

A New Bound of the $\mathcal{L}_2[0, T]$ -Induced Norm and Applications to Model Reduction¹

M. Sznaier*, A.C. Doherty†, M. Barahona†, H. Mabuchi†, J.C. Doyle‡

*Electrical Engineering, The Pennsylvania State University, University Park, PA 16802

†Department of Physics and ‡Control & Dynamical Systems,

California Institute of Technology, Pasadena, CA 91125

Abstract

We present a simple bound on the finite horizon $\mathcal{L}_2[0, T]$ -induced norm of a linear time-invariant (LTI), not necessarily stable system which can be efficiently computed by calculating the \mathcal{H}_∞ norm of a shifted version of the original operator. As an application, we show how to use this bound to perform model reduction of unstable systems over a finite horizon. The technique is illustrated with a non-trivial physical example relevant to the appearance of time-irreversible phenomena in statistical physics.

1 Introduction

Many problems of practical interest entail bounding the induced \mathcal{L}_2 norm of a not necessarily stable system over a finite time interval $[0, T]$. Examples include the worst-case identification of unstable plants [8], and obtaining low-order approximations to Hamiltonian systems common in physics such as large collections of coupled oscillators [1].

In principle, such a bound can be obtained by considering a slight extension of the Bounded Real Lemma, leading to a differential Riccati inequality or, equivalently, to a differential Linear Matrix Inequality. However, these differential inequalities are computationally expensive to solve, except in low dimensional cases.

In this paper, we present a simple new bound that can be readily obtained from the \mathcal{H}_∞ norm of a shifted version of the system. We can then use this bound to develop a framework for model reduction, over a finite interval, of non-Hurwitz LTI systems.

¹This work was supported in part by NSF, under grants ECS-9907051, ECS-0115946 and Caltech's Institute for Quantum Information, and by AFOSR, under MURI "Uncertainty Management in Complex Systems" and grant F49620-00-1-0020.

This framework is illustrated with an example from statistical physics showing that incomplete observation of microscopically time-reversible conservative laws can consistently lead to apparently irreversible macroscopic behavior.

The paper is organized as follows. Section 2 introduces notation and some preliminary results. In section 3 we present the derivation of the new bound. In Section 4 we illustrate how to apply this bound to model reduction of non-Hurwitz LTI systems. Finally, Section 5 contains some concluding remarks and points to some open research directions.

2 Preliminaries

2.1 Notation

\mathcal{L}_∞ denotes the Lebesgue space of complex valued matrix functions essentially bounded on the $j\omega$ axis, equipped with the norm $\|G(s)\|_\infty \doteq \text{ess sup}_\omega \bar{\sigma}(G(j\omega))$, where $\bar{\sigma}$ is the largest singular value. \mathcal{H}_∞ denotes the subspace of functions in \mathcal{L}_∞ with a bounded analytic continuation in $\mathcal{R}(s) \geq 0$.

$\mathcal{L}_2[0, T]$ denotes the space of vector valued real functions essentially bounded in the interval $[0, T]$, equipped with the norm $\|f\|_{\mathcal{L}_2[0, T]}^2 \doteq \int_0^T f'(t)f(t)dt$. Let \mathcal{L} represent the space of LTI, causal, bounded operators in $\mathcal{L}_2[0, T]$. The induced norm of an operator $M \in \mathcal{L}$ is given by

$$\|M\|_{\mathcal{L}_2[0, T], \text{ind}} \doteq \sup_{\|u\|_{\mathcal{L}_2[0, T]} \neq 0} \frac{\|Mu\|_{\mathcal{L}_2[0, T]}}{\|u\|_{\mathcal{L}_2[0, T]}}$$

It is a standard fact that the $\mathcal{L}_2[0, \infty)$ -induced norm of a LTI stable operator G coincides with the peak value of its frequency response, i.e.,

$$\|G\|_{\mathcal{L}_2[0, \infty), \text{ind}} = \|G\|_\infty.$$

Given a stable, finite dimensional operator $G: \mathcal{L}_2[0, \infty) \rightarrow \mathcal{L}_2[0, \infty)$ with a state space

realization $G = \left(\begin{array}{c|c} A & B \\ \hline C & \end{array} \right)$, its associated Hankel operator $\Gamma_G: \mathcal{L}_2(-\infty, 0] \rightarrow \mathcal{L}_2[0, \infty)$ is defined as:

$$\Gamma_G v = \begin{cases} \int_0^\infty C e^{A(t+\tau)} B v(\tau) d\tau & t \geq 0 \\ 0 & t < 0. \end{cases}$$

This operator can be thought of as mapping past inputs in $(-\infty, 0]$ to the corresponding output in $[0, \infty)$.

