SAND2009-0132C

A Framework for Reduced Order Modeling with Mixed Moment Matching and Peak Error Objectives

Keith R. Santarelli
Electrical and Microsystems Modeling
Sandia National Laboratories
Albuquerque, NM, USA
krsanta@sandia.gov

ABSTRACT

We examine a new method of producing reduced order models for LTI systems which attempts to minimize a bound on the peak error between the original and reduced order models subject to a bound on the peak value of the input. The method, which can be implemented by solving a set of linear programming problems that are parameterized via a single scalar quantity, is able to minimize an error bound subject to a number of moment matching constraints. Moreover, because all optimization is performed in the time-domain, the method can also be used to perform model reduction for infinite dimensional systems, rather than being restricted to finite order state space descriptions. We begin by contrasting the method we present here to two classes of standard model reduction algorithms, namely moment matching algorithms and singular-value-based methods. After motivating the class of reduction tools we propose, we describe the algorithm (which minimizes the L_1 norm of the difference between the original and reduced order impulse responses) and formulate the corresponding linear programming problem that is solved during each iteration of the algorithm. We then show how to incorporate moment matching constraints into the basic error bound minimization algorithm, and present an example which utilizes the techniques described herein. We conclude with some general comments for future work.

1. INTRODUCTION

The study of model order reduction (MOR) is a problem that has pervaded the engineering community for over thirty years. Stated simply, MOR attempts to replace a system description that is deemed “complex” by a simpler, approximate model that still accurately represents the salient features of the original system. The motivation for the inception of MOR tools from a simulation standpoint is clear: problems with fewer components, in general, take less time to simulate, so creating tools which reduce the size of a model without significantly sacrificing accuracy has great potential impact.

Much of the original work in MOR has roots in the systems and control community, with Moore’s work on principle component analysis [14] and Glover’s work on optimal model reduction in the Hankel norm [9] being the basis for a number of model reduction tools that are still used today. Outside of the realm of control, a great deal of attention has been placed on the development of MOR tools for simulation purposes (see, e.g., [2, 4, 5, 7, 8, 15, 16, 18, 19]). The paper by Gugercin et. al. [10] provides a comparison of the performance of several different linear model reduction techniques that are used today.

1.1 MOR for LTI Systems: Moment matching vs. Singular Values

As the focus of this paper revolves around MOR for LTI systems, we briefly review two of the main classes of model reduction methods for LTI systems, along with their associated benefits, as a means of motivating the particular problems and techniques that we investigate here. Two MOR methods for LTI systems that are popular in the literature today are methods which perform *moment matching* of transfer functions, and methods which compute *singular value decompositions* (SVD) of a linear operator that is associated with the state space description of the LTI system undergoing reduction. Moment matching methods operate by constraining either the value of the transfer function or some derivative (moment) of the transfer function to be the same for both the original and reduced order models at a specified set of frequencies (i.e., $G^{(m_{s_l})}(s_l) = G_r^{(m_{s_l})}(s_l)$, $m = 0, 1, \ldots$, $l = 1, 2, \ldots, L$, where $G(s)$ represents the transfer function of the original system, $G_r(s)$ represents the transfer function of the reduced order system, and $s_l \in \mathbb{C}$ represent N complex frequencies to be matched). One advantage of moment matching is that it can be used to preserve key frequency response characteristics between the original and reduced order systems. For instance, moment matching methods can be used to ensure that the DC gain for a reduced order system is the same as in the original system, an important property for systems which are primarily driven by step inputs. A disadvantage of these methods, however, is that, in general, they do not provide bounds on the error between the response of the original system and the response of the reduced order system for arbitrary inputs.

By contrast, SVD-based methods for model reduction do provide bounds on the error between the responses of the original and reduced order systems. Based upon computing the singular values of a joint controllability/observability measure, these methods produce a truncated state space description of the original system to serve as a reduced order approximation. When the inputs of interest are finite power signals, the outputs of the reduced order model are guaranteed to be "close" to the outputs of the original model in the sense that the power in the difference between the original system output and reduced system output is small. While such results provide a notion that the reduced order models are "good" for a wide range of inputs, classical SVD-based methods suffer from the fact that they do not incorporate moment matching constraints into the problem set-up. Hence, if exact matching of certain frequency response properties between the original and reduced order models is critical, SVD-based methods are typically not the method of choice.

