

Minimal subspace rotation on the Stiefel manifold for stabilization and enhancement of projection-based reduced order models for the compressible Navier–Stokes equations



Maciej Balajewicz^{a,*,1}, Irina Tezaur^{b,2}, Earl Dowell^{c,3}

^a Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA

^b Quantitative Modeling & Analysis Department, Sandia National Laboratories, P.O. Box 969, MS 9159, Livermore, CA, USA

^c Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC, USA

ARTICLE INFO

Article history:

Received 25 September 2015
Received in revised form 6 May 2016
Accepted 18 May 2016
Available online 25 May 2016

Keywords:

Projection-based reduced order model (ROM)
Proper Orthogonal Decomposition (POD)
Compressible flow
Stabilization
Trace minimization
Stiefel manifold

ABSTRACT

For a projection-based reduced order model (ROM) of a fluid flow to be stable and accurate, the dynamics of the truncated subspace must be taken into account. This paper proposes an approach for stabilizing and enhancing projection-based fluid ROMs in which truncated modes are accounted for *a priori* via a minimal rotation of the projection subspace. Attention is focused on the full non-linear compressible Navier–Stokes equations in specific volume form as a step toward a more general formulation for problems with generic non-linearities. Unlike traditional approaches, no empirical turbulence modeling terms are required, and consistency between the ROM and the Navier–Stokes equation from which the ROM is derived is maintained. Mathematically, the approach is formulated as a trace minimization problem on the Stiefel manifold. The reproductive as well as predictive capabilities of the method are evaluated on several compressible flow problems, including a problem involving laminar flow over an airfoil with a high angle of attack, and a channel-driven cavity flow problem.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

The past several decades have seen an exponential growth of computer processing speed and memory capacity. The massive, complex simulations that run on supercomputers allow exploration of fields for which physical experiments are too impractical, hazardous, and/or costly. A striking number of these fields require computational fluid dynamics (CFD) models and simulations. Accurate and efficient CFD simulations are critical to many defense, climate and energy missions, e.g., simulations aimed to design and qualify nuclear weapons components carried within an aircraft weapons bay; global climate simulations aimed to predict anticipated twenty-first century sea-level rise; aero-elastic simulations for optimal design of wind systems for power generation.

* Corresponding author.

E-mail address: mbalajew@illinois.edu (M. Balajewicz).

¹ Assistant Professor.

² Principal Member of Technical Staff.

³ William Holland Hall Professor.

Unfortunately, even with the aid of massively parallel next-generation computers, CFD simulations for applications such as these are still too expensive for real-time and multi-query applications such as uncertainty quantification (UQ), optimization and control design. Reduced order modeling is a promising tool for bridging the gap between high-fidelity, and real-time simulations/UQ. Reduced order models (ROMs) are derived from a sequence of high-fidelity simulations but have a much lower computational cost. Hence, ROMs can enable real-time simulations of complex systems for more rapid analysis, control and decision-making in the presence of uncertainty.

Most existing ROM approaches are based on projection. In projection-based reduced order modeling, the state variables are approximated in a low-dimensional subspace. There exist a number of approaches for calculating this low-dimensional subspace, e.g., Proper Orthogonal Decomposition (POD) [44,18], Dynamic Mode Decomposition (DMD) [37,41], balanced POD (BPOD) [34,49], balanced truncation [16,26], and the reduced basis method [38,47]. In all of these methods, a basis for the low-dimensional subspace is obtained from a basis for a higher-dimensional subspace through truncation – the removal of modes that are believed to be unimportant in representing a problem solution. Typically, the size of the reduced basis is chosen according to an energy criterion: modes with low energy are discarded, so that the reduced basis subspace is spanned by the highest energy modes. Although truncated modes are negligible from a data compression point of view, they are often crucial for representing solutions to the dynamical flow equations. Dynamics of the truncated orthogonal subspace must be taken into account for to ensure stability and accuracy of the ROM.

