Optimal Robust H_∞ Control

Stephen D. Patek
Michael Athans
Room 35-406
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139

33rd CDC

Keywords: H_∞, robust control, optimal control

Abstract

This paper briefly introduces and explores "optimal robust H_∞ control." The main idea is to obtain optimal *guaranteeable* H_∞ performance given uncertainty about the plant's state space parameters. The uncertainty can be either time-varying or time-invariant; both cases are treated separately. An example is presented which illustrates properties of the optimal robust designs.

1 Introduction

This paper is concerned with optimal robust H_∞ control design. The systems for which controllers will be designed are nominally LTI. Only parametric uncertainty is considered, with allowable perturbations being norm-bounded and either time-varying or time-invariant. The objective is to achieve robust stability and optimal guaranteeable H_∞ performance. This topic has recently been considered for the case of full-state feedback in [3].

This work began with the development of a design methodology for time-varying parametric uncertainty. [4] The result was RHINF (Robust H_∞), a robust analog of standard (nominal) H_∞ optimization. The goal was to explore the possibility of replacing standard H_∞ design with RHINF in DK iteration. [1] This would provide a less conservative design framework for problems with mixed LTI (unmodeled dynamic) and time-varying parametric uncertainty.

The original version of RHINF was conceived as a synthesis extension of the robustness analysis test in [6]. One particularly exciting aspect of the analysis test was that it claimed to be both necessary and sufficient with respect to real-valued parameter uncertainties. Unfortunately, while this claim is true, the fact that perturbations can be time-varying makes the robustness problem so demanding that the distinction between real and complex-valued perturbations is no longer significant. (Destabilizing time-varying perturbations in the proof of necessity can always be constructed as real-valued for any norm-bound on the size of allowable perturbations.) This is true in both the "quadratic stability" framework of [6] and in the more general input-output setting of [5].

The purpose of this paper is to report on some interesting aspects of "optimally robust" designs. The results are qualitative; most of the theoretical infrastructure is still classified as work in progress. Section 2 defines the optimal robust control problem and specifies the class of uncertain systems for which designs are to be produced. Section 3 develops the separate optimal robust control design methodologies (LTV-RHINF and LTI-RHINF) for design with respect to time-varying and time-invariant uncertainty, respectively. Since the optimization problems are non-convex, the resulting methodologies typically yield sub-optimal designs. This paper takes an initial look at what happens when these difficulties are ignored. Section 4 provides a numerical example that suggests some interesting properties that may hold in general. Section 5 contains a brief summary and suggestions for future work.

2 Set-up

The class of uncertain (LTV) "design-plants" for which optimally robust designs are to produced is described as follows,

$$\Pi = \left[\begin{array}{c|cc} A+\Delta A(t) & B_1+\Delta B_1(t) & B_2+\Delta B_2(t) \\ \hline C_1+\Delta C_1(t) & D_{11}+\Delta D_{11}(t) & D_{12}+\Delta D_{12}(t) \\ C_2+\Delta C_2(t) & D_{21}+\Delta D_{21}(t) & D_{22}+\Delta D_{22}(t) \end{array}\right] \tag{1}$$

The Δ-terms are constrained according to,

$$\begin{bmatrix} \Delta A(t) & \Delta B_1(t) & \Delta B_2(t) \\ \Delta C_1(t) & \Delta D_{11}(t) & \Delta D_{12}(t) \\ \Delta C_2(t) & \Delta D_{21}(t) & \Delta D_{22}(t) \end{bmatrix} = \begin{bmatrix} H_1 \\ H_2 \\ H_3 \end{bmatrix} \Delta(t) \begin{bmatrix} E_1 & E_2 & E_3 \end{bmatrix} \tag{2}$$

The H_i's and E_i's are assumed to have been determined through modeling and are constant-valued. $\Delta(t)$ parameterizes the time-varying, parametric uncertainty. All that is known about it is that it is Lebesgue integrable and has bounded Euclidean gain according to,

$$\bar{\sigma}[\Delta(t)] \leq \rho \tag{3}$$

The "level of uncertainty" ρ is nonnegative and is assumed to be known a priori. In this paper, no additional structure is imposed on $\Delta(t)$. (More generally, however, $\Delta(t)$ should be considered diagonal, with repeated scalars, to reflect the fact that uncertainty is parametric.) The nominal design-plant is LTI and can be recovered through $\Delta(t) = 0$. The notation for Π is an extension of that normally used in the literature: performance outputs $z(t)$ and measurement outputs $y(t)$ are mapped from the state and exogenous inputs through the second and third row partitions, respectively. Similarly, disturbance inputs $w(t)$ and control inputs $u(t)$ are mapped to the state and system outputs through the second and third column partitions, respectively. As a matter of notation, let $P(\Delta(t))$ represent an allowable design-plant corresponding to a particular realization of the uncertainty $\Delta(t)$, satisfying (3).

