On Ellipsoidal Bounds for Amplitude Control of LTI Systems

Technical Report of the ISIS Group
at the University of Notre Dame
ISIS-98-003
April 1998

Christopher J. Bett and Michael Lemmon
Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN 46556

Interdisciplinary Studies of Intelligent Systems

On Ellipsoidal Bounds for Amplitude Control of LTI Systems

Christopher J. Bett and Michael Lemmon*

Department of Electrical Engineering

University of Notre Dame

Notre Dame, IN 46556

April 1998

Abstract

One approach to estimating the peak-to-peak gain of an LTI system is to bound the reachable states of the system with an invariant ellipsoid. This report presents numerical evidence suggesting that the best ellipsoidal bound (the so-called $*$-norm) can be poor in the presence of dominant and/or unstable zero dynamics and/or repeated poles. In parallel, the report investigates the use of the system ∞-norm to evaluate the conservativeness of the $*$-norm.

1 Introduction

One approach to estimating the peak-to-peak gain of an LTI system is to bound the reachable states of the system with an invariant ellipsoid [1, 2, 4]. This approach is useful because, unlike the ℓ_1 methods [5], it can be applied directly to continuous-time systems and it guarantees an upper bound on controller dimension no greater than the order of the plant. (The maximum dimension of ℓ_1 controllers, on the other hand, is unbounded.) In addition to these advantages, the ellipsoid approach presents LMI based analysis and synthesis conditions. This offers efficient computation as well as the ability to combine with other performance measures [8].

In some cases, however, it has been shown [11] that even the best invariant ellipsoidal approximation to the reachable states (yielding the so-called $*$-norm [1]) can be arbitrarily conservative. This report presents an empirical investigation into cases where such amplitude bounds can be conservative. Numerical evidence is presented to suggest that the $*$-norm can be poor in the presence of dominant and/or unstable zero dynamics

*The authors gratefully acknowledge the partial financial support of the Army Research Office (DAAH04-95-1-0600, DAAH04-96-1-0134).

and/or repeated poles. In parallel, the report investigates the use of the system ∞-norm to evaluate the conservativeness of the $*$-norm via

$$\frac{\|H\|_\infty}{\|H\|_*} \leq \frac{\|H\|_{i\infty}}{\|H\|_*} \leq 1$$

For exposition, the LTI systems considered here are SISO and strictly proper. In addition, only analysis results are presented. It is expected that the insights gained here can be applied to control synthesis problems as well. Note that while the focus of this report is on cases where the $*$-norm is conservative, the numerical evidence presented here also implies that the $*$-norm can be nonconservative, and therefore useful.

The remainder of this report is organized as follows. Section 2 outlines notation and the basic notions associated with bounding reachable stated of the system with invariant ellipsoids. Section 3 briefly discusses the potential conservativeness of the $*$-norm. Section 4 presents numerical examples designed to lend insight to the problem of conservativeness and Section 5 discusses these insights and their implications.

2 Mathematical Preliminaries

2.1 Notation and definitions

Throughout the paper, $\|\cdot\|$ denotes the Euclidean vector norm. The $\mathcal{L}_\infty$ norm (peak norm) of a signal $f : \mathbb{R}^+ \to \mathbb{R}^n$ is defined as

$$\|f\|_\infty := \operatorname*{ess\,sup}_{t \geq 0} \|f(t)\|$$

The linear space $\mathcal{L}_\infty^n := \left\{ f : \mathbb{R}^+ \to \mathbb{R}^n \,|\, \|f\|_\infty < \infty \right\}$ denotes all finite-horizon infinity norm bounded signals; the set of all such signals bounded to the unit sphere is denoted $\mathcal{BL}_\infty^n$. When the dimension of the signal is either clear from the context or immaterial, the superscript will be dropped.

