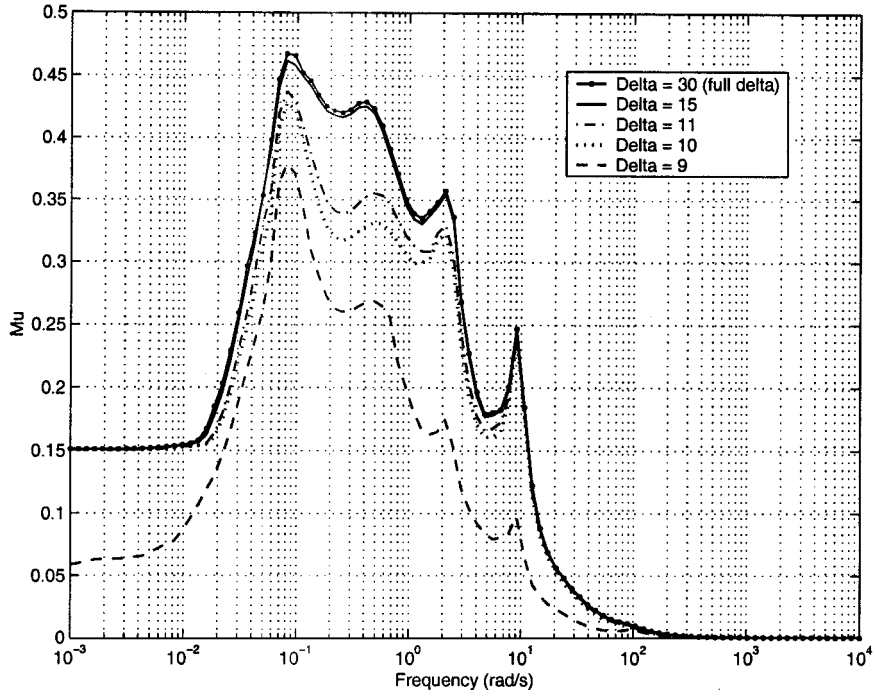

Fig. 7 Robustness of IFPC system to plant dynamics variations between 80 and 50 kn: μ upper bounds for different size Δ.

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

$$\min f(x) = \min_{\Delta_i \in \mathcal{R}, \omega \in \Omega} \bar{\sigma}(\Delta) \text{ subject to } |\det(I - \Delta M_{11}(j\omega)| \leq \epsilon \tag{13}$$

where ϵ is a user-defined parameter that can be used to tradeoff computation time vs tightness of the resulting lower bound. To locate the minimizing $\boldsymbol{x}$, it is common for optimization algorithms to consider the first two terms of the Taylor approximation of f at a candidate $\boldsymbol{x}$. This recasts the minimization as the well-known quadratic programming problem:

$$f(x) \approx \tfrac{1}{2} x^T H x + x^T g \tag{14}$$

where H is the symmetric matrix of the 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.[27,28] In this paper, commercially available optimization software[29,30] has been used to solve Eq. (14). Here, $f(\boldsymbol{x})$ is minimized on a two-dimensional subspace $S = \langle \boldsymbol{s}_1, s_2 \rangle$, where $\boldsymbol{s}_1$ is a vector in the steepest descent direction g so that the algorithm demonstrates fast convergence, whereas s_2 considers the approximate Newton direction, that is, $H \cdot s_2 = -g$, in an attempt to locate a global minimum. Note that f is minimized using a line search on S:

$$\min_{\alpha \in [0,\ldots,1]} f\colon \langle \alpha \boldsymbol{s}_1, (1-\alpha) s_2 \rangle \tag{15}$$

