

A FRAMEWORK FOR REDUCED ORDER MODELING WITH MIXED MOMENT MATCHING AND PEAK ERROR OBJECTIVES*

KEITH R. SANTARELLI†

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 with 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 prove that, for a certain class of LTI systems, the method we propose can be used to produce reduced order models of arbitrary accuracy even when the original system is infinite dimensional. We then show how to incorporate moment matching constraints into the basic error bound minimization algorithm, and present three examples which utilize the techniques described herein. We conclude with some comments on extensions to multi-input, multi-output systems, as well as some general comments for future work.

Key words. model reduction, L_1 norm, moment matching, simulation, Ritz approximation

AMS subject classifications. 00A69, 00A72, 34L10

DOI. 10.1137/090761136

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 [19] and Glover’s work on optimal model reduction in the Hankel norm [12] as 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: Bashir et al. [2] investigated a method of producing reduced order models for simulation when the initial condition is known to lie in a certain prespecified set; Gad and M. Nakhla [11] proposed a method for producing reduced order models for predicting the DC solution of large nonlinear circuits; N. Nakhla, M. Nakhla, and Achar [20] devised a method for model reduction of interconnect circuits; Rewiński and White [23] and Bond and Daniel [5] devised separate approaches to MOR of nonlinear circuits

*Received by the editors June 3, 2009; accepted for publication (in revised form) December 14, 2009; published electronically March 3, 2010.

<http://www.siam.org/journals/sisc/32-2/76113.html>

†Discrete Mathematics & Complex Systems Department, Sandia National Laboratories, Albuquerque, NM 87185 (krsanta@sandia.gov).

via piecewise linearization and projection; Roychowdhury [24] developed a method for the MOR of time-varying systems that has applications for modulation and sampling-type systems; Dong and Roychowdhury [8] developed a method of model reduction for nonlinear systems via representation by piecewise-polynomial functions; Feldmann and Freund [10] investigated a moment matching method for linear networks called the Padé via Lanczos method, and Odabasioglu, Celik, and Pileggi [21] developed a passivity-preserving MOR method for linear networks; Coelho, Phillips, and Silveira [7] investigated an optimization-based method where a reduced order model is found by solving a nonlinear least squares problem; Bui-Thanh, Willcox, and Ghattas [6] considered the problem of finding a reduced order model for a class of parameterized systems. The paper by Gugercin and Antoulas [13] 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* (SVDs) 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., $G^{(m_{s_l})}(s_l) = G_r^{(m_{s_l})}(s_l)$, $l = 1, 2, \dots, L$, 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. Hence, while moment matching methods provide a guarantee that the steady-state response will be the same for both the original and reduced order models for a finite set of *sinusoidal* inputs, there are typically no provable guarantees that the response of the reduced order system will be accurate at frequencies other than the matching frequencies.

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.¹ While such results

¹A similar statement exists when one considers inputs that are finite *energy* signals.

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, Daniel, and Silveira [22] 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, Antoulas, and Beattie [14] 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, $\text{Re}\{s_l\} < 0$.

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. To the best of the author’s knowledge, the result is the first of its kind and, hence, takes an important first step into investigating the problem of mixed moment matching/error-bounding reduction methods. 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. Figure 1.1 illustrates a hypothetical example where the spiky signal represents the output of an original full order system and the remaining 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 two 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 Figure 1.1 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 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 $y : [0, \infty) \rightarrow \mathbb{R}$, then the peak amplitude can be taken as the standard infinity norm:

$$(1.1) \quad \|y\|_\infty = \sup_{t \geq 0} |y(t)|.$$

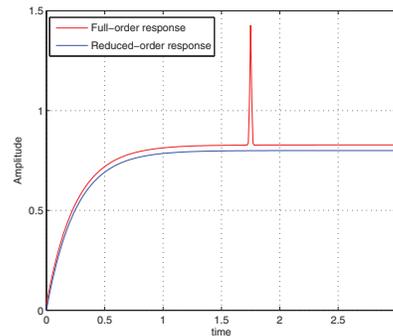


FIG. 1.1. Hypothetical responses of an original and reduced order system produced via an SVD-based method.

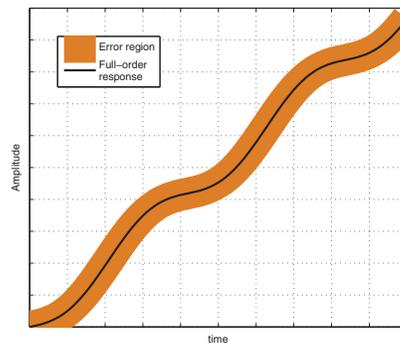


FIG. 1.2. Depiction of a full order output signal surrounded by an “error region.”

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 (1.1) that

$$(1.2) \quad |y(t) - y_r(t)| \leq \epsilon \quad \forall t \geq 0.$$

Figure 1.2 depicts the meaning of (1.2) graphically. In the figure, the black signal represents the response of the original system $y(t)$, and the surrounding area denoted “error region” represents a desired region in which one would like the response of a corresponding reduced order model $y_r(t)$ to lie. In the context of (1.2), the “height” of the error region at every given time t is 2ϵ , indicating the desire for $y_r(t)$ to be close to $y(t)$ uniformly over all times.

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 $L^\infty(\mathbb{R}^+)$

$$(1.3) \quad L^\infty(\mathbb{R}^+) = \left\{ u : [0, \infty) \rightarrow \mathbb{R} : \sup_{t \geq 0} |u(t)| < \infty \right\},$$

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

$$(1.4) \quad \|y - y_r\|_\infty \leq M \|u\|_\infty.$$

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, [17]) that the *smallest* value of M as given in (1.4) is the L_1 norm of the *error system* with impulse response $h(t) - h_r(t)$:

$$(1.5) \quad \|h - h_r\|_1 = \int_0^\infty |h(t) - h_r(t)| dt.$$

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 and Vidyasagar [9] first considered this problem in the context of some examples. In the discrete-time setting, Sebakhy and Aly [25] consider a simple form of impulse response truncation to minimize the l_1 norm of an error sequence ($\|e\|_1 = \sum_{k=1}^\infty |e_k|$). 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 [16]. While there are substantial differences with the class of problems being considered here as compared to [16], 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.

1.4. Document outline. Section 2 outlines the L_1 norm minimization algorithm and provides a complete characterization of the linear programming (LP) problem that is solved at each iteration. In section 3, we introduce the family of basis functions over which the minimization algorithm searches to produce a reduced order model with small L_1 error norm. Moreover, we show that this family of basis functions is sufficiently rich by proving that a certain large class of LTI systems can be approximated with arbitrary accuracy via an expansion of these basis functions. Section 4 describes how to incorporate moment matching constraints into the LPs to formulate mixed moment matching/norm minimization problems and formally proves that arbitrary accuracy via an expansion of basis functions is retained even in the presence of moment matching constraints. Section 5 summarizes the overall algorithm and describes some practical considerations in the problem of selecting an optimal basis from the family of basis functions under consideration. Section 6 illustrates the techniques described herein for three examples, two taken from the solution of a one-dimensional heat equation, and one taken from the circuits world. Section 7 briefly describes how to extend the methods here—designed only for single-input, single-output (SISO) systems—to a multi-input, multi-output (MIMO) generalization. We provide concluding remarks in section 8.

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, and show that this problem can be efficiently cast as the solution of a (relatively) small number of LPs.

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:

$$(2.1) \quad h_r(t) = \sum_{k=1}^N a_k g_k(t),$$

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

$$(2.2) \quad \begin{aligned} & \min \int_0^\infty z(t) dt, \\ & \text{subject to } z(t) \geq h(t) - \sum_{k=1}^N a_k g_k(t), \\ & \quad z(t) \geq - \left(h(t) - \sum_{k=1}^N a_k g_k(t) \right), \end{aligned}$$

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.² Note that (2.2) 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 (2.2) 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 $T > 0$,

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

where the inequality follows via repeated applications of the triangle inequality to $|h(t) - \sum_{k=1}^N a_k g_k(t)|$. By introducing the slack variables $w_k \geq |a_k|$ for $k = 1, 2, \dots, N$, (2.3) leads to the following LP:

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

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

$$\begin{aligned}
 (2.5) \quad \bar{h} &= \int_T^\infty |h(t)|dt, \\
 \beta_k &= \int_T^\infty |g_k(t)|dt.
 \end{aligned}$$

By virtue of (2.3), the minimal cost of the LP in (2.4) provides an upper bound for the minimal cost of the original infinite horizon LP of (2.2). 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 usually be chosen sufficiently large such that the effect of \bar{h} on the minimal cost in (2.4) is negligible).

²The reader unfamiliar with such arguments is referred to pp. 15–21 of [3] for a discussion of slack variables.