Given a linear operator, $H : \mathcal{L}_\infty \to \mathcal{L}_\infty$, the induced $\mathcal{L}_\infty$ norm of H is defined as

$$\|H\|_{i\infty} := \sup_{w \in \mathcal{BL}_\infty} \|Hw\|_\infty$$

Thus, the induced $\mathcal{L}_\infty$ norm of H can be seen as the maximum peak-to-peak gain of H. If H has an n-dimensional state space realization

$$\dot{x} = Ax + Bw \quad (1)$$

$$z = Cx \quad (2)$$

the induced $\mathcal{L}_\infty$ norm of H may also be described in terms of the set $\mathcal{X} \subset \mathbb{R}^n$ of all states reachable from the origin in finite time by some input, $w \in \mathcal{BL}_\infty$. In fact,

$$\|H\|_{i\infty} = \sup_{x \in \mathcal{X}} \|Cx\|$$

Any set $\mathcal{P} \subset \mathbb{R}^n$ which contains the origin and is invariant for H under $w \in \mathcal{BL}_\infty$ must necessarily contain $\mathcal{X}$. Thus, any such invariant provides an upper bound $\sup_{x\in\mathcal{P}} \|Cx\|$ on $\|H\|_{i\infty}$. The tightest such upper bound which can be obtained using invariant ellipsoids is defined as the $*$-norm of the system, H, and is denoted $\|H\|_*$ [1].

2.2 Lyapunov characterization of the $*$-norm

Assume that the system in (1-2) is stable and define $\kappa := -2\max(\Re(\text{spec}(A)))$. As shown in [1,4], the set of all invariant ellipsoids which bound the reachable states of the system in (1-2) subject to $w \in \mathcal{BL}_\infty$ may be characterized by positive definite solutions, Q, to the matrix inequality

$$QA' + AQ + \alpha Q + \frac{1}{\alpha}BB' \leq 0 \tag{3}$$

where $\alpha \in (0, \kappa)$ is a scalar. It is not difficult to show that for a fixed $\alpha \in (0, \kappa)$, the smallest invariant which may be obtained from (3) is given by the unique positive definite solution, Q_α, to the Lyapunov equation

$$Q_\alpha A' + AQ_\alpha + \alpha Q_\alpha + \frac{1}{\alpha}BB' = 0 \tag{4}$$

In other words, the smallest invariant ellipsoid for the system in (1-2) is found by solving the parameterized Lyapunov equation as α sweeps through $(0, \kappa)$. Thus, the $*$-norm may now be defined as

$$\|H\|_* := \inf_{\alpha\in(0,\kappa)} \|CQ_\alpha C'\|^{1/2} \tag{5}$$

where Q_α is the unique positive definite solution to (4).

3 Conservativeness of the Ellipsoidal Bound

Because the $*$-norm is characterized by the smallest ellipsoids bounding the reachable states of the system, it is therefore the least conservative upper bound on $\|H\|_{i\infty}$ which can be determined by invariant ellipsoids. The conservativeness of the bound, however, is problem dependent. In [11], it is shown that the bound can be arbitrarily conservative with respect to the induced $\mathcal{L}_\infty$ norm, although that finding fails to classify the types of systems for which the $*$-norm is conservative or nonconservative. In fact, within class of systems described in [11], there exist examples where the ellipsoidal approximation is very good.

In the remainder of this report, we present empirical results which show classes of systems where the ellipsoidal bound can be conservative. Specifically, we provide numerical and intuitive evidence to suggest that systems with

- zeros close to the $j\omega$-axis relative to the system poles,
- widely spaced allpass components, or
- repeated poles

will admit conservative ellipsoidal bounds on the reachable states of the systems. These results include examples drawn from the literature used to demonstrate *extreme* behavior between control designs when different signal norms are used to quantify system performance.