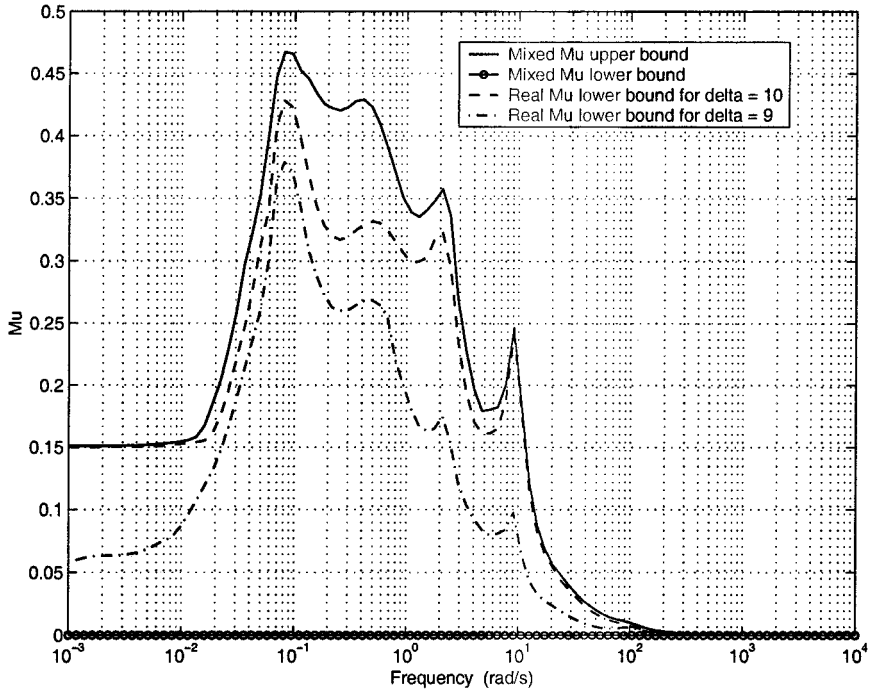

Fig. 8 Robustness of IFPC system to plant dynamics variations between 80 and 50 kn: comparison of μ lower bounds.

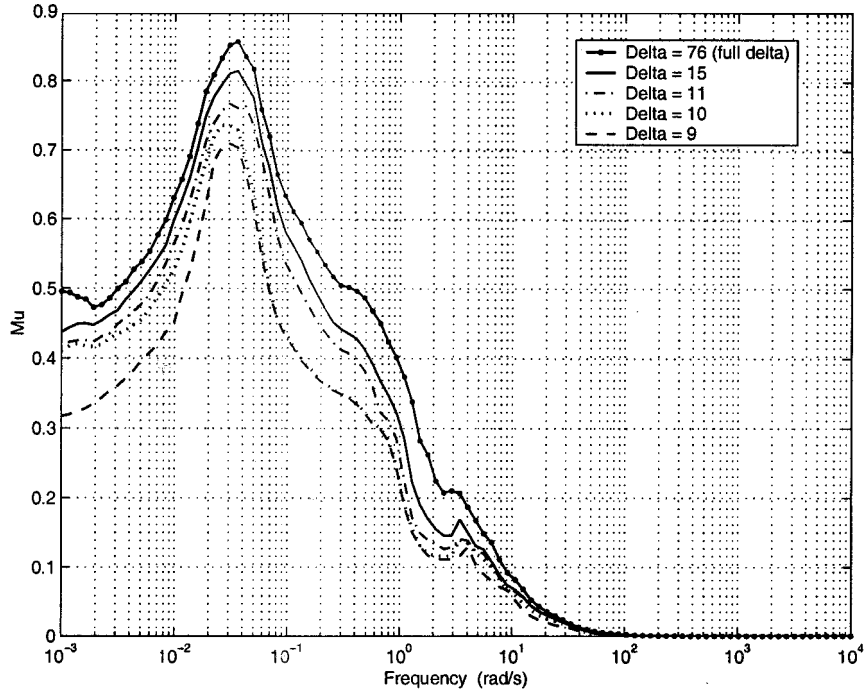

Fig. 9 Robustness of IFPC system to variations in mass and *x*c.g.: μ upper bounds for different size Δ.

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 destabilizing 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 destabilizing Δ will be computed. This offers a significant improvement over existing μ lower bound algorithms,[6] where convergence does not occur with real uncertainty. Best results are obtained from the optimization algorithm by initially using the maximum right singular vector associated with a SVD on M. Subsequently, the destabilizing Δ is used as the starting vector at the next frequency step for another iteration of the optimization if necessary. Because the search for a worst-case destabilizing Δ is nonconvex, local minima can occur. When a local minimum has been detected, the algorithm is then restarted using vectors from the SVD on M.