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 (2.4) via a Riemann sum leads to

$$(2.6) \quad \begin{aligned} \min \quad & \Delta \sum_{m=1}^M z_m + \sum_{k=1}^N \beta_k w_k, \\ \text{subject to} \quad & z_m \geq h_m - \sum_{k=1}^N a_k g_{km}, \\ & z_m \geq - \left(h_m - \sum_{k=1}^N a_k g_{km} \right), \\ & w_k \geq a_k, \\ & w_k \geq -a_k \end{aligned}$$

for all $k = 1, 2, \dots, N$ and $m = 1, 2, \dots, M$. 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 (2.4) and the approximate value in (2.6) 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 (2.2). The above LP can be written in multiple forms, and can be transformed into whatever form is most convenient for the particular software package that is used to provide a numerical solution.

2.2. L_1 norm minimization algorithm. The LP formulation of the last section begs the question: how does one choose the basis functions $g_k(t)$? First, recall that since we are trying to represent our approximate impulse response $h_r(t)$ as a finite dimensional model (meaning that the corresponding transfer function $H_r(s)$ is a rational function of s), $h_r(t)$ must be expressible as a linear combination of (possibly complex) exponential terms. This suggests that the functions $g_k(t)$ should involve “simple” linear combinations of exponential terms. Perhaps the simplest choice is $g_k(t) = \exp(-\alpha_k t)$, $\text{Re}\{\alpha_k\} \geq 0$ so that our approximation takes the form

$$(2.7) \quad h_r(t) = \sum_{k=1}^N a_k e^{-\alpha_k t}.$$

The problem of trying to choose the values of a_k and α_k to globally minimize $\|h - h_r\|_1$ is a nonconvex optimization problem, and hence is practically not solvable. If, however, we focus our attention on *local* minimizers, one naïve method of computing an upper bound on the global minimum is as follows. Whenever the values of α_k are *fixed*, the problem of approximating $h(t)$ via the $h_r(t)$ as in (2.7) is an LP. By gridding each value of $\alpha_k \in \mathbb{C}$ over some bounded region in the closed right half-plane, one could solve a sequence of LPs (one for each possible combination of grid points) and use the values of a_k and α_k which achieve the smallest cost over all the LPs that are solved.

While simple in concept, the above algorithm is computationally expensive since the number of LPs which must be solved grows exponentially with the order of approximation N . Indeed, if we grid each value of α_k using P points, we must solve a total of P^N LPs. For a value of $N = 10$, even using a coarse grid of $P = 10$ points per value of α_k results in 10^{10} LPs to be solved.

As an alternative to the above basis, consider the following choice:

$$(2.8) \quad g_k(t) = t^{k-1} e^{-\alpha t},$$

where the *single* parameter α again satisfies $\text{Re}\{\alpha\} \geq 0$. Such a choice for $g_k(t)$ results in an approximation of the form

$$(2.9) \quad h_r(t) = (a_1 + a_2 t + \dots + a_N t^{N-1}) e^{-\alpha t},$$

i.e., a *polynomial* in t multiplying a single decaying exponential term. The computational advantage of using such a basis is, in fact, quite large. For a given value of N , rather than having to grid N independent values and solving P^N LPs, one need only grid the single scalar variable α resulting in P LPs. Hence, again considering the case where $P = N = 10$, we reduce the number of LPs we need to solve from 10^{10} down to 10 by using the choice of g_k in (2.8). Note that we can generalize this idea to consider approximations of the form

$$(2.10) \quad h_r(t) = \sum_{j=1}^J p_j(t) e^{-\alpha_j t},$$

where the functions $p_j(t)$ are polynomials of fixed order with undetermined coefficients. Such an approximation would require J independent grids, and assuming that $J \ll N$, one still gains a large computational advantage over the original method since $P^J \ll P^N$.

With such a large savings in computation, it is natural to wonder whether the choice of (2.8) is somehow too narrow to accurately approximate a sufficiently rich class of signals. Fortunately, the answer to this question is no. Focusing on the case of approximation with just a single exponential parameter α , we show in the next section that there is a broad class of signals which can be well-approximated by expressions of the form shown in (2.9). More formally, we show that, under some mild assumptions, one can find a sequence of approximations of the form (2.9) for increasing N such that the L_1 norm of the error system $h - h_r$ converges to 0 as $N \rightarrow \infty$.

3. Convergence of approximations: Ritz basis. In this section, we prove that the approximations as determined by the choice of $g_k(t)$ in (2.8) converge in the L_1 norm to a given impulse response $h(t)$ under some mild assumptions. Since the main result of this section is an adaptation of existing results in functional analysis, we review those results first.

3.1. Convergence in L^2 : Ritz approximations. Let the space $L^2(\mathbb{R}^+)$ denote the set of Lebesgue measurable functions $f(t)$ defined for $t \geq 0$ such that the corresponding L_2 norm of f

$$(3.1) \quad \|f\|_2 = \left(\int_0^\infty |f(t)|^2 dt \right)^{\frac{1}{2}}$$

is finite. A well-studied problem in the systems and control community is that of approximating a given $f \in L^2(\mathbb{R}^+)$ via an expansion of the form $\sum a_k g_k(t)$, where

$G_k(s)$, the Laplace transform of $g_k(t)$, is a rational function of s . Convergence in these problems is naturally measured in terms of the L_2 norm; we say that the expansion converges if the sequence of partial sums

$$(3.2) \quad f_N(t) = \sum_{k=1}^N a_k g_k(t)$$

converges in the L_2 norm:

$$(3.3) \quad \lim_{N \rightarrow \infty} \|f - f_N\|_2 = 0.$$

A basis for $L^2(\mathbb{R}^+)$ is said to be *complete* if there exists a sequence of coefficients a_k in the expansion of (3.2) such that the $\|f - f_N\|_2$ converges to 0 for every $f \in L^2(\mathbb{R}^+)$. Complete bases, therefore, provide a set of elements that can well-approximate a wide range of functions. Perhaps the most popular basis for $L^2(\mathbb{R}^+)$ is the *Laguerre basis*, whose elements are described in the frequency domain via

$$(3.4) \quad G_k(s) = \frac{\sqrt{2\alpha}}{s + \alpha} \left(\frac{s - \alpha}{s + \alpha} \right)^{k-1}, \quad k = 1, 2, \dots,$$

where α satisfies $\text{Re}\{\alpha\} > 0$. A complete basis [18], the Laguerre basis is often the basis of choice for establishing theoretical statements due to the orthonormality of the basis functions:

$$\int_0^\infty g_j(t) \overline{g_k(t)} dt = \begin{cases} 0, & j \neq k, \\ 1, & j = k. \end{cases}$$

Because of this property, the coefficients a_k in the series expansion can be computed via a projection of f onto the corresponding basis functions $g_k(t)$.

While useful for theoretical statements, the Laguerre basis is less practical for numerical computation, as it can be empirically noted that LP solvers often run into numerical difficulties when dealing with the impulse responses $g_k(t)$ for large k . Moreover, easy-to-obtain upper bounds on the quantity

$$\beta_k = \int_T^\infty |g_k(t)| dt$$

(from (2.5) of the last section) are difficult to obtain and/or are very conservative, making for overinflated estimates of the minimal cost in the LP of (2.6). A more useful basis for our purposes here is a “deorthogonalized” version of the Laguerre basis known as the *Ritz basis*, whose elements are described in the frequency domain via

$$(3.5) \quad G_k(s) = \left(\frac{\alpha}{s + \alpha} \right)^k, \quad k = 1, 2, \dots$$

for $\text{Re}\{\alpha\} > 0$ with corresponding impulse responses

$$(3.6) \quad g_k(t) = \frac{\alpha^k}{(k-1)!} t^{k-1} e^{-\alpha t}, \quad k = 1, 2, \dots$$

Observe that the impulse responses of the Ritz basis vectors are scaled versions of the proposed basis vectors of the last section in (2.8). Also, note that when α is real, β_k can be calculated exactly as

$$(3.7) \quad \beta_k = e^{-\alpha T} \sum_{j=1}^k \frac{(\alpha T)^{j-1}}{(j-1)!} < 1 \quad \forall k \geq 1.$$

When α is complex, exact expressions for β_k are complicated, but simple upper bounds are readily obtainable. Indeed, if we denote $\alpha = -\alpha_r + j\omega$, where the real parameter α_r is positive, decomposition of the exponential terms into terms of the form $\exp(-\alpha_r t) \cos(\omega t)$ and $\exp(-\alpha_r t) \sin(\omega t)$ yields upper bounds on β_k by taking advantage of the fact that

$$(3.8) \quad \int_T^\infty |t^k e^{-\alpha_r t} \cos(\omega t + \phi)| dt \leq \int_T^\infty t^k e^{-\alpha_r t} dt$$

for any values of $\omega, \phi \in \mathbb{R}$.