In addition, we demonstrate the application of a fundamental signal norm bound in evaluating the conservativeness of the ellipsoidal bound for LTI systems described by (1-2). Combining the upper bound on $\|H\|_{i\infty}$ established so far with a well-known lower bound

$$\|H\|_{\infty} \leq \|H\|_{i\infty} \leq \|H\|_{*}$$

yields

$$\frac{\|H\|_{\infty}}{\|H\|_{*}} \leq \frac{\|H\|_{i\infty}}{\|H\|_{*}} \leq 1 \tag{6}$$

Both the ∞-norm and the $*$-norm of H are easily computed via state-space methods. The ratio on the left-hand side represents a worst-case measure of the conservativeness of the ellipsoidal bound. Unfortunately, the quantity can also be conservative since [3]

$$\frac{\|H\|_{i\infty}}{\|H\|_{\infty}} \leq 2n \tag{7}$$

In other words, if the induced $\mathcal{L}_{\infty}$ norm of H is much larger that the induced $\mathcal{L}_2$ norm, then the worst-case bound from (6) may imply poor performance of the ellipsoidal bound when, in fact, the ellipsoidal bound is adequate. Examples of systems where the bound in (7) is tight are discussed in [3].

4 Examples

We now demonstrate the use of the ellipsoidal bound on some numerical examples illustrating the classes of systems described in the previous section. Some examples were chosen from extreme cases in the literature demonstrating the difference between the induced $\mathcal{L}_{\infty}$ norm and the $\mathcal{H}_{\infty}$ norm approaches in order to demonstrate the utility of the bound in (6).

4.1 Example 1: Second-order system with a simple zero

Consider

$$H(s) = \frac{\lambda s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \tag{8}$$

describing a general second-order system with a real zero. For fixed ζ and ω_n, the zero location can be parameterized by $\lambda \in \mathbb{R}$. Figure 1 shows the performance of the ellipsoidal bound versus the zero location for $\omega_n = 10$ and $\zeta = 1.5$ (real pole locations $\{-26.2, -3.82\}$).

As demonstrated by Figure 1, the zero location greatly influences the ellipsoidal bound performance. For minimum phase zero locations, performance degrades as the zero moves closer to the $j\omega$-axis; the poorest

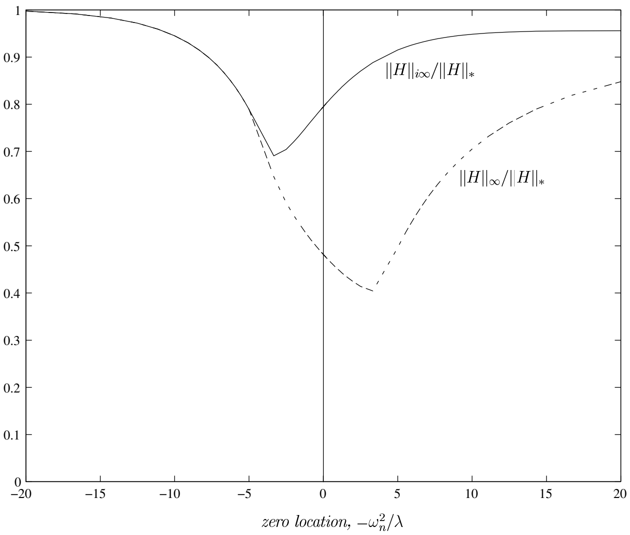

Figure 1: Ellipsoidal bound performance versus zero location for example 1

performance occurs when the zero is closer to the axis than the pole at -3.82. The same trend is observed for nonminimum phase zero locations.

These results are not surprising based on intuition that might be found in an undergraduate text on classical control systems. For systems described by (8), it is well-known that transient behavior is highly dependent on zero locations. To gain insight into how the ellipsoidal bound is affected, write a state space realization of (8) as

$$\dot{x} = \begin{bmatrix} -2\zeta\omega_n & -\omega_n^2 \\ 1 & 0 \end{bmatrix} x + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u \tag{9}$$

$$y = \begin{bmatrix} \lambda & \omega_n^2 \end{bmatrix} x \tag{10}$$

For fixed ω_n and ζ, the solutions Q_α to (4) vary with α, but are independent of λ. On the other hand, variation in λ over $\mathbb{R}$ causes $180°$ rotation of the output vector. Apparently, the phase variation associated with zeros close to the imaginary axis is related to poor alignment between the output vector and the (fixed) family of bounding ellipsoids $\{Q_\alpha\}$.