Application of the described procedure to the three robustness analysis problems described in Sec. II with $\epsilon = 1e-5$, produced the results shown in Figs. 11–13. From Figs. 11–13, we see that a tight lower bound is produced for all three problems over the frequency range of interest. Note again that the standard mixed μ lower bound software fails to converge for all three of these problems.

Computing times for the numerical optimizations involved in the described approach are a function of the problem size, ϵ, and the number of optimization restarts and also depend on internal algorithm settings. Table 2 shows average computing times per

frequency for each of the robustness analysis problems described in Sec. II assuming only one restart per frequency. The optimizations were carried out on a SPARC 10 operating at 120 MHz. Table 2 illustrates that problem size impacts directly on computing time. However the rate of increase is not as severe as that observed with existing lower bound algorithms for μ (Ref. 16).

Computational experience using exponential time lower bound algorithms for real μ on smaller systems suggests that existing upper bound software[6] yields tight upper bounds. Therefore, it is reasonable to conclude that where a lower bound on μ does not track the upper bound quite closely then the optimization algorithm is locating a local minimum. A number of additions to the basic algorithm have been considered in an attempt to improve the lower bound on μ at the expense of some increase in computing times:

1) Reduce ϵ in Eq. (13). For larger systems particularly, it is necessary to relax the ϵ parameter so that ΔM is not so close to being singular. A smaller ϵ reduces the gap between the upper and lower bounds on μ as the set of local minima satisfying Eq. (13) will now be smaller. Experience has shown that $\epsilon = 1e-5$ is a reasonable choice for this problem.

2) Consider frequencies that are close to a local minimum. It is a straightforward extra step to consider a region very close to a

Table 2 Computing times per frequency for numerical optimizations

Uncertainty description	Size of Δ	Average time, s
Variations in m and (xc.g.)	76	165
80–50 kn variations	30	8.2
80–120 kn variations	30	8.2

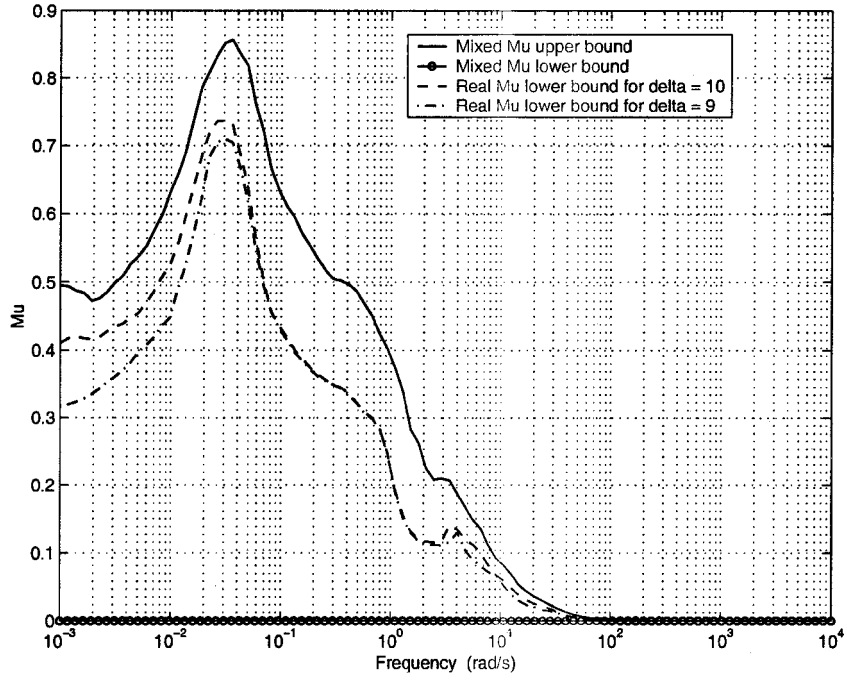

Fig. 10 Robustness of IFPC system to variations in mass and xc.g.: comparison of μ lower bounds.

