

An Optimization Framework for Goal-Oriented, Model-Based Reduction of Large-Scale Systems

K. Willcox, O. Ghattas, B. van Bloemen Waanders, B. Bader

Abstract—Optimization-ready reduced-order models should target a particular output functional, span an applicable range of dynamic and parametric inputs, and respect the underlying governing equations of the system. To achieve this goal, we present an approach for determining a projection basis that uses a goal-oriented, model-based optimization framework. The mathematical framework permits consideration of general dynamical systems with general parametric variations. The methodology is applicable to both linear and nonlinear systems and to systems with many input parameters. This paper focuses on an initial presentation and demonstration of the methodology on a simple linear model problem of the two-dimensional, time-dependent heat equation with a small number of inputs. For this example, the reduced models determined by the new approach provide considerable improvement over those derived using the proper orthogonal decomposition.

I. INTRODUCTION

Model reduction is a powerful tool that permits the systematic generation of cost-efficient representations of large-scale systems that, for example, result from discretization of partial differential equations (PDEs). The task of determining these representations may be posed as an optimization problem: determine the reduced model that provides the optimal representation (with respect to some measure) of the large-scale system behavior. For very large systems, determination of the best reduced model via direct optimization has not been pursued, due to challenges in solving the resulting optimization problem. Instead, several reduction methods have been developed that trade off optimality for tractability, and these have been applied in many different settings with considerable success, including controls, fluid dynamics, structural dynamics, and circuit design. However, a number of open issues remain with these methods, including the reliability of reduction techniques, guarantees associated with the quality of the reduced models, and the generation of reduced models that are suitable for optimal design, optimal control and inverse problem applications.

Recent advances in scalable algorithms for large-scale optimization of systems governed by PDEs have led to solution of problems with millions of state and optimization variables [1], [2]. The problem of determining a reduced model can be

cast in such a model-constrained optimization context. In particular, we consider a *goal-oriented* formulation in which the reduced model is chosen to optimally represent a particular output functional. Whereas other large-scale reduction methods, such as the proper orthogonal decomposition (POD), are purely data-driven and do not consider the underlying equations, our *model-based* optimization approach enforces the reduced-order governing equations as constraints. This improves on a data-driven approach by bringing additional knowledge of the reduced-order governing equations into the construction of the basis.

Most large-scale model reduction frameworks are based on a projection approach, which can be described in general terms as follows. Consider the general linear, time-invariant (LTI) dynamical system

$$M\dot{u} + Ku = f, \quad (1)$$

$$g = Cu, \quad (2)$$

with initial condition

$$u(0) = u_0, \quad (3)$$

where $u(t) \in \mathbb{R}^N$ is the system state, $\dot{u}(t)$ is the derivative of $u(t)$ with respect to time, and the vector u_0 contains the specified initial state. In general, we are interested in systems of the form (1) that arise from spatial discretization of PDEs. In this case, the dimension of the system, N , is very large and the matrices $M \in \mathbb{R}^{N \times N}$ and $K \in \mathbb{R}^{N \times N}$ result from the spatial discretization of the underlying governing equations. The vector $f(t) \in \mathbb{R}^N$ defines the input to the system and the matrix $C \in \mathbb{R}^{Q \times N}$ defines the Q outputs of interest, which are contained in the output vector $g(t)$.

A reduced-order model of (1)–(3) can be derived by assuming that the state $u(t)$ is represented as a linear combination of m basis vectors

$$\hat{u} = \Phi\alpha \quad (4)$$

where $\hat{u}(t)$ is the reduced model approximation of the state $u(t)$ and $m \ll N$. The projection matrix $\Phi \in \mathbb{R}^{N \times m}$ contains as columns the basis vectors ϕ_i , i.e., $\Phi = [\phi_1 \ \phi_2 \ \cdots \ \phi_m]$, and the vector $\alpha(t) \in \mathbb{R}^m$ contains the corresponding modal amplitudes. This yields the reduced-order model with state $\alpha(t)$ and output $\hat{g}(t)$

$$\hat{M}\dot{\alpha} + \hat{K}\alpha = \hat{f}, \quad (5)$$

$$\hat{g} = \hat{C}\alpha, \quad (6)$$

$$\hat{M}\alpha(0) = \Phi^T M u_0, \quad (7)$$

where $\hat{M} = \Phi^T M \Phi$, $\hat{K} = \Phi^T K \Phi$, $\hat{f} = \Phi^T f$, and $\hat{C} = C \Phi$.