If possible, it is clearly desirable to develop MOR tools which can both incorporate moment matching constraints into the reduction problem, and provide error bounds for general classes of inputs. To date, however, results that provide for mixed formulations which incorporate both error bounds and which simultaneously preserve general properties of the frequency response are limited. Phillips et. al. in [17] provide an algorithm which, while not able to preserve moment matching properties explicitly, does provide an SVD-based method that is guaranteed to preserve passivity of the reduced order model. Gugercin et. al. in [11] explain how the solution to a model reduction problem which minimizes the $\mathcal{H}_2$-norm of the corresponding error system is guaranteed to match moments at mirror images of the pole locations of the reduced order model (e.g., $G(-s_l) = G_r(-s_l)$ where $s_l \in \mathbb{C}$ is a pole of the reduced order model $G_r(s)$). This result is limited, however, since the matching frequencies cannot be chosen arbitrarily. Moreover, certain useful frequencies cannot be matched (such as frequencies along the imaginary axis), since the reduced order models are stable and, hence, $\mathrm{Re}\{s_l\} < 0$.

Some recent work by Astolfi in [1] considers a technique which can simultaneously match moments and produce small error bounds via the introduction of a free parameter into the state space description of the corresponding reduction problem. Nevertheless, when attempting to use model reduction tools for the inherent purpose of *simulation*, the error bounds produced by this tool— and the error bounds produced by *all* SVD-based reduction methods—are not the most desirable because of the *way* they measure error. One of the primary motivations of the work we present herein is that error is measured in a manner that is more useful for designers than the standard measures of error. We now present an example to illustrate the main issue along with a proposed resolution.

1.2 Measures of Error: Power vs. Peak Amplitude

Fig. 1 illustrates a hypothetical example where the red signal represents the output of an original full order system and the blue signal represents the output of a reduced order model that was created using an SVD-based technique. The moral of the example is this: an SVD-based method will consider the red and blue responses to be "close" because the power in the difference between the two signals is apparently small (note that the large spike in the full-order signal is very narrow and, hence, contributes very little energy). While such a measure of closeness may be appropriate for certain applications, if the signals depicted in Fig. 1 represent a critical parameter whose value should never exceed 1, then it is clear that the reduced order model does not adequately represent the original model since the response of the full-order

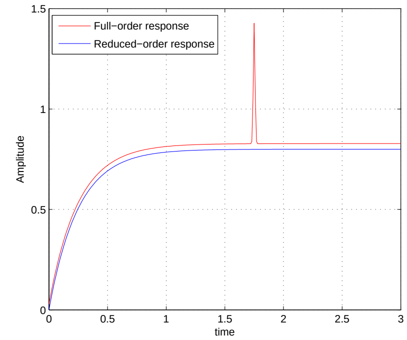

Figure 1: Hypothetical responses of an original and reduced order system produced via an SVD-based method.

system significantly exceeds 1 while the response of the reduced order system stays well below 1.

From a simulation perspective, a somewhat more useful notion of error can be measured in terms of *peak amplitude*. Formally, if we consider right-sided continuous-time signals $y : [0, \infty) \to \mathbb{R}$, then the peak amplitude can be taken as the standard infinity norm:

$$||y||_\infty = \sup_{t \geq 0} |y(t)|. \tag{1}$$

In the context of model reduction, if we define $y(t)$ as the response of an original system and $y_r(t)$ as the response of a reduced order system for an identical input $u(t)$, it is reasonable to desire that $||y - y_r||_\infty$ be a small quantity. Indeed, if for a particular pair $y(t)$ and $y_r(t)$ we define $\epsilon = ||y - y_r||_\infty$, then it immediately follows from the definition in Eqn. 1 that