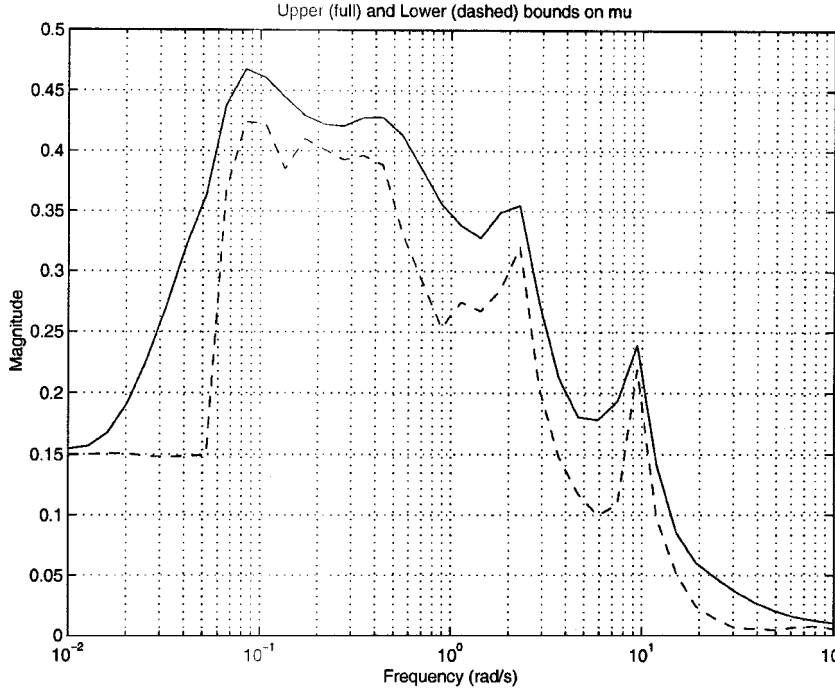

Fig. 11 Robustness of IFPC system to plant dynamics variations between 80 and 50 kn: μ lower bound from constrained optimization.

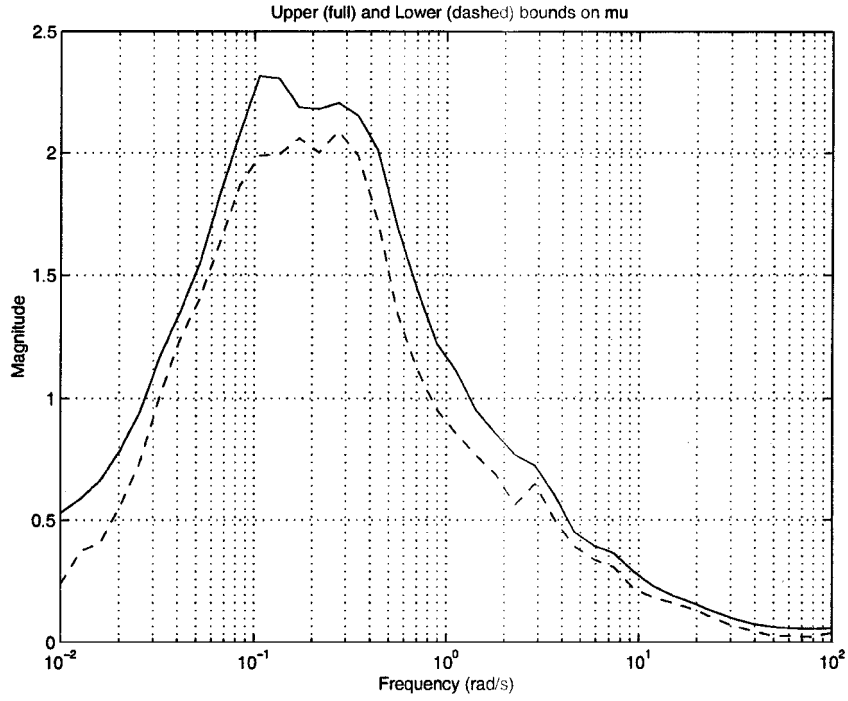

Fig. 12 Robustness of IFPC system to plant dynamics variations between 80 and 120 kn: μ lower bound from constrained optimization.