3.2. Remark: Ritz approximations for model reduction. Our main task in this paper is to produce finite order models which approximate a higher (possibly infinite) dimensional model. Finite truncations of Ritz approximations provide for reduced order models by approximating the original model by a finite order model with repeated poles, with the order of the reduced order model being equal to the term of *highest* degree in the truncation. While we shall not discuss this here, it is a relatively straightforward task in theory to convert truncations of Ritz approximations into finite order state space models, which is often a much more convenient form for simulation.

3.3. Ritz approximation convergence in L^1 . We denote by $L^1(\mathbb{R}^+)$ the set of Lebesgue measurable functions $f(t)$ defined for $t \geq 0$ for which $\|f\|_1$ is finite. In this section, we prove that a broad subset of $L^1(\mathbb{R}^+)$ can be well-approximated via a Ritz approximation such that the partial sums of the form (3.2) (where $g_k(t)$ represent the Ritz basis vectors of (3.6)) converge in the L_1 norm: $\|f - f_N\|_1 \rightarrow 0$ as $N \rightarrow \infty$. The specific subset of $L^1(\mathbb{R}^+)$ we consider is described in the following proposition, whose proof can be found in the appendix.

PROPOSITION 3.1. *Consider the set \mathcal{S} of functions $f \in L^2(\mathbb{R}^+)$ which satisfy the condition that $f(t) = O(t^{-\gamma})$ for $\gamma > 1$, where $f(t) = O(g(t))$ is equivalent to the existence of constants $C > 0, t_0 \geq 0$ such that*

$$|f(t)| \leq C|g(t)| \quad \forall t \geq t_0.$$

Then $\mathcal{S} \subset L^1(\mathbb{R}^+)$.

The set \mathcal{S} encompasses a wide range of functions that are interesting from an application standpoint, including all bounded functions that decay exponentially, and all bounded functions that decay polynomially with exponent strictly greater than 1. Note that \mathcal{S} is broader than either of these two common subclasses and includes, for instance, unbounded functions such as

$$f(t) = \begin{cases} t^{-\frac{1}{4}}, & 0 < t \leq 1, \\ 0, & \text{otherwise.} \end{cases}$$

The ultimate goal of this section is to prove that there exists a Ritz approximation for every $f \in \mathcal{S}$ which converges in the L_1 norm. We first prove this result for a particular subset of \mathcal{S} , and then we use this result to prove the result for all $f \in \mathcal{S}$.

PROPOSITION 3.2. Consider the subset of $\bar{\mathcal{S}} \subset \mathcal{S}$ defined via $f \in \bar{\mathcal{S}}$ iff $f \in \mathcal{S}$ and $f(t) = O(t^{-2})$. For every $f \in \bar{\mathcal{S}}$, there exists a Ritz approximation such that the sequence of partial sums

$$(3.9) \quad f_N = \sum_{k=1}^N a_k \frac{\alpha^k t^{k-1}}{(k-1)!} e^{-\alpha t}$$

satisfies $\|f - f_N\|_1 \rightarrow 0$ as $N \rightarrow \infty$.

Equipped with Proposition 3.2 (whose proof can be found in the appendix), we now prove the first main theoretical result of the paper.

THEOREM 3.3. For every $f \in \mathcal{S}$, there exists a Ritz approximation such that the sequence of partial sums of (3.9) satisfies $\|f - f_N\|_1 \rightarrow 0$ as $N \rightarrow \infty$.

Proof. First, note that for any function $\hat{f}_{\hat{N}}(t)$

$$(3.10) \quad \|f - f_N\|_1 = \left\| f - \frac{\hat{f}_{\hat{N}}}{t^{\frac{\gamma}{2} + 1}} + \frac{\hat{f}_{\hat{N}}}{t^{\frac{\gamma}{2} + 1}} - f_N \right\|_1 \leq \left\| f - \frac{\hat{f}_{\hat{N}}}{t^{\frac{\gamma}{2} + 1}} \right\|_1 + \left\| \frac{\hat{f}_{\hat{N}}}{t^{\frac{\gamma}{2} + 1}} - f_N \right\|_1.$$

Note that the first norm on the right-hand side of (3.10) satisfies

$$(3.11) \quad \left\| f - \frac{\hat{f}_{\hat{N}}}{t^{\frac{\gamma}{2} + 1}} \right\|_1 \leq \left\| \frac{1}{t^{\frac{\gamma}{2} + 1}} \right\|_2 \left\| (t^{\frac{\gamma}{2} + 1})f - \hat{f}_{\hat{N}} \right\|_2.$$

Since $\|(t^{\frac{\gamma}{2} + 1})^{-1}\|_2 \leq \sqrt{\gamma/(\gamma - 1)}$, the left-hand side of (3.11) can be made arbitrarily small for sufficiently large \hat{N} if the rightmost L_2 norm of (3.11) can be made arbitrarily small. Since $f(t) = O(t^{-\gamma})$, $(t^{\frac{\gamma}{2} + 1})f(t) = O(t^{-\frac{\gamma}{2}})$, and it follows that $(t^{\frac{\gamma}{2} + 1})f(t) \in L^2(\mathbb{R}^+)$. Hence, there exists a Ritz approximation $\hat{f}_{\hat{N}}(t)$ for which the left-hand side of (3.11) can be made arbitrarily small for \hat{N} sufficiently large.

Now, for every fixed \hat{N} , note that the Ritz approximation $\hat{f}_{\hat{N}}(t)$ is a finite sum of exponentially decaying terms. It follows that

$$\frac{\hat{f}_{\hat{N}}(t)}{t^{\frac{\gamma}{2} + 1}} = O(t^{-2}).$$

(Functions which decay exponentially decay polynomially for any rate, hence they decay at a rate of t^{-2} ; dividing an exponentially decaying function by $t^{\frac{\gamma}{2} + 1}$ does not change this fact.) Hence, the result of Proposition 3.2 applies, and there exists a Ritz approximation $f_N(t)$ such that the rightmost norm of (3.10) can be made arbitrarily small for N sufficiently large. It thus follows that there exists a Ritz approximation $f_N(t)$ such that $\|f - f_N\|_1$ can be made arbitrarily small for N sufficiently large. \square

Theorem 3.3 states that for any $f \in \mathcal{S}$, by taking N sufficiently large, one can always well-approximate f in the L_1 norm via a Ritz approximation for any value of the parameter α of (3.6). Hence, given a desired tolerance ϵ for which one desires $\|f - f_N\|_1 < \epsilon$, the process of finding some reduced order model which satisfies a given tolerance constraint is easy: pick some value of α , and keep increasing the order N until the desired error tolerance is achieved. It should be apparent, however, that certain values of the parameter α are better than others in the sense that a poor choice of α could lead to a very large value of N that is required to satisfy a given tolerance constraint which potentially defeats the point of model reduction. We discuss the issue of trying to find “good” choices of α in a later section.

4. 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

$$(4.1) \quad \frac{1}{m!}H^{(m)}(s_0) = \frac{1}{m!} \sum_{k=1}^N a_k G_k^{(m)}(s_0),$$

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. When using the Ritz basis, the moments of $G_k(s)$ can be calculated explicitly as

$$(4.2) \quad G_k^{(m)}(s) = \begin{cases} \left(\frac{\alpha}{s+\alpha}\right)^k, & m = 0, \\ \frac{(-1)^m i(i+1)\dots(i+m-1)\alpha^k}{m!} \left(\frac{1}{s+\alpha}\right)^{k+m}, & m \geq 1. \end{cases}$$

Hence, whenever the value of the parameter α in the Ritz approximations 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:

$$(4.3) \quad \begin{aligned} \min \quad & \Delta \sum_{m=1}^M z_m + \sum_{k=1}^N \beta_k w_k, \\ \text{subject to} \quad & z_m \geq h_m - \sum_{k=1}^N a_k g_{km}, \\ & z_m \geq -\left(h_m - \sum_{k=1}^N a_k g_{km}\right), \\ & w_k \geq a_k, \\ & w_k \geq -a_k, \\ & H^{(m_{s_l})}(s_l) = \sum_{k=1}^N a_k G_k^{(m_{s_l})}(s_l), \end{aligned}$$

where $s_l \in \mathbb{C}$, $l = 1, 2, \dots, L$, represent a set of (possibly repeated) frequencies for which we wish to match the m_{s_l} th moment of the original and reduced order models.

4.1. Convergence of Ritz approximations with moment matching constraints. We now prove that the addition of moment matching constraints does not affect our ability to well-approximate in the L_1 norm via Ritz approximations. In fact, as we show in the next two propositions, *every* well-defined moment of a function $f(t)$ with Laplace transform $F(s)$ has a Ritz approximation whose corresponding moment converges to the true moment of $F(s)$.

PROPOSITION 4.1. *For every $f(t) \in \mathcal{S}$ with Laplace transform $F(s)$, the following statements hold:*

1. *The zeroth order moment of $F(s)$ at frequency s_0 exists for all s_0 with $\text{Re}\{s_0\} \geq 0$. Moreover, any Ritz approximation $f_N(t)$ for which $\|f - f_N\|_1 \rightarrow 0$ also satisfies the condition*

$$(4.4) \quad F_N(s_0) \rightarrow F(s_0)$$

for all s_0 in the closed right half-plane, where $F_N(s)$ denotes Laplace transforms of $f_N(t)$.