$$|y(t) - y_r(t)| \leq \epsilon \qquad \forall t \geq 0. \tag{2}$$

1.3 Problem Formulation: L_1 Norm minimization

We now focus on formulating the formal problem to be investigated in this paper. Our focus is limited strictly to LTI systems, for which we wish to develop bounds of the following nature: if we denote by $L^\infty(\mathbb{R}^+)$

$$L^\infty(\mathbb{R}^+) = \left\{ u : [0, \infty) \to \mathbb{R} : \sup_{t \geq 0} |u(t)| < \infty \right\} \tag{3}$$

then for every input $u \in L^\infty(\mathbb{R}^+)$, we wish to find some (hopefully small) real number $M > 0$ such that

$$||y - y_r||_\infty \leq M ||u||_\infty. \tag{4}$$

If such a bound exists for an original system model and a reduced system model for *every* bounded input u, then the peak output of the error between the original and reduced model is always less than some multiple of the peak input value. In particular, due to the assumption of linearity, when $M < 1$, such a bound provides a guarantee that the pointwise error between $y(t)$ and $y_r(t)$ will never be more than a fixed percentage of the peak input value. When we denote by $h(t)$ the impulse response operator of the original system and by $h_r(t)$ the impulse response of the reduced order system, it is a well-known fact (see, for instance, [13]) that the *smallest* value of M as given in Eqn. 4 is the L_1 *norm* of the *error system* with impulse response $h(t) - h_r(t)$:

$$||h - h_r||_1 = \int_0^\infty |h(t) - h_r(t)|dt. \quad (5)$$

Hence, the problem of finding a reduced order model of a given LTI system for which the peak error between the original output and reduced order output is small can be posed in the following manner: for a given order N, find some choice of $h_r(t)$ of order N for which $||h - h_r||_1$ is small. Ideally, one would like to find that choice of $h_r(t)$ of order N such that the quantity $||h - h_r||_1$ is *minimized*, and that is the essential viewpoint that we take here. While the problem of finding that choice of $h_r(t)$ which *globally* minimizes the L_1 norm of the error system is nonconvex and intractable to compute from a practical perspective, we focus here on methods that search for local minimizers over a sufficiently rich set of choices for $h_r(t)$ so as to provide reduced order approximations that are both sufficiently accurate and computationally tractable.

The problem of producing reduced order models via minimization of the L_1 norm appears to have been seldom considered in the literature. El-attar et. al. first considered this problem in the context of some examples [6]. In the discrete-time setting, Sebakhy et. al. consider a simple form of impulse response truncation to minimize the l_1 norm of an error sequence ($||e||_1 = \sum_{k=1}^\infty |e_k|$) [22]. The closest work to the problem we consider here appears to be a result from the System Identification literature in which a reduced order model for a discrete-time system which minimizes the l_1 norm of an error metric is computed via a linear programming approach [12]. While there are substantial differences with the class of problems being considered here as compared to [12], the underlying technique of casting such problems as linear programs is the same. As we discuss in a later section, a major advantage of this approach is that mixed problems in which the L_1 norm of an error system is minimized subject to a set of moment matching constraints can be easily handled by our approach since the set of moment matching conditions can be cast as a set of linear constraints on a set of decision variables. Also, as a byproduct of our approach, the tools we develop here will be able to perform MOR for *infinite dimensional* systems, a stark contrast to standard moment matching and SVD-based tools which operate only on finite order state space descriptions.

The work we present here is a significantly abbreviated version of another document [21] which includes a formal mathematical discussion of associated convergence issues related to the algorithm we describe here, as well as an additional example, and extensions to multi-input, multi-output (MIMO) systems.

2. ALGORITHM FOR REDUCED ORDER MODELING VIA L_1 NORM MINIMIZATION

In this section, we describe a technique for computing reduced order models via an attempt to minimize the L_1 norm of the corresponding error system $h(t) - h_r(t)$. We first consider a relaxed problem in which the reduced order model is constrained to be a linear combination of a fixed set of basis functions and show that this problem can be cast as an LP. We then turn to the process of selecting an appropriate set of basis functions.