K. Willcox is with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA kwillcox@mit.edu

B. Bader and B. van Bloemen Waanders are with Sandia National Laboratories, Albuquerque, NM 87185, USA bwbader@sandia.gov, bartv@sandia.gov

O. Ghattas is with the Institute for Computational Engineering and Sciences and Departments of Geological Sciences, Mechanical Engineering, Biomedical Engineering, and Computer Sciences, University of Texas at Austin, Austin, Texas 78712, USA omar@ices.utexas.edu

Projection-based model reduction techniques seek to find a basis Φ so that the reduced system (5)–(7) provides an accurate representation of the large-scale system (1)–(3) over the desired range of inputs. An optimal reduced model can be defined as one that minimizes the H-infinity norm of the difference between the reduced and original system transfer functions; however, no polynomial-time algorithm is known to achieve this goal. Algorithms such as optimal Hankel model reduction [3], [4], [5] and balanced truncation [6] have been used widely throughout the controls community to generate suboptimal reduced models with strong guarantees of quality. These algorithms can be carried out in polynomial time; however, the computational requirements make them impractical for application to large systems such as those arising from the discretization of PDEs, for which system orders often exceed 10^4 .

Considerable effort has been applied in recent years towards development of algorithms that extend balanced truncation to large-scale LTI systems [7], [8], [9]; however, efficient algorithms for very large systems remain a challenge. The proper orthogonal decomposition (POD) [10], [11] has emerged as a popular alternative for reduction of very large dynamical systems; however, it lacks the quality guarantees of methods such as balanced truncation.

Effective model reduction methods for optimal design, optimal control and inverse problem applications remain a challenge. Approaches developed for dynamical systems, such as POD and Krylov-based methods, have been applied in an optimization context [12], [13], [14]; however, the number of parameters in the optimization application was small. A key challenge that must be addressed in order to provide optimization-ready reduced-order models is the need for the reduced models to capture variation over a parametric input space, which, for many optimization applications, will be of high dimension. In recent work for steady-state problems, methods are presented for constructing reduced models that are of guaranteed quality over a range of inputs via the use of error estimates and adaptivity [15].

In this paper, we formulate the problem of determining a projection basis using a goal-oriented, model-based optimization framework. The mathematical framework permits consideration of general dynamical systems with general parametric variations. The methodology is applicable to both linear and nonlinear systems and to systems with many input parameters. This paper focuses on an initial presentation and demonstration of the methodology on a simple model problem that is linear and has a small number of inputs. We propose an efficient solution strategy that borrows concepts from the POD and employs recent methods for optimization of systems governed by PDEs to make the approach tractable for large-scale problems. The paper is organized as follows. First, the general dynamical system framework with parametric variations is described. This is followed by a description of the goal-oriented basis optimization formulation and the proposed model reduction methodology. The approach is then demonstrated for a linear model problem that considers the unsteady two-dimensional heat equation

with parametrically varying boundary control inputs. Finally, we present conclusions and directions for future research.

II. DYNAMICAL SYSTEM FRAMEWORK

The standard LTI system framework is defined by (1)–(3). In this section, we present the more general case that includes parametric variation in the system.

A. Parametric input variations

We consider a finite set of instantiations of the governing equations (1)–(3) that could arise from variations in the coefficient matrices M and K , the input f , or the initial state u_0 . For example, where (1)–(3) represent a spatially discretized PDE, these variations stem from changes in the domain shape, boundary conditions, coefficients, initial conditions, or sources of the underlying PDEs. The general dynamical system for S different instances is thus written

$$M^k \dot{u}^k + K^k u^k = f^k, \quad k = 1, \dots, S \quad (8)$$

$$u^k(0) = u_0^k, \quad k = 1, \dots, S \quad (9)$$

$$g^k = C^k u^k, \quad k = 1, \dots, S \quad (10)$$

where the superscript k denotes the k th instance of the system, which has corresponding state $u^k(t)$ and output $g^k(t)$.

Using the projection framework described in the previous section, a reduced-order model of (8)–(10) is obtained as

$$\hat{M}^k \dot{\alpha}^k + \hat{K}^k \alpha^k = \hat{f}^k, \quad k = 1, \dots, S \quad (11)$$

$$\hat{g}^k = \hat{C}^k \alpha^k, \quad k = 1, \dots, S \quad (12)$$

$$\hat{M}^k \alpha(0)^k = \Phi^T M^k u_0^k, \quad k = 1, \dots, S \quad (13)$$

where $\hat{M}^k = \Phi^T M^k \Phi$, $\hat{K}^k = \Phi^T K^k \Phi$, $\hat{f}^k = \Phi^T f^k$, and $\hat{C}^k = C^k \Phi$.

B. Proper orthogonal decomposition

POD is a widely used approach to determine the reduced basis Φ . POD can be applied efficiently to large systems using the method of snapshots [10] as follows. Consider the collection of “snapshots”, $u^k(t_j)$, $j = 1, \dots, T$, $k = 1, \dots, S$, where $u^k(t_j) \in \mathbb{R}^N$ is the solution of the governing equations (8) at time t_j for parameter instance k . T time instants are considered for each parameter instance, yielding a total of ST snapshots. We define the snapshot matrix $U \in \mathbb{R}^{N \times ST}$ as

$$U = [u^1(t_1) \ u^1(t_2) \ \dots \ u^1(t_T) \ u^2(t_1) \ \dots \ \dots \ u^S(t_T)] \quad (14)$$

and we will refer to the i th column of U as the i th snapshot, denoted by U_i .