2. The m th order moments of $F(s)$ at frequency s_0 exists for all $m = 1, 2, \dots$ for all s_0 with $\text{Re}\{s_0\} > 0$. Moreover, any Ritz approximation $f_N(t)$ for which $\|f - f_N\|_1 \rightarrow 0$ also satisfies the condition

$$(4.5) \quad \frac{1}{m!} F_N^{(m)}(s_0) \rightarrow \frac{1}{m!} F^{(m)}(s_0)$$

for all s_0 in the open right half-plane.

The proof of this statement can be found in the appendix.

If $f(t)$ decays polynomially, higher order moments may not exist when $\text{Re}\{s_0\} = 0$ since $t^m f(t)$ may grow unboundedly as $t \rightarrow \infty$. If, however, $f(t)$ decays exponentially, all moments are well-defined on the $j\omega$ axis. The following proposition, whose proof can be found in the appendix, formalizes this statement.

PROPOSITION 4.2. Consider $f(t) \in L^2(\mathbb{R}^+)$ and $f(t) = O(e^{-\gamma t})$, $\gamma > 0$. The m th order moments of $F(s)$ (the Laplace transform of $f(t)$) at a frequency s_0 with $\text{Re}\{s_0\} = 0$ exist for all $m = 0, 1, \dots$. Moreover, any Ritz approximation which satisfies the condition $\|f - f_N\|_1 \rightarrow 0$ also satisfies the condition

$$(4.6) \quad \frac{1}{m!} F_N^{(m)}(s_0) \rightarrow \frac{1}{m!} F^{(m)}(s_0)$$

for all s_0 on the $j\omega$ axis.

Using Propositions 4.1 and 4.2, we can now prove that the addition of a finite number of moment matching constraints does not affect the ability of $\|f - f_N\|_1$ to converge to 0.

THEOREM 4.3. Consider $f(t) \in \mathcal{S}$. Subject to a finite number of well-defined moment matching constraints (as given by Propositions 4.1 and 4.2), there exists a Ritz approximation f_N for which $\|f - f_N\|_1 \rightarrow 0$.

Proof. From the results of Propositions 4.1 and 4.2, it follows that for every $\epsilon > 0$, there exists an N th order Ritz approximation $f_N(t) = \sum_{k=1}^N a_k g_k(t)$ with N sufficiently large such that

$$(4.7) \quad \int_0^\infty \left| f(t) - \sum_{k=1}^N a_k g_k(t) \right| dt \leq \epsilon,$$

$$(4.8) \quad \left| F^{(m_{s_l})}(s_l) - \sum_{k=1}^N a_k G^{(m_{s_l})}(s_l) \right| \leq \epsilon$$

for $l = 1, 2, \dots, L$, where the L lower inequalities represent moment matching constraints of order m_{s_l} at frequency s_l . For each L , define

$$(4.9) \quad \epsilon_l = F^{(m_{s_l})}(s_l) - \sum_{k=1}^N a_k G^{(m_{s_l})}(s_l).$$

Clearly, $|\epsilon_l| \leq \epsilon$ for $l = 1, 2, \dots, L$. Now, suppose there exists a sequence $\{\delta_k\}_{k=1}^N$ such that

$$(4.10) \quad \sum_{k=1}^N \delta_k G^{(m_{s_l})}(s_l) = \epsilon_l, \quad l = 1, 2, \dots, L$$

(we shall prove existence of such a sequence shortly). Then it follows that the Ritz approximation with coefficients $\tilde{a}_k = a_k + \delta_k$ satisfies each of the L moment matching constraints. To establish convergence of the Ritz approximation with coefficients \tilde{a}_k in the L_1 norm, we wish to show that

$$(4.11) \quad \int_0^\infty \left| f(t) - \sum_{k=1}^N \tilde{a}_k g_k(t) \right| dt < \lambda \epsilon$$

for some value of λ that does not depend on ϵ . We have

$$(4.12) \quad \left\| f(t) - \sum_{k=1}^N \tilde{a}_k g_k(t) \right\|_1 \leq \left\| f(t) - \sum_{k=1}^N a_k g_k(t) \right\|_1 + \left\| \sum_{k=1}^N \delta_k g_k(t) \right\|_1.$$

The first term on the right-hand side above is upper bounded by ϵ by assumption. Convergence, hence, reduces to showing that the second term on the right-hand side can be made sufficiently small.

Note that the coefficients δ_k satisfy a linear constraint of the form

$$(4.13) \quad A_N \bar{\delta} = \bar{\epsilon},$$

where $\bar{\delta} = [\delta_1 \ \delta_2 \ \dots \ \delta_N]$, and $\bar{\epsilon} = [\epsilon_1 \ \epsilon_2 \ \dots \ \epsilon_L]$. We assume that the L rows of A are linearly independent (otherwise, there is a redundant moment matching constraint that can be removed).

We now show that for any $N \geq L$, there exists $\bar{\delta}$ which satisfies (4.13) and which makes the rightmost term of (4.12) sufficiently small. Let $N_0 \geq L$. Then there exists L linearly independent columns of A_{N_0} . Consider a map $N_0(j)$ for $j = 1, 2, \dots, L$ such that the $L \times L$ matrix $A_{N_0}^L$ whose j th column is the $N_0(j)$ th column of A_{N_0} is invertible. Similarly, let $\bar{\delta}^L$ be defined via $\bar{\delta}_j^L = \delta_{N_0(j)}$. Then the system of equations $A_{N_0}^L \bar{\delta}^L = \bar{\epsilon}$ is solvable for $\bar{\delta}^L$ and can be written explicitly as $\bar{\delta}^L = (A_{N_0}^L)^{-1} \bar{\epsilon}$. It follows that $\|\bar{\delta}^L\|_\infty \leq \epsilon \|(A_{N_0}^L)^{-1}\|_\infty$. Now, observe that the vector $\bar{\delta}$ with

$$(4.14) \quad \bar{\delta}_k = \begin{cases} \bar{\delta}_k^L, & k = N_0(j), \quad j = 1, 2, \dots, L, \\ 0, & \text{otherwise} \end{cases}$$

satisfies (4.13) for every $N \geq N_0$. Hence, it follows that there exists a choice of the δ_k 's such that

$$(4.15) \quad \int_0^\infty |\delta_k g_k(t)| dt \leq \epsilon \|(A_{N_0}^L)^{-1}\|_\infty \sum_k \int_0^\infty |g_k(t)| dt,$$

where the sum on the right-hand side is taken over only those k which can be represented as $k = N_0(j)$, $j = 1, 2, \dots, L$. Since the sum is finite, convergence follows. \square

5. Summary of algorithm and practical considerations. We begin by briefly summarizing the steps of the algorithm we propose for producing reduced order models with small L_1 error norm subject to moment matching constraints.

User-specified data. To begin, the user must select an order N for the reduced order model, and a set of values \mathcal{A} for the parameter α in the Ritz approximation of (3.6). Additionally, the user must select a horizon time T over which the L_1 norm will be approximated, along with an appropriate grid spacing Δ with which to sample the time axis. Smaller Δ will, obviously, provide better Riemann sum approximations to the desired integrals, and larger T will provide less conservative upper bounds on the L_1 norm computation (reflected by the fact that the upper bound coefficients β_k of (2.5) are monotonically decreasing functions of T). Note that, once Δ and T have been selected, the samples h_m of the original impulse response $h(t)$ in the LP formulation of (4.3) are automatically determined.

In addition to the above quantities, the user must also specify a (possibly empty) set of frequencies s_l and corresponding moments m_{s_l} to be matched (i.e., we compute $H^{(m_{s_l})}(s_l)$).

Quantities to compute for each LP iteration. For each value of $\alpha \in \mathcal{A}$, one must compute several quantities in order to set up the corresponding LP. First, one must compute the samples g_{km} of the Ritz basis vectors $g_k(t)$ for $k = 1, 2, \dots, N$. Also, the quantities β_k must be computed as well. Finally, one must compute the values $G^{(m_{s_l})}(s_l)$ that are necessary for moment matching constraints.

With all of the above quantities in place, one can loop over all $\alpha \in \mathcal{A}$ and solve each corresponding LP. The value of α which minimizes the cost of (4.3) yields the minimal upper bound on $\|h - h_r\|_1$, and the coefficients a_k together with the transfer functions $g_k(t)$ for this value of α determine an N th order model $h_r(t)$ which achieves the minimal upper bound.