The design objective is to find an LTI feedback controller, $K: y(t) \to u(t)$, which yields robust stability and optimal guaranteeable H_∞ performance in the presence of allowable perturbations. The problem can be rephrased as,

$$\min_K \{\gamma : \|F_l[P(\Delta(t)), K]\|_\infty \leq \gamma, \ \forall \ \Delta(t) \ \text{allowable}\} \tag{4}$$

where $F_l[P(\Delta(t)), K]$ is the closed loop map corresponding to K and a particular realization of the uncertainty $\Delta(t)$. (Strictly speaking, the H_∞ norm does not exist for time-varying systems. The use of this notation in the time-varying case refers to the induced L_2 norm of the zero-state response of the system.) Of course, K must be nominally internally stabilizing.

3 Optimal robust control design

This section begins with a brief review of some familiar analysis tests for robust H_∞ performance. These tests, which invariably manifest themselves as "scaled" variations on the small gain theorem, are then used to develop algorithms for optimally robust H_∞ control design.

3.1 Analysis for a desired performance level γ

The goal here is to be able to answer the following yes or no question: *Given a compensator K, is the resulting (uncertain) closed loop system robustly stable with a guaranteed H_∞ performance level of γ?*

One approach to answering this question is to apply the small gain theorem directly. This requires that the uncertain system be represented as the feedback connection of the uncertainty Δ and the nominal closed loop corresponding to $P(0)$ and K. This is a familiar manipulation which results in a new design-plant based on Π with extra disturbance inputs and extra performance outputs (relating to how parameter perturbations affect the plant's dynamics). The new design-plant, denoted P_s, is deterministic and takes the form,

$$P_s = \left[\begin{array}{c|c|c} A & \begin{bmatrix} \rho H_1 & \gamma^{-1}B_1 \end{bmatrix} & B_2 \\ \hline \begin{bmatrix} E_1 \\ C_1 \end{bmatrix} & \begin{bmatrix} 0 & \gamma^{-1}E_2 \\ \rho H_2 & \gamma^{-1}D_{11} \end{bmatrix} & \begin{bmatrix} E_3 \\ D_{12} \end{bmatrix} \\ C_2 & \begin{bmatrix} \rho H_3 & \gamma^{-1}D_{21} \end{bmatrix} & D_{22} \end{array}\right] \tag{5}$$

Note that this representation normalizes the "level of uncertainty" by incorporating ρ as a scaling on the uncertainty input. Let the new normalized perturbation be denoted $\Delta_n(t)$. The desired level of performance is similarly normalized by incorporating γ^{-1} as a scaling on the original disturbance input. Let $\Delta_{aug}(t)$ be the diagonal augmentation of $\Delta_n(t)$ and a similar fictitious perturbation block $\Delta_f(t)$ for performance:

$$\Delta_{aug}(t) = \begin{bmatrix} \Delta_n(t) & 0 \\ 0 & \Delta_f(t) \end{bmatrix} \tag{6}$$

By thinking of $\Delta_{aug}(t)$ as an arbitrary operator with unity induced L_2 norm, the small gain theorem gives a sufficient condition for a "yes" answer to the robustness question. Namely, the closed loop system with K is robustly input-output stable and has a guaranteeable H_∞ (i.e. zero-state L_2 induced) performance level of γ if,

$$||F_l[P_s, K]||_\infty < 1 \tag{7}$$

This is a sufficient condition for robust performance because the augmented uncertainty block used in the above argument has structure not accounted for by the small gain theorem.

Reducing conservatism for time-varying uncertainty One way to reduce conservatism in the above is to apply a set of "invariant" transformations to the set of allowable perturbations. For the case of time-varying uncertainty, this amounts to small gain analysis using "constant D-scales". The constant scalings are used to account for the block-diagonal structure of the augmented perturbation block. Thus, a less conservative test for robustness is as follows: If there exists $\varepsilon > 0$ such that,

$$|| \begin{bmatrix} \sqrt{\varepsilon}I & 0 \\ 0 & I \end{bmatrix} F_l[P_s, K] \begin{bmatrix} \sqrt{\varepsilon^{-1}}I & 0 \\ 0 & I \end{bmatrix} ||_\infty < 1 \tag{8}$$

then the closed loop system with K is robustly input-output stable and has a guaranteeable H_∞ (i.e. zero-state L_2 induced) performance level of γ.