Figure 1 also depicts the lower bound on the conservativeness of the $*$-norm from (6). Clearly, the bound is quite good while until the zero moves close to the $j\omega$-axis, when the bound deteriorates quickly. The lower bound is also observed to be poor for nonminimum phase zero locations. This trend was observed in general.

As a final remark, we note that for $\lambda > 0$, the system in (8) is passive. Thus, the ellipsoidal bound can be conservative even for passive systems, which are considered to have good peaking properties [7, 9].

4.2 Example 2: Sandwiched pole/zero configurations

In this example, we continue to investigate the notion that the ellipsoidal bound can be poor when a zero lies close to the $j\omega$-axis relative to the poles of the system. To emphasize this case, we consider a systems with sandwiched pole/zero configurations which consist of simple zeros sandwiched by a pair of simple poles except, possibly, for the zero closest to the $j\omega$-axis.

First, consider

$$H(s) = \frac{\prod_{k=1}^{m}(s/\lambda^{f_{3k-2}}+1)}{\prod_{k=1}^{m}(s/\lambda^{f_{3k-1}}+1)\prod_{k=1}^{m-1}(s/\lambda^{f_{3k}}+1)} \tag{11}$$

for real $\lambda > 1$, fixed positive integer, m, and a sequence $\{f_i\}$ defined by

$$f_1 = 1; \qquad f_i = f_{i-1} + \nu, \;\; i \geq 2 \tag{12}$$

with $\nu > 0$. Equation (11) characterizes a sandwiched pole/zero configuration where the zero closest to the $j\omega$-axis which is sandwiched by a simple pole and the $j\omega$-axis. Using intuition gained in example 1, we expect the ellipsoidal bound to be conservative. Such pole/zero configurations were used in [10] to demonstrate extreme behavior between optimal $\mathcal{L}_1$ and $\mathcal{H}_\infty$ controllers, so we expect the lower bound from (6) to be poor as well.

For a fixed m and λ, the spacing of the system time constants can be controlled by ν. Note that the pole and zero closest to the $j\omega$-axis remain fixed for all values of ν. Figure 2 illustrates the conservativeness of the $*$-norm as a function of ν for $m = 3$ and $\lambda = 3$. Performance degrades as ν increases, corresponding to greater distance between the zero at $-\lambda = -1$ and the pole at $-\lambda^{f_1} = -(3^{\nu+1})$. The lower bound from (6) also deteriorates as ν increases. These observations therefore reinforce the intuition gained from the previous example.

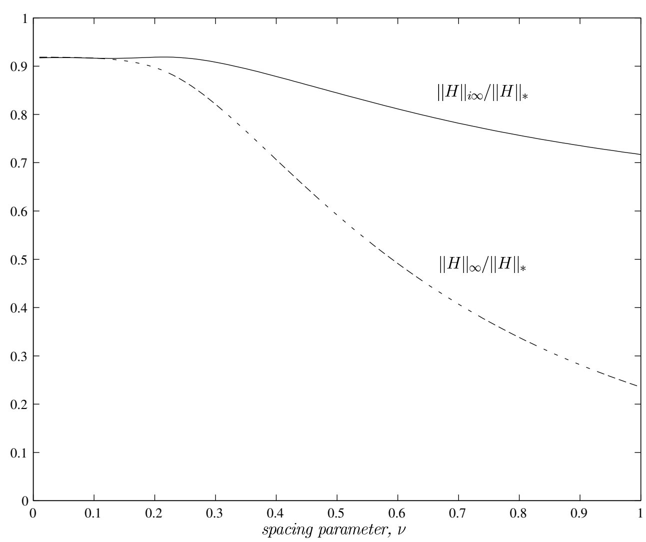

Figure 2: Ellipsoidal bound performance for sandwich configuration with zero closest to $j\omega$-axis.