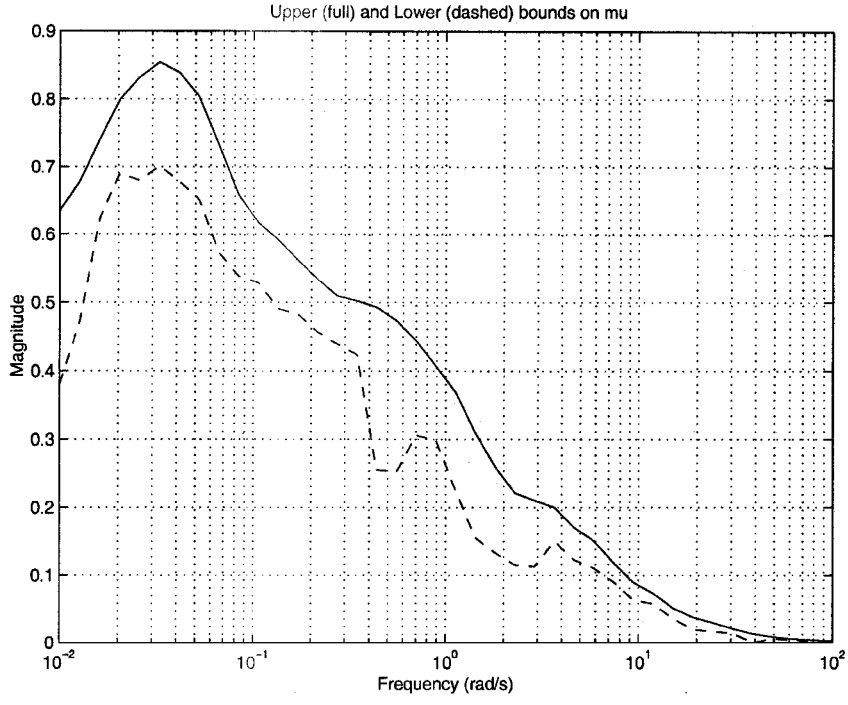

Fig. 13 Robustness of IFPC system to variations in mass and *x*c.g.: μ lower bound from constrained optimization.

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. Although useful at times, computational experience has yielded no real trend in problem matrices that may provide poor local minima.

3) Exploit good starting Δ. A double sweep over the frequency range under consideration has also been considered, which uses a list of the best Δ, that is, those that close the gap between the upper and lower bound on μ to within a certain user-defined level. These Δ are used for a second pass over frequencies where the lower bound is not so good. Computational experience indicates that the likely benefit from this second sweep is quite modest, thus indicating that the worst-case Δ for this problem varies quite significantly as a function of frequency. Figure 14 shows the small improvement that is achieved despite a near doubling of computing times involved.

4) Consider different starting vectors for the optimization. The lower bound on μ determined with a purely random starting vector is compared with an SVD vector restart in Fig. 15. It is clear that using the SVD restart for the optimization results in a consistently better μ lower bound. Computational experience suggests that an SVD restart offers the best lower bound on μ via constrained optimization.

5) Edit internal optimization settings. Table 3 shows the limited improvement that is obtained by an order of magnitude increase in the number of function evaluations that are allowed during the minimization of Eq. (13) at a particular frequency. At this point the algorithm has begun to stitch, that is, the improvement per extra function evaluation becomes vanishingly small. Further relaxation of the algorithm exit criteria for this problem exacerbates the stitching problem to such an extent that the code can break down.

Clearly, the existence and location of local minima through the use of nonlinear optimization-based approaches of the type described here means that the computation of actual rather than approximate worst-case real destabilising Δ remains a problem. Interior trust region methods[31] have been proposed as one way of improving on line search methods so as to further close the gap between the upper and lower bounds on real μ. These ideas are currently being applied by the authors to improve the global convergence properties of the approach described earlier.