2.1 Relaxation: Approximation via a Fixed Basis

At the heart of the algorithm we propose is an approximation scheme where the reduced order model is constrained to be a linear combination of a *fixed* set of functions:

$$h_r(t) = \sum_{k=1}^N a_k g_k(t) \quad (6)$$

where $g_k(t)$, $k = 1, 2, \dots, N$, represent a set of fixed, known functions with finite L_1 norm, and where the parameters $a_k \in \mathbb{R}$ represent a set of decision parameters that we wish to select to make $||h - h_r||_1$ as small as possible. As we show here, this problem can be cast as an LP that can be solved using existing software packages. The reader unfamiliar with linear programming is referred to [3] for an excellent introduction to the subject.

To begin, note that the problem of minimizing $||h - h_r||_1$ is equivalent to:

$$\begin{aligned} \min \quad & \int_0^\infty z(t)dt \quad (7) \\ \text{subject to} \quad & z(t) \geq h(t) - \sum_{k=1}^N a_k g_k(t) \\ & z(t) \geq -\left(h(t) - \sum_{k=1}^N a_k g_k(t)\right) \end{aligned}$$

since the two inequality constraints are equivalent to $z(t) \geq |h(t) - h_r(t)|$, and the choice of $z(t)$ which minimizes the integral expression must achieve this inequality with equality. Note that Eqn. 7 represents an infinite dimensional LP with decision variables a_k and $z(t)$ for all $t \geq 0$. In order to solve this LP, we must resolve two issues: first, the infinite dimensional LP must be replaced by an appropriate finite dimensional LP to fit the form of standard LP solvers. This will be achieved by gridding the real time axis in an appropriate manner. A second issue arises from the fact that the horizon in Eqn. 7 is infinite. In practice, it is possible to solve a finite horizon LP whose optimal solution is an upper bound for the optimal solution of the original infinite horizon problem. We deal with the second of these issues first.

To begin, note that for any $T > 0$

$$\begin{aligned} \int_0^\infty z(t)dt &= \int_0^T z(t)dt + \int_T^\infty z(t)dt \quad (8) \\ &\leq \int_0^T z(t)dt + \int_T^\infty |h(t)|dt + \sum_{k=0}^N |a_k| \int_T^\infty |g_k(t)|dt. \end{aligned}$$

By introducing the slack variables $w_k \geq |a_k|$ for $k = 1, 2, \dots, N$, Eqn. 8 leads to the following LP:

$$\begin{aligned} \min \quad & \int_0^T z(t)dt + \bar{h} + \sum_{k=1}^N \beta_k w_k \quad (9) \\ \text{subject to} \quad & z(t) \geq h(t) - \sum_{k=1}^N a_k g_k(t) \\ & z(t) \geq -\left(h(t) - \sum_{k=1}^N a_k g_k(t)\right) \\ & w_k \geq a_k \\ & w_k \geq -a_k \end{aligned}$$

where T is a specified horizon, $k = 1, 2, \ldots, N$, and where

$$\bar{h} = \int_T^\infty |h(t)| dt \quad (10)$$
$$\beta_k = \int_T^\infty |g_k(t)| dt.$$

By virtue of Eqn. 8, the minimal cost of the LP in Eqn. 9 provides an upper bound for the minimal cost of the original infinite horizon LP of Eqn. 7. Note that for any given choice of $h(t)$, the quantity $\bar{h}$ is a constant, and hence may be removed from the cost function (in practice, T can always be chosen sufficiently large such that the effect of $\bar{h}$ on the minimal cost in Eqn. 9 is negligible).

Now, to relax the infinite dimensional LP to a finite dimensional version, we introduce a grid on the time axis. While there are many ways to do this, here we consider the simplest method of imposing a grid that is uniformly spaced over the horizon length T. If we let Δ represent the sampling interval, and define $z_m = z(m\Delta)$, $h_m = h(m\Delta)$, $g_{km} = g_k(m\Delta)$, and $M = \lfloor T/\Delta \rfloor$, then an approximation of the integral in Eqn. 9 via a Riemann sum leads to:

$$\begin{aligned} \min \quad & \Delta \sum_{m=1}^{M} z_m + \sum_{k=1}^{N} \beta_k w_k && (11) \\ \text{subject to} \quad & z_m \geq h_m - \sum_{k=1}^{N} a_k g_{km} \\ & z_m \geq -\left(h_m - \sum_{k=1}^{N} a_k g_{km} \right) \\ & w_k \geq a_k \\ & w_k \geq -a_k \end{aligned}$$