The results of [5] (extended to continuous-time systems and applied to robust performance analysis) seem to suggest that this test is both necessary and sufficient. "Necessity" would mean that if there is no $\varepsilon > 0$ such that the condition in (8) is true, then a structured performance-degrading perturbation can be constructed which necessarily keeps the closed loop system from meeting specifications. Unfortunately, such degrading perturbations typically have to be constructed as LTV dynamic systems. This means that "necessity" cannot apply to the case of time-varying *parametric* perturbations. Fortunately, however, the fact that performance-degrading perturbations can always be constructed as real-valued means that there is no additional conservatism associated with the fact that complex-valued perturbations are allowable. (Complex perturbations are allowable because the only active constraint on the perturbations is that they must be norm-bounded.)

The case that $H_2 = 0$ and $E_2 = 0$ in equation (2) is considered within the context of "quadratic stability" in [6]. (In that article, D_{11} is further constrained to equal zero. The extension to non-zero D_{11} follows straightforwardly.) The main result is that the uncertain closed loop system is (robustly) "quadratically stable with disturbance rejection γ" if and only if there exists $\varepsilon > 0$ such that (8) holds. To be "quadratically stable with disturbance rejection γ" means that there exists a positive semi-definite matrix Q such that,

1. robust closed loop stability can be proven using a single quadratic Lyapunov function of the form $V(x(t)) = x^T(t)Qx(t)$ with Q time-invariant, and

2. robust performance can be established via manipulation of square-integral functions parameterized by Q (time-invariant).

To use Lyapunov theory to prove the internal stability of LTV systems, one must generally make use of *time-varying* quadratic Lyapunov functions. Moreover, it is important to note that robust performance could have been established through manipulations of square-integrable functions parameterized by a *time-varying* positive semi-definite matrix $Q(t)$. Thus, there is a lot of conservatism built into the definition of "quadratic stability with disturbance rejection γ" . Also, the test for "quadratic stability with disturbance rejection γ" fails to be "necessary" in the traditional sense (where, if the robustness condition fails, an allowable de-robustifying perturbation can be constructed).

Reducing conservatism for LTI uncertainty When allowable parameter perturbations are known to be time-invariant, the conservatism of the small-gain test can be reduced through application of the structured singular value, μ. [2] To do this, the augmented perturbation block in (6) must first be constrained to be stable, linear, and time-invariant. Let S be the "block structure" associated with the newly constrained augmented perturbation block. Then, the closed loop system with K is robustly input-output stable and has a guaranteeable H_∞ performance level of γ if,

$$\mu_S(F_l[P_s, K])(j\omega) < 1, \quad \forall\ \omega \tag{9}$$

Considering the simple structure of this μ-test, μ_s can be evaluated explicitly as,

$$\mu_S(F_l[P_s, K])(j\omega) = \inf_{\varepsilon>0} \bar{\sigma} \left[\begin{bmatrix} \sqrt{\varepsilon}I & 0 \\ 0 & I \end{bmatrix} F_l[P_s, K](j\omega) \begin{bmatrix} \sqrt{\varepsilon^{-1}}I & 0 \\ 0 & I \end{bmatrix} \right] \tag{10}$$

In this case, an (infimizing) ε-scaling is evaluated for each frequency ω.

Robustness tests using the structured singular value are both necessary and sufficient when the LTI perturbations are dynamic (i.e. complex-valued in the frequency domain). When perturbations are parametric (that is, non-dynamic and *real-valued* in the frequency domain) μ robustness tests become only sufficient. This is because robustness is guaranteed with respect to a larger class of perturbations than which actually exist. To make up for this, the additional stipulation that Δ_{aug} be real-valued must be made. The resulting theory, representing a significantly more detailed and complex level of analysis, is known as "mixed μ" [7] and is ignored for the remainder of this paper.

3.2 Optimal robust design

The design algorithms discussed in this paper can be decomposed into two parts: an outer loop where a sequence of desired levels of robust performance γ are chosen and an inner loop where it is determined whether a given level of performance can actually be achieved robustly. (See Figure 1.)The outer loop corresponds exactly to "γ-iteration", commonly used in implementations of state-space H_∞ design algorithms. The inner loop corresponds to robustness-measure optimization for a given desired level of robust performance. For the case of time-invariant perturbations, this inner loop takes the form of "μ-synthesis" via DK iteration. The resulting sub-optimal design algorithm is referred to as LTI-RHINF henceforth.

