SAND REPORT

SAND2010-8815
Unlimited Release
Printed December 2010

Solution Methods for Very Highly Integrated Circuits

Yao Chen, Ting Mei, Ryan Nong, Keith Santarelli,
Heidi Thornquist, and Ray Tuminaro

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed
and operated by Sandia Corporation, a wholly owned subsidiary of
Lockheed Martin Corporation, for the U.S. Department of Energys
National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online ordering:	http://www.ntis.gov/ordering.htm

SAND2010-8815
Unlimited Release
Printed December 2010

Solution Methods for Very Highly Integrated Circuits

Yao Chen, Ting Mei, Ryan Nong, Keith Santarelli,
Heidi Thornquist, and Ray Tuminaro

Sandia National Laboratories
P.O. Box 5800
Mail Stop 0316
Albuquerque, NM 87185-0316

Abstract

While advances in manufacturing enable the fabrication of integrated circuits containing tens-to-hundreds of millions of devices, the time-sensitive modeling and simulation necessary to design these circuits poses a significant computational challenge. This is especially true for mixed-signal integrated circuits where detailed performance analyses are necessary for the individual analog/digital circuit components as well as the full system. When the integrated circuit has millions of devices, performing a full system simulation is practically infeasible using currently available Electrical Design Automation (EDA) tools. The principal reason for this is the time required for the nonlinear solver to compute the solutions of large linearized systems during the simulation of these circuits.

The research presented in this report aims to address the computational difficulties introduced by these large linearized systems by using Model Order Reduction (MOR) to (i) generate specialized preconditioners that accelerate the computation of the linear system solution and (ii) reduce the overall dynamical system size. MOR techniques attempt to produce macromodels that capture the desired input-output behavior of larger dynamical systems and enable substantial speedups in simulation time. Several MOR techniques that have been developed under the LDRD on "Solution Methods for Very Highly Integrated Circuits" will be presented in this report. Among those presented are techniques for linear time-invariant dynamical systems that either extend current approaches or improve the time-domain performance of the reduced model using novel error bounds and a new approach for linear time-varying dynamical systems that guarantees dimension reduction, which has not been proven before. Progress on preconditioning power grid systems using multi-grid techniques will be presented as well as a framework for delivering MOR techniques to the user community using Trilinos and the Xyce circuit simulator, both prominent world-class software tools.

Contents

1 Introduction

While advances in manufacturing enable the fabrication of integrated circuits containing tens-to-hundreds of millions of devices, the time-sensitive modeling and simulation necessary to design these circuits poses a significant computational challenge. This is especially true for mixed-signal integrated circuits where detailed performance analyses are necessary for the individual analog/digital circuit components as well as the full system. When the integrated circuit has millions of devices, performing a full system simulation is practically infeasible using currently available Electrical Design Automation (EDA) tools. The principal reason for this is the time required for the nonlinear solver to compute the solutions of large linearized systems during the simulation of these circuits.

The research presented in this report aims to address the computational difficulties introduced by these large linearized systems by using Model Order Reduction (MOR) to (i) generate specialized preconditioners that accelerate the computation of the linear system solution and (ii) reduce the overall dynamical system size. MOR techniques attempt to produce macromodels that capture the desired input-output behavior of larger dynamical systems and enable substantial speedups in simulation time. Several MOR techniques that have been developed under the LDRD on ”Solution Methods for Very Highly Integrated Circuits” will be presented in this report. Among those presented are techniques for linear time-invariant (LTI) dynamical systems that either extend current approaches or improve the time-domain performance of the reduced model using novel error bounds and a new approach for linear time-varying dynamical systems that guarantees dimension reduction, which has not been proven before. Progress on preconditioning power grid systems using multi-grid techniques will be presented as well as a framework for delivering MOR techniques to the user community using Trilinos and the Xyce circuit simulator, both prominent world-class software tools.

1.1 Analog Circuit Simulation

In analog circuit simulation, the circuit is represented as a system of coupled DAE's, which are obtained from the enforcement of Kirchhoff's current and voltage laws (KCL and KVL, respectively) across an electrical network. The resulting system of differential algebraic equations (DAE) has the following form:

$$f(x(t)) + \frac{dq(x(t))}{dt} = b(t). \tag{1}$$

Simulation of this transient equation results in linear systems of the form:

$$(G + Q/\delta t)\delta x = (b - f)/\delta t \tag{2}$$

involving the conductance matrix $G(t) = \frac{df}{dx}(x(t))$, and the capacitance matrix $Q(t) = \frac{dq}{dx}(x(t))$. In general, there are several mathematical formulations that may be used to produce this DAE system, but in practice, nearly all circuit simulators such as SPICE [15], use the “modified KCL” formulation [23]. This is also the formulation used by **Xyce** [12].

1.2 Model Order Reduction

The goals of any MOR technique are to (i) automatically and efficiently generate an inexpensive macromodel that captures the desired input-output behavior of the larger dynamical system and (ii) enable substantial speedups in simulation time. While these techniques have been successfully used in industry on linear dynamical systems, MOR techniques for parametrized and weakly/strongly nonlinear systems are an active area of research. The vast majority of current MOR methods are projection based, meaning that a macromodel of the large-scale dynamical system is generated by projecting it onto some low-dimensional subspace.

Projection based MOR methods generate their subspace using either a moment matching based method (Krylov subspace methods) or SVD based method (balanced realization, proper orthogonal decomposition). Moment matching methods use Krylov subspace algorithms to generate subspaces that include a certain number of moments from the transfer function of the original dynamical system. The resulting macromodels lack provable error bounds and any guarantee of preserving stability or passivity of the original model. Nonetheless, moment matching algorithms like Padé via Lanczos (PVL [5]), Passive Reduced-order Interconnect Macromodeling Algorithm (PRIMA [17]), and Structure-Preserving Reduced-Order Interconnect Macromodeling (SPRIM [6]) are popular in the electrical engineering community. Truncated Balanced Realization (TBR) methods generate macromodels that have provable error bounds and guaranteed preservation of stability when the dynamical system is linear time-invariant. However, the numerical cost of generating the macromodel is high due to the solution of Lyapunov equations, which makes TBR methods impractical for large-scale dynamical systems. Proper Orthogonal Decomposition (POD) computes the projection subspace from time or frequency domain solutions, called snapshots, of the original dynamical system. The macromodel generated using POD is not guaranteed to preserve stability.

Nonlinear MOR techniques usually approximate the nonlinear system using linearization, polynomial (Taylor) expansion, or functional (Volterra) expansion and then apply moment matching or SVD-based methods on the approximation to compute a macromodel. A natural consequence of using linearization or expansion series on a nonlinear model is that these approaches generate macromodels that are only valid within a region of the initial operating point. Therefore, this approach can only be applied to weakly nonlinear dynamical systems and the resulting macromodels are not suitable for large input disturbances.

Simulating very highly integrated circuits, where the number of devices is in the tens-to-hundreds of millions, requires a novel application of MOR methods. Generating a macromodel of the whole circuit may not be computationally efficient, especially if the full system simulation of this circuit is part of a design phase where it is likely to change. Commercial EDA tools use linear MOR techniques to expedite the modeling of circuit interconnects, which can be done automatically from the circuit level descriptions before the full system simulation is performed. Using linear/nonlinear MOR techniques on the individual devices or logically associated groups of devices in a circuit is a novel extension of the current use of MOR in circuit simulation. This approach would reduce the system of coupled DAEs to a system of macromodels and has the potential of being done automatically using the circuit level descriptions. Furthermore, this approach allows different, device-specific MOR techniques to be combined in the full circuit simulation as well as the reuse of macromodels for future design cycles. In the end, this approach reduces the overall dimension of the dynamical system but the simulation of this system will still require the use of parallel solution methods, which in turn will necessitate efficient and scalable preconditioned iterative linear solvers.