Table 3 Effect of increasing maximum number of function evaluations

ω	Quick μ computing time, s	Slow μ computing time, s	%, improvement
0.01	0.3807	0.3807	0
0.0127	0.4719	0.5154	9
0.0160	0.6241	0.6287	2
0.0203	0.6900	0.6900	0

V. Conclusions

New tools were presented for the computation of tight lower bounds on the structured singular value μ. The tools are designed to overcome current computational problems associated with generating lower bounds on μ for high-order plants subject to purely real parametric uncertainty. The first approach uses μ sensitivities to reduce the order of the real uncertainty matrix Δ, so that exponential time lower bound algorithms can be applied. The second approach formulates the search for a worst-case real destabilizing perturbation as a constrained nonlinear optimization problem. Both approaches were applied to the problem of analyzing the stability robustness properties of an IFPC system for an experimental V/STOL aircraft configuration. Currently available software tools for calculating lower bounds on real μ are shown to fail for this problem, whereas both new approaches deliver tight bounds over the frequency range of interest. The two approaches described in this paper are of course complementary: The tightest lower bound will be the maximum of that produced by each method at each frequency.

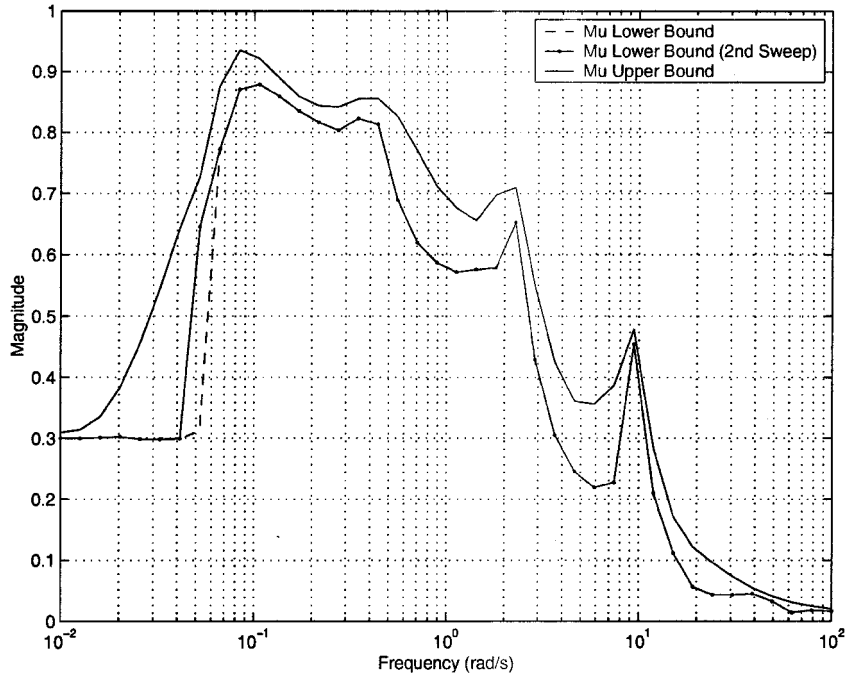

Fig. 14 Improvements in μ lower bound from constrained optimization using the second sweep technique.

