



Brief Paper

A fast algorithm for the computation of an upper bound on the μ -norm[☆]

Craig T. Lawrence^{a,1}, André L. Tits^{b,*}, Paul Van Dooren^{c,2}

^aALPHATECH, Inc., 2101 Wilson Blvd., Arlington, VA 22201, USA

^bDepartment of Electrical and Computer Engineering and Institute for Systems Research, University of Maryland, College Park, MD 20742, USA

^cUniversité Catholique de Louvain, CESAME, Bâtiment Euler, B-1348 Louvain-la-Neuve, Belgium

Received 19 May 1998; revised 8 December 1998; received in final form 24 June 1999

Abstract

A fast algorithm for the computation of the optimally frequency-dependent scaled \mathcal{H}_∞ -norm of a finite-dimensional LTI system is presented. It is well known that this quantity is an upper bound to the “ μ -norm”; furthermore, it was recently shown to play a special role in the context of slowly time-varying uncertainty. Numerical experimentation suggests that the algorithm generally converges quadratically. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Numerical algorithms; Numerical methods; Robust control; Robust stability; Structured singular value; Uncertain linear systems; Uncertainty

1. Introduction

In the context of robust control analysis and synthesis a quantity of great interest is the *structured singular-value norm*, or μ -norm, of the system. Consider a feedback connection of a continuous-time system with real coefficients as in Fig. 1. Let $P(s) = C(sI - A)^{-1}B$ be an $m \times m$ stable transfer function matrix and let $\Delta(s)$ be a *structured* perturbation constrained to lie in the set

$$\mathcal{R}(\Delta) \triangleq \{ \Delta \in \mathcal{R}_S : \Delta(s_0) \in \Delta \ \forall s_0 \in \bar{\mathbb{C}}_+ \}$$

where \mathcal{R}_S is the set of all real-rational, proper, stable, $m \times m$ transfer matrices and the uncertainty set

$$\Delta \triangleq \{ \text{diag}[\Delta_1, \dots, \Delta_F] : \Delta_j \in \mathbb{C}^{m_j \times m_j} \}$$

is defined for integers m_1, \dots, m_F . Let $\sigma_1(\cdot)$ denote the maximum singular value of its matrix argument and define the complex *structured singular value* for a constant matrix $M \in \mathbb{C}^{m \times m}$ as (see Zhou, Doyle & Glover, 1995 for a complete discussion of the structured singular value)

$$\mu_\Delta(M) \triangleq \frac{1}{\min\{\sigma_1(\Delta) : \Delta \in \Delta, \det(I - M\Delta) = 0\}}$$

unless $\det(I - M\Delta) \neq 0$ for all $\Delta \in \Delta$, in which case $\mu_\Delta(M) \triangleq 0$. Finally, define the “ μ -norm”³ as

$$\|P\|_\Delta \triangleq \sup_{\omega \in \mathbb{R}} \mu_\Delta(P(j\omega)).$$

It has been shown that the computation of $\mu_\Delta(M)$ is NP-hard (see Toker & Özbay, 1995), in consequence no efficient algorithms are likely to exist for its computation.

³ Actually, it is not a norm.

[☆]This paper has appeared, in slightly different form, in the Proceedings of the 1996 IFAC World Congress (Lawrence, Tits & Van Dorren, 1996). This paper was recommended for publication in revised form by Associate Editor P-O. Gutman under the direction of Editor T. Basar.

* Corresponding author. Tel.: + 1-301-405-3669; fax: + 1-301-405-6707.

E-mail address: andre@eng.umd.edu (A.L. Tits)

¹ Supported in part by NSF's Engineering Research Center Program, under grant NSFD-CDR-88-03012.

² Supported by the Belgian Programme on Interuniversity Poles of Attraction, initiated by the Belgian State, Prime Minister's Office for Science, Technology and Culture. The scientific responsibility rests with the authors.




Fig. 1. Standard framework for robustness analysis.

In practice, a standard upper bound is used in its place. This upper bound is computed as follows. Define

$$\mathcal{D} \triangleq \{\text{diag}[d_1 I_{m_1}, \dots, d_{F-1} I_{m_{F-1}}, I_{m_F}]: 0 < d_j \in \mathbb{R}\}.$$

Then, for a constant complex matrix M , an upper bound for $\mu_\Delta(M)$ is

$$\hat{\mu}_\Delta(M) \triangleq \inf_{D \in \mathcal{D}} \sigma_1(DMD^{-1}).$$

Substituting this into the expression for $\|P\|_\Delta$, an upper bound on the μ -norm of the system $P(s)$ is obtained as the optimally frequency-dependent scaled \mathcal{H}_∞ -norm⁴

$$\|P\|_{\hat{\mu}} \triangleq \sup_{\omega \in \mathbb{R}} \hat{\mu}_\Delta(P(j\omega)).$$

It is well-known that if $P(s)$ is stable, then $\|P\|_\Delta < 1$ is necessary and sufficient (thus $\|P\|_{\hat{\mu}} < 1$ is sufficient) for uniform robust stability of the $P - \Delta$ loop (Fig. 1) for any linear time-invariant structured $\Delta(s)$ of L_2 -gain no greater than one (see, e.g., Corollary 3 in Tits & Balakrishnan, 1998). Further, it has recently been shown that $\|P\|_{\hat{\mu}} < 1$ is necessary and sufficient for uniform robust stability of the $P - \Delta$ loop for any linear, arbitrarily slowly time-varying structured Δ of L_2 -gain no greater than one (Pooila & Tikku, 1995).