1.3 Preconditioning

The parallel performance of full system simulations is hindered by the lack of efficient and scalable preconditioners for the iterative linear solver. Although specific features of large scale circuit models are exploited by current preconditioners, the present state-of-the-art eventually treats the resulting algebraic problem as a single large system. The algorithmic complexity of these preconditioners prohibits massive scalability because eventually the cost of applying the preconditioner dominates the total solution time. Scalable preconditioning techniques ultimately require some multi-level concepts, like those used by multi-grid methods.

2 MOR for LTI Systems

Performing model order reduction on linear time-invariant systems essentially starts with the state space realization of the original LTI system $\mathbf{\Sigma} \equiv (\mathbf{C}, \mathbf{G}, \mathbf{B}, \mathbf{L})$ as in

$$\begin{array}{l} \mathbf{C}\frac{d\mathbf{x}}{dt} = -\mathbf{G}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \\ \mathbf{y}(t) = \mathbf{L}^T\mathbf{x}(t), \end{array} \tag{3}$$

where $\mathbf{C} \in \mathbb{R}^{n\times n}$, $\mathbf{G} \in \mathbb{R}^{n\times n}$, $\mathbf{B} \in \mathbb{R}^{n\times p}$ and $\mathbf{L} \in \mathbb{R}^{n\times p}$ and $\mathbf{x}(t)$ is the state, $\mathbf{u}(t)$ the input and $\mathbf{y}(t)$ the output of the system. Also, n is the size of the original system and p the number of inputs (outputs). If $p = 1$, then (3) is referred to as a single-input-single-output (SISO) system and, if $p > 1$, it is a multiple-input-multiple-output (MIMO) system.

For model order reduction, one constructs two projection matrices $\mathbf{W} \in \mathbb{R}^{n\times k}$ and $\mathbf{V} \in \mathbb{R}^{n\times k}$ such that $\mathbf{W}^T\mathbf{V} = \mathbf{I}_k$, where k is the desired size of the reduced system $(k \ll n)$. The reduced system is now $\hat{\mathbf{\Sigma}} \equiv (\hat{\mathbf{C}}, \hat{\mathbf{G}}, \hat{\mathbf{B}}, \hat{\mathbf{L}})$ governed by the following set of first-order LTI differential equations

$$\begin{array}{l} \hat{\mathbf{C}}\frac{d\hat{\mathbf{x}}}{dt} = -\hat{\mathbf{G}}\hat{\mathbf{x}}(t) + \hat{\mathbf{B}}\mathbf{u}(t) \\ \hat{\mathbf{y}}(t) = \hat{\mathbf{L}}^T\hat{\mathbf{x}}(t), \end{array} \tag{4}$$

where $\hat{\mathbf{C}} = \mathbf{W}^T\mathbf{C}\mathbf{V}$, $\hat{\mathbf{G}} = \mathbf{W}^T\mathbf{G}\mathbf{V}$, $\hat{\mathbf{B}} = \mathbf{W}^T\mathbf{B}$, $\hat{\mathbf{L}}^T = \mathbf{L}^T\mathbf{V}$. For more details, we refer the readers to [1], [7] and the references therein.

The frequency input-output relationships of the original (3) and reduced (4) systems are determined by their corresponding transfer functions:

$$\begin{aligned} \mathbf{H}(s) &= \mathbf{L}^T(s\mathbf{C} + \mathbf{G})^{-1}\mathbf{B}, \\ \hat{\mathbf{H}}(s) &= \hat{\mathbf{L}}^T(s\hat{\mathbf{C}} + \hat{\mathbf{G}})^{-1}\hat{\mathbf{B}}. \end{aligned} \tag{5}$$

2.1 Moment Matching Scheme

Consider the Laurent expansions of the transfer functions (5) of the original and reduced systems about a given point $s_0 \in \mathbb{C}$ as follows:

$$\begin{aligned} \mathbf{H}(s_0 + \sigma) &= \eta_0 + \eta_1\sigma + \eta_2\sigma + \dots, \\ \hat{\mathbf{H}}(s_0 + \sigma) &= \hat{\eta}_0 + \hat{\eta}_1\sigma + \hat{\eta}_2\sigma + \dots, \end{aligned}$$

where η_i and $\hat{\eta}_i$ are the moments of $\mathbf{\Sigma}$ and $\hat{\mathbf{\Sigma}}$ at s_0 respectively. The moment matching based methods aim to compute a reduced system $\hat{\mathbf{\Sigma}}$, with a certain number of moments matching those of the original system $\mathbf{\Sigma}$, i.e.,

$$\eta_i = \hat{\eta}_i, \quad i = 1, \dots, l,$$

for some $l \ll n$. Note that l does not necessarily equal k/p. In the survey performed during the first year of the LDRD [16], two moment matching techniques were considered: PRIMA and RKS. Since these techniques can be implemented iteratively, they are quite numerically efficient. However, global error bounds are not available.

PRIMA

The Passive Reduced-Order Interconnect Macromodeling Algorithm (PRIMA) is proposed by Odabasioglu et al. [17] in 1998. The algorithm utilizes the block Arnoldi procedure. Note that for PRIMA, $\mathbf{W} = \mathbf{V}$. The resulting reduced system $\hat{\mathbf{\Sigma}}$ is proven to be passive and hence, stable. The number of matched moments is equal to the desired size of the reduced system $\hat{\mathbf{\Sigma}}$ divided by the number of inputs, i.e., $l = k/p$. While no global error bound is available, Heres [10] provides some heuristic considerations for error control of PRIMA in his Ph.D. thesis. Based on the numerical experiments presented in the survey [16], the quality of the error estimation depends very highly on the interpolation points and hence, it is local.

RKS

The Rational Krylov Subspace (RKS) method for model order reduction is proposed by Skoogh [28] and is based on the rational Krylov algorithm by Ruhe [24]. The rational Krylov algorithm is a generalization of the standard Arnoldi and Lanczos methods. The advantage of the rational Krylov algorithm is that it provides the flexibility of choosing a set of m different interpolation points ($m \leq k/p$). The reduced system $\hat{\mathbf{\Sigma}}$ matches $l = k/p$ moments of the original system $\mathbf{\Sigma}$ at these interpolation points. Reduced models resulting from RKS are not guaranteed to be passive and stable, and also no global error bound is available.

2.2 Balanced Truncation Reduction

The balanced truncation reduction is classified as an SVD-based scheme. The scheme constructs the reduced model $\hat{\mathbf{\Sigma}}$ based on the Hankel singular values of the original system $\mathbf{\Sigma}$. For the LTI system $\mathbf{\Sigma}$ as in (3), the Hankel singular values can be computed by solving the following two generalized Lyapunov equations for the system Grammians $\mathcal{P}$ and $\mathcal{Q}$:

$$\begin{aligned} \mathbf{G}\mathcal{P}\mathbf{C}^T + \mathbf{C}\mathcal{P}\mathbf{G}^T &= \mathbf{B}\mathbf{B}^T \\ \mathbf{G}^T\mathcal{Q}\mathbf{C} + \mathbf{C}^T\mathcal{Q}\mathbf{G} &= \mathbf{L}\mathbf{L}^T. \end{aligned} \tag{6}$$

Then the Hankel singular values of $\mathbf{\Sigma}$ are $\sigma_i(\mathbf{\Sigma}) = \sqrt{\lambda_i(\mathcal{P}\mathcal{Q})}, i = 1, \ldots, n$, the square roots of the eigenvalues of the product of the system Grammians. The projection matrices can be computed using the system Grammians and the reduced system $\hat{\mathbf{\Sigma}}$ results balanced. In addition, the reduced system $\hat{\mathbf{\Sigma}}$ has the following guaranteed properties: (a) stability is preserved, and (b) global error bounds exist in Hankel-norm approximation and they can be computed as follows:

$$\sigma_{k+1} \leq \|\mathbf{\Sigma} - \hat{\mathbf{\Sigma}}\|_\infty \leq 2(\sigma_{k+1} + \ldots + \sigma_n),$$

where k is the desired size of the reduced system $\hat{\mathbf{\Sigma}}$.