Let $\Gamma_G^*: \mathcal{L}_2[0, \infty) \rightarrow \mathcal{L}_2(-\infty, 0]$ denote the adjoint operator of Γ_G . The Hankel singular values σ_i^H of G are defined as the square roots of the eigenvalues of the operator $\Gamma_G^* \Gamma_G$. A well known result (see for instance [6], Chapter 6) establishes that these eigenvalues coincide with the eigenvalues of $W_c W_o$, the product of the controllability and observability Gramians of G .

2.2 Preliminary Results

In this section, we provide some preliminary results that relate the \mathcal{L}_2 -induced norm of a finite dimensional LTI operator to the existence of positive definite solutions to an algebraic or a differential matrix inequality. These results will be used to obtain a bound on the $\mathcal{L}_2[0, T]$ -induced norm.

Lemma 1 (Bounded Real) Consider a strictly proper, finite dimensional, LTI, stable system G with state space realization:

$$G = \left(\begin{array}{c|c} A & B \\ \hline C & \end{array} \right).$$

Then the following statements are equivalent:

1. The $\mathcal{L}_2[0, \infty)$ -induced gain is bounded by $\gamma > 0$: $\|G\|_{\mathcal{L}_2[0, \infty), \text{ind}} < \gamma$
2. The following LMI admits a positive definite solution $X > 0$:

$$\begin{bmatrix} A'X + XA + C'C & XB \\ B'X & -\gamma^2 I \end{bmatrix} < 0 \quad (1)$$

Lemma 2 (Bounded Real, Differential version) Consider a strictly proper, finite dimensional, not necessarily stable LTI system G with state space realization:

$$G = \left(\begin{array}{c|c} A & B \\ \hline C & \end{array} \right).$$

Assume that the following differential matrix inequality admits a positive definite solution

$X(t), \forall t \in [0, T]$:

$$\begin{bmatrix} A'X + XA + \dot{X} + C'C & XB \\ B'X & -\gamma^2 I \end{bmatrix} < 0. \quad (2)$$

Let $u \in \mathcal{L}_2[0, T]$ denote an arbitrary input and z the corresponding output. Then the following holds:

$$\int_0^T z' z dt < \gamma^2 \int_0^T u' u dt \quad (3)$$

Corollary 1 If the inequality (2) holds, then $\|G\|_{\mathcal{L}_2[0, T], \text{ind}} < \gamma$.

3 A Simple Bound of the $\mathcal{L}_2[0, T]$ -induced norm

In this section, we show that a bound on the $\mathcal{L}_2[0, T]$ -induced norm can be obtained by simply computing the \mathcal{H}_∞ norm of a shifted version of the system under consideration.

Formally, the $\mathcal{L}_2[0, T]$ -induced norm of a given LTI operator G is equivalent to the $\mathcal{L}_2[0, \infty)$ induced norm of a time-varying system with convolution kernel:

$$G(t, \tau) \doteq W(t)G(t - \tau), \quad (4)$$

where $W(\cdot)$ is a step window of the form

$$W(t) = \begin{cases} 1 & 0 \leq t \leq T \\ 0 & \text{otherwise} \end{cases}$$

However, as mentioned above, there are no computationally efficient procedures for performing this calculation. To circumvent this difficulty, one can approximate the step window $W(t)$ by an exponential window of the form e^{-at} , where the time constant a is such that $e^{-at} \ll 1$ for $t > T$. The advantage of this approach is that the resulting kernel, $e^{-at}G(t)$ can be associated with a new LTI operator, whose frequency response is a shifted version of the frequency response of the original system. If a is chosen such that this new LTI operator is stable, computing its $\mathcal{L}_2[0, \infty)$ -induced norm (i.e., its \mathcal{H}_∞ norm) is now a standard problem. These observations motivate the following result:

Theorem 1 Consider a strictly proper, finite dimensional, LTI, (not necessarily stable system) G with state space realization

$$G = \left(\begin{array}{c|c} A & B \\ \hline C & \end{array} \right).$$

If there exists $a > 0$ such that

$$G_a = \left(\begin{array}{c|c} A - aI & B \\ \hline C & \end{array} \right)$$

is stable, with $\|G_a\|_\infty < \gamma$, then the following bound holds:

$$\|G\|_{\mathcal{L}_2[0,T], \text{ind}} < \gamma e^{aT}. \quad (5)$$

Proof: Since by hypothesis $\|e^{aT}G_a\|_\infty < \gamma e^{aT}$, it follows from Lemma 1 that there exists $X_a \geq 0$ such that:

$$\begin{bmatrix} A'_a X_a + X_a A_a + e^{aT} C' C e^{aT} & X_a B \\ B' X_a & -\gamma^2 e^{2aT} I \end{bmatrix} < 0 \quad (6)$$

where $A_a = A - aI$. Define $X(t) = e^{-2at} X_a$ and consider $t \in [0, T]$. Multiplying (6) by $e^{-2at} I_{(n+m) \times (n+m)}$ we have:

$0 >$

$$\begin{bmatrix} A'X + XA + \dot{X} + e^{2a(T-t)} C' C & XB \\ B' X & -\gamma^2 e^{2a(T-t)} I \end{bmatrix} >$$

$$\begin{bmatrix} A'X + XA + \dot{X} + C' C & XB \\ B' X & -\gamma^2 e^{2aT} I \end{bmatrix}$$

where the last inequality relies on $a > 0$ and $t \leq T$. The proof now follows immediately from Lemma 2.

4 Application: Model Reduction of Unstable Systems

In this section, we use the bound (5) to solve the problem of model reduction of non-stable systems over a finite horizon. This problem is relevant, for instance, in the context of classical physics, where fundamental models arise from Hamiltonians with all the eigenvalues of A purely imaginary. In this framework, the algorithm developed here, with upper bounds on the approximation error, provides an alternative to uncontrolled, formal approximation procedures currently used in the physics literature.

4.1 A Simple Algorithm for Model Reduction in $\mathcal{L}_2[0, T]$

Consider a finite dimensional LTI system $G(s)$ with McMillan degree n . A well known result [4] states that a rank r approximation G_r to a stable G can be obtained by considering a balanced realization of $G(s)$ and discarding the states associated with the smallest $n - r$ Hankel singular values $\sigma_i^H, i = r + 1, \dots, n$. The corresponding approximation error is then bounded by $\|G - G_r\|_\infty \leq 2 \sum_{k=r+1}^n \sigma_k^H$.

On the other hand, few results are available for the case where G is not necessarily Hurwitz and it is only desired to find approximations (in the \mathcal{L}_2 -induced sense) over a finite time interval. Model reduction of unstable systems in the \mathcal{L}_∞ (Hankel) sense has been addressed in [9]. However, contrary to the stable case, \mathcal{L}_∞ approximation error bounds do not lead to bounds in the $\mathcal{L}_2[0, T]$ sense when G is non Hurwitz. In principle, model reduction in the $\mathcal{L}_2[0, T]$ sense can be accomplished by treating the system as LTV (setting $A(t), B(t), C(t), D(t) = 0$ for $t > T$) and using the results in [7] or [5]. However, they both pose practical difficulties: the former requires solving two differential Lyapunov inequalities subject to an additional constraint on the structure of the solutions, while the latter applies to discrete-time systems, with no continuous-time counterpart presently available.

These problems can be circumvented using the following simple algorithm, motivated by Theorem 1:

Algorithm 1

- 0.- Take as inputs a state space realization $G(s) = C(sI - A)^{-1}B + D$, and a number $a \in \mathcal{R}^+$ such that $G(s + a) \in \mathcal{H}_\infty$.
- 1.- Find a stable reduced order approximation $G_{r,a} \doteq \left(\begin{array}{c|c} A_r & B_r \\ \hline C_r & D_r \end{array} \right)$ to the shifted system $G_a \doteq \left(\begin{array}{c|c} A - aI & B \\ \hline C & D \end{array} \right)$.
- 2.- Use the system $G_r \doteq \left(\begin{array}{c|c} A_r + aI & B_r \\ \hline C_r & D_r \end{array} \right)$ as an approximation to G in the interval $[0, T]$.

Remark 1 Since G_a is stable, an approximation $G_{r,a}$ can be readily obtained using standard model reduction techniques. In particular, using balanced truncations leads to the error bound: $\|G_a - G_{r,a}\|_\infty \leq 2 \sum_{i=r+1}^n \sigma_{a,i}^H$, where $\sigma_{a,i}^H$ denotes the Hankel singular values of G_a (ordered in decreasing order). Now, from Theorem 1 it follows that

$$\|G - G_r\|_{\mathcal{L}_2[0,T], \text{ind}} \leq 2e^{aT} \sum_{i=r+1}^n \sigma_{a,i}^H. \quad (7)$$

Remark 2 From equation (7) it follows that the proposed algorithm gives error bounds comparable to those of the stable LTI case when $a \sim 1/T$. An intuitive explanation of this relationship is given in the next section.