Algorithms for the efficient computation of $\hat{\mu}_\Delta(P(j\omega))$, for given ω , have long been available. In fact, given ω ,

$$\hat{\mu}_\Delta(P(j\omega))^2 = \inf_{\alpha \geq 0, D \in \mathcal{D}} \{\alpha: P(j\omega)DP(j\omega)^* - \alpha D < 0\}, \quad (1)$$

which is a linear matrix inequality (LMI) problem. Efficient algorithms exist for obtaining global solutions to such problems, e.g., Boyd, El Ghaoui, Feron and Balakrishnan (1994). Note that minimizers (or approximate minimizers) D for $\sigma_1(DMD^{-1})$ are related to the minimizers (or approximate minimizers) \tilde{D} for (1) by $D = \tilde{D}^{1/2}$. $\|P\|_{\hat{\mu}}$ is usually computed via a “frequency sweep”, i.e. choose a set of frequencies $\{\omega_1, \dots, \omega_N\}$ and use the approximation

$$\|P\|_{\hat{\mu}} \approx \max\{\hat{\mu}_\Delta(P(j\omega_1)), \dots, \hat{\mu}_\Delta(P(j\omega_N))\}.$$

The drawbacks of this approach are obvious. First, a large number of $\hat{\mu}_\Delta(\cdot)$ computations are required.

Second, an upper bound on $\|P\|_\Delta$ is not necessarily obtained. Finally, the result can be arbitrarily bad, i.e. it is difficult to bound the error.

Another approach that has been used to compute an upper bound to the μ -norm is based on the Main Loop Theorem (see, e.g., Packard & Doyle, 1993) and the extension to μ of the Maximum Modulus Theorem (see, e.g., top of p. 1201 in Tits & Fan, 1995). In the discrete-time case, given $P(z) = C(zI - A)^{-1}B + D$, with A stable, and a positive scalar γ , it can be seen that the μ -norm of P is less than γ if and only if $\mu_{\Delta_p}(M(\gamma)) < 1$, where

$$M(\gamma) \triangleq \begin{bmatrix} A & B \\ \frac{1}{\gamma}C & \frac{1}{\gamma}D \end{bmatrix}$$

and Δ_p is an “augmented” uncertainty structure given by

$$\Delta_p \triangleq \{\text{diag}[\delta I_n, \Delta]: \delta \in \mathbb{C}, \Delta \in \Lambda\}.$$

The original idea is due to Doyle and Packard (1987). The continuous-time case can be reduced to the discrete-time case by means of a bilinear transformation (see Section 10.2 in Zhou et al., 1995 for details). Repeated evaluation of the upper bound $\hat{\mu}_{\Delta_p}(\cdot)$ according to a bisection search over γ yields an upper bound on the μ -norm of P . On the down side, note that this upper bound is generally less tight than that obtained by gridding, because the augmented uncertainty structure involves an additional block which, to make things worse, is of the “repeated” type. The bisection search can be done away with, as shown by Ferreres and Fromion (1997), by invoking the “skewed μ ” proposed and studied by Fan and Tits (1992).⁵ Computation then entails the solution of a single LMI constrained quasi-convex generalized eigenvalue minimization problem (GEVP; see Boyd et al., 1994).⁶ The approach proposed in this paper improves on the scheme of Ferreres and Fromion (1997) in that it computes the same (tighter) upper bound as the gridding approach. As far as computational cost is concerned, the tradeoff is that of a sequence of LMIs (μ upper bound computations with respect to the original, non-augmented uncertainty structure) versus a single GEVP with a larger number of variables and constraints. Due to the “repeated” block, the complexity of the latter is strongly affected by the dimension of the state space.

Recently, an efficient algorithm has been proposed (Boyd & Balakrishnan, 1990; Bruinsma & Steinbuch,

⁵ In Ferreres and Fromion (1997) the continuous-time case is dealt with directly, by limiting the frequency range to a compact interval and using a “repeated real scalar” block in the augmentation. This idea is due to Sideris (1992).

⁶ To make the optimization problem of Fan and Tits (1992) into a standard GEVP, a symmetric $n \times n$ matrix of slack variables must be introduced; see p. 8–33 in Gahinet et al. (1995).

⁴ Again, this is not a norm.

1990) for the computation of the \mathcal{H}_∞ -norm

$$\|P\|_\infty = \sup_{\omega \in \mathbb{R}} \sigma_1(P(j\omega)). \quad (2)$$

This algorithm makes use of the well-known fact that a given scalar $\xi > 0$ is a singular value of $C(j\omega - A)^{-1}B$ if and only if $j\omega$ is an eigenvalue of the related Hamiltonian matrix

$$H(\xi, A, B, C) \triangleq \begin{bmatrix} A & \xi^{-1}BB^* \\ -\xi^{-1}C^*C & -A^* \end{bmatrix}.$$

Given any $\xi \in (0, \|P\|_\infty)$, a set of frequency intervals may thus be computed where maximizers for (2) are known to lie. At step k , ξ is selected as $\sigma_1(P(j\omega_{k-1}))$ and ω_k is chosen as the mid-point of the largest among the intervals thus determined. A quadratic rate of convergence ensues.

In this paper, using the idea just outlined as a stepping stone, an algorithm is constructed for the fast computation of $\|P\|_{\hat{\mu}}$. (A similar algorithm can be used for the computation of the real stability radius; see Sreedhar, Van Dooren & Tits, 1996.)

2. Key ideas

Let

$$\hat{\mu}(\omega) \triangleq \inf_{D \in \mathcal{D}} \sigma_1(DP(j\omega)D^{-1}).$$

The goal of the algorithm is to maximize $\hat{\mu}(\omega)$ over $\omega \in \mathbb{R}$. Define

$$\hat{\mu}^* \triangleq \sup_{\omega \in \mathbb{R}} \hat{\mu}(\omega)$$

and

$$\omega^* \triangleq \operatorname{argmax}_{\omega \in \mathbb{R}} \hat{\mu}(\omega)$$

assuming for now that such a unique maximizer exists. Finally, for a fixed $D \in \mathcal{D}$, define the curve

$$v_D(\omega) \triangleq \sigma_1(DP(j\omega)D^{-1}).$$

While reading the following, it may be helpful to refer to Fig. 2. The main idea of the algorithm is as follows (more details are given in Section 3). At iteration k , suppose ξ_k is the best-known lower bound to $\hat{\mu}^*$ thus far, and let ω_k be the current trial frequency. Suppose further that ω^* is known to lie in a certain open set Ω_k .⁷ Compute

$$D_{k+1} = \operatorname{argmin}_{D \in \mathcal{D}} \sigma_1(DP(j\omega_k)D^{-1})$$




Fig. 2. An example of the various definitions.