The POD basis vectors are chosen to be the orthonormal set that maximizes the following cost [16]:

$$\phi = \arg \max_{\varphi} \frac{\langle |(\varphi, u)|^2 \rangle}{(\varphi, \varphi)}, \quad (15)$$

where (u, ϕ) denotes the scalar product of the basis vector with the field $u(t)$ evaluated over the domain, and $\langle \cdot \rangle$ represents a time-averaging operation. In the case of the discrete snapshots contained in U , (15) is maximized when

the m basis vectors are chosen to be the first m left singular vectors of U . For a fixed basis size, the POD basis therefore minimizes the error between the original snapshots and their representation in the reduced space defined by

$$E = \sum_{k=1}^S \sum_{j=1}^T [u^k(t_j) - \tilde{u}^k(t_j)]^T [u^k(t_j) - \tilde{u}^k(t_j)], \quad (16)$$

where $\tilde{u}^k(t_j) = \Phi \Phi^T u^k(t_j)$. This error is equal to the sum of the singular values corresponding to those singular vectors not included in the basis

$$E = \sum_{i=m+1}^{ST} \sigma_i, \quad (17)$$

where σ_i is the i^{th} singular value of U .

The POD is an optimal basis in the sense that it minimizes the data reconstruction error given by (16); however, it is important to note that this optimality applies only to the representation of a known state solution $u^k(t_j)$ in the reduced space, i.e. \tilde{u} is computed as $\tilde{u}^k(t_j) = \Phi \Phi^T u^k(t_j)$, not by solution of the reduced model ($\tilde{u} \neq \hat{u}$). Therefore, the error expression does not apply to the resulting POD reduced-order model (5). In particular, the error expression yields no rigorous information regarding the accuracy of the solution of the reduced model and thus whether \hat{u} is a good approximation of u . Moreover, the POD basis does not account for the system outputs, although methods to augment the standard approach have been proposed that use adjoint information [17], [18]. In addition, because no information regarding the governing equations is included in the POD process, the POD basis does not properly reflect the fact that the snapshots $u^k(t_j)$ are associated with different parametric instances of the system.

In the following section we present an alternative method to determine the reduced-space basis. This method seeks to minimize an error similar in form to (16); however, we will improve upon the POD, first, by minimizing the error in the outputs (as opposed to states) and, second, by imposing additional constraints that $\hat{u}^k(t)$ should result from satisfying the reduced-order governing equations for each parameter instance k .

III. OPTIMIZED REDUCED-ORDER BASIS

A. Constrained optimization formulation for projection basis

We pose the problem of selecting the basis Φ as a goal-oriented optimization problem that seeks to minimize the difference between the full-space and reduced-order output solution over a selected set of inputs and the interval $[0, t_f]$, subject to satisfying the underlying governing equations. The problem of determining the optimal basis, $\Phi \in \mathbf{R}^{N \times m}$, can be written as

$$\begin{aligned} \min_{\Phi, \alpha} \mathcal{G} &= \frac{1}{2} \sum_{k=1}^S \int_0^{t_f} (g^k - \hat{g}^k)^T (g^k - \hat{g}^k) dt \\ &+ \frac{\beta}{2} \sum_{j=1}^m (1 - \phi_j^T \phi_j)^2 \end{aligned} \quad (18)$$

subject to

$$\Phi^T M^k \Phi \dot{\alpha}^k + \Phi^T K^k \Phi \alpha^k = \Phi^T f^k, \quad k = 1, \dots, S \quad (19)$$

$$\Phi^T M^k \Phi \alpha^k(0) = \Phi^T M^k u_0^k, \quad k = 1, \dots, S \quad (20)$$

$$\hat{g}^k = C^k \Phi \alpha^k, \quad k = 1, \dots, S. \quad (21)$$

In the case of a linear relationship between outputs and state as in (10), the objective function can be written

$$\begin{aligned} \mathcal{G} &= \frac{1}{2} \sum_{k=1}^S \int_0^{t_f} (u^k - \hat{u}^k)^T H^k (u^k - \hat{u}^k) dt \\ &+ \frac{\beta}{2} \sum_{j=1}^m (1 - \phi_j^T \phi_j)^2, \end{aligned} \quad (22)$$

where $H^k = C^k T C^k$ can be interpreted as a weighting matrix that defines the states relevant to the specified output. While the first term in the objective function (22) has similarities with that minimized by the POD, given by (16), there are two important distinctions to note. First, the goal-oriented nature of the formulation (22) focuses on reduction of the error for a particular output functional rather than for the general state vector. Second, through the optimization constraints (19)–(21), the general optimization approach requires satisfaction of the reduced-order governing equations to compute \hat{u} . The error minimized by the optimization approach is thus tied rigorously to the reduced-order model, whereas the POD is based purely on snapshot data. In both cases, however, the definition of the error is limited to a discrete set of observations.