4.2 A Nontrivial Example: Coupled-oscillator Models of a Thermal Bath

To illustrate its usefulness, we have applied our general framework to a widely studied model (see for example [2, 10]) that describes the dynamics of a harmonic oscillator coupled to a thermal bath— itself modelled as a very large collection of harmonic oscillators with a distribution of frequencies. This model, shown schematically in Figure 1(a), has been proposed as a particularly simple example of the means by which microscopic conservative, time-reversible laws can give rise to irreversible, dissipative macroscopic behavior [2, 10].

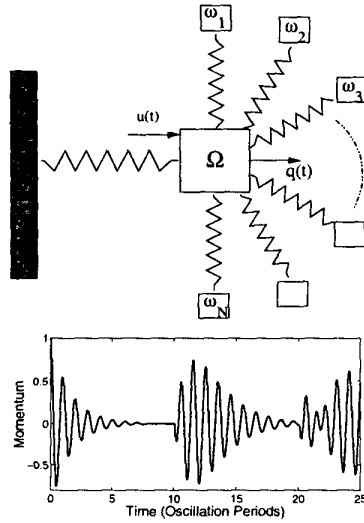


Figure 1: (top) Harmonic oscillator in a thermal bath. (bottom) Impulse response for $\Omega = 1, \gamma = 0.1, N = 100, \omega_c = 10$.

A state-space realization of this model is given by:

$$\begin{aligned} \dot{q} &= p/m \\ \dot{p} &= -m\Omega^2 q + \sum_{i=1}^N g_i (\omega_i q_i - g_i q/m_i) + u(t) \\ \dot{q}_i &= p_i/m_i \\ \dot{p}_i &= -m_i \omega_i^2 q_i + g_i \omega_i q \\ y &= p, \end{aligned} \quad (8)$$

where the variables q, p, q_i, p_i describe, respectively, the position and momentum of the harmonic oscillator with frequency Ω , and of the N bath oscillators with frequencies ω_i . The generality of (8) resides in the freedom to choose the distribution of ω_i and the coupling strengths g_i . Here, we consider a standard ohmic bath with frequencies evenly spaced in the interval $(0, \omega_c]$, such that $\omega_i = i\omega_c/N$

with ω_c a cut-off frequency, and with constant coupling $g_i^2/m_i = 4m\Gamma\omega_c/\pi N$. In this case, standard results [2, 10] establish that in the limit of $N, \omega_c \rightarrow \infty$, the harmonic oscillator should behave like a damped oscillator with frequency Ω and damping rate Γ . However, it is worth emphasizing that most approximation techniques through which this result is derived in the physics literature are *ad hoc*, with no error bounds and often requiring formal limits with some infinite scaling.

In contrast, we approach this problem as that of obtaining low-order approximations, over a finite time horizon, of an input-output system with high-dimensional, linear Hamiltonian dynamics. Indeed, the impulse response shown in Figure 1(b) strongly suggests that the observed dynamics of the system at early times may be very close to that of a damped oscillator even for a moderate number of bath modes. However, since the model is Hamiltonian, the corresponding matrix A has all its eigenvalues on the $j\omega$ -axis, and thus standard infinite horizon model reduction techniques cannot be directly applied. On the other hand, any $a > 0$ renders the shifted system A_a stable. This allows the application of Theorem 1, bound (7), and Algorithm 1 to obtain reduced order models of (8), with guaranteed approximation error bounds, over a finite time horizon T .