(assume for simplicity that the minimum is achieved), and let

$$\xi'_{k+1} \triangleq \hat{\mu}(\omega_k) = v_{D_{k+1}}(\omega_k).$$

If $\xi'_{k+1} > \xi_k$, take $\xi_{k+1} = \xi'_{k+1}$ as the new estimate of $\hat{\mu}^*$ (as is the case in Fig. 2), otherwise keep the old estimate (i.e. $\xi_{k+1} = \xi_k$). Note that

$$\hat{\mu}(\omega) \leq v_{D_{k+1}}(\omega) \quad \forall \omega \in \mathbb{R}, \quad (3)$$

with equality at $\omega = \omega_k$, since by definition $\hat{\mu}(\omega)$ is the lower envelope of the family of curves $\{v_D(\omega) : D \in \mathcal{D}\}$. Now note that $v_{D_{k+1}}(\omega)$ is the maximum singular value curve for the transfer function $D_{k+1}P(s)D_{k+1}^{-1}$ at $s = j\omega$. Thus, as suggested in Section 1, the ideas developed for the computation of the \mathcal{H}_∞ -norm may be applied to locate the open set of frequencies, call it Ω'_{k+1} , in which $v_{D_{k+1}}(\omega) > \xi_{k+1}$ (in Fig. 2, $\Omega'_{k+1} \approx (1, 5.1 \times 10^4)$). In view of (3) and since ξ_{k+1} is a lower bound to $\hat{\mu}^*$, it is clear that (i) if $\Omega'_{k+1} = \emptyset$, then $\hat{\mu}^* = \xi_{k+1}$ and (ii) if $\Omega'_{k+1} \neq \emptyset$, then $\omega^* \in \Omega'_{k+1}$. In the latter case it follows that $\omega^* \in \Omega_{k+1} \triangleq \Omega_k \cap \Omega'_{k+1}$ (in Fig. 2, it is assumed that $\Omega_k \supset \Omega'_{k+1}$, thus $\Omega_{k+1} = \Omega'_{k+1}$). Now, the next frequency ω_{k+1} must be selected. Several possibilities come to mind. For example, ω_{k+1} could simply be chosen as the mid-point of the largest interval contained in Ω_{k+1} . Since, for all k , $\Omega_{k+1} \not\ni \omega_k$, the size of the largest interval in Ω_k would go to zero as $k \rightarrow \infty$ and convergence of ξ_k to $\hat{\mu}^*$ would ensue (see Section 4 below). This, however, results in a relatively slow rate of convergence. More sophisticated alternatives are considered in the next section.

A rough outline of the algorithm described above is as follows.

Algorithm

Data: $P(s) = C(sI - A)^{-1}B$, Δ , $\varepsilon > 0$.

Initialization: $k = 0$, pick $\omega_0 \geq 0$, $\xi_0 = 0$, $\Omega_0 = (0, \infty)$.

⁷ Below, Ω_k is restricted to be a subset of $(0, \infty)$ for all k . Thus, in the case that $\omega^* = 0$, clearly ω^* will not be in Ω_k . It is readily checked that this will not affect any of the results.

- Step 1:** Compute D_{k+1}, ζ_{k+1} .
- Step 2:** Compute Ω_{k+1} .
- Step 3:** Compute ω_{k+1} .
- Step 4:** $k \leftarrow k + 1$. If an appropriate stopping criterion is satisfied, stop. Otherwise, go back to Step 1.

3. Details of the algorithm

The computation of Ω'_{k+1} is now considered in greater detail. The pure imaginary eigenvalues of the Hamiltonian matrix

$$H_k = H(\zeta_{k+1}, A, BD_{k+1}^{-1}, D_{k+1}C)$$

tell us the frequencies at which *one* of the singular values of $D_{k+1}P(j\omega)D_{k+1}^{-1}$ takes the value ζ_{k+1} . That is, if $j\tilde{\omega}$ is an imaginary eigenvalue of H_k , then

$$\sigma_r(D_{k+1}P(j\tilde{\omega})D_{k+1}^{-1}) = \zeta_{k+1}$$

for some $r \in \{1, \dots, m\}$ (see Fig. 3). Of course, the frequencies of interest are those for which $r = 1$. Suppose the Schur decomposition of the Hamiltonian matrix H_k has been performed and the $\tilde{\omega}$'s have been identified. Let $\tilde{\omega}_1, \dots, \tilde{\omega}_q$ be the nonnegative $\tilde{\omega}$'s, sorted in increasing order, and let $\lambda_\ell = j\tilde{\omega}_\ell, \ell = 1, \dots, q$. In Boyd and Balakrishnan (1990), the authors suggest determining those frequencies corresponding to the maximum singular value by eliminating the frequencies for which

$$\zeta_{k+1}^2 I - D_{k+1}^{-1}P^*(j\omega)D_{k+1}^2P(j\omega)D_{k+1}^{-1}$$

is not positive semi-definite. This is a rather costly computation, though. An alternative scheme is as follows. The pure imaginary simple eigenvalues of the Hamiltonian may be differentiated with respect to the scalar parameter ζ via the expression

$$\frac{\partial \lambda_\ell}{\partial \zeta}(\zeta_{k+1}) = \frac{u_\ell^*(\partial H / \partial \zeta)(\zeta_{k+1}, A, BD_{k+1}^{-1}, D_{k+1}C)v_\ell}{u_\ell^*v_\ell},$$

$$\ell = 1, \dots, q,$$

(see, e.g., Theorem IV.2.3 in Stewart and Sun (1990)) where v_ℓ and u_ℓ^* are respectively right and left eigenvectors of H_k corresponding to the imaginary eigenvalue λ_ℓ ; these few vectors can be computed from the Schur decomposition at little extra cost, e.g., using ‘‘inverse iteration’’. Since $\lambda_\ell = j\tilde{\omega}_\ell$,

$$\frac{\partial \zeta}{\partial \omega}(\tilde{\omega}_\ell) = \left(-j \frac{\partial \lambda_\ell}{\partial \zeta}(\zeta_{k+1}) \right)^{-1},$$

which is the slope of the corresponding singular-value curve for the matrix $D_{k+1}P(j\omega)D_{k+1}^{-1}$ at the frequency $\omega = \tilde{\omega}_\ell$ and is a real number for every simple eigenvalue on the imaginary axis. This is all the information needed to compute Ω'_{k+1} .