It is the time-varying case that is somewhat new and interesting. The robustness test of (8) says that a performance level of γ (as well as stability) can be guaranteed as long as there exists a constant "D-scale" (parameterizedformance level of γ (as well as stability) can be guaranteed as long as there exists a constant "D-scale" (parameterized by ε) that keeps the H_∞ norm of the *scaled* closed loop system less than unity. Thinking of this H_∞ norm as a measure of robustness (for a given desired performance level γ), a brute force approach is taken to optimize this quantity. The optimization is computed with respect to both K (stabilizing) and $\varepsilon > 0$. The key step here is to note that standard H_∞ optimal design can be employed to get an optimal K for any given value of $\varepsilon > 0$. This requires that ε be incorporated into the augmented scaled design-plant P_s. Assuming that these optimal closed loop H_∞ norms are convex in ε, a straightforward

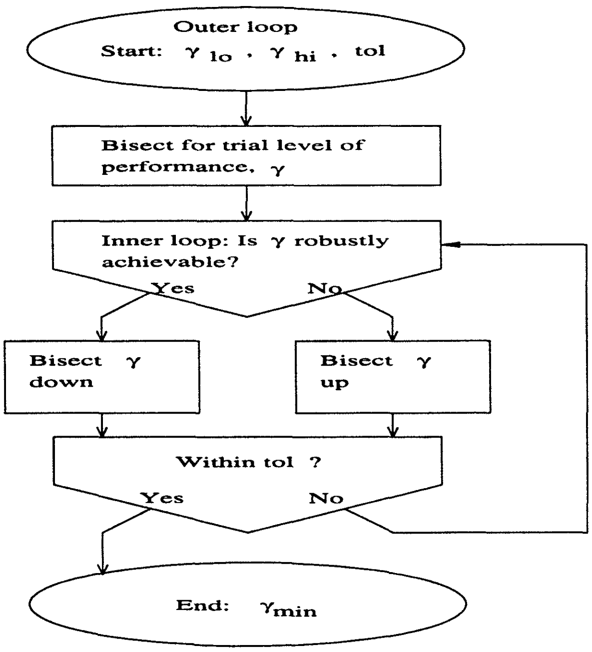

Notes:

1. For LTI-RHINF the inner loop is evaluated via DK iteration using standard (complex) μ.
2. For LTV-RHINF the inner loop is evaluated via a brute force search over ε . (For each ε , an optimal K can be computed using nominal H-infinty design tools.)

Figure 1: Flow chart for optimal robust design.

search for the global robustness-measure-optimizing K can be implemented. The resulting design algorithm, incorporating this search as the "inner loop", is referred to henceforth as LTV-RHINF.

LTI-RHINF and LTV-RHINF have been implemented in MATLAB. This code was used in generating the optimally robust designs for the mass-spring system in the following subsection. Additional examples of RHINF designs can be found in [4].

4 An example

Consider the mass-spring system illustrated in Figure 2. Control design for this system is a well known problem that is often used as a benchmark-test for robust design methodologies. The object is to regulate the position x_2 of the right-hand mass in the presence of measurement "noise" v and disturbance forces d. A control force u may be applied to the left-hand mass based on the corrupted measurement $y = x_2 + v$. The spring's modulus C_s is uncertain and is allowed to be time-varying such that $.5 < C_s(t) < 2.0$. The state space equations for this system are easy to derive and can be found elsewhere in the literature.

One way to put this problem into the optimal-robust-control framework is to say (arbitrarily) that C_s is nominally constant with $C_{s,nom} = 1.25$. Thus, allowable perturbations $\Delta(t)$ to $C_{s,nom}$ must satisfy $|\Delta(t)| \le \rho$, where $\rho = .6C_{s,nom}$. To keep the H_∞ optimization problem from being singular, the performance vector z must account for both deviations in x_2 and control useage u. Moreover, the overall disturbance vector w must account for the effects of both v and w. In this example, $z = (u, x_2)^T$ and $w = (v, d)^T$. More generally, various weights can be incorporated into the various inputs and outputs to reflect a priori information about disturbances and performance specifications. With all of this said, the control problem can easily be cast into the form of Equations (1-3).