In Figure 2, we present a plot of the Hankel singular values of the shifted system G_a for several values of N . It shows that the system has only two significant Hankel singular values with the remaining ones typically several orders of magnitude smaller and tailing off rapidly. This is so even for moderate numbers of oscillators, far from any infinite N limiting behavior commonly invoked in the physics literature. In essence, this is an algorithmic derivation of the fact that the best model (in a rigorous sense) of the dynamics over a finite horizon is a damped oscillator. This is corroborated in Figure 3, which shows the approximation error incurred when using a second order (dissipative) approximation in $[0, T]$, and in Figure 4, which compares the impulse responses of system (8) with $N = 100$ modes and its second order approximation. The reduced model we obtain for these parameters has the following state-space realization:

$$G_r = \left(\begin{array}{cc|c} -0.30 & -1.02 & 1.01 \\ 1.02 & 0.10 & -0.05 \\ \hline 1.01 & 0.05 & \end{array} \right) \quad (9)$$

and even for this relatively small number of oscilla-

tors the error over a time horizon T is bounded by

$$\|G - G_r\|_{\mathcal{L}_2[0,T], \text{ind}} \leq 0.12e^{0.1T}. \quad (10)$$




Figure 2: Hankel singular values for the shifted coupled oscillator system. The circles, crosses, plus signs, stars, squares and diamonds correspond to $N = 100, 150, 200, 250, 350, 500$, respectively. The other parameters are $\Omega = 1$, $\Gamma = 0.1$, $a \propto N^{-1/3}$ and $\omega_c \propto N^{1/3}$.

These results are quite independent of the particular choice of parameters. Consistent numerical experiments indicate that the approximation error is small as long as $a \simeq 1/T$. This is intuitive: if the time horizon T is to lead to significant simplification of the dynamics, the weighting $\exp(-aT)$ should ensure that times beyond T do not contribute significantly to the norm. We also investigated the approximation error of the second order reduced models as N, ω_c and T are varied. It is expected that if the above approximation scheme is valid, then by increasing the cut-off frequency and the number of oscillators it should be possible to consider arbitrarily long time horizons. Indeed, this is the case, as shown in Figure 3 where we plot the bounds on the approximation error as $N, T, \omega_c \rightarrow \infty$ such that $\omega_c/Na \rightarrow 0$ and $a\omega_c/\Omega^2$ is held constant. We remark that, by using the bound on the induced norm (5), it is possible to show analytically that, with these scalings, the error in replacing the full system by the second order Langevin equation may be made arbitrarily small.

5 Conclusions

In this paper we have presented a simple bound on the $\mathcal{L}_2[0, T]$ -induced norm and exploited it to perform model reduction of not necessarily Hurwitz systems over a finite time interval. A salient feature of the proposed algorithm is its modest computational complexity, similar to that of obtaining

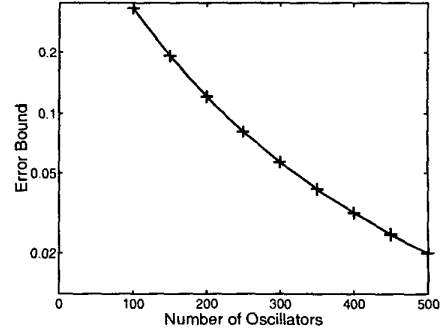


Figure 3: Upper bound on the approximation error as a function of the number of oscillators.




Figure 4: Impulse response of the coupled oscillator system (8) with $\Omega = 1, \Gamma = 0.1, N = 100, \omega_c = 10$. Dashed: full system, Dotted: Low order approximation with $a = 0.1$ ($T = 10$).

reduced order models of stable LTI plants of comparable size.

These results were illustrated with a non trivial physical example showing how microscopic reversible dynamics can lead to macroscopic dissipative-like behavior. From this perspective, the origin of dissipation is no particular mystery: it arises as a parsimonious description of incomplete observations of the dynamics over a finite horizon. Important to the above analysis is that the small number of degrees of freedom of the system that can be driven and observed result in a state space model with strongly observable and strongly controllable subspaces that are nearly orthogonal. Thus, in a systematic search for simple descriptions of finite horizon LTI systems, dissipative dynamics will usually arise, unless conservation of energy is artificially enforced. In fact, such a conservative description will typically be of higher order. Remarkably, the algorithmic results above are robust to changes in the parameters of the system (8) (such as the damping Γ , the cut-off frequency ω_c and the number of bath modes N), the horizon T and the shifting factor a , as long as some general scaling ratios hold.

Indeed, a key feature of our approach is that it offers the possibility of unravelling the dependence of the approximation error on all of those parameters, an analysis hitherto beyond the abilities of currently used statistical physics techniques.