Fig. 3. Sample singular-value plot.

Indeed, define

$$s_\ell \triangleq \frac{\partial \zeta}{\partial \omega}(\tilde{\omega}_\ell), \quad \ell = 1, \dots, q.$$

For the example in Fig. 3, the sign of s_ℓ is marked at each intersection point. As $P(s)$ is strictly proper, $v_{D_{k+1}}(\omega) \rightarrow 0$ as $\omega \rightarrow \infty$. Thus, as long as $\zeta_{k+1} > 0$,⁸ $\tilde{\omega}_q$ must correspond to the largest singular value, i.e. $v_{D_{k+1}}(\tilde{\omega}_q) = \zeta_{k+1}$, and Ω_{k+1} will be bounded. Let $c_\ell, \ell = 1, \dots, q, c_\ell \in \{1, \dots, m\}$, denote the number of the singular-value curve corresponding to $\tilde{\omega}_\ell$. That is,

$$\sigma_{c_\ell}(D_{k+1}P(j\tilde{\omega}_\ell)D_{k+1}^{-1}) = \zeta_{k+1}.$$

Thus, $c_q = 1$. The rest of the c_ℓ 's may be assigned by sweeping to the left from $\tilde{\omega}_{q-1}$ and using the following rules:⁹

- If $s_{\ell+1} < 0$ and $s_\ell < 0$ then $c_\ell = c_{\ell+1} + 1$.
- If $s_{\ell+1} < 0$ and $s_\ell > 0$ then $c_\ell = c_{\ell+1}$.
- If $s_{\ell+1} > 0$ and $s_\ell < 0$ then $c_\ell = c_{\ell+1}$.
- If $s_{\ell+1} > 0$ and $s_\ell > 0$ then $c_\ell = c_{\ell+1} - 1$.

Now define

$$\{\hat{\omega}_1, \dots, \hat{\omega}_r\} \triangleq \{\tilde{\omega}_\ell: c_\ell = 1\},$$

ordered so that $\hat{\omega}_1 < \dots < \hat{\omega}_r$. Clearly, $r \leq q$. In the obvious manner define $\hat{s}_\ell, \ell = 1, \dots, r$ as the slopes of the maximum singular value curve computed at $\hat{\omega}_\ell$. If $\hat{s}_1 < 0$,

⁸ Certainly true if $\hat{\mu}(\omega_0) > 0$, which will hold in all but a few pathological cases.

⁹ The algorithm as stated does not handle the case when $s_\ell = 0$ or the eigenvalue curve is non-differentiable. This happens when multiple eigenvalues are found on the imaginary axis. A simple way to get around the difficulty in such cases is to increase the current level ζ_k slightly.

then

$$\Omega'_{k+1} = (0, \hat{\omega}_1) \cup (\hat{\omega}_2, \hat{\omega}_3) \cup \dots \cup (\hat{\omega}_{r-1}, \hat{\omega}_r).$$

If $\hat{s}_1 > 0$, then

$$\Omega'_{k+1} = (\hat{\omega}_1, \hat{\omega}_2) \cup \dots \cup (\hat{\omega}_{r-1}, \hat{\omega}_r).$$

To help clarify the ideas, consider the example in Fig. 3 where $q = 6$. Letting $c_6 = 1$, we begin the left sweep. As $s_6 < 0$ and $s_5 < 0$, $c_5 = c_6 + 1 = 2$. Then, since $s_5 < 0$ and $s_4 > 0$, $c_4 = c_5 = 2$. Next, since $s_4 > 0$ and $s_3 > 0$, $c_3 = c_4 - 1 = 1$. Proceeding in this manner, we find $\{c_1, \dots, c_6\} = \{1, 1, 1, 2, 2, 1\}$. Thus, $r = 4$ and $\{\hat{\omega}_1, \dots, \hat{\omega}_4\} = \{\hat{\omega}_1, \hat{\omega}_2, \hat{\omega}_3, \hat{\omega}_6\}$. Finally, since $\hat{s}_1 = s_1 > 0$, $\Omega'_{k+1} = (\hat{\omega}_1, \hat{\omega}_2) \cup (\hat{\omega}_3, \hat{\omega}_6)$.

As mentioned in Section 2, one way to select the next trial point ω_{k+1} is to pick the mid-point of the largest interval in Ω_{k+1} . A better approach, yielding faster convergence, is to use some form of interpolation scheme in the final iterations. In particular, suppose $(\omega_{k_1}, \zeta_{k_1+1})$ and $(\omega_{k_2}, \zeta_{k_2+1})$ are the two most recent trial points for which $\hat{\mu}(\omega_{k_i}) = \zeta_{k_i+1}$, $i = 1, 2$. Then, for $k = k_i$, $i = 1, 2$, $\zeta_{k+1} = \zeta_{k_i+1}$. Thus ω_k is one of the $\tilde{\omega}_\ell$'s for the $(k + 1)$ st iteration, say $\omega_k = \tilde{\omega}_{\ell_k}$, and the corresponding c_{ℓ_k} is 1, i.e.,

$$\hat{\mu}(\omega_k) = \sigma_1(D_{k+1}P(j\omega_k)D_{k+1}^{-1}) = \zeta_{k+1}.$$