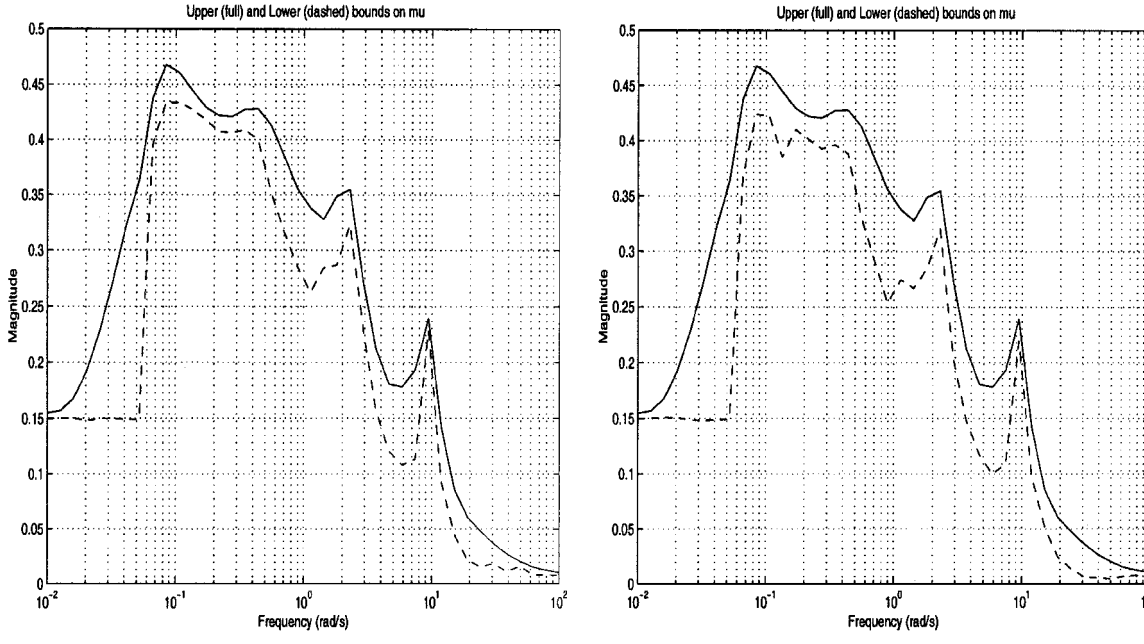

Fig. 15 Comparison of SVD (left) vs random (right) restart on μ lower bound obtained from constrained optimization.

References

[1]Zames, G., "On the Input–Output Stability of Time-Varying Non-Linear Feedback Systems," Pts. 1 and 2, *IEEE Transactions on Automatic Control*, Vol. AC-11, No. 2/3, 1966, pp. 228–238, 465–476.

[2]Doyle, J., "Analysis of Feedback Systems with Structured Uncertainty," *IEE Proceedings on Control Theory and Applications, Part D*, Vol. 129, No. 6, 1982, pp. 242–250.

[3]Packard, A., and Doyle, J., "The Complex Structured Singular Value," *Automatica*, Vol. 29, No. 1, 1993, pp. 71–109.

[4]Braatz, R., Young, P., Doyle, J., and Morari, M., "Computational Complexity of μ Calculation," *IEEE Transactions on Automatic Control*, Vol. AC-39, No. 5, 1994, pp. 1000–1002.

[5]Ferreres, G., *A Practical Approach to Robustness Analysis with Aeronautical Applications*, Kluwer Academic, New York, 1999, Chap. 1.

[6]Balas, G. J., Doyle, J. C., Glover, K., Packard, A., and Smith, R., *μ-Analysis and Synthesis Toolbox*, MathWorks, Natick, MA, 1995, Chap. 4.

[7]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*, Vol. 19, No. 3, 1991, pp. 271–280.

[8]Fan, M., Tits, A., and Doyle, J., "Robustness in the Presence of Mixed Parametric Uncertainty and Unmodelled Dynamics," *IEEE Transactions on Automatic Control*, Vol. AC-36, No. 1, 1991, pp. 25–38.

[9]Safonov, M., and Lee, P., "A Multiplier Method for Computing Real Multivariable Stability Margins," *Proceedings of the IFAC World Congress*, Vol. 1, International Federation of Automatic Control, Oxford, 1993, pp. 275–278.

[10]Packard, A., and Doyle, J., "A Power Method for the Structured Singular Value," *Proceedings of the American Control Conference*, Vol. 2, Inst. of Electrical and Electronics Engineers, New York, 1988, pp. 1213–1218.

[11]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, No. 6, 1995, pp. 573–590.

[12]Gately, S. L., Bates, D. G., Hayes, M. J., and Postlethwaite, I., "Robustness Analysis of Integrated Flight and Propulsion Control Systems for V/STOL Aircraft," European Control Conf., 2001 (submitted for presentation).

[13]Packard, A., and Pandey, P., "Continuity Properties of the Real/Complex Structured Singular Value," *IEEE Transactions on Automatic Control*, Vol. 38, No. 3, 1993, pp. 415–428.