Now, consider a similar pole/zero configuration with

$$H(s) = \frac{\prod_{k=1}^{m}(s/\lambda^{f_{3k-1}}+1)}{\prod_{k=1}^{m}(s/\lambda^{f_{3k-2}}+1)\prod_{k=1}^{m-1}(s/\lambda^{f_{3k}}+1)} \tag{13}$$

and f_i as defined in (12). Here, the zero closest to the $j\omega$-axis is sandwiched by a pair of simple poles, so we expect the ellipsoidal bound to improve over the previous sandwich configuration. Figure 3 illustrates the conservativeness of the $*$-norm as a function of ν for $m = 3$ and $\lambda = 3$. As expected, the ellipsoidal bound is quite good now that the zero is no longer close to the $j\omega$-axis relative to the system poles. We also observe that the ∞-norm bound from (6) is tight for all values of ν.

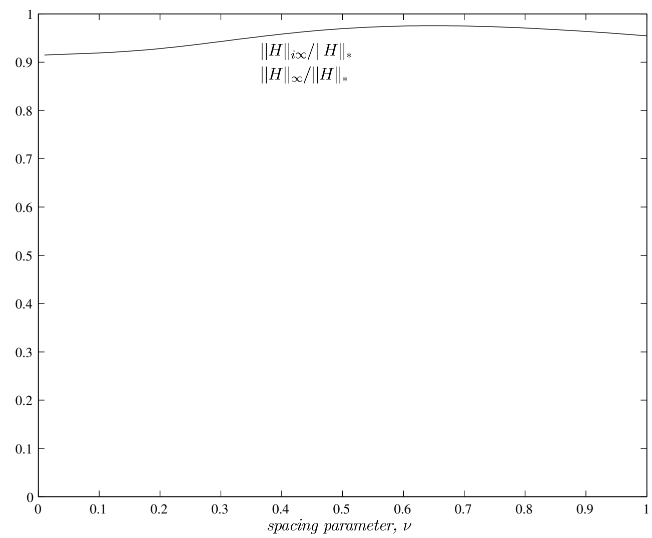

Figure 3: Ellipsoidal bound performance for sandwich configuration with pole closest to $j\omega$-axis.

4.3 Example 3: The Venkatesh example

We now consider the class of systems used in [11] to demonstrate arbitrary conservativeness of the ellipsoidal bound. Let

$$H(s) = \frac{1}{\sqrt{\lambda}} \frac{s-\lambda}{(s+\lambda)(s+\frac{1}{\lambda})} \tag{14}$$

for real $\lambda > 0$. For $\lambda < 1$, H contains a nonminimum phase zero close to the $j\omega$-axis relative to the system poles. Based on previous examples, we expect the ellipsoidal bound to be poor for $\lambda < 1$. Figure 4 shows the performance of the ellipsoidal bound and the associated ∞-norm bound on conservativeness from (6).

Clearly, Figure 4 supports the claim that the $*$-norm can be arbitrarily conservative. Note, however, that the bound is only conservative for small values of λ ($\lambda < 1$), as expected. For $\lambda > 1$, the zero is far from the imaginary axis relative to the pole at $-1/\lambda$ and the $*$-norm bound is actually quite good. Following the trend observed in example 1, the ∞-norm bound from (6) is seen to be conservative when the nonminimum phase zero location has the same or lower order of magnitude as the pole locations of the system.