Since the $\hat{\mu}(\cdot)$ curve lies entirely below the $\sigma_1(D_{k+1}P(j\cdot)D_{k+1}^{-1})$ curve, for all such k , whenever both curves are differentiable at ω_k ,

$$\frac{\partial \hat{\mu}}{\partial \omega}(\omega_k) = \frac{\partial}{\partial \omega} \sigma_1(D_{k+1}P(j\omega_k)D_{k+1}^{-1}) = \frac{\partial \zeta}{\partial \omega}(\tilde{\omega}_{\ell_k}) = s_{\ell_k}.$$

Using the available values of $\hat{\mu}$ and of its derivative at ω_{k_1} and ω_{k_2} , a cubic or rational (specifically, quadratic over linear) function (Barzilai & Ben-Tal, 1982) may be passed through the two points and ω_{k+1} may be taken as the maximum of the interpolating function, subject to the constraint $\omega_{k+1} \in \bar{\Omega}_{k+1}$ (closure of Ω_{k+1}).

Finally, a stopping criterion which guarantees that the algorithm terminates with an estimate that is within prescribed $\varepsilon > 0$ of $\|P\|_{\hat{\mu}}$ is as follows. At iteration k , given ζ_k , D_k , and Ω_k , define the set

$$S_k \triangleq \{\omega: v_{D_k}(\omega) > \zeta_k + \varepsilon\}.$$

Like Ω'_{k+1} , S_k may be computed using an appropriate Hamiltonian eigen-decomposition. Since $\omega^* \in \Omega_k$, if $\Omega_k \cap S_k = \emptyset$, then $\omega^* \notin S_k$, hence

$$\zeta_k < \|P\|_{\hat{\mu}} \leq \zeta_k + \varepsilon.$$

4. Convergence

Theorems 1 and 2 below are proved for the case when the mid-point rule is used throughout to select the next trial point ω_{k+1} .

Assumption 1. The choice of ω_0 is such that

$$\hat{\mu}(\omega_0) > 0.$$

Assumption 1 is necessary in order for the algorithm to be well-defined. Typically, $\omega_0 = 0$ will satisfy this assumption; if not, a random search will provide a suitable ω_0 as long as $P(s)$ is not identically 0. Under this assumption, Ω_1 is bounded and, in view of the fact that $\omega_k \notin \Omega_{k+1}$, a simple induction argument shows that the length of the largest interval in Ω_k tends to zero as $k \rightarrow \infty$.

Theorem 1. The sequence $\{\zeta_k\}_{k \in \mathbb{N}}$ converges to $\hat{\mu}^*$ as $k \rightarrow \infty$.

Proof. By contradiction. Suppose $\zeta_k \rightarrow \hat{\mu}^*$. Since the sequence $\{\zeta_k\}_{k \in \mathbb{N}}$ is monotone non-decreasing, and since $\zeta_k \leq \hat{\mu}^*$ for all k , $\zeta_k \rightarrow \zeta^*$ for some $\zeta^* < \hat{\mu}^*$. Thus, $E \triangleq \{\omega: \hat{\mu}(\omega) > \zeta^*\}$ is not empty and, clearly, $E \subset \Omega_k$ for all k . Since $\hat{\mu}(\cdot)$ is continuous (see Bercovici, Foias & Tannenbaum, 1990), E contains a non-trivial interval S . This contradicts the fact that the size of the largest interval in Ω_k goes to 0.

For this preliminary convergence analysis, consider now instead the problem of maximizing $\tilde{\mu}(\omega)$, where

$$\tilde{\mu}(\omega) \triangleq \min\{\sigma_1(DP(j\omega)D^{-1}): D \in \mathcal{D}, \underline{b}I \leq D^2 \leq \bar{b}I\} \quad (4)$$

with $\bar{b} \geq \underline{b} > 0$ arbitrary prescribed numbers. Note that $\tilde{\mu}(\omega)$ is an upper bound to $\hat{\mu}(\omega)$, that it can be efficiently computed, and that it can be made arbitrarily close to $\hat{\mu}(\omega)$, uniformly over $\omega \in \mathbb{R}$, by selecting $\underline{b} > 0$ small enough and \bar{b} large enough (Lee & Tits, 1993). Let

$$\tilde{\mu}^* \triangleq \sup_{\omega \in \mathbb{R}} \tilde{\mu}(\omega).$$

It is not difficult to show that Theorem 1 still holds when $\hat{\mu}(\omega)$ is replaced with $\tilde{\mu}(\omega)$ and $\hat{\mu}^*$ is replaced with $\tilde{\mu}^*$.

Theorem 2. Assume that ω^* is the unique global maximizer of $\tilde{\mu}(\omega)$. Then $\omega_k \rightarrow \omega^*$ as $k \rightarrow \infty$.

Proof. By contradiction. Let \mathcal{H} be an infinite index set and suppose that $\omega_k \rightarrow \tilde{\omega}$ on \mathcal{H} , with $\tilde{\omega} \neq \omega^*$ (since Ω_1 is bounded, such \mathcal{H} exists for some $\tilde{\omega}$ in the closure of Ω_1). By continuity of $\tilde{\mu}(\cdot)$ and uniqueness of the global maximizer ω^* , there exists $\delta > 0$ such that, for k large enough, $k \in \mathcal{H}$,

$$\tilde{\mu}(\omega_k) < \tilde{\mu}^* - \delta.$$

Let D_k be the minimizer of (4) at ω_k . Since $\{D_k\}$ lies in a compact set, there exists an infinite index set $\mathcal{H}' \subset \mathcal{H}$ such that $D_k \rightarrow D^*$ on \mathcal{H}' . By uniform continuity on compact sets,

$$v_{D_k}(\omega) < \tilde{\mu}^* - \delta,$$

for all $k \in \mathcal{K}'$, in some fixed neighborhood V of $\tilde{\omega}$. On the other hand, in view of Theorem 1, $\xi_k > \tilde{\mu}^* - \delta$ for k large enough. Thus, for $k \in \mathcal{K}'$ large enough, V is taken out of Ω_k in Step 2 of the algorithm, which contradicts the fact that $\omega_k \rightarrow \tilde{\omega}$ on \mathcal{K} .