Despite the advantageous properties, MOR via balanced truncation is not very attractive due to its computational requirements. The reason is in directly solving the two Lyapunov equations (6); as n gets large, the complexity in computation and storage required is prohibitive. A number of efforts have been made to solve the Lyapunov equations iteratively, which were not examined in the survey [16], instead the following technique is considered: GSHSR. This technique is a generalization of balanced truncation model order reduction to descriptor systems.

GSHSR

The Generalized Schur-Hammarling Square Root (GSHSR) method is proposed by Mehrmann and Stykel [14]. The essence of the algorithm is to decouple the descriptor system (3) into its proper and improper portions and then reduce each of the portions separately. In addition to truncating the states that are difficult to control and/or to observe, GSHSR also removes those that are uncontrollable and/or unobservable. The algorithm utilizes a collection of solvers for (generalized) Lyapunov and (generalized) Sylvester matrix equations. Specifically, as suggested in [14], to solve the generalized Sylvester equations, the generalized Schur method by Kågström and Westin [11] is used. The upper triangular Cholesky factors of the solutions of the generalized Lyapunov equations can be determined without computing the solutions themselves by using the generalized Hammarling method by Hammarling [9] and Penzl [18].

As mentioned above, GSHSR is a balanced truncation method. Therefore, there exist global error bounds for the approximation $\hat{\mathbf{\Sigma}}$. Since the error bounds are in terms of the Hankel singular values of $\mathbf{\Sigma}$, the quality of the approximation depends on the decay of the Hankel singular values. In other words, GSHSR is effective in model order reduction if the

Hankel singular values of Σ decay rapidly. In fact, this is a common feature of all of the balanced truncation techniques.

2.3 Optimal $\mathcal{H}_2$

The optimal $\mathcal{H}_2$ is classified as an SVD-based scheme in [1]. The reason is that it solves the following model order reduction problem: Given a stable system $\mathbf{\Sigma}$, an approximation $\hat{\mathbf{\Sigma}}$ is sought to satisfy the following conditions:

$$\sigma_{k+1}(\mathbf{\Sigma}) \leq \|\mathbf{\Sigma} - \hat{\mathbf{\Sigma}}\|_{\mathcal{H}_2} \leq \epsilon < \sigma_k(\mathbf{\Sigma}).$$

Therefore, the construction of the reduced system $\hat{\mathbf{\Sigma}}$ in this framework can mimic the procedure as presented in Section 2.2. However, in 2008, Gugercin et al. [7] observe the equivalence of the local optimality conditions for the model order reduction problem in the two different frameworks: interpolation-based and Lyapunov-based. This result gives birth to a new direction using the interpolation properties to construct an approximation $\hat{\mathbf{\Sigma}}$ without solving the two Lyapunov equations. In the survey [16], IRKA, a technique that takes advantage of this result, is extended for descriptor systems.

IRKA

The Iterative Rational Krylov Algorithm (IRKA) is proposed by Gugercin et al. [7]. The original algorithm as shown in [7] is proposed to work with non-descriptor systems. As mentioned above, by using the interpolation properties, constructing an approximation $\hat{\mathbf{\Sigma}}$ can be done iteratively and without solving the two Lyapunov equations.

Despite these computational advantages, IRKA does not guarantee stability for the approximation $\hat{\mathbf{\Sigma}}$. The success of the iteration depends on the convergence of shifts, which is unpredictable and depends on initial guesses. In addition, for MIMO descriptor systems where ill-conditioned generalized eigenvalue problems have to be solved at each iteration, shifts at infinity have to be taken care of. Essentially, one needs to remove the subspace corresponding to the shifts at infinity from the computation. Our current approach is to identify the unwanted subspace at each iteration and then replace it by some random subspace of the same dimension. This subspace replacement approach retains the size of the successive reduced systems. However, the collection of shifts keeps getting polluted by the replacements, which makes it very hard to achieve any convergence in shifts.

One resolution may be to remove the unwanted subspace without replacement. A drawback of this subspace removal approach is that the sizes of the successive reduced systems get smaller and smaller, reducing the system to an unacceptable size before any convergence of shifts can be observed. Another potential approach to resolving the issue of shifts at infinity may be to decouple the proper and improper portions of the descriptor system (3) similar to the approach in GSHSR and then reduce each of the portions separately.

With the current implementation with the subspace replacement approach, the IRKA algorithm for MIMO descriptor systems exhibits very unpredictable behavior with regards to the convergence of shifts, stability and formation of reduced systems. For more information on the extension of IRKA to descriptor systems and the results of this extension, see the survey paper [16].

2.4 Mixed Moment Matching and Peak Error Objectives

If possible, it is clearly desirable to develop MOR techniques 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 [19] 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 [7] 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., $\mathbf{H}(-s_l) = \hat{\mathbf{H}}(-s_l)$ where $s_l \in \mathbb{C}$ is a pole of the reduced order model $\hat{\mathbf{H}}(s)$). 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 [2] 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. The result was 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 performed on this topic is that error is measured in a manner that is more useful for designers than the standard measures of error. Consider an example to illustrate the main issue along with a proposed resolution.

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

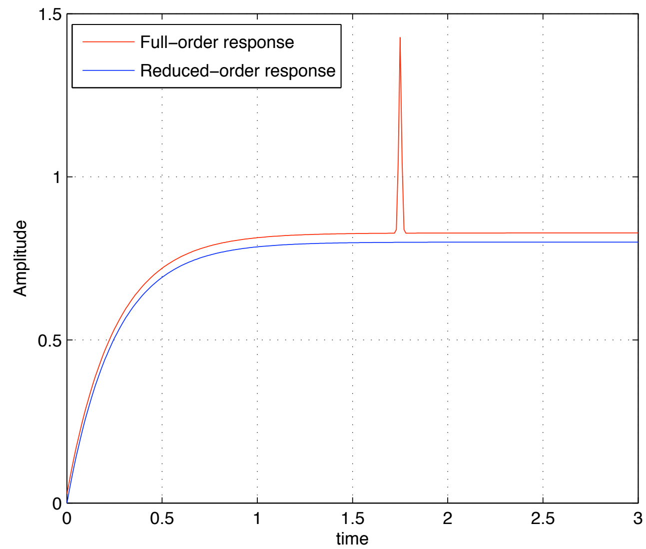

Figure 1. Hypothetical responses of an original and reduced order system produced via an SVD-based method.

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, consider right-sided continuous-time signals $\mathbf{y} : [0, \infty) \to \mathbb{R}$, then the peak amplitude can be taken as the standard infinity norm:

$$||\mathbf{y}||_\infty = \sup_{t \geq 0} |\mathbf{y}(t)|. \tag{7}$$

In the context of model reduction, define $\mathbf{y}(t)$ as the response of an original system and $\hat{\mathbf{y}}(t)$ as the response of a reduced order system for an identical input $\mathbf{u}(t)$, it is reasonable to desire that $||\mathbf{y} - \hat{\mathbf{y}}||_\infty$ be a small quantity. Indeed, if for a particular pair $\mathbf{y}(t)$ and $\hat{\mathbf{y}}(t)$ define $\epsilon = ||\mathbf{y} - \hat{\mathbf{y}}||_\infty$, then it immediately follows from the definition in Eqn. 7 that

$$|\mathbf{y}(t) - \hat{\mathbf{y}}(t)| \leq \epsilon \quad \forall t \geq 0. \tag{8}$$

Fig. 2 depicts the meaning of Eqn. 8 graphically. In the figure, the black signal represents the response of the original system $\mathbf{y}(t)$, and the surrounding area denoted “error region” represents a desired region in which one would like the response of a corresponding reduced order model $\hat{\mathbf{y}}(t)$ to lie. In the context of Eqn. 8, the “height” of the error region at every given time t is 2ϵ, indicating the desire for $\hat{\mathbf{y}}(t)$ to be close to $\mathbf{y}(t)$ uniformly over all times.