The second term in (22) is a regularization term that penalizes the deviation of the length of the basis vectors from unity and β is a constant weighting. This regularization acts only in the null space of the projected Hessian matrix of the first term of (22). Therefore, the reduced output approximation, \hat{g} , is unaffected by the regularization term, yet the conditioning of the optimization problem is improved. Note, however, that there remains a null space of the projected Hessian matrix that admits arbitrary rotations of the basis vectors; the optimization method chosen to solve (18)–(21) should therefore be tolerant of singular projected Hessian matrices. It is also important to note that the optimization problem (18)–(21) is nonlinear and nonconvex; therefore, there is no guarantee that a purely local optimization method will converge to the global optimum. Therefore, generating the initial guess is very important; strategies to address this issue will be discussed in the next section.

B. Optimality conditions and the reduced gradient

The optimality conditions for the system (18)–(21) can be derived by defining the Lagrangian functional

$$\mathcal{L}(\Phi, \alpha^k, \lambda^k, \mu^k) = \frac{1}{2} \sum_{k=1}^S \int_0^{t_f} (u^k - \Phi \alpha^k)^T H^k (u^k - \Phi \alpha^k) dt$$

$$\begin{aligned}
& + \frac{\beta}{2} \sum_{j=1}^m (1 - \phi_j^T \phi_j)^2 \\
& + \sum_{k=1}^S \int_0^{t_f} \lambda^{kT} (\Phi^T M^k \Phi \dot{\alpha}^k + \Phi^T K^k \Phi \alpha^k - \Phi^T f^k) dt \\
& + \sum_{k=1}^S \mu^{kT} (\Phi^T M^k \Phi \alpha(0) - \Phi^T M^k u_0^k), \quad (23)
\end{aligned}$$

where $\lambda^k = \lambda^k(t) \in \mathbb{R}^m$ and $\mu^k \in \mathbb{R}^m$ are Lagrange multipliers (also known as adjoint state variables) that respectively enforce the state ODE system and initial conditions for the k th sample. The optimality system can be derived by taking variations of the Lagrangian with respect to the adjoint, state, and basis vector variables.

Setting the first variation of the Lagrangian with respect to λ^k to zero and arguing that the variation of λ^k is arbitrary in $[0, t_f]$, and setting the derivative of the Lagrangian with respect to μ^k to zero, simply recovers the state equation and initial conditions (19)–(20).

Setting the first variation of the Lagrangian with respect to the α^k to zero, and arguing that the variation of α^k is arbitrary in $[0, t_f]$, at $t = 0$, and at $t = t_f$, yields the adjoint equation, final condition and definition of μ

$$-\Phi^T M^k \Phi \dot{\lambda}^k + \Phi^T K^{kT} \Phi \lambda^k = \Phi^T H^k (u^k - \Phi \alpha^k), \quad k = 1, \dots, S \quad (24)$$

$$\lambda^k(t_f) = 0, \quad k = 1, \dots, S \quad (25)$$

$$\mu^k = \lambda^k(0), \quad k = 1, \dots, S. \quad (26)$$

Note that, without loss of generality, M is assumed to be a symmetric matrix.

Taking the derivative of the Lagrangian with respect to the basis vector variables Φ yields the following matrix equation

$$\begin{aligned}
\delta \mathcal{L}_\Phi & = \sum_{k=1}^S \int_0^{t_f} (H^k \Phi \alpha^k \alpha^{kT} - H^k u^k \alpha^{kT}) dt \\
& + \beta \Phi \text{diag}(1 - \phi_i^T \phi_i) \\
& + \sum_{k=1}^S \int_0^{t_f} [M^k \Phi (\lambda^k \dot{\alpha}^{kT} + \dot{\alpha}^k \lambda^{kT}) \\
& + K^{kT} \Phi \lambda^k \alpha^{kT} + K^k \Phi \alpha^k \lambda^{kT} - f^k \lambda^{kT}] dt \\
& - \sum_{k=1}^S M^k u_0^k \mu^{kT} = 0. \quad (27)
\end{aligned}$$

The combined system (19)–(20), (24)–(26), and (27) represents the first-order Karush-Kuhn-Tucker optimality conditions for the optimization problem (18)–(21).

To solve the constrained optimization problem (18)–(21), we choose to eliminate the state variables α^k and state equations (19) and solve an equivalent unconstrained optimization problem in the Φ variables. Efficient solution of this unconstrained optimization problem requires a gradient-based method, which requires function and gradient evaluations. In this case, the gradient of the unconstrained objective with

respect to Φ is given by $\delta \mathcal{L}_\Phi$ when the α^k satisfy the state equations and (λ^k, μ^k) satisfy the adjoint equations. The procedure to compute the gradient of \mathcal{G} for any value of Φ can therefore be summarized as follows. First, solve the state equations (19)–(20) to determine $\alpha^k(t)$. Second, solve the adjoint equations (24)–(26) to determine $\lambda^k(t)$ and μ^k . Finally, use the computed α^k , λ^k , and μ^k in (27) to determine the gradient.