5. Numerical experiments

The algorithm was implemented¹⁰ in MATLABTM. In the implementation, a mid-point rule is used until Ω_k is reduced to just one interval and enough information has been accumulated to compute an interpolating function. All numerical experiments were run on a Sun UltraSparc 10, 300 MHz machine with 128 Meg of RAM, running Solaris 2.5 operating system.

The performance of the new algorithm was first compared to that of the “skewed μ ” approach mentioned in the introduction. For the former, function `mu` of the μ -Analysis and Synthesis Toolbox (Balas, Doyle, Glover, Packard & Smith, 1993) was used for computing ξ_{k+1} and D_{k+1} in Step 1. The GEVP in the latter was solved using the LMI Control Toolbox Version 1.0.4 (Gahinet, Nemirorski, Laub & Chilali 1995). Tolerance parameters were set so as to obtain roughly the same accuracy with both approaches; in particular, option C was used in `mu`. On 100 randomly generated¹¹ stable systems with 10 variables, 5 inputs, 5 outputs and 4 complex uncertainty block problems, there were 5 instances where the new approach yielded a μ -norm upper bound 1% or more lower than the upper bound obtained with the skewed- μ approach; in one of these instances, the difference was about 33%. The computation time was always less with the new algorithm, often by more than an order of magnitude.

To illustrate the behavior of the new algorithm, we then applied it to an example taken from Balas et al. (1993). The system is a model of an experimental aircraft (the HIMAT vehicle) with an \mathcal{H}_∞ sub-optimal pitch axis controller. Uncertainty and performance weightings are included in the model, for a total of 16 state variables. The resultant system is not strictly proper, hence a state space realization will have a non-zero feed-through matrix E . Redefine the Hamiltonian matrix as

$$H(\xi, A, B, C, E) \triangleq \begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} + \begin{bmatrix} B & 0 \\ 0 & -C^* \end{bmatrix} \times \begin{bmatrix} -E & \xi I \\ \xi I & -E^* \end{bmatrix}^{-1} \begin{bmatrix} C & 0 \\ 0 & B^* \end{bmatrix}, \tag{5}$$

and make sure that the initial frequency ω_0 is such that $\hat{\mu}(\omega_0) > \sigma_1(D)$. Then the algorithm can handle such

problems with no other changes. The uncertainty structure Δ consists of two 2×2 full blocks (robust performance problem). Fig. 4 shows a plot of $\hat{\mu}_\Delta(P(j\omega))$ (in this case, equal to $\mu_\Delta(P(j\omega))$) as a function of frequency.

The algorithm may now be applied to compute $\|P\|_{\hat{\mu}}$. In Table 1 the progress of the algorithm from iteration to iteration is shown. All quantities are as defined earlier. An asterisk in the column labeled I indicates that rational interpolation was used to generate ω_k . If Ω_k contains multiple intervals, or if the algorithm has not yet accumulated enough information, the mid-point rule is used to generate ω_k . It appears from the table that the algorithm exhibits quadratic convergence. Note that, in the process of computing $\hat{\mu}^*$, $\hat{\mu}$ has been evaluated at only a few points, which demonstrates great computational savings in comparison with a grid search. The total computation time was about 2.5 s. In comparison, on the same problem, the “skewed μ ” approach took over 50 s. Interestingly, it yielded the same value for $\hat{\mu}^*$.

6. Extension to mixed- μ

The algorithm is readily extended to compute an upper bound on the mixed- μ norm, that is, the μ -norm for systems with mixed dynamic and real parametric uncertainty. (See Feron (1997) for a sophisticated grid-based approach to this computation, accounting for possible discontinuities.) Two expressions for the widely used “D-G” upper bound $\hat{\mu}_\Delta(M)$ are as follows (see, e.g., Chapter 18 in Zhou & Doyle, 1998)

$$\hat{\mu}_\Delta(M) = \inf_{\substack{D \in \mathcal{D}, G \in \mathcal{G} \\ \xi > 0}} \left\{ \xi \cdot \bar{\sigma} \left(\left(\frac{DMD^{-1}}{\xi} - jG \right) \times (I + G^2)^{-1/2} \right) \leq 1 \right\}, \tag{6}$$




Fig. 4. μ -plot for the closed-loop system.

¹⁰ The implementation is available from the authors.

¹¹ MATLAB's function `sysrand` was used.

Table 1
Results for example system

k	ω_{k-1}	ξ_k	$ \omega_{k-1} - \omega^* $	$ \xi_k - \hat{\mu}^* $	I
1	0	0.7959766	$8.0 \times 10^{+1}$	8.9×10^{-1}	
2	3790.8315	0.7959766	$3.7 \times 10^{+4}$	8.9×10^{-1}	
3	367.31532	0.7959766	$2.9 \times 10^{+2}$	8.9×10^{-1}	
4	129.91042	1.3439461	$5.0 \times 10^{+1}$	3.4×10^{-1}	
5	64.955208	1.6041191	$1.6 \times 10^{+1}$	7.9×10^{-2}	
6	76.939264	1.6791253	3.5×10^0	3.7×10^{-3}	*
7	80.240400	1.6827812	2.1×10^{-1}	1.3×10^{-5}	*
8	80.450844	1.6827940	5.6×10^{-4}	9.5×10^{-11}	*

$$\hat{\mu}_\Delta(M) = \inf_{\substack{D \in \mathcal{D}, G \in \mathcal{G} \\ \alpha \geq 0 \\ -\alpha^2 D \leq 0}} \{ \alpha: M^*DM + j(GM - M^*G) \} \quad (7)$$

where \mathcal{D} and \mathcal{G} are determined by the structure. Expression (7) is an efficiently solvable LMI problem (see Boyd et al., 1994) and, given $\alpha, \tilde{D}, \tilde{G}$ satisfying the constraint in (7), the values $\xi = \alpha, D = \tilde{D}^{1/2}$, and $G = (1/\alpha)\tilde{D}^{-1/2}\tilde{G}\tilde{D}^{-1/2}$ satisfy the constraint in (6) for the same M (Zhou & Doyle, 1998). Moreover, given ξ, D , and G , it follows from (6), with $M = P(j\omega)$, that $\hat{\mu}_\Delta(P(j\omega)) \leq \xi$ for all ω such that

$$\bar{\sigma}(F(j\omega)) \leq 1, \quad (8)$$

with

$$F(j\omega) \triangleq \left(\frac{DP(j\omega)D^{-1}}{\xi} - jG \right) (I + G^2)^{-1/2}. \quad (9)$$