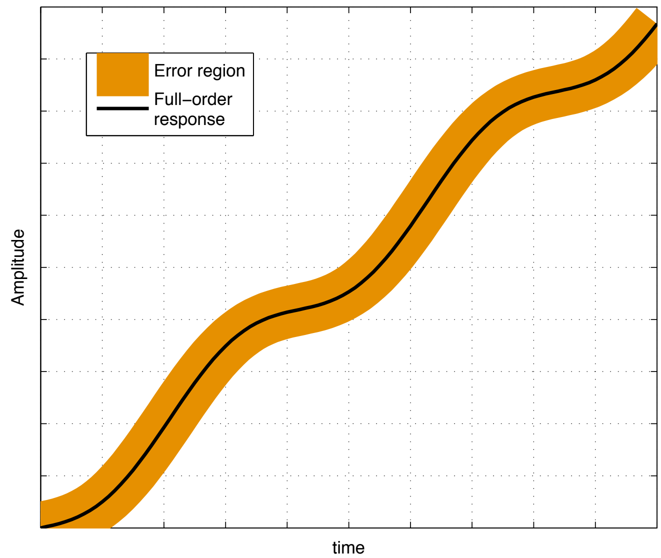

Figure 2. Depiction of full-order output signal surrounded by an "error region".

Problem Formulation: L_1 Norm minimization

As part of the LDRD, for LTI systems, it was endeavored to develop bounds of the following nature: denote by $L^\infty(\mathbb{R}^+)$

$$L^\infty(\mathbb{R}^+) = \left\{ \mathbf{u} : [0, \infty) \rightarrow \mathbb{R} : \sup_{t \geq 0} |\mathbf{u}(t)| < \infty \right\} \tag{9}$$

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

$$||\mathbf{y} - \hat{\mathbf{y}}||_\infty \leq M||\mathbf{u}||_\infty. \tag{10}$$

If such a bound exists for an original system model and a reduced system model for *every* bounded input $\mathbf{u}$, 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 $\mathbf{y}(t)$ and $\hat{\mathbf{y}}(t)$ will never be more than a fixed percentage of the peak input value. When the impulse operator of the original system is denoted by $\mathbf{h}(t)$ and the impulse operator of the reduced order system is denoted by $\hat{\mathbf{h}}(t)$, it is a well-known fact (see, for instance, [13]) that the *smallest* value of M as given in Eqn. 10 is the L_1 *norm* of the *error system* with impulse response $\mathbf{h}(t) - \hat{\mathbf{h}}(t)$:

$$||\mathbf{h} - \hat{\mathbf{h}}||_1 = \int_0^\infty |\mathbf{h}(t) - \hat{\mathbf{h}}(t)|dt. \tag{11}$$

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 $\hat{\mathbf{h}}(t)$ of order N for which $||\mathbf{h}-\hat{\mathbf{h}}||_1$ is small. Ideally, one would like to find that choice of $\hat{\mathbf{h}}(t)$ of order N such that the quantity $||\mathbf{h}-\hat{\mathbf{h}}||_1$ is *minimized*, and that is the essential viewpoint that taken here. While the problem of finding that choice of $\hat{\mathbf{h}}(t)$ which *globally* minimizes the L_1 norm of the error system is nonconvex and intractable to compute from a practical perspective, the research presented focuses on methods that search for local minimizers over a sufficiently rich set of choices for $\hat{\mathbf{h}}(t)$ 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 [4]. 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 ($||\mathbf{e}||_1 = \sum_{k=1}^{\infty} |\mathbf{e}_k|$) [27]. The closest work to the problem considered 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 [8]. While there are substantial differences with the class of problems being considered here as compared to [8], the underlying technique of casting such problems as linear programs is the same.

During the course of this LDRD, a new technique was developed for computing reduced order models via an attempt to minimize the L_1 norm of the corresponding error system $\mathbf{h}(t) - \hat{\mathbf{h}}(t)$. A major advantage of this technique 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 this 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 this approach, the technique 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. For more details on the L_1 norm based approach for reducing LTI systems and computational results, the reader is referred to [25, 26] and the references therein.

2.5 Integrating Reduced Models

Integrating the reduced model (4) into a larger circuit requires an admittance matrix to be generated from the original circuit (3). The admittance, or y-parameter, matrix provides the relationship between the input voltage and output current at any *port*, input or output node, of the original circuit. It is obtained by attaching a voltage source to every port of the original circuit so that every port becomes both an input and output node. The state space realization $(\mathbf{C}, \mathbf{G}, \mathbf{B}, \mathbf{L})$ of this modified circuit is then used, instead of the original, to generate the reduced model. As a result of this requirement for integration, even if the original circuit is SISO, the modified circuit used to generate the reduced model is always MIMO. This constraint will be taken into account during the presentation of the model order reduction schemes. For more details on reduced model integration, the reader is referred to [17] and the references therein.

3 MOR for LPTV Systems

Linear periodic time-varying (LPTV) systems arise in many electrical systems. Examples of LPTV systems include switched capacitor circuits, mixers, etc. For such systems, macromodels are normally constructed manually, which is heuristic and time consuming.

We present a novel algorithm which automatically abstracts the macromodels from the SPICE-level circuit description. Unlike previous LPTV macromodel techniques which use a time-invariant projection and produce a time-invariant reduced system, our method uses a time-varying projection, resulting in reduced time-varying systems which have the same form as the original system. This provides a key advantage during the simulation of the reduced systems: our method results in an LPTV reduced system which directly produces the "real" solution of the system. In previous approaches, the simulation is first performed on the reduced LTI system and some post-processing is required in order to obtain the "real" solution of the system.

Since the background for LPTV systems is not as commonly known as that of LTI systems, we will present some background material. This will motivate the discussion of the deficiencies in the previous techniques and the advantages of our new technique.

3.1 LPTV Systems Background

Consider a system driven by a large periodic signal $b_l(t)$ and a small signal $u(t)$ to produce an output $w(t)$. For simplicity, we assume that both $u(t)$ and $w(t)$ are scalars. The system can be described by the differential algebraic equations (DAEs)

$$\begin{aligned} \frac{dq(y)}{dt} + f(y) = &\quad b_l(t) + Bu(t), \\ w(t) = &\quad d^T y(t), \end{aligned} \tag{12}$$

where $y(t)$ is a vector of circuit unknowns (node voltages and branch currents), B and d are incidence vectors that capture the connection of the input to the output for the circuit.

To more conveniently derive the transfer function of LPTV systems, first recall the MPDE [3, 22] formulation of Eqn. 12, which separates the input and system time scales. This is given by

$$\begin{aligned} \left[\frac{\partial}{\partial t_1} + \frac{\partial}{\partial t_2}\right] q(\hat{y}(t_1,t_2)) + f(\hat{y}(t_1,t_2)) = &\quad b_l(t_2) + Bu(t_1), \\ \hat{w}(t_1,t_2) = &\quad d^T \hat{y}(t_1,t_2), \end{aligned} \tag{13}$$

where $\hat{y}(t_1,t_2)$ and $\hat{w}(t_1,t_2)$ are the bivariate forms of $y(t)$ and $w(t)$ in Eqn. 12.