for all $k = 1, 2, \ldots, N$ and $m = 1, 2, \ldots, M$. Here we assume that the value of Δ is taken sufficiently small (corresponding to a fine grid) so that the difference between the true value of the integral in Eqn. 9 and the approximate value in Eqn. 11 is negligible. As before, the decision variables a_k provide the relative weights for each basis function $g_k(t)$ in our approximation $h_r(t)$, and the auxiliary parameters w_k and z_m determine an upper bound on the minimal L_1 norm to the original problem of Eqn.7.

2.2 L_1 Norm Minimization Algorithm

The LP formulation of the last section begs the question: how does one choose the basis functions $g_k(t)$? The choice we use in this paper is given by the following:

$$g_k(t) = t^k e^{-\alpha t} \quad (12)$$

where α is a parameter with $\mathrm{Re}\{\alpha\} > 0$. Noting that the Laplace transform $G_k(s)$ of $g_k(t)$ is given by

$$G_k(s) = k! \left(\frac{1}{s+\alpha} \right)^{k+1}, \quad (13)$$

we see that choosing the basis $g_k(t)$ of Eqn. 12 corresponds to finding a reduced order model with repeated poles. The basic algorithm that we use, then, is the following: we impose a grid on the right half plane $\mathrm{Re}\{\alpha\} > 0$ and solve the corresponding LP for each value of α in the grid. The value of α which yields the minimal upper bound on $||h - h_r||$ is chosen as the optimal value, and the minimizing values of a_k for this value of α determine the impulse response of the reduced order model.

While the choice of $g_k(t)$ may appear restrictive, we discuss in [21] how a very large subset of impulse responses $h(t)$ can be well-approximated via such a basis for an *arbitrary* value of α with $\mathrm{Re}\{\alpha\} > 0$.

3. ADDITION OF MOMENT MATCHING CONSTRAINTS

We now turn to the incorporation of moment matching constraints into the L_1 minimization algorithm discussed in Section 2. Note that for any fixed basis choice $g_k(t)$, a moment matching constraint of order m at a frequency s_0 takes the form

$$\frac{1}{m!} H^{(m)}(s_0) = \frac{1}{m!} \sum_{k=1}^{N} a_k G_k^{(m)}(s_0) \quad (14)$$

where $H^{(m)}(s)$ and $G_k^{(m)}(s)$ represent the $m-$th derivatives of the Laplace transforms of the original impulse response $h(t)$ and basis functions $g_k(t)$, respectively. The moments of $G_k(s)$ of Eqn. 13 can be calculated explicitly as

$$G_k^{(m)}(s) = \begin{cases} \left(\frac{\alpha}{s+\alpha}\right)^k & m = 0 \\ \frac{(-1)^m i(i+1)\cdots(i+m-1)\alpha^k}{m!} \left(\frac{1}{s+\alpha}\right)^{k+m} & m = 1, 2, \ldots \end{cases} . \quad (15)$$

Hence, whenever the value of the parameter α in Eqn. 12 is fixed, each moment matching constraint is a linear equality constraint on the decision variables a_k and can be added as an additional constraint to the corresponding LP formulation.

4. PRACTICAL CONSIDERATIONS

The success of the above algorithm largely hinges on the ability to select a good value of the parameter α. This is directly correlated to the choice of the grid set, which we denote here via the symbol $\mathcal{A}$. In what follows, we focus on the case where α is a real parameter, though appropriate modifications can be made in the case that α is complex.

Perhaps the simplest way of selecting the set $\mathcal{A}$ is to uniformly grid the real axis:

$$\mathcal{A} = \{\alpha : \alpha = j\Delta_a, \quad j = 1, 2, \ldots, J\} \quad (16)$$

where $\Delta_a > 0$ and J is a user-specified constant. It is clear that smaller choices of Δ_a and larger choices of J provide a finer grid of the real axis and, hence, should produce smaller upper bounds on the minimal value of $||h - h_r||_1$.