C. Basis computation

The formulation defined by equations (18)–(21) provides a mathematical definition of the desired optimal basis; however, in practice this optimization problem may not be tractable for large-scale problems. First, we may not be able to afford storage of the entire time history for the full model, which leads us to adopt a snapshot-based approach. As in the POD, the time integrals in (18) are replaced by a summation over a finite number of discrete time instants. Our method therefore requires a priori computation of a set of high-fidelity solutions over a pre-determined set of time instants and input parameter values.

Second, even with this simplification, the number of optimization variables is equal to mN — the desired number of basis functions multiplied by the length of each basis vector — where for many applications $N \geq O(10^6)$. Therefore, it will be assumed that each basis vector can be represented as a linear combination of snapshots:

$$\phi_j = \sum_{i=1}^{ST} \gamma_i^j U_i \quad j = 1, \dots, m \quad (28)$$

where the coefficients γ_i^j are the variables in the modified optimization problem. This assumption reduces the number of optimization variables from mN to mST , where, for large-scale applications, typically $ST \ll N$. As a consequence, neither the gradient computation nor the optimization step computation (which dominate the cost of an optimization iteration) scale with the full system size N . Moreover, using adjoints ensures that the gradient computation requires just $2S$ reduced model solutions per optimization iteration, as opposed to mST as with direct sensitivities. The assumption that the basis vectors can be represented as a linear combination of snapshots is motivated by the singular value decomposition (SVD) theory underlying the POD, for which the relation (28) is exact (this is equivalent to solving the inner versus the outer SVD problem). As will be demonstrated in the following results section, in the case of the optimal basis formulation, numerical experiments suggest that this formulation is still capable of yielding very good results.

Equation (28) can be written in matrix form as

$$\Phi = U\Gamma, \quad (29)$$

where γ_i^j is the ij th element of $\Gamma \in \mathbb{R}^{ST \times m}$. Gradients of the objective function with respect to Γ are related simply to gradients with respect to Φ by

$$\frac{\partial \mathcal{L}}{\partial \Gamma} = U^T \frac{\partial \mathcal{L}}{\partial \Phi}. \quad (30)$$

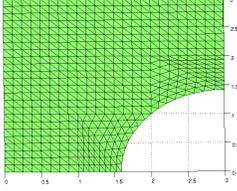


Fig. 1. Problem domain and boundary conditions: Neumann on right side, Dirichlet on all other boundaries.

The modified optimization formulation offers no guarantees of convexity and the choice of initial guess for the basis is thus very important. In this paper, we present two possible strategies. The first is to use the POD basis as an initial starting point. Since a snapshot set is required anyway, the additional cost of computing the POD basis is small. A second strategy is to employ continuation on the basis dimension. In this approach, the initial guess for the case of m basis vectors is chosen to be the solution of the optimization problem for $m-1$ basis vectors plus an arbitrary m th vector. This iterative procedure can be initialized at any value $m \geq 1$ with the POD basis vectors as an initial guess on the first iteration.

IV. MODEL PROBLEM AND RESULTS

A. Model problem description

Results are presented for a model problem that considers the two-dimensional time-dependent heat equation. In this case, the PDE is given by

$$\frac{\partial \bar{u}}{\partial t} - k \nabla^2 \bar{u} = 0 \text{ in } \Omega \quad (31)$$

$$\bar{u} = \bar{u}_c \text{ on } \Gamma \quad (32)$$

$$\bar{u} = \bar{u}_0 \text{ in } \Omega \text{ for } t = 0 \quad (33)$$

where $\bar{u}(x, y, t)$ is the temperature field defined on the domain Ω , $\bar{u}_c(x, y)$ is the boundary control function (which is assumed to be constant in time) applied on the boundary Γ , and $\bar{u}_0(x, y)$ is the given initial temperature field. The output of interest is the temperature over a specified sub-region of the domain.

A modification of the finite element formulation from [19] is used to discretize the problem in space, yielding a dynamical system of the form (8)–(10), where $u^k(t)$ represents the spatially discretized temperature field corresponding to forcing input f^k , and $g^k(t)$ contains those elements of u^k that lie within the specified region of interest. Figure 1 shows the specific domain Ω that was used, which is discretized with triangular elements. Results will be shown for a discretization containing a total of $N = 480$ temperature unknowns. The specified initial condition is $u = 0$ at $t = 0$, and the selected time integration scheme is an implicit Euler method with a constant time step over the time interval $[0, T]$. The boundary control is applied on $\Gamma_c = \{(0, y) : 0 \leq y \leq 3\}$, i.e., Dirichlet control on the left boundary of the domain. Neumann boundary conditions are specified on $\Gamma_N = \{(3, y) : 1.5 \leq y \leq 3\}$ and remaining part of the boundary, Γ_D , is fixed with zero Dirichlet conditions.