Assume that $y^*(t_2)$ is the periodic steady state solution of Eqn. 13 when $u(t_1) = 0$. Linearizing this MPDE around $y^*(t_2)$, we obtain

$$\begin{aligned} \frac{\partial C(t_2)\hat{x}(t_1,t_2)}{\partial t_1} + \frac{\partial C(t_2)\hat{x}(t_1,t_2)}{\partial t_2} + G(t_2)\hat{x}(t_1,t_2) = &\quad Bu(t_1), \\ \hat{z}(t_1,t_2) = &\quad d^T \hat{x}(t_1,t_2). \end{aligned} \tag{14}$$

Here, $C(t_2) = (\partial q(\hat{y})/\partial \hat{y})|_{y^*}$ and $G(t_2) = (\partial f(\hat{y})/\partial \hat{y})|_{y^*}$ are periodic time-varying matrices. $\hat{x}(t_1, t_2)$ and $\hat{z}(t_1, t_2)$ are the small signal versions of $\hat{y}(t_1, t_2)$ and $\hat{w}(t_1, t_2)$, respectively.

Performing a Laplace transform with respect to t_1, we further obtain

$$\begin{aligned} sC(t_2)\hat{X}(s,t_2) + & \frac{\partial C(t_2)\hat{X}(s,t_2)}{\partial t_2} + G(t_2)\hat{X}(s,t_2) \\ & = BU(s), \\ \hat{Z}(s,t_2) & = d^T \hat{X}(s,t_2), \end{aligned} \tag{15}$$

where $\hat{X}(s,t_2)$ and $\hat{Z}(s,t_2)$ are the transformed variables.

If we define the differential operator

$$\frac{D}{dt_2}[.] = \frac{\partial [C(t_2).]}{\partial t_2}, \tag{16}$$

then the time-varying transfer function can be written in operator form

$$\begin{aligned} H(s,t_2) & = \frac{\hat{Z}(s,t_2)}{U(s)} \\ & = d^T \left[sC(t_2) + \frac{D}{dt_2}[.] + G(t_2) \right]^{-1} B. \end{aligned} \tag{17}$$

An LPTV system can be converted to an artificial LTI system by discretizing the periodic time variation using a finite basis [21, 20]. For example, we can expand the t_2 dependence in Eqn. 15 using a time domain Backward Euler finite difference basis. Define the following long vectors

$$\begin{aligned} \vec{X}_{TD}(s) = & \quad [\hat{X}_0(s)^T, \hat{X}_1(s)^T, \cdots, \hat{X}_{N-1}(s)^T]^T, \\ \vec{Z}_{TD}(s) = & \quad [\hat{Z}_0(s)^T, \hat{Z}_1(s)^T, \cdots, \hat{Z}_{N-1}(s)^T]^T, \\ \vec{B}_{TD}(s) = & \quad [B^T, B^T, \cdots, B^T]^T. \end{aligned} \tag{18}$$

and

$$D = \begin{bmatrix} d & & & \\ & d & & \\ & & \ddots & \\ & & & d \end{bmatrix}. \tag{19}$$

Here $\hat{X}_i(s) = \hat{X}(s,t_2)|_{t_2=t_i}$ and $\hat{Z}_i(s) = \hat{Z}(s,t_2)|_{t_2=t_i}$. $t_2 \in [0, T_2]$. N is the number of sample points.

Then a time domain matrix form of Eqn. 15 is given by

$$\begin{aligned} \left[s\mathcal{C}_{TD} + \mathcal{J}_{TD} \right] \vec{X}_{TD}(s) & = \vec{B}_{TD} U(s), \\ \vec{Z}_{TD}(s) & = \mathcal{D}^T \vec{X}_{TD}(s), \end{aligned} \tag{20}$$

where

$$\begin{aligned}
\mathcal{J}_{TD} &= \mathcal{G}_{TD} + \Delta \mathcal{C}_{TD} \\
\mathcal{C}_{TD} &= \begin{bmatrix} C_0 & & & \\ & C_1 & & \\ & & \ddots & \\ & & & C_{N-1} \end{bmatrix}, \\
\mathcal{G}_{TD} &= \begin{bmatrix} G_0 & & & \\ & G_1 & & \\ & & \ddots & \\ & & & G_{N-1} \end{bmatrix}, \\
\Delta &= \begin{bmatrix} \frac{1}{\delta_1} I & & & -\frac{1}{\delta_1} I \\ -\frac{1}{\delta_2} I & \frac{1}{\delta_2} I & & \\ & & & \\ & & -\frac{1}{\delta_{N-1}} I & \frac{1}{\delta_{N-1} I} \end{bmatrix},
\end{aligned} \tag{21}$$

$C_i = C(t_2)|_{t_2=t_i}$, $G_i = G(t_2)|_{t_2=t_i}$, and $\delta_i = t_{2,i} - t_{2,i-1}$. A convenient matrix representation for the time-varying transfer function 17 can be written if you define

$$\vec{H}_{TD}(s) = [H_0(s)^T, H_1(s)^T, \cdots, H_{N-1}(s)^T]^T, \tag{22}$$

where $H_i(s) = H(s, t_2)|_{t_2=t_i}$. This is given by

$$\vec{H}_{TD}(s) = \mathcal{D}^T \left[s\mathcal{C}_{TD} + \mathcal{J}_{TD} \right]^{-1} \vec{B}_{TD}. \tag{23}$$

Then 20 is an LTI system and can be reduced by LTI MOR techniques, like block-Krylov methods[17]. The transfer function 23 can be written in the form

$$\vec{H}_{TD}(s) = \mathcal{L}^T \left[I + s\mathcal{A} \right]^{-1} \mathcal{R}, \tag{24}$$

where

$$\mathcal{L} = \mathcal{D}, \mathcal{R} = \mathcal{J}_{TD}^{-1} \vec{B}_{TD}, \mathcal{A} = \mathcal{J}_{TD}^{-1} \mathcal{C}_{TD}. \tag{25}$$

Then Eqn. 25 can be used to generate reduced-order models using block-Krylov methods. For example, applying the block Arnoldi algorithm with matrices $\mathcal{A}$ and $\mathcal{R}$, we obtain a orthogonal projection matrix V_q and a block Hessenberg matrix T_q. The transfer function of the reduced-order model is given by

$$\vec{H}q(s) = \mathcal{L}^T V_q \left[I_q + sT_q \right]^{-1} V_q^T \mathcal{R}. \tag{26}$$

It can be shown that $\vec{H}q(s)$ approximates $\vec{H}_{TD}(s)$.

The matrix-based MOR methods described above rely on a fixed, a-priori discretization of the time-varying differential operators to convert an LPTV system into an LTI one. For the operator-based MOR methods, that use the transfer function 17, the discretization basis can be changed dynamically during the model-order reduction process. This is achieved by modifying the internals of Krylov-subspace methods to use general function space operators, instead of matrices. If a fixed discretization basis is used throughout the whole

Arnoldi process, then discretizing the differential operator before the Arnoldi process (the matrix-based MOR methods) and discretizing the operator during the Arnoldi process (the operator-based MOR methods) produce the same results. For example, if the Backward Euler finite difference basis is used in both methods to produce a q-th Krylov subspace, the projection matrix from both approaches can be written as

$$V_q = \begin{bmatrix} V_0 \\ V_1 \\ \vdots \\ V_{N-1} \end{bmatrix}, \tag{27}$$

where $V_i = V(t_2)|_{t_2=t_i}$ and N is the number of time points during discretization. Note that the long vector V_q has the same form as $\vec{X}_{TD}(s)$, $\vec{Z}_{TD}(s)$ and $\vec{B}_{TD}(s)$ in Eqn. 18.

3.2 Previous LPTV MOR Approaches

The previous approaches impose several limitations on the model reduction process. First, the time-varying structure of the system is not preserved. The original LPTV system is first converted to an LTI system and then reduced to an LTI system. Simulations are performed on the reduced system and additional post-processing is required to convert the LTI solutions to LPTV solutions.