While simple, the above brute-force method can be computationally expensive if the user tries to search for a relatively tight upper bound on the minimal value of $||h - h_r||_1$. While the number of LPs which are solved grow linearly with J, it is typical to refine a grid by dividing the value of Δ_a by a particular value, i.e., by replacing Δ_a by $\Delta_a/2$. Assuming even that the maximal value of $\alpha \in \mathcal{A}$ does not increase during grid refinement, this causes the corresponding value of J to grow exponentially with successive refinements, eliminating some of the benefits of parameterizing the basis vectors $g_k(t)$ via a single scalar.

A simple heuristic approach which bypasses some of the above difficulty is the following: initially impose a coarse grid and refine the grid until one is fairly confident that the sampling is sufficiently fine to be indicative of the true behavior of

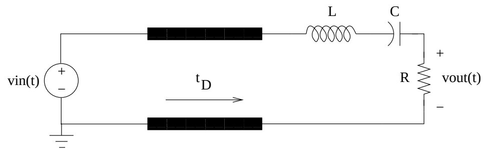

Figure 2: Block diagram of circuit with an RLC bandpass filter and ideal transmission line.

the minimal cost (this can be done, for instance, by examining a graph of the minimal cost at the sample points in the current grid). Once the grid is determined to be sufficiently fine, one can locate an interval around which a minimizing value of α appears to lie and then refine the grid *only* in this interval. When appropriately carried out, such a procedure can only guarantee convergence to *some* local minimum, rather than the minimal value on the interval $I = [\Delta_a, J\Delta_a]$ (the smallest interval containing the original grid). Still, it has been empirically observed in multiple examples that the minimal value of the cost function tends to vary slowly and with few changes in monotonicity, so that carrying out a procedure in this manner is likely to converge to the minimum on I for many problem instances.

5. EXAMPLE: BANDPASS FILTER WITH TIME DELAY

Fig. 2 depicts a circuit consisting of an RLC bandpass filter, along with an ideal transmission line. The transmission line is modeled mathematically as a pure time delay t_D. For a unit time delay, and for the values $R = 2$, $L = 1$, $C = 1/10001$, the input-output transfer function of this circuit is given by

$$H(s) \triangleq \frac{V_{\text{out}}(s)}{V_{\text{in}}(s)} = e^{-s}\frac{2s}{(s+1)^2 + 10,000} \tag{17}$$

with corresponding impulse response

$$h(t) = \begin{cases} e^{-(t-1)}\left(2\cos 100(t-1) - 0.02\sin 100(t-1)\right) & t \geq 1 \\ 0 & t < 1 \end{cases} . \tag{18}$$

It is clear that $H(s)$ is infinite-dimensional due to the presence of the term $\exp(-s)$, and our goal here is to find a finite dimensional approximation with small error norm $||h - h_r||_1$ subject to the additional constraint that the $H(s)$ and $H_r(s)$ match exactly at the resonant frequency of the RLC filter, i.e. that $H(100j) = H_r(100j)$. Because of the highly oscillatory nature of the impulse response, approximating the original $h(t)$ by an approximation whose parameter α is real is unwise, since this would require a very high order polynomial multiplying $\exp(-\alpha t)$ to match $h(t)$ with any reasonable amount of accuracy. Through trial and error, it was quickly discovered that using a value of $\alpha = -\alpha_r + 100j$, where α_r is a positive real parameter, appears to yield the best results, an unsurprising phenomenon since the system naturally oscillates at 100 rad/sec.

Using a 12th order model, we find that a value of $\alpha_r = 3.25$ yields an error norm upper bound of $||h - h_r||_1 \leq 0.297$. The impulse response of the original and reduced order models is shown in Fig. 3. While it is difficult to resolve finer features in this graph, observe qualitatively that $h_r(t)$ is small on the interval $0 \leq t \leq 1$ where $h(t)$ is identically 0, and quickly “catches up” to the oscillatory portion of $h(t)$ for $t \geq 1$.