Snapshots were generated by performing a time simulation of the system under different forcing conditions. The forcing was generated by applying a temperature distribution along the boundary Γ_c in Figure 1. For the results presented here, the forcing functions considered were parameterized using sinusoidal distributions with successively higher spatial frequency. Snapshots were generated over $S = 5$ instances of the control parameter forcing with $T = 20$ time instants for each parameter instance. Using the optimization formulation (18)–(21), we seek to select the m basis functions that minimize the error defined by (18) while satisfying the reduced-order state equations for each control instance. The basis functions are assumed to be a linear combination of available snapshots, hence there are $mST = 100m$ design variables in the optimization problem. The state and adjoint equations each consist of m uncoupled ODE systems of dimension $ST = 100$. An implicit backward Euler scheme is used to discretize the ODEs (note that the adjoint equations are marched backward in time). The resulting fully discrete systems are lower and upper tridiagonal for the state and adjoint, respectively, and thus can be solved very efficiently.

B. Optimized basis performance

For the first set of results, the output is defined to be the temperature over a strip of the domain in the region $0.5 < x < 1.0$, $0.5 < y < 2.5$, yielding an output vector of size $Q = 47$. To determine the goal-oriented basis, (18)–(21) were solved by using (27) to compute analytical gradients and employing an unconstrained optimization algorithm that uses a trust-region-based Newton method [20]. Figure 2 shows the resulting objective function values for bases ranging in size from $m = 1$ to $m = 10$. Figure 2 also shows the evaluation of (22) for the POD bases over this range of m . It can be seen clearly that the optimized basis outperforms the POD in all cases, particularly when m is small.

In order to provide a quantitative metric by which to judge the performance of the optimized basis, balanced truncation was applied to this problem. The problem was converted to standard LTI form by considering each parametric forcing function as an independent input. Figure 2 shows the evaluation of (22) for truncated balanced models of size $m = 1$ through $m = 10$. It can be seen that the optimized basis provides a substantial improvement over POD when both are compared to the results of balanced truncation. It is also important to note that balanced truncation uses both a left and a right projection basis, and thus has twice as many degrees of freedom as the goal-oriented optimized basis.

C. Comparison with POD

A significant advantage of the goal-oriented approach is that the basis can be optimized with respect to a particular output functional, whereas the POD seeks to minimize the reconstruction error over all states. Several different output definitions were considered in order to gain insight to the optimized basis results.

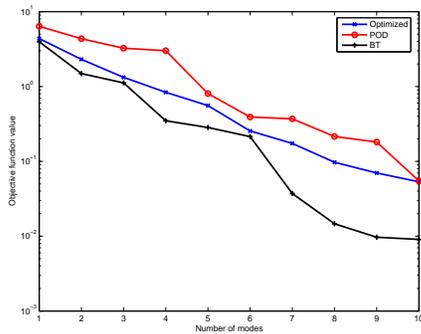


Fig. 2. The error (22) versus number of modes for the goal-oriented optimized basis, the POD basis, and balanced truncation.

If the output considered is to minimize the error of state prediction over the entire domain, that is, $H^k = 1$ in (22), then the goal-oriented approach seeks to minimize the same error as the POD. However, it is important to note again the difference in the representation of the term \tilde{u}^j , which for POD is computed directly from the known solution w^j , i.e. $\tilde{u}^j = \Phi\Phi^T w^j$. In this sense the POD is a purely data-based method that does not account in any way for the underlying governing equations. In contrast, our method computes \hat{u}^j in (22) by requiring the solution to satisfy the governing equations in the reduced-order space.

Results for this case are shown in the first row of Table I. Using the POD basis as an initial guess, the optimizer is able to make almost no improvement in the objective function. As shown in Table I, the reduction in the error is just 1%. For different values of S , T and m , the POD basis is found to be almost optimal with respect to state reconstruction error for this example. Due to the symmetry properties of the system (M and K are symmetric matrices), any congruent basis transformation, such as the POD, is guaranteed to preserve the stability of the system. Thus it is to be expected that the POD should perform well on this heat conduction example. As the results show, the additional error from solution of the governing equations in the reduced space is not significant in this case. In more complicated examples where this error becomes significant, the optimized basis might be expected to provide an advantage over the POD even in terms of full state reconstruction. This is particularly true for non-symmetric systems, such as those representing the Euler equations, for which the POD basis can routinely produce unstable reduced-order models.