[14]Barmish, B. R., Khargonekar, P. P., and Shi, Z., "Robustness Margin Need Not Be a Continuous Function of the Problem Data," *Systems and Control Letters*, Vol. 15, No. 1, 1990, pp. 91–98.

[15]Jones, R., "Structured Singular Value Analysis for Real Parameter Variations," *Proceedings of the AIAA Conference on Guidance, Navigation, and Control*, Vol. 2, AIAA, Washington, DC, 1987, pp. 1424–1432.

[16]Dailey, R., "A New Algorithm for the Real Structured Singular Value," *Proceedings of the American Control Conference*, Vol. 3, Inst. of Electrical and Electronics Engineers, New York, 1990, pp. 3036–3040.

[17]DeGaston, R., and Safanov, M., "Exact Calculation of the Multiloop Stability Margin," *IEEE Transactions on Automatic Control*, Vol. AC-33, No. 2, 1988, pp. 156–171.

[18]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, Asian Control Council, Tokyo, 1997, pp. 847–850.

[19]Hyde, R. A., *$\mathcal{H}_\infty$ Aerospace Control Design—A VSTOL Flight Application*, Springer-Verlag, New York, 1995, Chap. 1.

[20]Hayes, M. J., Bates, D. G., Gately, S. L., and Postlethwaite, I., "Stability Robustness Analysis of an Integrated Flight and Propulsion Controller Using the Structured Singular Value," *Proceedings of the 38th IEEE Conference on Decision and Control*, Vol. 5, Inst. of Electrical and Electronics Engineers, New York, 1999, pp. 4529, 4530.

[21]Bates, D. G., Gatley, S. L., Postlethwaite, I., and Berry, A. J., "Design and Piloted Simulation of a Robust Integrated Flight and Propulsion Controller," *Journal of Guidance, Navigation, and Control*, Vol. 23, No. 2, 2000, pp. 269–277.

[22]McFarlane, D., and Glover, K., "A Loop Shaping Design Procedure Using $\mathcal{H}_\infty$ Synthesis," *IEEE Transactions on Automatic Control*, Vol. AC-36, No. 6, 1992, pp. 759–769.

[23]Varga, A., Looye, G., Moorman, D., and Grubel, G., "Automated Generation of LFT-Based Parametric Uncertainty Descriptions from Generic Aircraft Models," GARTUER Technical-Publ., TP-088-36, GARTEUR, 1997.

[24]Morton, B., and McAfoos, R., "A μ-Test for Real Parameter Variations," *Proceedings of the American Control Conference*, Vol. 1, Inst. of Electrical and Electronics Engineers, New York, 1985, pp. 135–138.

[25]Braatz, R. D., and Morari, M., "μ-Sensitivities as an Aid for Robust Identification," *Proceedings of the American Control Conference*, Vol. 1, Inst. of Electrical and Electronics Engineers, New York, 1991, pp. 231–236.

[26]Douglas, J., and Athans, M., "The Calculation of μ-Sensitivities," *Proceedings of the American Control Conference*, Vol. 1, Inst. of Electrical and Electronics Engineers, New York, 1995, pp. 437–441.

[27]Fletcher, R., *Practical Methods of Optimization*, 2nd ed., Wiley, New York, 1987, Chap. 6.

[28]Gill, P. E., Murray, W., and Wright, M. H., *Practical Optimization*, Academic, New York, 1981, Chap. 1.

[29]Branch, M. A., and Grace, A., *MATLAB Optimization Toolbox User's Guide*, MathWorks, Natick, MA, 1996, Chap. 3.

[30]Coleman, T. F., and Li, Y., "A Reflective Newton Method for Minimizing a Quadratic Function Subject to Bounds on Some of the Variables," *SIAM Journal on Optimization*, Vol. 6, No. 7, 1996, pp. 1040–1058.

[31]Coleman, T. F., and Li, Y., "An Interior, Trust Region Approach for Nonlinear Minimisation Subject to Bounds," *SIAM Journal on Optimization*, Vol. 6, No. 4, 1996, pp. 418–445.