For linear systems, a variety of techniques for generating low-dimensional projection-based ROMs with rigorous stability guarantees and accuracy bounds are available [16,26,1,22]. Equivalent results are lacking for nonlinear systems. Traditionally, low-dimensional ROMs of fluid flows have been stabilized and enhanced using empirical turbulence models. In this approach, the nonlinear dynamics of the truncated subspace are modeled using additional constant and linear terms in the ROM system of ordinary differential equations (ODEs) [2,33,11,13]. More recently, nonlinear eddy-viscosity models have also been proposed [29,27,19,31]. One downside of turbulence models is that they destroy consistency between the Navier–Stokes partial differential equations (PDEs) and the ODE system of the ROM. Accurately identifying and matching free coefficients of the turbulence models is another challenge. Moreover, these methods are usually limited to the incompressible Navier–Stokes equations.

Consider, for concreteness, the POD/Galerkin approach to model reduction applied to the incompressible Navier–Stokes equations. For these equations, the natural choice of inner product for the Galerkin projection step of the model reduction procedure is the L^2 inner product. This is because, in these models, the solution vector is taken to be the velocity vector \mathbf{u} , so that $\|\mathbf{u}\|_2$ is a measure of the global kinetic energy in the domain, and the POD modes optimally represent the kinetic energy present in the ensemble from which they were generated. The same is not true for the compressible Navier–Stokes equations. Hence, if a compressible fluid ROM is constructed in the L^2 inner product (a common choice of inner product in projection-based model reduction), the ROM solution may not satisfy the conservation relation implied by the governing equations, and may exhibit non-physical instabilities [36].

Unfortunately, ROM instability is a real problem for many compressible flow problems: as demonstrated in [8,6,21], a compressible fluid POD/Galerkin ROM might be stable for a given number of modes, but unstable for other choices of basis size. Several researchers have proposed ways to circumvent this difficulty through the careful construction of an energy-based inner product for the projection step of the model reduction. Rowley et al. [36] show that Galerkin projection preserves the stability of an equilibrium point at the origin if the ROM is constructed in an energy-based inner product. Barone et al. [6], Kalashnikova et al. [21] demonstrate that a symmetry transformation leads to a stable formulation for a Galerkin ROM for the linearized compressible Euler equations and nonlinear compressible Navier–Stokes equations with solid wall and far-field boundary conditions. Serre et al. [42] propose applying the stabilizing projection developed by Barone et al. [6], Kalashnikova et al. [21] to a skew-symmetric system constructed by augmenting a given linear system with its adjoint system. The downside to these methods is that they are inherently embedded methods: access to the governing PDEs and/or the code that discretizes these PDEs is required.

Other ROM approaches, e.g., the Gauss–Newton with Approximated Tensors (GNAT) method of Carlberg et al. [9], have better stability properties, as they formulate the ROM at the fully discrete level. The drawback of this approach is that an additional layer of approximation – usually called hyper-reduction – is required to gain computational speed-up. Moreover, the approach lacks stability guarantees for low-dimensional expansions.

In this paper, a stabilization and enhancement approach to ROMs for the compressible Navier–Stokes equations is developed. The approach is an extension of the methodology developed in [3,5] specifically for the incompressible Navier–Stokes equations. The specific volume (ζ -) form of the compressible Navier–Stokes equations is utilized. Since these equations have polynomial (quadratic) nonlinearities, the Galerkin projection can be computed offline, once and for all; no hyper-reduction is required [20]. Unlike traditional eddy-viscosity-based stabilization methods, the proposed approach requires no additional empirical turbulence modeling terms – truncated modes are accounted for *a priori* via a minimal rotation of projection subspace. The method is also non-intrusive, as it operates only on the matrices and tensors defining a ROM ODE system, which are stabilized through the offline solution of a small trace minimization problem on the Stiefel manifold. The proposed new approach can be interpreted as a combination of several previously developed ideas. Following Iollo et al. [20], we propose to stabilize and enhance projection-based ROMs by modifying the projection subspace in order to capture more of the low-energy, but high dissipative scales of the flow solution. Similarly to Amsallem and Farhat [1], a rotation of the projection subspace is used to achieve this goal. Specifically, a larger set of basis is linearly superimposed to provide a smaller set of basis that generate a stable and accurate ROM. Finally, in the spirit of most previously proposed eddy-viscosity-based

turbulence models, the eigenvalues of the linear part of the Galerkin system of ODEs are used as a proxy to guide the stabilization algorithm.