Table I shows the results for other outputs corresponding to various specified $x-y$ regions (and thus different weightings H in the objective function). Note that the POD basis is computed in the standard way and thus is insensitive to the choice of output functional. The values in the column \mathcal{G}_{pod} represent the standard POD basis evaluated using the criterion defined by (22) for each different instance of H (i.e. the metric \mathcal{G}_{pod} is case-dependent). It can be seen that by defining an output functional, the goal-oriented basis can yield substantial improvements in errors over the POD basis. It should be emphasized that our method does not simply

TABLE I

COMPARISON OF OPTIMIZATION RESULTS. THE OBJECTIVE FUNCTION GIVEN BY (22) IS EVALUATED FOR THE OPTIMIZED BASIS (\mathcal{G}_{opt}) AND THE POD BASIS (\mathcal{G}_{pod}).

| Min. error over | S | T | m | \mathcal{G}_{opt} | \mathcal{G}_{pod} |
|------------------------|-----|-----|-----|---------------------|---------------------|
| All states | 5 | 20 | 5 | 28.9829 | 29.2762 |
| $x = 0.625, y = 0.625$ | 3 | 20 | 5 | 2.9038e-3 | 0.01066 |
| $0.5 < x, y < 1$ | 5 | 20 | 5 | 0.01282 | 0.1932 |
| $0.5 < x, y < 1$ | 5 | 20 | 5 | 0.5555 | 0.8062 |

“ignore” states that lie outside of the region of interest, since \hat{u}^j is computed by solving the reduced-order equations over the entire domain. Therefore the basis must represent *all* states – but the optimization formulation allows the basis energy to be focused appropriately to achieve the desired objective. One might draw conceptual parallels between this approach and a posteriori error estimates to manage grid adaptivity.

Figure 3 shows the output errors in the case of an output functional defined over the region $0.5 < x < 1, 0.5 < y < 1$. Each plot in the figure corresponds to one of the nine grid points that lie within the region of interest (for clarity, just four of the points are shown). The first $T = 20$ snapshots correspond to the first instance of control forcing, the second $T = 20$ correspond to the second instance, and so on. The figure shows that for almost every snapshot in the ensemble, the optimized basis results in a more accurate prediction of the temperature at the point of interest. In many cases, the error is reduced by almost an order of magnitude.

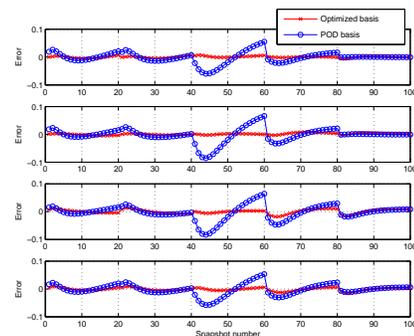


Fig. 3. Error in temperature prediction for each snapshot using POD and optimized basis. The optimized basis was selected so as to minimize the error over the region $0.5 < x < 1, 0.5 < y < 1$. Errors are shown for four of the nine points contained within this region.

The reduced output errors shown in Figure 3 come at a cost. Figure 4 shows the norm of the errors computed over the entire domain for each snapshot. In order to reduce the errors at the specified points, the optimized basis yields less accurate predictions for other states. However, it is again important to note that this trade-off in accuracy is done in a systematic way using both the governing equations and the defined output functional. According to the optimization result, the larger errors observed in other areas of the domain are compatible with the task of reducing the error in the region of interest.

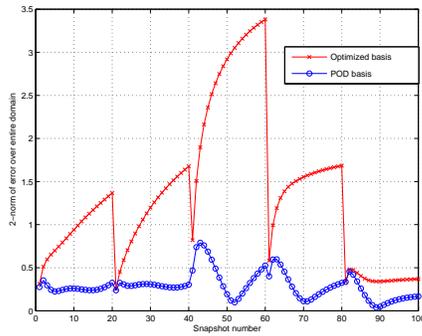


Fig. 4. Norm of the error in temperature prediction over the entire domain for each snapshot using POD and optimized basis. The optimized basis was selected so as to minimize the error over the region $0.5 < x < 1, 0.5 < y < 1$.

Figure 5 shows an example solution for one particular snapshot. The snapshot chosen corresponds to the first time instant in the third control function (snapshot number 41 in Figures 3 and 4). The output of interest is the temperature over the region $0.5 < x < 1, 0.5 < y < 1$, indicated by the bold square in the figure. As can be seen from Figures 3 and 4, this snapshot corresponds to a case where the output error is substantially reduced by the optimized basis. The reductions in output error are approximately an order of magnitude – for example, in the bottom plot in Figure 3 the error in the output for this snapshot is reduced in magnitude from -0.0209 to 0.0036 . In contrast, the norm of the error over the domain is substantially increased from 0.4675 to 0.8201 . By comparing the plots in Figure 5, this effect can be clearly seen. The optimized basis result indicates that the large errors near the control boundary are acceptable (in fact optimal) if one is concerned only with predicting the temperature within the indicated region.

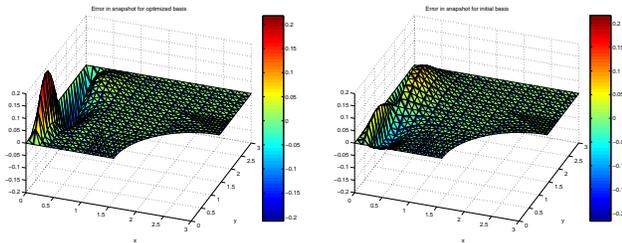


Fig. 5. Reduced-order model temperature prediction for snapshot number 41 using optimized (left) and POD (right) basis. The optimized basis was selected so as to minimize the error over the region $0.5 < x < 1, 0.5 < y < 1$.