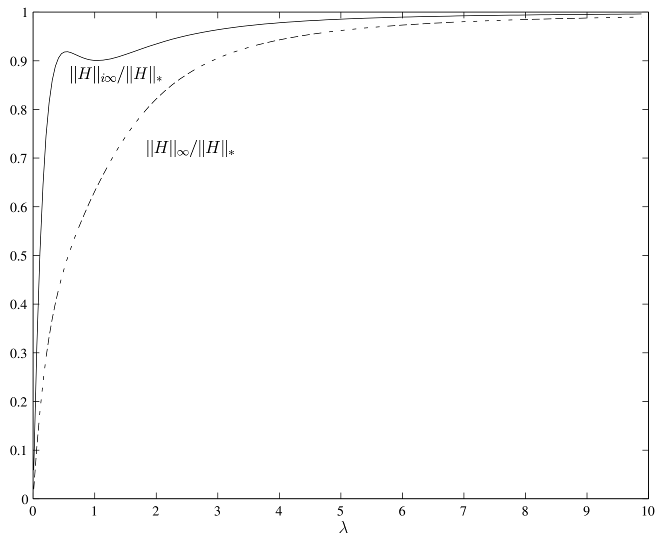

Figure 4: Ellipsoidal bound performance for Venkatesh example

We remark that the class of systems described by (14) is important since it represents a stable system with a first-order Taylor model of a time delay. In light of this observation, a class of systems for which the ellipsoidal bound is conservative can also be seen as systems with time delays which dominate the dynamical behavior of the system. Such systems require special care, so it is no surprise that the simple ellipsoidal bound can be quite conservative here.

4.4 Example 4: Allpass Systems

In example 3, we observed that allpass terms can result in poor ellipsoidal bounds, especially when the allpass terms are at low frequencies relative to the remaining system dynamics. In this example, we show that additional, widely spaced allpass terms can further degrade the ellipsoidal bounds. Consider the allpass system

$$H(s) = \frac{1}{(s/\lambda^{f_{m+1}} + 1)} \prod_{k=1}^{m} \frac{(s - \lambda^{f_k})}{(s + \lambda^{f_k})} \tag{15}$$

for real $\lambda > 1$, fixed positive integer, m, and a sequence $\{f_i\}$ as defined in (12). Note that we have modified the system with an additional lowpass term at high frequencies to make the system strictly proper. For a fixed m and λ, the spacing of the allpass terms can be controlled by ν. Figure 5 shows that the ellipsoidal bound becomes extremely conservative as the spacing (ν) increases. Note that the ∞-norm bound is poor for all values of ν.

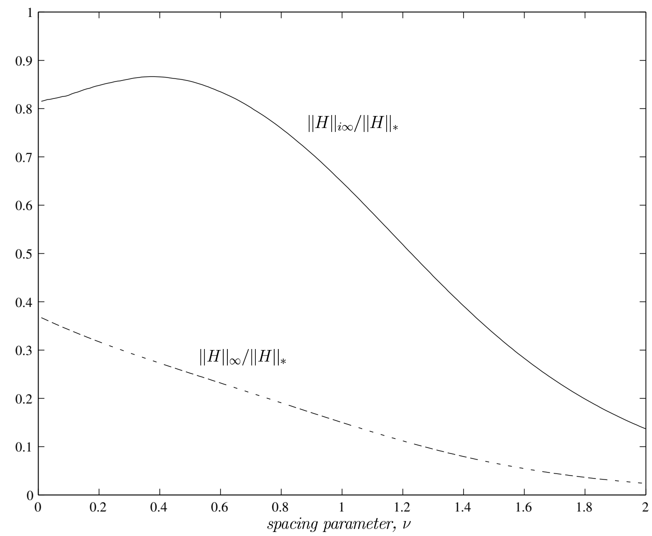

Figure 5: Ellipsoidal bound performance for versus allpass term spacing in example 4

4.5 Example 5: Repeated poles

Finally, we demonstrate that systems with repeated poles can also exhibit conservative $*$-norm bounds. Consider

$$H(s) = \prod_{k=1}^{m} \frac{1}{(s/\lambda + 1)} \tag{16}$$

for fixed real $\lambda > 0$ and positive integer, m. The transfer function in (16) describes systems with poles of multiplicity m at $-\lambda$. Note that $\|H\|_\infty = \|H\|_{i\infty} = 1$, so that the bound from (6) is tight. Ellipsoidal bound performance versus multiplicity is shown in Figure 6.