5.1. Selecting the set \mathcal{A} . The success of the above algorithm largely hinges on the ability to select a good value of the Ritz approximation parameter α which is directly correlated to the choice of the grid set \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:

$$(5.1) \quad \mathcal{A} = \{\alpha : \alpha = j\Delta_a, \quad j = 1, 2, \dots, J\},$$

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 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 guarantee convergence only 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. In fact, while we shall not do so here, one can actually prove that the minimal cost varies *continuously* as a function of α , which seemingly indicates further promise for convergence to a minimum on I , but which does *not*, unfortunately, provide a provable guarantee.

6. Examples. In this section, we apply the techniques described in the prior sections to three infinite dimensional examples. In each case, we compare the performance of the reduced order models we obtain via mixed L_1 norm minimization/moment matching to an existing MOR method.

6.1. Example: One-dimensional heat equation. Consider a semi-infinite rod described by the half-line $x \geq 0$, and let $u(x, t)$ denote the temperature of the rod at position x at time t . The evolution of the temperature distribution can be described via the *one-dimensional heat equation*, given by

$$(6.1) \quad u_t - u_{xx} = 0.$$

Suppose that the temperature at $x = 0$ is controlled via an actuator, and that we are interested in computing the corresponding temperature at $x = 1$ for all $t \geq 0$. Such a problem is of interest in the case when one wishes to control the temperature of the rod at a certain point but can only apply heat at a different location. Subject to the initial value constraint $u(x, 0) = 0, x \geq 0$ and the boundary constraint, $u(\infty, t) = 0, t \geq 0$, this corresponds to finding the transfer function and/or impulse response operator from $u(0, t)$ to $u(1, t)$. A simple calculation shows that the transfer function $H(s)$ is given by

$$(6.2) \quad H(s) \triangleq \frac{U(1, s)}{U(0, s)} = e^{-\sqrt{s}},$$

where $U(0, s)$ and $U(1, s)$ represent Laplace transforms of $u(0, t)$ and $u(1, t)$, respectively. The corresponding impulse response $h(t)$ is given by

$$(6.3) \quad h(t) = \frac{1}{\sqrt{4\pi t^3}} e^{-\frac{1}{4t}}, \quad t > 0.$$

It is apparent that the transfer function $H(s)$ is infinite dimensional since it is not a rational function of s . Also, observe that $h(t)$ decays as $1/t^{1.5}$, which satisfies the condition of Theorem 3.3 so that $h(t)$ can be well-approximated in the L_1 norm via a Ritz approximation. Suppose that we are interested in producing a finite order approximation $h_r(t)$ to the original impulse response $h(t)$ such that $\|h - h_r\|_1$ is small, and also such that the DC gain of the original and reduced order models are equal (i.e., $H(0) = H_r(0) = 1$).

Using the procedure described in the paper, we found a 10th order model $H_r(s)$ with a Ritz parameter value of $\alpha = 0.5$. A plot of the original impulse response $h(t)$ and the reduced order impulse response $h_r(t)$ is shown in Figure 6.1. This value of α was found to produce an upper bound on the error norm $\|h - h_r\|_1$ of 0.206. Hence

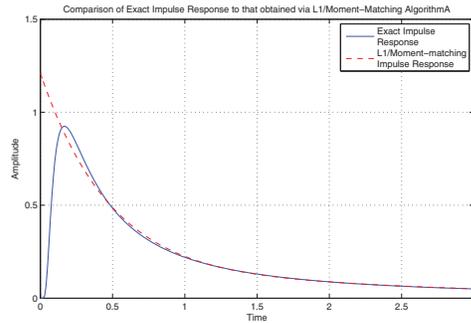


FIG. 6.1. Impulse responses of the original and reduced order systems using the mixed L_1 /moment matching algorithm.

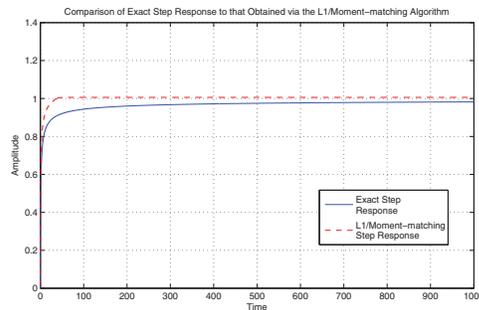


FIG. 6.2. Step responses of the original and reduced order systems using the mixed L_1 /moment matching algorithm.

for any bounded input $x(t)$, the outputs of the original model $y(t)$ and the reduced order model $y_r(t)$ satisfy the bound

$$(6.4) \quad |y(t) - y_r(t)| \leq 0.206 \|x\|_\infty.$$

In particular, when $\|x\|_\infty = 1$, we find that $|y(t) - y_r(t)| \leq 0.206$ for *any* bounded input with peak magnitude of 1.

The step responses of both the original and reduced order system are depicted in Figure 6.2. Observe that the reduced order step response approaches the final steady-state value of 1 much more quickly than the step response of the original system. This should be expected since the original impulse response $h(t)$ of (6.3) decays polynomially as $1/t^{1.5}$, whereas the Ritz approximation decays exponentially. Nevertheless, the steady-state values are both equal to 1 by design, since we imposed the constraint that $H(0) = H_r(0)$. Moreover, (6.4) guarantees that the approximate response will never deviate from the true response by more than 0.206. Indeed, one can verify by examining the graph of Figure 6.2 that the maximum deviation between the two step responses is roughly $0.09 < 0.206$.

As a comparison, we also computed a 10th order reduced order model via a moment matching method. Specifically, we represented $H_r(s) = p(s)/q(s)$, where $q(s) = (s + 0.5)^{10}$ (hence using the same pole locations as we used above for the L_1 /moment matching technique), and we chose the numerator $p(s)$ such that $H(s) = H_r(s)$ at

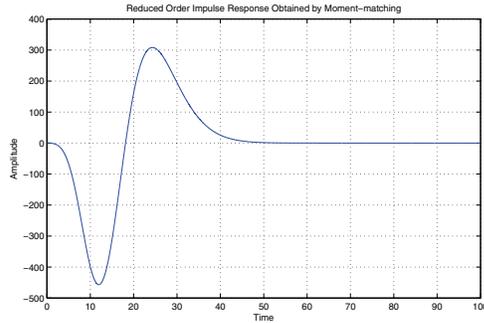


FIG. 6.3. Impulse response of the reduced order system obtained via moment matching.

the frequencies $s = 0, j, 2j, 3j, 4j,$ and $5j$. In other words, we found a reduced order model whose frequency response matches the original frequency response at several points in a low frequency range. The reduced order impulse response that is obtained via this method is shown in Figure 6.3. While the steady-state value of the step response of this reduced order system is equal to the steady-state value of the original system, it is clear that the step responses will deviate *very* significantly during the transient phase. One finds that the step response of this approximate system dips to a value of approximately -3800 at $t \approx 18$ s. Indeed, the value of $\|h - h_r\|_1$ for this reduced order model is roughly 7600!

6.2. Example: One-dimensional heat equation with two repeated poles.

This next example is a continuation of the last example where we now find a reduced order model with not one but two repeated poles. More formally, we consider a reduced order model $h_r(t)$, which can be written in the form

$$(6.5) \quad h_r(t) = p_1(t)e^{-\alpha_1 t} + p_2(t)e^{-\alpha_2 t},$$

with $\text{Re}\{\alpha_1\}, \text{Re}\{\alpha_2\} > 0$, and where $p_1(t)$ and $p_2(t)$ are polynomials of order $N_1 - 1$ and $N_2 - 1$, respectively. In the frequency domain, this can be expressed equivalently as

$$(6.6) \quad H_r(s) = \frac{Q_1(s)}{(s + \alpha_1)^{N_1}} + \frac{Q_2(s)}{(s + \alpha_2)^{N_2}},$$

where $Q_1(s)$ and $Q_2(s)$ are polynomials of degree less than or equal to N_1 and N_2 , respectively. It is clear from both (6.5) and (6.6) that the order of the resulting approximation is equal to $N_1 + N_2$.

Some comments are in order. First, note that considering reduced order models of the forms (6.5) and (6.6) does not require additional convergence theorems to validate this choice as a “good” form. For instance, if one takes $Q_1(s) = 0$ in (6.6), the model reduces to the single- α model for which all of the prior theoretical results on convergence have been established. Hence, since existence of a finite order model which well-approximates the original impulse response $h(t)$ can be guaranteed when $Q_1(s) = 0$, it can clearly also be guaranteed when this restriction is lifted.