V. CONCLUSIONS

The goal-oriented, model-based optimization approach presented here provides a general framework for construction of reduced models, and is particularly applicable for optimal design, optimal control and inverse problems. The optimization approach provides significant advantages over the POD by allowing the projection basis to be targeted to output functionals, by providing a framework in which to

consider multiple parameter instances, and by incorporating the reduced-order governing equations as constraints in the basis derivation.

REFERENCES

- [1] V. Akcelik, G. Biros, and O. Ghattas, "Parallel multiscale Gauss-Newton-Krylov methods for inverse wave propagation," in *Proceedings of SC2002*, 2002.
- [2] V. Akcelik, J. Bielik, G. Biros, I. Epanomeritakis, A. Fernandez, O. Ghattas, E. Kim, J. Lopez, D. O'Hallaron, T. Tu, and J. Urbanic, "Terascale forward and inverse earthquake modeling," in *Proceedings of SC2003*, 2003.
- [3] V. Adamjan, D. Arov, and M. Krein, "Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur-Takagi problem," *Math. USSR Sbornik*, vol. 15, pp. 31–73, 1971.
- [4] M. Bettayeb, L. Silverman, and M. Safonov, "Optimal Approximation of Continuous-Time Systems," *Proceedings of the 19th IEEE Conference on Decision and Control*, vol. 1, December 1980.
- [5] S.-Y. Kung and D. Lin, "Optimal Hankel-norm model reductions: Multivariable systems," *IEEE Transactions on Automatic Control*, vol. AC-26, no. 1, pp. 832–52, August 1981.
- [6] B. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Transactions on Automatic Control*, vol. AC-26, no. 1, pp. 17–31, August 1981.
- [7] D. Sorensen and A. Antoulas, "The sylvester equation and approximate balanced reduction," *Linear Algebra and Its Applications*, Fourth Special Issue on Linear Systems and Control, edited by Blondel et al., 351–352, pp. 671–700, 2002.
- [8] J. Li and J. White, "Low rank solution of Lyapunov equations," *SIAM Journal on Matrix Analysis and Applications*, vol. 24, no. 1, pp. 260–280, 2002.
- [9] S. Gugercin and A. Antoulas, "A survey of model reduction by balanced truncation and some new results," *International Journal of Control*, vol. 77, pp. 748–766, 2004.
- [10] L. Sirovich, "Turbulence and the dynamics of coherent structures. Part 1 : Coherent structures," *Quarterly of Applied Mathematics*, vol. 45, no. 3, pp. 561–571, October 1987.
- [11] P. Holmes, J. Lumley, and G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge, UK: Cambridge University Press, 1996.
- [12] L. Daniel, O. Siang, L. Chay, K. Lee, and J. White, "Multiparameter moment matching model reduction approach for generating geometrically parameterized interconnect performance models," *Trans. on Computer Aided Design of Integrated Circuits*, vol. 23, no. 5, pp. 678–693, May 2004.
- [13] M. Hinze and S. Volkwein, "Proper orthogonal decomposition surrogate models for nonlinear dynamical systems: Error estimates and sub-optimal control," Technische Universitat Dresden, Tech. Rep. SFB609-Preprint-27-2004, 2004.
- [14] K. Kunisch and S. Volkwein, "Control of Burgers' equation by reduced order approach using proper orthogonal decomposition," *Journal of Optimization Theory and Applications*, vol. 102, pp. 345–371, 1999.
- [15] C. Prud'homme, D. Rovas, K. Veroy, Y. Maday, A. Patera, and G. Turinici, "Reliable real-time solution of parameterized partial differential equations: Reduced-basis output bound methods," *Journal of Fluids Engineering*, vol. 124, pp. 70–80, 2002.
- [16] G. Berkooz, P. Holmes, and J. Lumley, "The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows," *Annual Review of Fluid Mechanics*, vol. 25, pp. 539–575, 1993.
- [17] S. Lall, J. Marsden, and S. Glavaski, "A subspace approach to balanced truncation for model reduction of nonlinear control systems," *International Journal on Robust and Nonlinear Control*, vol. 12, no. 5, pp. 519–535, 2002.
- [18] K. Willcox and J. Peraire, "Balanced model reduction via the proper orthogonal decomposition," *AIAA Journal*, vol. 40, no. 11, pp. 2323–30, November 2002.
- [19] J. Albery, C. Carstensen, and S. Funken, "Remarks around 50 lines of matlab: Short finite element implementation," *Numerical Algorithms*, vol. 20, pp. 117–137, 1999.
- [20] T. Coleman and Y. Li, "An interior, trust region approach for nonlinear minimization subject to bounds," *SIAM Journal on Optimization*, vol. 6, pp. 418–445, 1996.