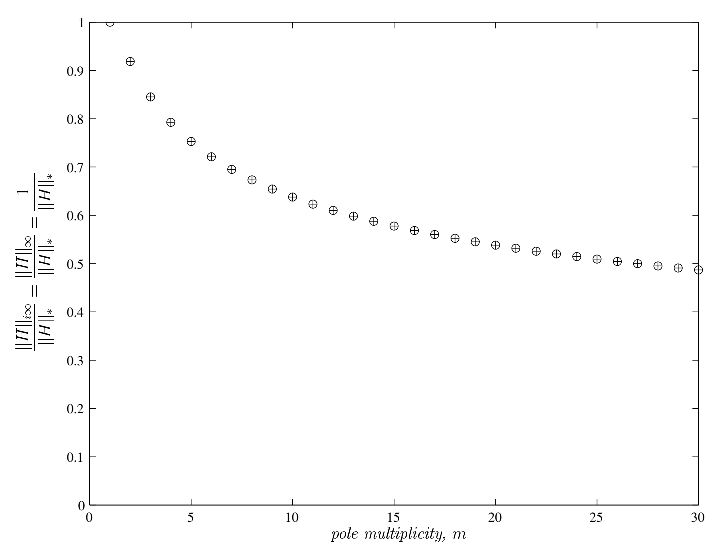

Figure 6: Ellipsoidal bound performance versus pole multiplicity

Figure 6 demonstrates that performance degrades with increased multiplicity. We remark that the dip in Figure 4 seen at $\lambda = 1$ alludes to a similar conclusion. This result is likely due to the fact that repeated poles can lead to defective eigensystems [12] (i.e. systems of generalized eigenvalues) which will affect solution of the Lyapunov equation in (4). Similar notions appear in other areas of control theory, e.g. pole placement algorithms.

5 Discussion

The examples presented in the previous section provide empirical evidence for cases where ellipsoidal methods used to bound peak-to-peak system gains may be conservative. The results motivate engineering guidelines warranting additional or alternative analysis efforts if the system under study falls into a class described in this report. This additional effort will appear in either a more detailed investigation of system performance, or in the use of more complicated analysis methods, such as the ℓ_1 approach [5].

While the results consider system analysis only, parallels can likely be drawn for control synthesis. Because conservativeness of the ellipsoidal approach has been demonstrated in practical problems, the methods are more likely to be useful in mixed performance problems. In other words, the $*$-norm should not be used as performance criteria without additional constraints even if the $*$-norm serves as an upper bound on those constraints, e.g. the $\mathcal{H}_\infty$ norm. Such problems actually occur more naturally since disturbance inputs and uncertainty are rarely caputured by a single mathematical model. These ideas are stated in [5, 10], in the context of ℓ_1 methods. Note, however, that the $*$-norm approach appears to be more amenable to mixed performance problems due to the LMI nature of analysis and synthesis constraints.

An important question arises in characterizing the conservativeness of the $*$-norm as a bound on the peak-to-peak gain of the system. In this report, we investigated the application of the fundamental bound stated in (6) to this problem. The numerical evidence included here suggests that (6) is only useful if $\|H\|_\infty/\|H\|_*$ is close to unity, proving that the $*$-norm is nonconservative. Otherwise, as the examples indicate, the ratio can be a poor approximation to $\|H\|_{i\infty}/\|H\|_*$. Unfortunately, it is unclear if one can do better than this ratio, short of numerical integration which will become extremely inefficient for MIMO systems of high dimension.