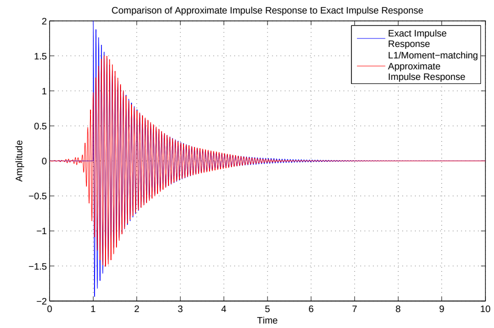

Figure 3: Impulse response of original model and reduced order model produced via mixed L_1/moment matching algorithm.

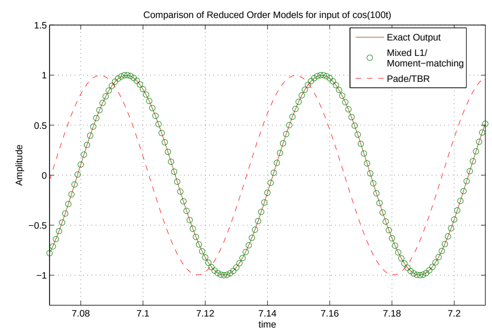

Figure 4: Comparison of responses of original model, L_1/moment matching reduced order model, and Padé/TBR reduced order model for the input $\cos(100t)$.

As a comparison, we created an alternative reduced order model using the following technique: the time delay in Eqn. 17 was approximated via a high order (50th order) Padé approximation, and the resulting system was reduced to a 12th order system using a Truncated Balanced Realization (TBR) algorithm. Since TBR operates on finite-dimensional LTI systems, one must first approximate the non-rational portion by a rational approximation before applying the algorithm.

Performing the above process, we found that the reduced order model obtained via this alternate method produces an L_1 error norm $||h - h_r||_1 = 8.889$, more than an order of magnitude larger than the reduced order model obtained via the mixed L_1/moment matching algorithm. Also for comparison, we computed the response of the original system, reduced order system obtained via L_1/moment matching, and the reduced order system obtained via the Padé approximation and TBR for the input $\cos 100t$. The responses we obtained are depicted in Fig. 4. Observe that the responses of the original model and reduced order model obtained via mixed L_1/moment matching track each other exactly. Such tracking is guaranteed from the moment matching constraint $H(100j) = H_r(100j)$. By contrast, the response of the system produced via the Padé approximation and TBR exhibits a phase lag.

6. CONCLUSION

We have introduced a new framework for model order reduction of LTI systems that is well-suited for simulation purposes. The framework, which can preserve key frequency characteristics of the original model while simultaneously minimizing a bound on the "closeness" of the original and reduced order responses in a point-wise sense, can be implemented efficiently using a relatively small number of user-specified iterations, as demonstrated by two specific examples.

One specific direction in which the work described here can be extended involves investigating methods which produce tighter error bounds. While the bounds provided here are *least* upper bounds in the sense that there always exists an input for which the corresponding difference in outputs will achieve the L_1 norm error bound with equality, such inputs are typically very specific and not encountered in typical application. Methods of constraining the set of inputs to be less "diabolical" via the addition of additional constraints on the input (e.g., bounds on the *derivative* of the input) or via some sort of weighting procedure are desirable.

Finally, while gridding the set $\mathcal{A}$ is not too daunting of a task, it is of interest to examine whether there are intelligent of methods of selecting good values of the parameter α to circumvent gridding entirely. A recent publication [20] examines the use of Laguerre expansions for a model reduction problem which attempts to minimize a quadratic cost function, and computation of the optimal value of α is observed to converge in a process that requires very few iterations. While the problem considered in [20] is significantly different from what we consider here, it would be interesting to investigate whether there exists a class of similar problems which can be used as a heuristic for determining good values of α for the given L_1 setting.

7. REFERENCES

[1] Astolfi A. A new look at model reduction by moment matching for linear systems, *Proc. 46th Conf. Decision and Control*, 2367-2372, New Orleans, 2007.

[2] Bashir O, Willcox K, Van Bloemen Waanders B, Hill J and Ghattas O. Reduced-order model construction for high-dimensional systems, *CSRI Summer Proc. 2006*, 251-263, Sandia National Laboratories, 2006.