Moreover, during the process of converting an LPTV system to an LTI one, the size of the system is enlarged by N, which is the number of discretization points. If we denote the size of the original LPTV system by n, then the size of the LTI system after discretization is $n \times N$, or nN. For example, in Eqn. 20, $\vec{X}_{TD}(s)$, $\vec{Z}_{TD}(s)$ and $\vec{B}_{TD}(s)$ are vectors of size $nN \times 1$. $\mathcal{C}_{TD}$ and $\mathcal{J}_{TD}$ are square matrices of size $nN \times nN$ and $\mathcal{D}$ is a rectangular matrix of size $nN \times N$. The LTI MOR techniques are performed on this enlarged system. As a result, the reduced LTI system can potentially have a size of up to nN in order to accurately approximate the original system. In other words, the reduced system can have a bigger size than that of the original system. If we denote the size of the reduced system by q, then $q \leq nN$. However, it is not guaranteed $q \leq n$. As demonstrated in section 3.4, the size of reduced system generated using previous approaches is larger than the original system size for some examples.

3.3 A Novel Structure Preserving MOR Technique for LPTV Systems

In this section, we present a structure preserving MOR technique for LPTV systems. This novel approach not only preserves the time-varying structure of the original system during the model reduction process, but also guarantees the reduced system is smaller than the original system.

The Algorithm

Instead of converting an LPTV system to an LTI one, the system and projection matrix have a time-varying structure. The block vector V_q in Eqn. 27 can be written in a time-varying form as $V(t_2)$ in which $V(t_2)$ evaluated at $t_2 = t_i$ corresponds to the i-th block of V_q. Using this time-varying projection matrix, $V(t_2)^T$, to multiply the first row in the time-varying system 15 and applying the change of variable $\hat{X}(t_1,t_2) = V(t_2)\hat{X}_q(t_1,t_2)$, we obtain

$$\begin{aligned} V(t_2)^T\tfrac{\partial C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_1} + \quad & V(t_2)^T\tfrac{\partial C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_2} + V(t_2)^T G(t_2)V(t_2)\hat{x}_q(t_1,t_2) \\ & = V(t_2)^T Bu(t_1), \\ \hat{z}_q(t_1,t_2) \quad & = d^T V(t_2)\hat{x}_q(t_1,t_2). \end{aligned} \tag{28}$$

Since $V(t_2)$ is not dependent on t_1,

$$V(t_2)^T\frac{\partial C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_1} = \frac{\partial V(t_2)^T C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_1}. \tag{29}$$

Applying the chain rule, gives

$$V(t_2)^T\frac{\partial C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_2} = \frac{\partial V(t_2)^T C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_2} - \frac{\partial V(t_2)^T}{\partial t_2}C(t_2)V(t_2)\hat{x}_q(t_1,t_2). \tag{30}$$

Inserting Eqn. 29 and 30 into Eqn. 28 yields

$$\begin{aligned} & \tfrac{\partial V(t_2)^T C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_1} + \tfrac{\partial V(t_2)^T C(t_2)V(t_2)\hat{x}_q(t_1,t_2)}{\partial t_2} \\ & + V(t_2)^T G(t_2)V(t_2)\hat{x}_q(t_1,t_2) - \tfrac{\partial V(t_2)^T}{\partial t_2}C(t_2)V(t_2)\hat{x}_q(t_1,t_2) = V(t_2)^T Bu(t_1), \\ & \hat{z}_q(t_1,t_2) = d^T V(t_2)\hat{x}_q(t_1,t_2). \end{aligned} \tag{31}$$

Using the definitions

$$\begin{aligned} \hat{C}(t_2) &= V(t_2)^T C(t_2)V(t_2), \\ \hat{G}(t_2) &= V(t_2)^T G(t_2)V(t_2) - \tfrac{\partial V(t_2)^T}{\partial t_2}C(t_2)V(t_2), \\ \hat{B}(t_2) &= V(t_2)^T B, \\ \hat{d}(t_2) &= V(t_2)^T d, \end{aligned} \tag{32}$$

the reduced system 31 can be rewritten as

$$\begin{aligned} \tfrac{\partial \hat{C}(t_2)\hat{x}_q(t_1,t_2)}{\partial t_1} + \tfrac{\partial \hat{C}(t_2)\hat{x}_q(t_1,t_2)}{\partial t_2} + \hat{G}(t_2)\hat{x}_q(t_1,t_2) &= \hat{B}(t_2)u(t_1), \\ \hat{z}_q(t_1,t_2) &= \hat{d}(t_2)^T\hat{x}_q(t_1,t_2). \end{aligned} \tag{33}$$

Note that the reduced system 33 has exactly the same form as the original MPDE system 14. Therefore, we can write the DAE form of the reduced system 33 as

$$\begin{aligned} \tfrac{d\hat{C}(t)x_q(t)}{dt} + \hat{G}(t)x_q(t) &= \hat{B}(t)u(t), \\ z_q(t) &= \hat{d}(t)^T x_q(t). \end{aligned} \tag{34}$$

Here $x_q(t)$ is a vector of size q. The q-th order reduced system 34 is an LPTV system, which is the same as the original system. Simulations are performed on this reduced system to directly generate LPTV solutions. No additional post-processing is required. Following the same procedure as in 3.1, we can write the time-varying transfer function of the reduced system in operator form

$$H_q(s,t_2) = \hat{d}(t_2)^T \left[s\hat{C}(t_2) + \frac{\partial \left[\hat{C}(t_2). \right]}{\partial t_2} + \hat{G}(t_2) \right]^{-1} \hat{B}(t_2). \tag{35}$$

Properties of this Technique

One of the main advantages of this method is that the time-varying structure is preserved throughout the whole reduction process, leading to an LPTV reduced system. By using a time-varying projection $V(t_2)$, the separation of the LPTV system time scale t_2 and signal time scale t_1 is preserved during the reduction process, which is not true for the previous matrix based approaches. Preserving the time-varying structure also led to another desirable property: the size of the reduced system is never bigger than the original system, i.e., $q \leq n$.

To understand the property of guaranteed dimension reduction, we first apply the Laplace transform on t_1 and write the matrix equivalent form of 28

$$\begin{aligned} \left[sV_{TD}^T \mathcal{C}_{TD} V_{TD} + V_{TD}^T \mathcal{J}_{TD} V_{TD} \right] \vec{X}_q(s) &= V_{TD}^T \vec{B}_{TD} U(s), \\ \vec{Z}_q(s) &= \mathcal{D}^T V_{TD} \vec{X}_q(s). \end{aligned} \tag{36}$$

where $\mathcal{C}_{TD}$ and $\mathcal{J}_{TD}$ are square matrices of size $nN \times nN$, defined in 21. $\vec{B}_{TD}(s)$ is a vector of size $nN \times 1$ and $\mathcal{D}$ is a rectangular matrix of size $nN \times N$, defined in 18 and 19, respectively. Vectors $\vec{X}_q(s)$ and $\vec{Z}_q(s)$ are of size $qN \times 1$, defined similarly to $\vec{X}_{TD}(s)$ and $\vec{Z}_{TD}(s)$ in 18. The projection matrix V_{TD} is a matrix of size $nN \times qN$ defined as

$$V_{TD} = \begin{bmatrix} V_0 & & & \\ & V_1 & & \\ & & \ddots & \\ & & & V_{N-1} \end{bmatrix} \tag{37}$$

where $V_i = V(t_2)|_{t_2=t_i}$. Note the blocks V_i in V_{TD} are the same as those in projection matrix 27 from previous methods. However, the blocks are arranged differently as a result of using a time-varying projection. This special structure preserves the time-varying structure during the reduction process.