As the examples indicate, performance of the $*$-norm bound depends heavily on zero locations. The insights revealed by the numerical examples rely on the fact that only SISO systems were considered. Generalizing these notions to MIMO systems is nontrivial, however, because the concept of zeros becomes more complicated. To generalize the insights obtained here as well as provide a quantifiable test for conservativeness, we might study the relationship between the eigenstructure of the system and the eigenstructure of the zero dynamics [6]. In light of the evidence presented in this report, the $*$-norm will be conservative for systems which posess dominant and/or unstable zero dynamics.

Finally, we remark that evaluation of the $*$-norm amounts to finding the controllability grammian, Q_α, of the matrix pair $(A + \frac{1}{2}\alpha I, \frac{1}{\sqrt{\alpha}}B)$, i.e.

$$Q_\alpha = \int_0^\infty \frac{e^{\alpha t}}{\alpha} e^{At} BB' e^{A't} dt$$

The invariant ellipsoid bounding the system states is therefore intimately related to the way system inputs map to left eigenvectors of A. On the other hand, (5) implies that

$$\|H\|_* = \inf_{\alpha \in (0,\kappa)} \int_0^\infty \frac{e^{\alpha t}}{\alpha} C e^{At} BB' e^{A't} C' dt$$

which relates the right eigenvectors of A to the outputs of the system. This simple interpretation suggests that clues to the conservativeness of the $*$-norm lie in the interaction of the input and output matrices with the eigenstructure of A.

6 Conclusions

This report has described, in empirical fashion, cases where invariant ellipsoid based bounds on the peak-to-peak gain of an LTI can be conservative. Numerical evidence was presented to suggest that the best ellipsoidal bound (the $*$-norm) can be poor in the presence of dominant and/or unstable zero dynamics and/or repeated poles. In parallel, the report has investigated the use of the system ∞-norm to evaluate the conservativeness of the $*$-norm via

$$\frac{\|H\|_\infty}{\|H\|_*} \leq \frac{\|H\|_{i\infty}}{\|H\|_*} \leq 1$$

We conclude that unless the bound proves that the $*$-norm is nonconservative, it offers little more than the intuition.

References

[1] J. Abedor, K. Nagpal, and K. Poolla. A linear matrix inequality approach to peak-to-peak gain minimization. *International Journal of Robust and Nonlinear Control*, 6:899–927, 1996.

[2] C.J. Bett. *Bounded Amplitude Multiple Agent Control.* PhD thesis, University of Notre Dame, Notre Dame, IN, 1998.

[3] S. Boyd and J. Doyle. Comparison of peak and RMS gains for discrete-time systems. *Systems & Control Letters*, 9:1–6, 1987.

[4] S. Boyd, L. El Gaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory.* Society for Industrial and Applied Mathematics, 1994.

[5] M.A. Dahleh and I.J. Diaz-Bobillo. *Control of uncertain systems: a linear programming approach.* Prentice-Hall, 1995.

[6] A. Isidori. *Nonlinear Control Systems.* Springer-Verlag, 2nd edition, 1989.

[7] P. Kokotović. The joy of feedback: Nonlinear and adaptive. *IEEE Control Systems Magazine*, 12(3):7–17, June 1992.

[8] C. Scherer, P. Gahinet, and M. Chilali. Multiobjective output-feedback control via LMI optimization. *IEEE Transactions on Automatic Control*, 42(7):896–911, July 1997.

[9] H.J. Sussmann and P.V. Kokotović. The peaking phenomenon and the global stabilization of nonlinear systems. *IEEE Transactions on Automatic Control*, 36(4):424–440, April 1991.

[10] S.R. Venkatesh and M.A. Dahleh. Examples of extreme cases in ℓ_1 and H_∞ optimization. *Systems & Control Letters*, 23:229–236, 1994.

[11] S.R. Venkatesh and M.A. Dahleh. Does star norm capture ℓ_1 norm? In *Proceedings of the American Control Conference*, pages 944–945, Seattle, WA, June 1995.

[12] J.H. Wilkinson. *The Algebraic Eigenvalue Problem.* Oxford University Press, 1965.