The motivation for considering a model of the forms (6.5) and (6.6) lies in that, while in theory one can always find a reduced order model by considering a single Ritz expansion (i.e., a single value of the parameter α), in practice, the order N which may

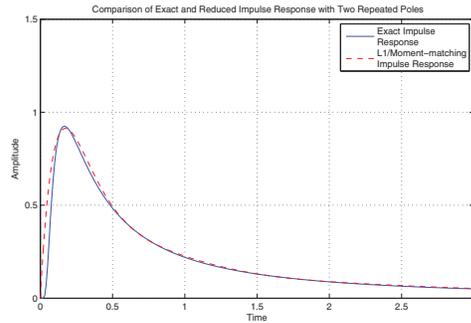


FIG. 6.4. Impulse response of the reduced order system obtained using two repeated poles.

be required to obtain a reasonably small approximation error may be prohibitively large. Thought about in another way, if the desired order of a reduced order model N is fixed, it may be possible to lower the approximation error by considering models which have multiple repeated poles. An intuitive interpretation of this idea, in the context of the impulse response $h(t)$ of (6.3) (depicted in Figure 6.1), is as follows. It is clear that for sufficiently large times t , $h(t) \approx 1/\sqrt{4\pi t^3}$, which decays monotonically and without oscillation. Hence, it is reasonable to expect that, for large times, $h(t)$ is well-approximated by a low frequency signal. On the other hand, at times close to 0, $h(t)$ is not monotonic (the “hump” centered close to $t = 0.25$), and it is reasonable to expect that some higher frequency component may be needed to well-approximate $h(t)$ in this region.

As in the last section, we search for a 10th order model (i.e., $N_1 + N_2 = 10$) which attempts to minimize the L_1 norm of the error approximation subject to the DC moment matching constraint $H_r(0) = H(0) = 1$. Experimentation with different values of N_1 and N_2 shows that $N_1 = 3$ (for the low frequency component) and $N_2 = 7$ (for the high frequency component) are good values relative to the other integer combinations. Gridding the corresponding values of α_1 and α_2 on the real line, we find that the values $\alpha_1 = 0.1$ and $\alpha_2 = 3.25$ achieve an L_1 error norm upper bound of 0.13, less than $2/3$ of the bound we found using a single value of α . The original and reduced order impulse responses obtained using this two repeated pole approximations are shown in Figure 6.4.

At this point, it is natural to question whether the use of a few repeated poles—as opposed to considering techniques which allow the pole locations to all be completely different—is constraining in terms of the minimum achievable L_1 norm. As we saw above, considering a model with two distinct, real poles decreased the L_1 norm by roughly 33%. By how much could the L_1 norm be reduced if we considered 3, 4, or up to 10 distinct pole locations in the form of the reduced order model? While a formal answer to this is currently unknown, we present the following example (provided by one of the anonymous reviewers) as a means of providing intuition on this topic. If one ignores the moment matching constraint as an explicit constraint in the design procedure, then one can use the vector-fitting method of [15] (a method which is also capable of dealing with infinite dimensional systems, albeit completely in the frequency domain) to obtain a 10th order reduced model, as well. One of the advantages to this method is that the pole locations are unconstrained and are determined by the vector-fitting procedure. The resulting reduced order impulse response

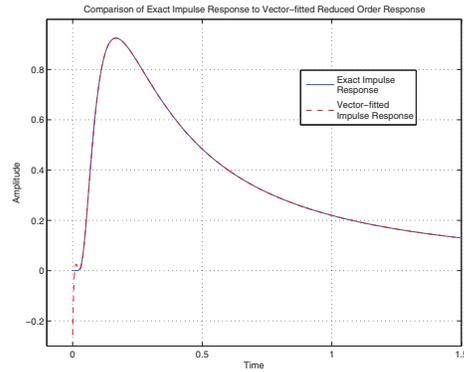


FIG. 6.5. Impulse response of the reduced order system obtained using the vector-fitting method of [15].

obtained by using this method is as shown in Figure 6.5. The corresponding upper bound on the L_1 error norm in this case can be computed to be roughly 0.114, a close to 13% improvement over the two-repeated-pole approximation we consider above.

The “punchline” of the above comparison is mixed: on one hand, it is clear that there is some room for improvement by considering additional poles in the form of the approximation, a result which should not be surprising. On the other hand, if the result of this particular example is indicative of the typical performance gain obtained by considering unrestricted pole locations (a very open question), then the utility of expanding the number of distinct poles in the approximation scheme may not be of great practical use. For one, as mentioned previously, the computational complexity involved in solving the corresponding sequence of LPs grows exponentially with the number of distinct poles. From a different perspective, however, if one decides to “jump ship” and use different methods entirely—such as the vector-fitting method shown above, since it produces a smaller L_1 error norm in this example—then one must abandon the idea of considering mixed problems which consider both error norms and moment matching constraints. Indeed, the vector-fitting method of [15] cannot incorporate moment matching constraints directly and, hence, the quality of the moments of the resulting reduced order model cannot be guaranteed. This, in turn, relates to the question of what the “best” MOR method is for a particular problem and that, in turn, is often a function of the user. For this particular example, it is arguable that a DC moment matching constraint is quite important since, in practice, step inputs are very common in heat applications and, hence, accuracy of the corresponding steady-state solutions to step inputs is critical. For another application with the same transfer function, however, matching the DC gain may not be as important and other methods may be more appropriate.

6.3. Example: Bandpass filter with time delay. Figure 6.6 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

$$(6.7) \quad H(s) \triangleq \frac{V_{\text{out}}(s)}{V_{\text{in}}(s)} = e^{-s} \frac{2s}{(s+1)^2 + 10,000},$$

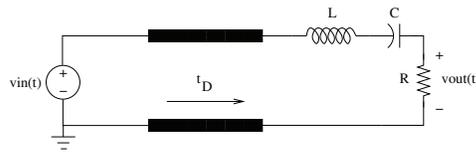


FIG. 6.6. Schematic of a circuit with an RLC bandpass filter and an ideal transmission line.

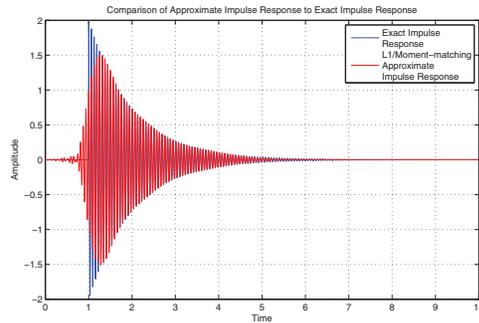


FIG. 6.7. Impulse responses of the original model and the reduced order model produced via the mixed L_1 /moment matching algorithm.

with corresponding impulse response

$$(6.8) \quad h(t) = \begin{cases} e^{-(t-1)} (2 \cos 100(t-1) - 0.02 \sin 100(t-1)), & t \geq 1, \\ 0, & t < 1. \end{cases}$$

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 a Ritz 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 Figure 6.7. 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 $t \geq 1$.

As a comparison, we created an alternative reduced order model using the following technique: the time delay in (6.7) 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. The essential reasoning behind such an approach is that TBR algorithms are designed to produce small error norms, where error is measured in terms of power instead of peak amplitude. Intu-

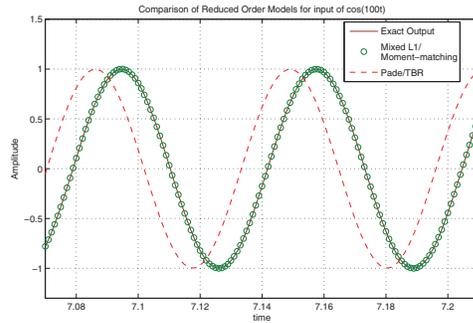


FIG. 6.8. Comparison of responses of the original model, L_1 /moment matching reduced order model, and Padé/TBR reduced order model for the input $\cos 100t$.

itively speaking, by comparing the behavior of the minimal L_1 norm approaches here to an SVD-based method, we are trying to assess whether SVD-based methods produce small L_1 norms “automatically” without explicitly incorporating such measures into their cost criterion. Since TBR operates on finite dimensional LTI systems, one must first approximate the nonrational 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 Figure 6.8. 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.

One of the primary reasons as to why the TBR algorithm fails on this particular example is due to the fact that the 50th order Padé approximation does not well-approximate the time delay over the bandwidth of interest. By examining the Bodé phase curves of a true unit time delay and a 50th order Padé approximation, one will find that the two are in close agreement up to approximately 70 rad/sec, beyond which the two phase curves begin to vary significantly. Because the bandpass filter has a resonant frequency at 100 rad/sec, the Padé approximation is not sufficient to capture the behavior of the system at the set of frequencies where it is most “active.” While this may appear to be sheer negligence, this example was constructed to illustrate a “soft” point: the reason that a Padé approximation of order higher than 50 was not used to approximate the time delay before applying TBR was that the MATLAB routine that was used to find the Padé approximation (the standard `padé` command in the LTI toolbox) suffered numerical instabilities when the order was increased significantly beyond this range (the author noticed that the resulting TBR reduced order model was unstable starting with Padé approximations of order 55). While it is possible to circumvent this issue by developing more numerically stable Padé approximation and TBR software routines, the *necessity* of having to create

new computational machinery to handle this issue is clearly undesirable, particularly when there is another method that circumvents entirely the need to find a Padé approximation in the first place.