LTI-RHINF and LTV-RHINF optimal designs for this problem were computed using MATLAB. To summarize the results:

1. The nominal H_∞ design is not robustly stable— even with respect to time-invariant perturbations. Nominal performance level equals $\gamma_{H_\infty} = 3.1126$. (Compensator is 4^{th} order, i.e. has four states.)

2. The LTV-RHINF design is robust with respect to allowable *time-varying* perturbations. It is robustly stable and yields a guaranteeable performance level of $\gamma_{\text{LTV-RHINF}} = 28.04 \pm .01$. (Compensator is

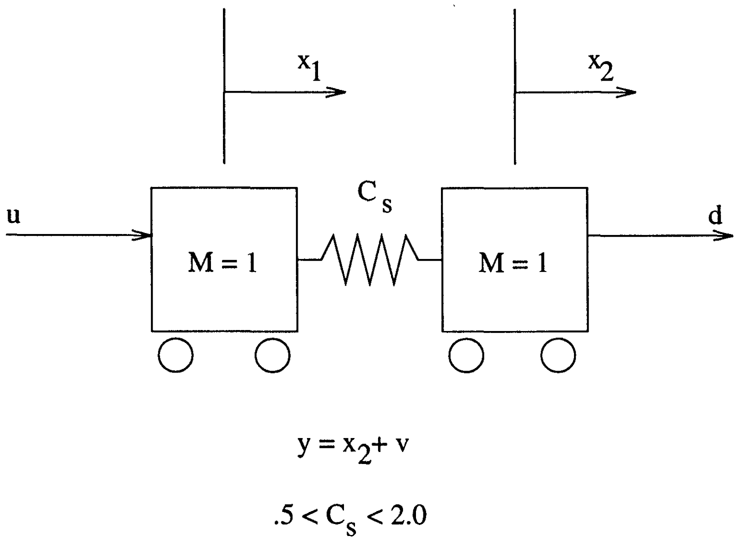

Figure 2: Schematic for the mass-spring system.

4^{th} order.)

3. The LTI-RHINF design is robust with respect to norm-bounded *time-invariant* perturbations. For these perturbations, it is robustly stable and yields a guaranteeable performance level of $\gamma_{\text{LTI-RHINF}} = 27.65 \pm .05$. (Compensator is 10^{th} order.)

Results for time-invariant perturbations are summarized in Figure 3. The traces in Figure 3 show resultant closed-loop H_∞ performance for constant-valued spring constant perturbations. (These are often referred to as "performance buckets".) The poor robustness of the nominal H_∞ design is clear from the fact that its bucket traces blow-up for values of C_s not far from $C_{s,nom}$. It is interesting to note that, at least with respect to time-invariant perturbations, the LTI-RHINF and LTV-RHINF designs yield very similar performance. (Does this say anything significant about the underlying problem? The fact that spring constant perturbations *can* be time-varying seems to be unimportant for this performance metric. It happens to be true that LTI-RHINF and LTV-RHINF designs have very similar frequency responses.) Perhaps the most interesting aspect of these traces is the fact that, for both robust designs, the guaranteeable levels of performance are actually achieved for every possible constant-valued perturbations. This shows up in the Figure 3 as performance buckets with perfectly flat bottoms. This means every possible constant-valued perturbation is a worst-case perturbation. (No allowable value for C_s is any worse or better than any other allowable value.) That this is true for the LTV-RHINF design is especially strange since (presumably) time-varying perturbations are "worse" than time-invariant ones.

5 Discussion

The simple example in the previous section has pointed out a number of interesting peculiarities:

1. Insignificant gap between LTI-RHINF and LTV-RHINF optimal performance.

2. Worst-case performance achieved for every allowable constant-valued perturbation (for both robust designs).

3. Worst-case performance for the LTV-RHINF design achieved by a time-invariant perturbation.

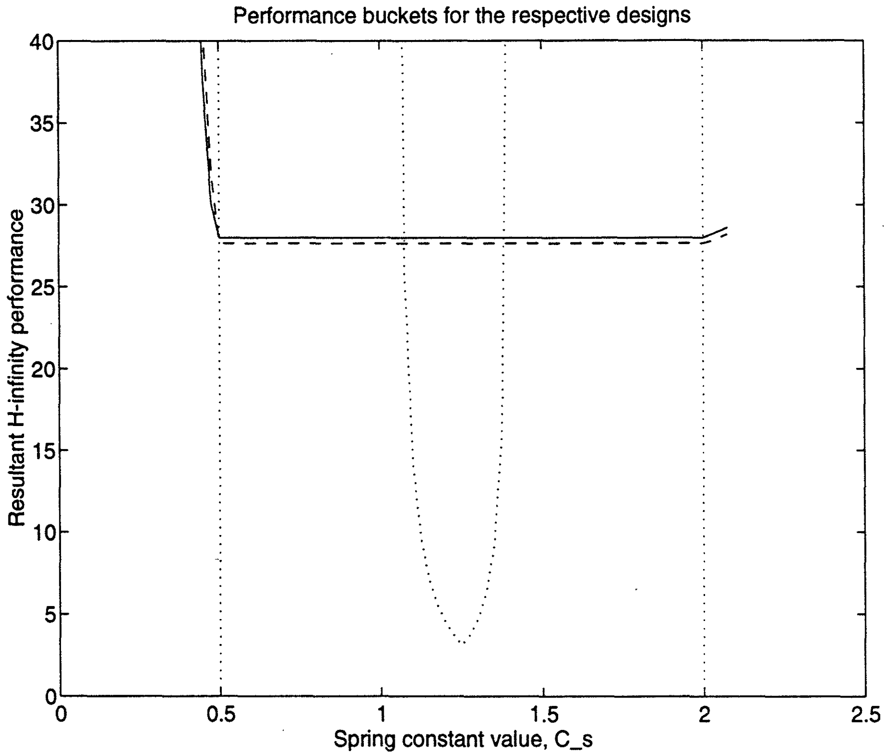

Figure 3: Performance-bucket plots for the example designs. Legend: solid $\leftrightarrow$ RHINF, dashed $\leftrightarrow$ LTI-RHINF, dotted $\leftrightarrow$ H_∞.

Some of these features may actually hold in general for optimally robust designs. One that definitely does not is Feature 1. Other examples can be found [4] in which significant gaps between LTI-RHINF and LTV-RHINF performance exist. In general, the amount of gap is problem (and problem formulation) dependent.

To address Features 2 and 3, it is the author's opinion that these may actually hold in general and are related to the fact that H_∞ optimal designs yield closed loop frequency responses that are "all-pass". ("All-pass" refers to the flat magnitude responses seen in H_∞ optimal closed loop transfer functions $T_{zw}(j\omega)$. One-block H_∞ problems, where only sensitivity is weighted, have optimal closed loop maps that are truly all-pass. More generally, H_∞ problems that are not singular have optimal closed loop maps that are "all-pass" up to some break-frequency where competing cost factors force the closed loop transfer to roll-off.) This conjecture seems plausible considering the state-space μ-tests discussed in [2], where the Laplace variable s can be broken out of a transfer function and treated as a perturbation via an LFT. With this in mind, a strong basis for verifying the conjecture may come through studying properties of the following optimization problem:

$$\inf_{\gamma>0} \inf_{Q,D\in D_s} \bar{\sigma}\left[\begin{bmatrix} D^{\frac{1}{2}} & 0 \\ 0 & \gamma^{-1}I \end{bmatrix}(T_1+T_2QT_3)\begin{bmatrix} D^{-\frac{1}{2}} & 0 \\ 0 & I \end{bmatrix}\right] \tag{11}$$

where T_1, T_2, T_3 are fixed matrices, Q is a matrix of appropriate dimension, and D_s is a prescribed set of positive definite matrices. Much of the machinery for such an analysis has already been developed and appears in [3] among other places.

References

[1] J. Doyle. Structured uncertainty in control systems design. *Proc. 24th IEEE Conf. Decision and Control*, pages 260–265, 1985.

[2] A. Packard and J. Doyle. The complex structured singular value. *Automatica*, 29(1):71–110, January 1992.

[3] A. Packard, K. Zhou, P. Pandey, J. Leonhardson, and G. Balas. Optimal, constant I/O similarity scaling for full-information and state-feedback control problems. *Systems and Control Letters*, 19:271–280, 1992.

[4] S. Patek. Robust H_∞ control via quadratic stabilization. Master's thesis, Massachussetts Institute of Technology, Cambridge, MA, 1993.

[5] J. Shamma. Robustness analysis for time-varying systems. *Proc. 31st IEEE Conf. Decision and Control*, 1992.

[6] L. Xie, M. Fu, and C. de Souza. H_∞ control and quadratic stabilization of systems with parameter uncertainty via output feedback. *IEEE Transactions on Automatic Control*, 37(8):1253–1256, August 1992.

[7] P. Young, M. Newlin, and J. Doyle. μ analysis with real parametric uncertainty. *Proc. 30th IEEE Conf. Decision and Control*, pages 1251–1256, 1991.