When V_{TD} becomes square matrix, i.e., $q = n$, the transfer function of reduced system is

$$\begin{aligned} \vec{H_{qN}}(s) &= \mathcal{D}^T V_{TD} \left[sV_{TD}^T \mathcal{C}_{TD} V_{TD} + V_{TD}^T \mathcal{J}_{TD} V_{TD} \right]^{-1} V_{TD}^T \vec{B}_{TD} \\ &= \mathcal{D}^T V_{TD} \left[V_{TD}^T (s\mathcal{C}_{TD} + \mathcal{J}_{TD}) V_{TD} \right]^{-1} V_{TD}^T \vec{B}_{TD} \\ &= \mathcal{D}^T V_{TD} V_{TD}^{-1} \left[s\mathcal{C}_{TD} + \mathcal{J}_{TD} \right]^{-1} (V_{TD}^T)^{-1} V_{TD}^T \vec{B}_{TD} \\ &= \mathcal{D}^T \left[s\mathcal{C}_{TD} + \mathcal{J}_{TD} \right]^{-1} \vec{B}_{TD}. \end{aligned} \tag{38}$$

Therefore, the transfer function of reduced system becomes the same as that of the original system 23. In other words, we need a maximum of nN columns in V_{TD} to accurately approximate $\vec{H}_{TD}(s)$ i.e., $qN \leq nN$. Therefore, for our method

$$q \leq n. \tag{39}$$

3.4 LPTV Examples

In this section, we apply the structure preserving MOR technique to several LPTV systems and demonstrate the advantage of this method over previous methods.

A Simple Upconverter

A simple upconverter from [21] is shown in Fig. 3. It consists of a low-pass filter, an ideal mixer, and two bandpass filter stages. The purpose of showing this simple example is to validate the structure preserving method since the transfer function of this example can be analytically derived.

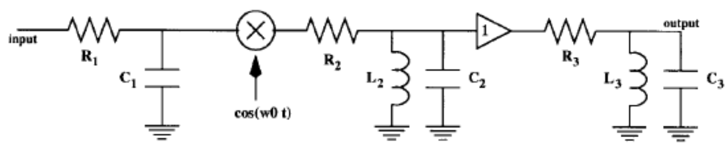

Figure 3. A simple upconverter.

Harmonic balance (HB) is used to discretize the LPTV system and consider the time-varying transfer function at $t_2 = 0$. Fig. 4 shows the comparison of the results from the original system and the reduced system using the structure preserving method. As can be seen in Fig. 4, both transfer functions match well across a broad frequency range.

A Double-Balanced Mixer

A double-balanced mixer from [29] is shown in Fig. 5. The mixer uses a local oscillator frequency of 10kHZ.

The comparison of results from the original system, the reduced system using the structure preserving method, and the reduced system using previous methods are in Fig. 6-7. We consider the second harmonic of the time-varying transfer function $H_2(s)$. The original system is an LPTV system of size 5. Using previous methods, which produce a reduced LTI system, a reduced system size of at least 6 is needed to accurately approximate the original system, as can be seen in Fig. 6. On the other hand, the reduced LPTV system from the structure preserving method with a size of 2 produces results that match well the results from the original system, as can be seen in Fig. 7.

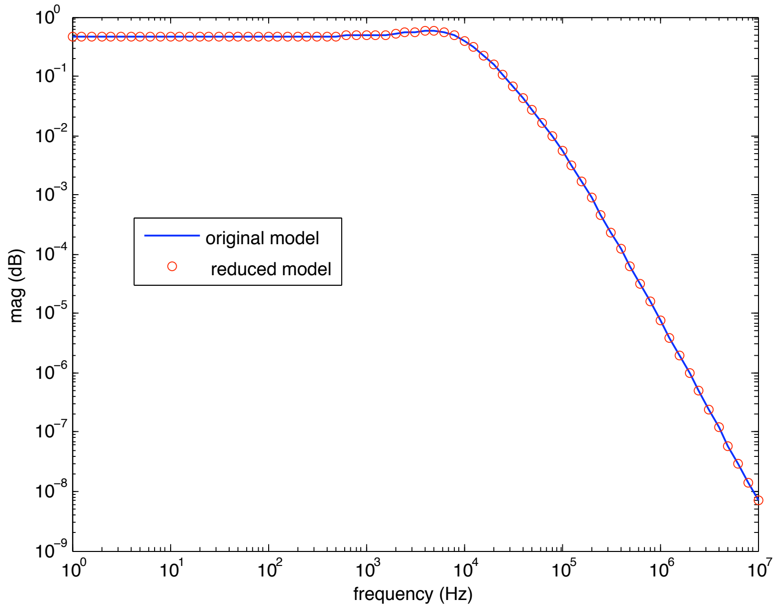

Figure 4. $H_0(s)$ comparison: the original system vs. the reduced system using the structure preserving method.

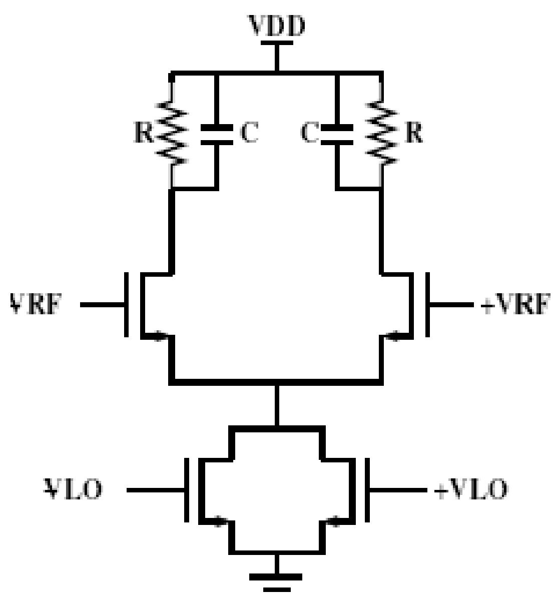

Figure 5. A double-balanced mixer.

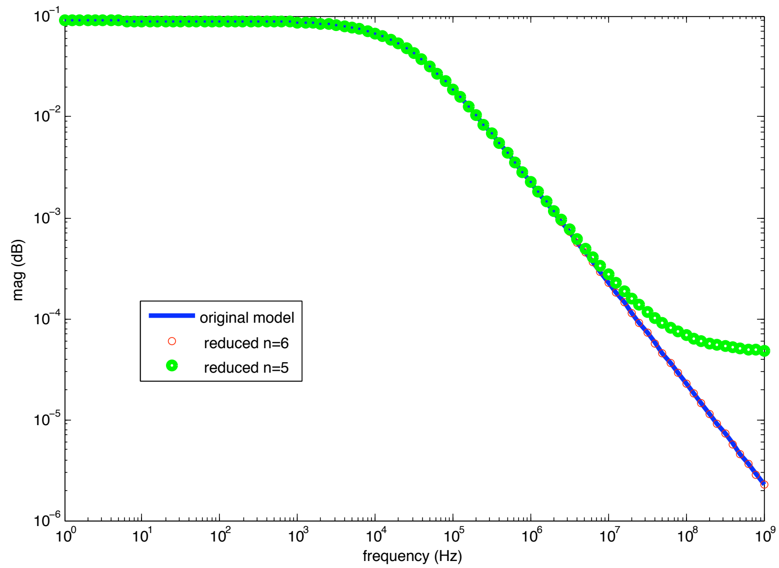

Figure 6. $H_2(s)$ comparison: the original system vs. the reduced system from a previous method.