6.4. General remarks. The computation for both of the examples presented above was performed in MATLAB using the freeware `sedumi` as the core LP solver engine. For a given order of approximation N , each example required solving for no more than 10–15 values of the Ritz parameter α before the improvement in the L_1 error norm ceased to be significant.

The order N in each of these examples was chosen somewhat arbitrarily. While larger N and longer horizon times T did indicate that the upper bounds on the error norms were strictly decreasing, the run time of each LP became significantly slower beyond certain thresholds since the computer system would run out of cache memory. The value of N was, therefore, kept not-too-large in order to be able to carry out the necessary computations quickly. It is an important area of future research to devise computationally efficient ways to cast and solve the core linear programs that are solved at each iteration of the algorithm.

While perhaps diabolical to say, it is the hope of the author that by reading this Example section, the reader is left with more questions than answers. The main goal of this exposition is to provide a new perspective on ways to assess the quality of a reduced order model (measured here as peak-to-peak error), and to show that there exists a rudimentary method of posing such problems in a manner which is computationally tractable and which possesses good theoretical properties (to show that the method is “not bad”). Clearly, it will be important to determine when the method described herein has practical benefit for particular classes of applications, but the examples provided here do not provide detailed comparisons. Rather, they serve to demonstrate some of the *potential* benefits and drawbacks of the proposed method as a means of motivating future research.

7. Extensions to multi-input, multi-output systems. While we have focused our discussion around SISO systems, the techniques discussed here can be readily applied to MIMO systems as well. The changes to the algorithms in moving from SISO to MIMO are primarily centered in the formulation of the corresponding LP problems, so we touch upon that briefly here.

A MIMO LTI system with impulse response matrix $\mathcal{H}(t)$ has L_1 norm given by

$$(7.1) \quad \|\mathcal{H}\|_1 = \max_i \sum_j \int_0^\infty |h^{ij}(t)| dt,$$

where $h^{ij}(t)$ represents the (i, j) th entry of $\mathcal{H}(t)$. Again, assuming a fixed set of basis functions $g_k(t)$, our task is to find a matrix $\mathcal{H}_r(t)$ with (i, j) th entry $h_r^{ij}(t)$ of the form

$$(7.2) \quad h_r^{ij}(t) = \sum_{k=1}^{N_{ij}} a_k^{ij} g_k(t)$$

such that $\|\mathcal{H} - \mathcal{H}_r\|_1$ is minimized. We can cast this minimization problem in the

following manner:

$$\begin{aligned}
 (7.3) \quad & \min C, \\
 & \text{subject to } C \geq c_i, \\
 & c_i = \sum_j c_{ij}, \\
 & c_{ij} = \int_0^\infty \left| h^{ij}(t) - \sum_{k=1}^{N_{ij}} a_k^{ij} g_k(t) \right| dt.
 \end{aligned}$$

Each c_{ij} represents the cost functional for a single SISO L_1 norm minimization problem. Hence, one can introduce slack variables and set up moment matching constraints for each c_{ij} , as presented in the prior sections.

8. 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 pointwise sense, can be implemented efficiently using a relatively small number of user-specified iterations, as demonstrated by three specific examples.

The primary concern of the methods shown here is likely to be the issue of gridding the parameter space to find a local minimum. While this is computationally tractable in low-dimensional parameter spaces, it is desirable to utilize methods which eliminate exhaustive gridding procedures whenever possible. Some preliminary work to this effect has recently been considered in [4], where a nonlinear programming (NLP) formulation (in which the value of the Ritz parameter α is a decision parameter) is examined. Because the value of α in this NLP is a decision parameter, explicit gridding is eliminated, but convergence of the corresponding algorithm to a local minimum is no longer guaranteed. The manuscript [4] examines the behavior of the corresponding NLP for multiple examples.

Appendix. Proofs of technical statements.

Proof of Proposition 3.1. We must show that every $f \in \mathcal{S}$ satisfies $f \in L^1(\mathbb{R}^+)$. To begin, note that if $\|f\|_1$ exists, then

$$(A.1) \quad \|f\|_1 = \int_0^{t_0} |f(t)|dt + \int_{t_0}^\infty |f(t)|dt,$$

where t_0 is defined as in Proposition 3.1. By virtue of the Cauchy–Schwartz inequality, the first integral in (A.1) satisfies

$$(A.2) \quad \int_0^{t_0} |f(t)|dt \leq \sqrt{t_0} \left(\int_0^{t_0} |f(t)|^2 \right)^{\frac{1}{2}} < \infty,$$

since $f \in L^2(\mathbb{R}^+)$. The second integral in (A.1) satisfies

$$(A.3) \quad \int_{t_0}^\infty |f(t)|dt \leq \int_{t_0}^\infty C t^{-\gamma} dt = \frac{C t_0^{1-\gamma}}{\gamma - 1},$$

since $\gamma > 1$. Since both quantities on the right-hand side of (A.1) are bounded above, $f \in L^1(\mathbb{R}^+)$.

Proof of Proposition 3.2. We first introduce the notation $\hat{f}_N(t) = (\alpha t + 1)f_N(t)$, where α is as in (3.9). For every $N \in \mathbb{Z}^+$, as a result of the Cauchy–Schwartz inequality,

$$(A.4) \quad \|f - f_N\|_1 = \left\| f - \frac{\hat{f}_N}{\alpha t + 1} \right\|_1 \leq \left\| \frac{1}{\alpha t + 1} \right\|_2 \left\| (\alpha t + 1)f - \hat{f}_N \right\|_2.$$

Notice that since $\|(\alpha t + 1)^{-1}\|_2 = 1/\sqrt{\alpha}$, convergence in the L_1 norm will follow if $\hat{f}_N(t)$ converges to $(\alpha t + 1)f(t)$ in the L_2 norm. Indeed, since $f(t) = O(t^{-2})$, $(\alpha t + 1)f(t) = O(t^{-1})$, and, hence, $(\alpha t + 1)f(t) \in L^2(\mathbb{R}^+)$. Therefore, it follows from the standard results on L_2 theory presented at the beginning of section 3 that there exists a Ritz approximation $\hat{f}_N(t)$ which converges to $(\alpha t + 1)f(t)$ in the L_2 norm.

What remains to be shown is that $f_N(t)$, itself, can be written as a Ritz approximation. If we assume that $f_N(t)$ has a Ritz approximation for the same value of α in (3.9) as the Ritz approximation $\hat{f}_N(t)$, and if we denote the corresponding expansion coefficients of $f_N(t)$ and $\hat{f}_N(t)$ by a_k and \hat{a}_k , respectively, then multiplying the series expansion of $f_N(t)$ by $\alpha t + 1$ and equating coefficients yields the relationship

$$(A.5) \quad \hat{a}_k = a_k + k a_{k-1}.$$

Since (A.5) implies that $\hat{a}_1 = a_1$, the above relationship can be solved recursively for a_k via $a_k = \hat{a}_k - k a_{k-1}$, which shows that $f_N(t)$ can be expressed as a Ritz approximation.

Proof of Proposition 4.1. We prove only the second item as the proof of the first item is similar. In what follows, we scale all moments by $m!$ for notational simplicity. To begin, note that the scaled m th moment at a frequency s_0 can be written as

$$(A.6) \quad F^{(m)}(s_0) = (-1)^m \int_0^\infty t^m f(t) e^{-s_0 t} dt.$$

Hence, the m th order moment at frequency s_0 is well-defined if the integral on the right-hand side of (A.6) converges. We have

$$(A.7) \quad \left| \int_0^\infty t^m f(t) e^{-s_0 t} dt \right| \leq \sup_{t \geq 0} |t^m e^{-s_0 t}| \int_0^\infty |f(t)| dt,$$

because $\operatorname{Re}\{s_0\} > 0$, $t^m e^{-s_0 t}$ is bounded. Moreover, $f(t) \in L^1(\mathbb{R}^+)$ via Proposition 3.1; hence, the moment is well-defined.

To prove convergence of the Ritz approximation moments to the true moment, let $f_N(t)$ be any Ritz approximation which converges in the L_1 norm. Now

$$(A.8) \quad \begin{aligned} \left| F^{(m)}(s_0) - F_N^{(m)}(s_0) \right| &= \left| \int_0^\infty t^m (f(t) - f_N(t)) e^{-s_0 t} dt \right| \\ &\leq \sup_{t \geq 0} |t^m e^{-s_0 t}| \|f(t) - f_N(t)\|_1. \end{aligned}$$

Hence, convergence of the moment follows via convergence of $\|f - f_N\|_1$.

Proof of Proposition 4.2. Again, we scale all moments by $m!$ for notational simplicity. By the assumed exponential decay of f , there exists a constant C such that $|f(t)| \leq Ce^{-\gamma t}$ for all $t \geq 0$. Hence, with $\text{Re}\{s_0\} = 0$

$$(A.9) \quad \left| F^{(m)}(s_0) \right| = \left| \int_0^\infty t^m f(t) e^{-s_0 t} dt \right| \leq C \int_0^\infty t^m e^{-\gamma t} dt.$$

Since the rightmost integral converges for all $m \geq 0$, we conclude that the m th order moments exist and are well-defined at all frequencies along the $j\omega$ axis.