Note that, given a realization (A, B, C, E) for $P(j\omega)$, $F(j\omega)$ can be expressed as $C_F(j\omega I - A)^{-1}B_F + E_F$, where $B_F = \xi^{-1/2}BD^{-1}(I + G^2)^{-1/2}$, $C_F = \xi^{-1/2}DC$, and $E_F = (\xi^{-1}DED^{-1} - jG)(I + G^2)^{-1/2}$ are obtained by inspection from (9). Thus, the end-points of intervals where (8) holds correspond to imaginary eigenvalues of $H = H(1, A, B_F, C_F, E_F)$, where H is as in (5). When $E = 0$ (no feedthrough term in P), this reduces to¹²

$$H = \begin{bmatrix} A & 0 \\ 0 & -A^* \end{bmatrix} + \frac{1}{\xi} \begin{bmatrix} jBD^{-1}GDC & BD^{-2}B^* \\ -C^*D(I + G^2)DC & jC^*DGD^{-1}B^* \end{bmatrix}.$$

With this in hand, the algorithm discussed in the previous sections is readily extended to the mixed- μ case.

However, the $\hat{\mu}(\cdot)$ curve does not enjoy the same regularity properties as in the purely complex case. As a result, the interpolation rules for updating ω_{k+1} may not be appropriate and quadratic convergence may be lost. In the extreme case where the supremum of $\hat{\mu}(\omega)$ is strictly greater than the supremum of its lower envelope (function whose epigraph is the closure of that of $\hat{\mu}(\omega)$), convergence of ξ_k to $\hat{\mu}^*$ (Theorem 1) may even be lost unless frequency points where $\hat{\mu}$ is discontinuous (including possibly $\omega = \infty$) are checked separately, and the initial value of ξ_k is set accordingly. This can happen even in the trivial case of “real μ ” (real uncertainty) of a scalar transfer function such as $p(s) = 1/(s^3 + (3/2)s^2 + s + 1)$, since $\hat{\mu}(p(j0)) = 1, \mu(p(\pm j)) = 2$, and $\hat{\mu}(p(j\omega)) = 0$ for $\omega \notin \{0, \pm 1\}$. Indeed, the real μ of a scalar is equal to its upper bound $\hat{\mu}$ and is the absolute value of the scalar if the scalar is real, and zero otherwise.

7. Concluding remarks

The algorithm just outlined in this note can be refined in various ways. For instance, in the process of minimizing $v_D(\omega_k)$ to evaluate $\hat{\mu}(\omega_k)$, if it is found that the minimum value is less than ξ_k , then it is unnecessary to compute it (or D_k) with great accuracy.

The new algorithm can be extended to discrete-time systems. The Hamiltonian eigenvalue problem $H(\xi, A, B, C)$ is replaced by a symplectic eigenvalue problem $S(\xi, A, B, C)$ which will have an eigenvalue $e^{j\omega}$ on the unit circle if and only if ξ is a singular value of $C(e^{j\omega}I - A)^{-1}B$. We refer the reader to Hinrichsen and Son (1989), for the appropriate formulas. Alternatively, a corresponding continuous-time quadruple (A, B, C, D) can be obtained by means of a bilinear transformation.

References

Balas, G. J., Doyle, J. C., Glover, K., Packard, A., & Smith, R. (1993). *μ -analysis and synthesis toolbox: user's guide*. MUSYN Inc. and The MathWorks, Inc.


¹² Once ξ in Section 5 of Lawrence et al. (1996) is changed to ξ^2 (to match the notation used in the other sections of Lawrence et al. (1996)), the G matrix here corresponds to G/ξ in Section 5 of Lawrence et al. (1996). The Hamiltonian matrix obtained here is equivalent to that given in Lawrence et al. (1996); the latter may be obtained from the former by a similarity transformation $S^{-1}HS$ with $S = \text{diag}[I, -\xi I]$.