[3] Bertsimas D and Tsitsiklis J. *Introduction to Linear Optimization*, Athena Scientific, 1997.

[4] Bond B and Daniel L. Stabilizing schemes for piecewise-linear reduced order models via projection and weighting functions, *Proc. 2007 Int. Conf. Computer-aided Design*, 860-867, San Jose, 2007.

[5] Dong N and Roychowdhury J. General-purpose nonlinear model-order reduction using piecewise-polynomial representations, *IEEE Trans. Computer-aided Design of Integrated Circuits and Systems,* **(27(2)**:249-264, 2008.

[6] El-Attar RA and Vidyasagar M. Order reduction by l_1 and l_∞ norm minimisation, *IEEE Trans Auto. Control*, **AC-23**:731-734, 1978.

[7] Feldmann P and Freund RW. Efficient linear circuit analysis by Padé approximation via the Lanczos process, *IEEE Trans. Computer-aided Design* , **14**:639-649, 1995.

[8] Gad E and Nakhla M. Model reduction for DC solution of large nonlinear circuits, *Proc. 1999 Int. Conf. Computer-aided Design*, 376-379, San Jose, 1999.

[9] Glover K. All optimal hankel-norm approximations of linear multivariable systems and their $L^\infty-$error bounds, *Int. J. Control*, **39**:1115-1193, 1984.

[10] Gugercin S and Antoulas AC. A comparative study of 7 algorithms for model reduction, *Proc. 39th Conf. Decision and Control*, 2367-2372, Sydney, 2000.

[11] Gugercin S, Antoulas AC and Beattie CA. H2 model reduction for large-scale linear dynamical systems, *SIAM J. Matrix Analysis and Applications*, **30(2)**:609-638, 2008.

[12] Hakvoort R. Worse-case system identification in l_1: error bounds, optimal models and model reduction, *Proc. 31st Conf. Decision and Control*, 499-504, Tucson, 1992.

[13] Khalil H. *Nonlinear Systems*, Prentice-Hall, second edition, 1996. %bibitemknockaertKnockaert L. On orthonormal Münt-Laguerre Filters,

[14] Moore BC. Principal component analysis in linear systems: controllability, observability and model reduction, *IEEE Trans. Auto. Control*, **AC-26**:17-32, 1981.

[15] Nakhla N, Nakhla M and Achar R. Spare and passive reduction of massively coupled large multiport interconnects, *Proc. 2007 Int. Conf. Computer-aided Design*,622-626, San Jose, 2007.

[16] Odabasioglu A, Celik M and Pileggi L. PRIMA: passive reduced-order interconnect macromodeling algorithm, *IEEE Trans. Computer-aided Design of Integrated Circuits and Systems*, **17(8)**:645-654, 1998.

[17] Phillips J, Daniel L and Silveira LM. Guaranteed passive balancing transformations for model order reduction, *IEEE Trans. Computer-aided Design of Integrated Circuits and Systems*, **22(8)**: 1027-1041, 2003.

[18] Rewienski M and White J. A trajectory piecewise-linear approach to model order reduction and fast simulation of nonlinear circuits and micromachined devices, *Proc. 2001 Int. Conf. Computer-aided Design*, 252-257, San Jose, 2001.

[19] Roychowdhury J. Reduced-order modeling of time-varying systems, *IEEE Trans. Circuits and Systems–II: Analog and Digital Signal Processing*,**46(10)**:1273-1288, 1999.

[20] Salimbahrami B, Eid R and Lehmann B. On the choice of an optimal interpolation point in Krylov-based order reduction, *Proc. of the 47th Conf. on Decision and Contr.*, 4209-4214, Cancun, 2008.

[21] Santarelli K. A Framework for Reduced Order Modeling with Mixed Moment Matching and Peak Error Objectives. Technical Report, Sandia National Laboratories, Dec. 2008.

[22] Sebakhy OA and Aly MN. Discrete-time model reduction with optimal zero locations by norm minimisation, *IEEE Proc. Control Theory Appl.*, **145(6)**:499-506, 1998.