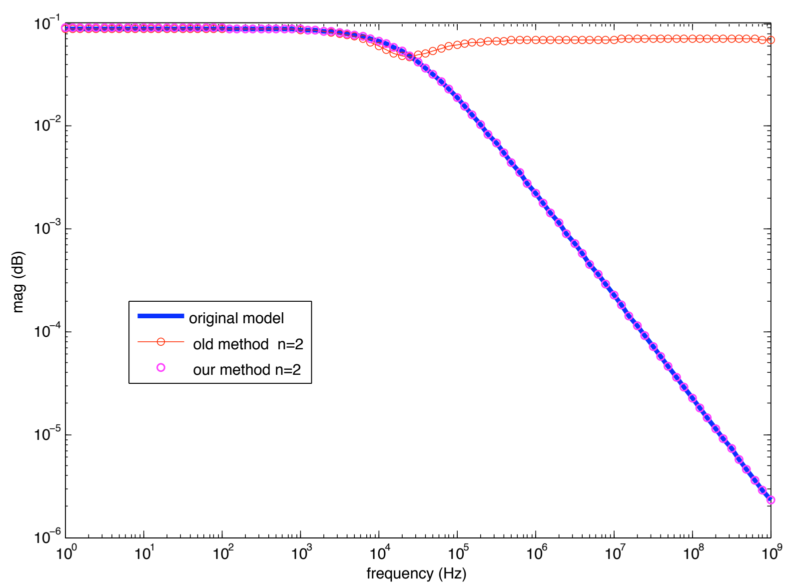

Figure 7. $H_2(s)$ comparison: the original system, the reduced system from a previous method and the reduced system using the structure preserving method.

4 Conclusions

In this paper we presented research in novel techniques to address the computational difficulties introduced by these large linearized systems by using Model Order Reduction (MOR) to (i) generate specialized preconditioners that accelerate the computation of the linear system solution and (ii) reduce the overall dynamical system size. MOR techniques attempt to produce macromodels that capture the desired input-output behavior of larger dynamical systems and enable substantial speedups in simulation time. Several MOR techniques that have been developed under the LDRD on ”Solution Methods for Very Highly Integrated Circuits” were presented in this report. Among those presented were techniques for linear time-invariant dynamical systems that either extend current approaches or improve the time-domain performance of the reduced model using novel error bounds and a new approach for linear time-varying dynamical systems that guarantees dimension reduction, which has not been proven before. Progress on preconditioning power grid systems using multi-grid techniques will be presented as well as a framework for delivering MOR techniques to the user community using Trilinos and the Xyce circuit simulator, both prominent world-class software tools.

References

[1] A. C. Antoulas, D. C. Sorensen, and S. Gugercin. A survey of model reduction methods for large-scale systems. *Structured Matrices in Mathematics, Computer Science and Engineering*, Vol. I. Comtemporaty Mathematics Series, 280:193–219, 2001.

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

[3] H. G. Brachtendorf, G. Welsch, R. Laur, and A. Bunse-Gerstner. Numerical steady state analysis of electronic circuits driven by multi-tone signals. *Electrical Engineering*, 79:103–112, 1996.

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

[5] Peter Feldmann and Roland W. Freund. Efficient linear circuit analysis by Padé approximation via the Lanczos process. *IEEE Transactions on Computer-Aided Design*, 14:137–158, 1993.

[6] Roland W. Freund. SPRIM: structure-preserving reduced-order interconnect macromodeling. *Proceedings of the 2004 IEEE/ACM International Conference on Computer-Aided Design*, pages 80–87, 2004.

[7] S. Gugercin, A. C. Antoulas, and C. Beattie. $\mathcal{H}_2$ model reduction for large-scale linear dynamical systems. *SIAM Journals on Matrix Analysis and Applications*, 30:609–638, 2008.

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

[9] S. J. Hammarling. Numerical solution of the stable, non-negative definite Lyapunov equation. *IMA Journal of Numerical Analysis*, 2:303–323, 1982.

[10] P. J. Heres. *Robust and Efficient Krylov Subspace Methods for Model Order Reduction*. PhD thesis, Eindhoven University of Technology, 2005.

[11] B. Kågström and L. Westin. Generalized Schur methods with condition estimators for solving the generalized Sylvester equation. *IEEE Transactions on Automatic Control*, 34:745–751, 1989.

[12] Eric R. Keiter, Scott A. Hutchinson, Robert J. Hoekstra, Lon J. Waters, and Thomas V. Russo. Xyce parallel electronic simulator design: Mathematical formulation, version 2.0. Technical Report SAND2004-2283, Sandia National Laboratories, Albuquerque, NM, June 2004.

[13] H. Khalil. *Nonlinear Systems*. Prentice-Hall, second edition, 1996.

[14] V. Mehrmann and T. Stykel. Balanced truncation model reduction for large-scale systems in descriptor form. *Dimension Reduction of Large-Scale Systems, Lect. Notes Comput. Sci. Eng.*, 45:83–115, 2005.

[15] L. W. Nagel. Spice 2, a computer program to simulate semiconductor circuits. Technical Report Memorandum ERL-M250, 1975.

[16] R. Nong and H. Thornquist. A survey of model order reduction methods for LTI systems in descriptor form. *CSRI Summer Proceedings 2008*, pages 14–32, 2008.

[17] A. Odabasioglu, M. Celik, and L. T. Pileggi. PRIMA: Passive reduced-order interconnect macromodeling algorithm. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17:645–654, 1998.

[18] T. Penzl. Numerical solution of generalized Lyapunov equations. *Advances in Computational Mathematics*, 8:33–48, 1998.

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

[20] J. R. Phillips. Projection-based approaches for model reduction of weakly nonlinear, time-varying systems. *IEEE Trans. CAD*, 22(2):171187, 2003.

[21] J. Roychowdhury. Reduced-order modelling of time-varying systems. *IEEE Trans. Ckts. Syst. II: Sig. Proc.*, 46(10):12731288, 1999.

[22] J. Roychowdhury. Analyzing circuits with widely separated time scales using numerical PDE methods. *IEEE Transactions on Circuits and Systems — I: Fundamental Theory and Applications*, 48(5):578–594, 2001.

[23] C. W. Ho A. E. Ruehli and P. A. Brennan. The modified nodal approach to network analysis. *IEEE Trans. Circuits Systems*, 22:505–509, 1988.

[24] A. Ruhe. Rational Krylov algorithms for nonsymmetric eigenvalue problems. ii. matrix pairs. *Linear Algebra and Its Applications*, 197, 198:283–295, 1994.

[25] Keith R. Santarelli. A framework for reduced order modeling with mixed moment matching and peak error objectives. Technical Report SAND2009-0196, Sandia National Laboratories, Albuquerque, NM, January 2009.

[26] Keith R. Santarelli. A framework for reduced order modeling with mixed moment matching and peak error objectives. *SIAM J. Sci. Comput.*, 32(2):745–773, 2010.

[27] O.A. Sebakhy and M.N. Aly. Discrete-time model reduction with optimal zero locations by norm minimisation. *IEEE Proc. Control Theory Appl.*, 145(6):499–506, 1998.

[28] D. Skoogh. A rational Krylov method for model order reduction. *Blåserien*, 47, 1998.

[29] Z. Chen Z. Zhang and J. Lau. A 900mhz cmos balanced harmonic mixer for direct conversion receivers. *Proc. IEEE Radio and Wireless Conference (RAWCON)*, pages 219–222, 2000.

DISTRIBUTION:

10 MS 0316
Eric R. Keiter, 01445

1 MS 0316
Robert J. Hoekstra, 01427

1 MS 0316
Richard Schiek, 01445

1 MS 0316
Heidi Thornquist, 01445

1 MS 0316
Ting Mei, 01445

1 MS 1110
Raymond S. Tuminaro, 01442

1 MS 0123
Donna L. Chavez, 01011

1 MS 0899
Technical Library (electronic copy), 9536

1 MS 0612
Review & Approval Desk, for DOE/OSTI, 9612