- Barzilai, J., & Ben-Tal, A. (1982). Nonpolynomial and inverse interpolation for line search: synthesis and convergence rates. *SIAM Journal of Numerical Analysis*, 19(6), 1263–1277.
- Bercovici, H., Foias, C., & Tannenbaum, A. (1990). Structured interpolation theory. In I. Gohberg (Ed.), *Extension and interpolation of linear operators and matrix functions* Operator Theory, Advances and Applications, Vol. 47. Boston: Birkhauser (pp. 195–220).
- Boyd, S. P., & Balakrishnan, V. (1990). A regularity result for the singular values of a transfer matrix and a quadratically convergent algorithm for computing its L_∞ -norm. *Systems and Control Letters*, 15, 1–7.
- Boyd, S. P., El Ghaoui, L., Feron, E., & Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. Philadelphia: SIAM.
- Bruinsma, N. A., & Steinbuch, M. (1990). A fast algorithm to compute the \mathcal{H}_∞ -norm of a transfer function matrix. *Systems and Control Letters*, 14(5), 287–293.
- Doyle, J. C., & Packard, A. (1987). Uncertain multivariable systems from a state space perspective. Proceedings of the American control conference (pp. 2147–2152). Minneapolis.
- Fan, M. K. H., & Tits, A. L. (1992). A measure of worst-case H_∞ performance and of largest acceptable uncertainty. *Systems and Control Letters*, 18, 409–421.
- Feron, E. (1997). A more reliable robust stability indicator for linear systems subject to parametric uncertainties. *IEEE Transactions on Automatic Control*, 42(9), 1326–1330.
- Ferreres, G., & Fromion, V. (1997). Computation of the robustness margin with the skewed μ tool. *Systems and Control Letters*, 32, 193–202.
- Gahinet, P., Nemirovski, A., Laub, A. L., & Chilali, M. (1995). *LMI control toolbox — for use with MATLAB*. The MathWorks, Inc.
- Hinrichsen, D., & Son, N. K. (1989). The complex stability radius of discrete-time systems and symplectic pencils. *Proceedings of the Conference on Decision and Control*, Tampa (pp. 2265–2270).
- Lawrence, C. T., Tits, A. L., & Van Dooren, P. (1996). A fast algorithm for the computation of an upper bound on the μ -norm. *Proceedings of the IFAC World Congress*, San Francisco, Vol. H. (pp. 59–64).
- Lee, L., & Tits, A. L. (1993). On continuity/discontinuity in robustness indicators. *IEEE Transactions on Automatic Control*, AC-38(10), 1551–1553.
- Packard, A., & Doyle, J. C. (1993). The complex structured singular value. *Automatica*, 29, 71–109.
- Poola, K., & Tikku, A. (1995). Robust performance against time-varying structured perturbations. *IEEE Transactions on Automatic Control*, AC-40(9), 1589–1602.
- Sideris, A. (1992). Elimination of frequency search from robustness tests. *IEEE Transactions on Automatic Control*, AC-37(10), 1635–1640.
- Sreedhar, J., Van Dooren, P., & Tits, A. L. (1996). A fast algorithm to compute the real structured stability radius. In R. Jeltsch, M. Mansour (Eds.), *Stability theory: proceedings of centenary conference*, Ticino, Switzerland, May 21–26, (pp. 219–230). Boston: Birkhäuser.
- Stewart, G. W., & Sun, J. (1990). *Matrix perturbation theory*. New York: Academic Press.
- Tits, A. L., & Balakrishnan, V. (1998). Small- μ theorems with frequency-dependent uncertainty bounds. *Mathematics of Control Signals and Systems*, 11(3), 220–243.
- Tits, A. L., & Fan, M. K. H. (1995). On the small- μ theorem. *Automatica*, 31(8), 1199–1201.
- Toker, O., & Özbay, H. (1995). On the NP-hardness of the purely complex μ computation, analysis/synthesis, and some related problems in multi-dimensional systems. *Proceedings of the American Control Conference* (pp. 447–451).
- Zhou, K., & Doyle, J. C. (1998). *Essentials of robust control*. Englewood Cliffs, NJ: Prentice-Hall.

Zhou, K., Doyle, J. C., & Glover, K. (1995). *Robust and optimal control*. Englewood Cliffs, NJ: Prentice Hall.



Craig T. Lawrence received the B.S. degree in Electrical Engineering and the B.S. degree in Mathematics in 1993, the M.S. degree in Electrical Engineering in 1996, and the Ph.D. in 1998, all at the University of Maryland, College Park. Since graduating, he has been with ALPHATECH, Inc., in Arlington, Virginia working on algorithms for intelligence, surveillance, and reconnaissance. His research interests are in the areas of optimization and robust control theory.



André L. Tits was born in Verviers, Belgium on April 13, 1951. He received the 'Ingenieur Civil' degree from the University of Liege, Belgium and the M.S. and Ph.D. degrees from the University of California, Berkeley, all in Electrical Engineering, in 1974, 1979, and 1980, respectively. Since 1981, Dr. Tits has been with the University of Maryland, College Park. Currently, he is a Professor of Electrical Engineering and he holds a permanent joint appointment with the Institute for Systems Research. He has held visiting

positions at the University of California, Berkeley, at the Lund Institute of Technology, at INRIA, at the Catholic University of Louvain at Louvain-la-Neuve, Belgium and at the Australian National University.

Dr. Tits received a 1985 NSF Presidential Young Investigator Award. He is a Fellow of the Institute of Electrical and Electronics Engineers. He is an Associate Editor of *Automatica* and the Editor for Technical Notes and Correspondence of the *IEEE Transactions on Automatic Control*. His main research interests lie in various aspects of numerical optimization, optimization-based system design and robust control with emphasis on numerical methods.



Paul M. Van Dooren received the engineering degree in computer science and the doctoral degree in applied sciences, both from the Katholieke Universiteit te Leuven, Belgium, in 1974 and 1979, respectively. He held research and teaching positions at the Katholieke Universiteit te Leuven (1974–1979), the University of Southern California (1978–1979), Stanford University (1979–1980), the Australian National University (1984), Philips Research Laboratory Belgium (1980–1991), the University of Illinois at Urbana-Champaign (1991–1994), and the Université Catholique de Louvain (1980–1991, 1994–now) where he is currently a professor of Mathematical Engineering.

Dr. Van Dooren received the Householder Award in 1981 and the Wilkinson Prize of Numerical Analysis and Scientific Computing in 1989. He is an Associate Editor of *Journal of Computational and Applied Mathematics*, *Numerische Mathematik*, *SIAM Journal on Control and Optimization*, *Applied Mathematics Letters*, *Linear Algebra and its Applications*, *Journal of Numerical Algorithms*, the *Electronic Transactions of Numerical Analysis*, *Applied and Computational Control, Signals and Systems*, *Mathematics of Control, Signals and Systems* and editor in chief of *SIAM Journal on Matrix Analysis and Applications*. His main interests lie in the areas of numerical linear algebra, systems and control theory, digital signal processing, and parallel algorithms. He is currently the Chairman of the SIAM Special Interest Group on Linear Algebra. He has published over 60 journal papers and 80 refereed conference papers.