Finally, we would like to briefly comment on some interesting open issues. The first concerns the conservativeness of the bound (5). It can be shown that this bound is tight in some cases. Consider the case of a time delay $G_\tau(s) = e^{-s\tau}$, $\tau < T$. Clearly, $\|G_\tau\|_{\mathcal{L}_2[0,T], \text{ind}} = 1$. On the other hand, for any $a > 0$,

$$e^{aT} \|G_a(s)\|_\infty = e^{aT} \|e^{-a\tau} e^{-s\tau}\|_\infty = e^{a(T-\tau)}$$

and direct application of Theorem 1 yields:

$$\|G\|_{\mathcal{L}_2[0,T], \text{ind}} \leq e^{a(T-\tau)}.$$

Hence the bound can be made arbitrarily tight by taking $\tau \rightarrow T$. In other cases, however, there is a non-zero gap. Consider the case $G(s) = 1/s$. It can be shown that the worst case signal is $u = \cos(\pi t/2T)$ and, therefore, $\|G\|_{\mathcal{L}_2[0,T], \text{ind}} = 2T/\pi$. Now, application of Theorem 1 leads to:

$$e^{aT} \|G_a\|_\infty = e^{aT} \left\| \frac{1}{s+a} \right\|_\infty = \frac{1}{a} e^{aT}.$$

Optimizing over a to obtain the tightest bound yields $aT = 1$ and $e^{aT} \|G_a\|_\infty = T \cdot e$. Therefore, in this case the ratio of the best bound to the actual norm is $\pi e/2 \sim 4.27$. Clearly, it would be of interest to obtain conditions relating the size of the gap to properties of the plant or, alternatively, to identify classes of plants where the bound is tight.

The second open issue is related to finding the value of a that yields the tightest bound for a given plant. Note that from the maximum modulus theorem it follows that $\|G(s)_a\|_\infty$ is a non-increasing function of a . Thus, in general, the bound (5) is the product of an increasing (e^{aT}) and a decreasing function of a and there is usually an optimal value of the parameter. In simple cases such as the integrator considered above, this value can be found by solving an optimization problem. However, it is not known at this point whether this problem is tractable in more complicated cases. A similar situation arises in connection with the bound (7): typically $\sum_{i=\tau+1}^n \sigma_{a,i}^H$ is a decreasing function of a , and thus, for given G and τ there exists a value of a that minimizes the right hand side of the bound.

Finally, in the case of model reduction of stable systems it is well known that the approximation error is

bounded below by $\|G - G_\tau\|_\infty \geq \sigma_{\tau+1}^H$. No comparable lower bound is available at this time for the approximation error when using Algorithm 1 to perform model reduction over a finite horizon.

Acknowledgements

The authors are indebted to Pablo Parrilo, Aristo Asimakopoulou, Henrik Sandberg and Alexandre Megretski for discussions on model reduction of unstable systems and its application to problems in classical physics.

References

- [1] A.C. Doherty, M. Barahona, J. Doyle, H. Mabuchi, and M. Sznaier, "Robustness and Dynamics in the Quantum Classical Transition", Sixth SIAM Conference on Applications of Dynamical Systems, Snowbird, UT, May 23 2001.
- [2] G. W. Ford and M. Kac, "On the quantum Langevin equation," *J. Stat. Phys.*, 46, 803-810, 1987.
- [3] P. Gahinet and P. Apkarian, "A Linear Matrix Inequality Approach to \mathcal{H}_∞ Optimization Problems," *Int. J. Robust and Nonlinear Ctrl.*, 4, 421-448, 1994.
- [4] K. Glover, "All Optimal Hankel Norm Approximations of Linear Multivariable Systems and Their \mathcal{L}^∞ Error Bounds," *Int. J. Control*, 39, 1984, pp. 1115-1193.
- [5] S. Lall, C. Beck, and G. Dullerud, "Guaranteed Error Bounds for Model Reduction of Linear Time-Varying Systems," *Proc. 1998 ACC*, Philadelphia, PA, 1998, pp. 634-638.
- [6] R.S. Sánchez Peña and M. Sznaier, *Robust Systems: Theory and Applications*, Wiley, New York, 1998.
- [7] H. Sandberg and A. Rantzer, "An Upper Error Bound for Model Reduction of Linear Time Variable Systems," *preprint*, 2001.
- [8] M. Sznaier and C. Mazzaro, "Worst Case Identification of Unstable Plants," *submitted*, 2001.
- [9] K. Zhou, G. Salomon, and E. Wu, "Balanced Realization and Model Reduction for Unstable Systems," *Int. J. Robust and Nonlinear Ctrl.*, 9, 183-198, 1999.
- [10] R. Zwanzig, "Nonlinear generalized Langevin equations," *J. Stat. Phys.*, 9, 215-220, 1973