To prove convergence of the moments, if we denote $D = |F^{(m)}(s_0) - F_N^{(m)}(s_0)|$, observe that for any $\epsilon > 0$

$$(A.10) \quad \begin{aligned} D &= \left| \int_0^\infty t^m (f(t) - f_N(t)) e^{-s_0 t} dt \right| \\ &= \left| \int_0^\infty t^m e^{-\epsilon t} (f(t) e^{\epsilon t} - f_N(t) e^{\epsilon t}) e^{-s_0 t} dt \right| \\ &\leq \sup_{t \geq 0} |t^m e^{-\epsilon t}| \int_0^\infty |f(t) e^{\epsilon t} - f_N(t) e^{\epsilon t}| dt. \end{aligned}$$

The supremum in the last line above exists for every $m \geq 0$ for any $\epsilon > 0$. Hence, convergence reduces to finding a choice of ϵ sufficiently small such that $\|f e^{\epsilon t} - f_N e^{\epsilon t}\|_1 \rightarrow 0$. Note that for any $\epsilon > 0$, we have

$$(A.11) \quad \begin{aligned} \int_0^\infty |f(t) e^{\epsilon t} - f_N(t) e^{\epsilon t}| dt &\leq \int_0^\infty |f(t) - f_N(t)| dt + \\ &\leq \int_0^\infty |f(t) - f_N(t)| (e^{\epsilon t} - 1) dt. \end{aligned}$$

The leftmost integral above can be made arbitrarily small for a sufficiently large choice of N by the assumed convergence of the Ritz approximation in the L_1 norm. To prove that the rightmost integral can be made arbitrarily small, note that

$$(A.12) \quad \begin{aligned} \|(f - f_N)(e^{\epsilon t} - 1)\|_1 &= \int_0^T |f(t) - f_N(t)| (e^{\epsilon t} - 1) dt \\ &\quad + \int_T^\infty |f(t) - f_N(t)| (e^{\epsilon t} - 1) dt \end{aligned}$$

for any $T > 0$. Now, since $f(t) = O(e^{-\gamma t})$ and $f_N(t)$ is a finite sum of exponentially decaying functions, there exists $\gamma' > 0$ such that $f(t) - f_N(t) = O(e^{-\gamma' t})$. Hence $(f(t) - f_N(t))(e^{\epsilon t} - 1) = O(e^{(-\gamma' + \epsilon)t})$, which decays exponentially for ϵ sufficiently small. Hence, by choosing T sufficiently large, the rightmost integral in (A.12) can be made arbitrarily small for N sufficiently large. Now, observe that

$$(A.13) \quad \int_0^T |f(t) - f_N(t)| (1 - e^{\epsilon t}) dt \leq |e^{\epsilon T} - 1| \int_0^T |f(t) - f_N(t)| dt.$$

Since $\lim_{\epsilon \rightarrow 0} e^{\epsilon T} = 1$ for every fixed T , one can choose ϵ sufficiently to make the right-hand side of (A.13) arbitrarily small for sufficiently large N .

REFERENCES

- [1] A. ASTOLFI, *A new look at model reduction by moment matching for linear systems*, in Proceedings of the IEEE 46th Conference on Decision and Control, New Orleans, 2007, pp. 2367–2372.
- [2] O. BASHIR, K. WILLCOX, O. GHATTAS, B. VAN BLOEMEN WAANDERS, AND J. HILL, *Hessian-based model reduction for large-scale systems with initial condition inputs*, Internat. J. Numer. Methods Engrg., 73 (2008), pp. 844–868.
- [3] D. BERTSIMAS AND J. TSITSIKLIS, *Introduction to Linear Optimization*, Athena Scientific, Nashua, NH, 1997.
- [4] P. BHANSALI AND K. R. SANTARELLI, *L_1 -MOR: An automated model order reduction framework based on L_1 norm and moment-matching constraints*, to appear in CSRI Summer Proceedings 2009, Sandia National Laboratories, 2009.
- [5] B. BOND AND L. DANIEL, *Stabilizing schemes for piecewise-linear reduced order models via projection and weighting functions*, in Proceedings of the 2007 IEEE/ACM International Conference on Computer-Aided Design, San Jose, CA, 2007, pp. 860–867.
- [6] T. BUI-THANH, K. WILLCOX, AND O. GHATTAS, *Model reduction for large-scale systems with high-dimensional parametric input space*, SIAM J. Sci. Comput., 30 (2008), pp. 3270–3288.
- [7] C. COELHO, J. PHILLIPS, AND L. SILVEIRA, *Robust rational function approximation algorithm for model generation*, in Proceedings of the 36th IEEE/ACM Design Automation Conference, New Orleans, 1999, pp. 207–212.
- [8] N. DONG AND J. ROYCHOWDHURY, *General-purpose nonlinear model-order reduction using piecewise-polynomial representations*, IEEE Trans. Comput. Aided Des. Integr. Circuits Syst., 27 (2008), pp. 249–264.
- [9] R. A. EL-ATTAR AND M. VIDYASAGAR, *Order reduction by l_1 and l_∞ norm minimisation*, IEEE Trans. Automat. Control, 23 (1978), pp. 731–734.
- [10] P. FELDMANN AND R. W. FREUND, *Efficient linear circuit analysis by Padé approximation via the Lanczos process*, IEEE Trans. Comput. Aided Des. Integr. Circuits Syst., 14 (1995), pp. 639–649.
- [11] E. GAD AND M. NAKHLA, *Model reduction for DC solution of large nonlinear circuits*, in Proceedings of the 1999 IEEE/ACM International Conference on Computer-Aided Design, San Jose, CA, 1999, pp. 376–379.
- [12] K. GLOVER, *All optimal Hankel-norm approximations of linear multivariable systems and their L^∞ -error bounds*, Internat. J. Control, 39 (1984), pp. 1115–1193.
- [13] S. GUGERCIN AND A. C. ANTOULAS, *A comparative study of 7 algorithms for model reduction*, in Proceedings of the 39th IEEE Conference on Decision and Control, Sydney, 2000, pp. 2367–2372.
- [14] S. GUGERCIN, A. C. ANTOULAS, AND C. A. BEATTIE, *H_2 model reduction for large-scale linear dynamical systems*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 609–638.
- [15] B. GUSTAVSEN AND A. SEMLYEN, *Rational approximation of frequency domain responses by vector fitting*, IEEE Trans. Power Delivery, 14 (1999), pp. 1052–1061.
- [16] R. HAKVOORT, *Worst-case system identification in l_1 : Error bounds, optimal models and model reduction*, in Proceedings of the 31st IEEE Conference on Decision and Control, Tucson, AZ, 1992, pp. 499–504.
- [17] H. KHALIL, *Nonlinear Systems*, 2nd ed., Prentice-Hall, Englewood cliffs, NJ, 1996.
- [18] L. KNOCKAERT, *On orthonormal Müntz-Laguerre filters*, IEEE Trans. Signal Process., 49 (2001), pp. 790–793.
- [19] B. C. MOORE, *Principal component analysis in linear systems: Controllability, observability and model reduction*, IEEE Trans. Automat. Control, 26 (1981), pp. 17–32.
- [20] N. NAKHLA, M. NAKHLA, AND R. ACHAR, *Sparse and passive reduction of massively coupled large multiport interconnects*, in Proceedings of the 2007 IEEE/ACM International Conference on Computer-Aided Design, San Jose, 2007, pp. 622–626.
- [21] A. ODABASIOGLU, M. CELIK, AND L. PILEGGI, *PRIMA: Passive reduced-order interconnect macromodeling algorithm*, IEEE Trans. Comput. Aided Des. Integr. Circuits Syst., 17 (1998), pp. 645–654.
- [22] J. PHILLIPS, L. DANIEL, AND L. M. SILVEIRA, *Guaranteed passive balancing transformations for model order reduction*, IEEE Trans. Comput. Aided Des. Integr. Circuits Syst., 22 (2003), pp. 1027–1041.
- [23] M. REWIENSKI AND J. WHITE, *A trajectory piecewise-linear approach to model order reduction and fast simulation of nonlinear circuits and micromachined devices*, in Proceedings of the 2001 IEEE/ACM International Conference on Computer-Aided Design, San Jose, 2001, pp. 252–257.

- [24] J. ROYCHOWDHURY, *Reduced-order modeling of time-varying systems*, IEEE Trans. Circuits Systems II: Analog Digital Signal Process., 46 (1999), pp. 1273–1288.
- [25] O. A. SEBAKHY AND M. N. ALY, *Discrete-time model reduction with optimal zero locations by norm minimisation*, IEE Proc. Control Theory Appl., 145 (1998), pp. 499–506.