The remainder of this paper is organized as follows. In § 2, the standard projection-based model reduction approach is outlined in the context of the specific volume form of the compressible Navier–Stokes equations. Also overviewed is the POD method for constructing an optimal reduced basis, and some eddy-viscosity based closure models used for accounting for modal truncation. The proposed methodology of “rotating” the projection subspace into a more dissipative regime to better resolve the small, energy dissipative scales of the flow is detailed in § 3. Here, the approach is formulated mathematically as a constrained optimization problem on the Stiefel manifold. In § 4, the performance of the proposed method is evaluated on several compressible flow problems, including a problem involving a laminar flow over an airfoil with a high angle of attack, and a channel-driven cavity flow problem. Finally, conclusions are offered in § 5.

2. Projection-based model reduction for nonlinear compressible flow

In this section the standard projection-based model reduction approach is laid out. In § 2.1 the approach for a general nonlinear system is presented while in § 2.2 the approach is applied to the compressible Navier–Stokes equations. The POD method for calculating a reduced basis using a set of snapshots from a high-fidelity simulation is outlined in § 2.3, followed by a brief overview of eddy-viscosity-based closure models that account for modes truncated in the application of the POD method (§ 2.4).

2.1. Nonlinear projection-based model order reduction

Consider a dynamical system of the form:

$$\frac{d}{dt} \mathbf{w} = \mathbf{F}(\mathbf{w}), \quad (1)$$

where \mathbf{F} is the propagator in H , a Hilbert space. In fluid flows, the state variable $\mathbf{w} = \mathbf{w}(\mathbf{x}, t) \in H$ depends on space $\mathbf{x} \in \Omega$, Ω being the flow domain, and time $t \in [0, T]$, T representing the period of integration. Then, the propagator \mathbf{F} contains spatial derivatives. The associated Hilbert space of square-integrable functions $L^2(\Omega)$ is equipped with the standard inner product for its elements $\mathbf{v}, \mathbf{w} \in L^2(\Omega)$, defined by:

$$(\mathbf{v}, \mathbf{w})_{\Omega} := \int_{\Omega} \mathbf{v} \cdot \mathbf{w} \, d\mathbf{x}. \quad (2)$$

In the Galerkin ROM approach, the governing variable, $\mathbf{w}(\mathbf{x}, t)$ is discretized using basis functions (modes) $\{\mathbf{w}_i(\mathbf{x})\}_{i=1}^n \in H$ with corresponding mode coefficients $\{a_i(t)\}_{i=1}^n$

$$\mathbf{w}(\mathbf{x}, t) \approx \mathbf{w}_0(\mathbf{x}) + \mathbf{w}^{[1..n]}(\mathbf{x}, t) := \mathbf{w}_0(\mathbf{x}) + \sum_{i=1}^n a_i(t) \mathbf{w}_i(\mathbf{x}), \quad (3)$$

where $\mathbf{w}_0(\mathbf{x})$ denotes the (steady) mean flow.

In the method of lines, the modes \mathbf{w}_i are known *a priori* and the goal is to find mode coefficients a_i that satisfy the differential equation (1). In general, the modes \mathbf{w}_i can be chosen in a number of ways. In the context of spectral methods in CFD for example, the basis vectors are usually analytical functions, e.g. trigonometric functions or Chebyshev polynomials. The advantage of these functions is that their spatial derivatives have analytical representations and numerically efficient algorithms such as the Fast Fourier Transform (FFT) can be utilized. In the context of ROMs, the spatial basis functions are usually derived *a posteriori* from a snapshot of a solution data set, like the Proper Orthogonal Decomposition (POD) [18] or Dynamic Mode Decomposition (DMD) [37,41]. Attention is restricted here to modes computed using the POD method (detailed in § 2.3), but it is noted that the methods proposed here hold for any choice of reduced basis. The reason for the choice of the POD reduced basis is two-fold. First, the POD is a widely used approach for computing efficient bases for fluid dynamical systems. Moreover, ROMs constructed via the POD/Galerkin method lack in general an *a priori* stability guarantee (meaning POD/Galerkin ROMs would benefit from ROM stabilization approaches such as the one developed herein).

The mode coefficients in (3) a_i are chosen to minimize the residual of the Galerkin expansion

$$\left(\mathbf{w}_i, \frac{d}{dt} \mathbf{w}^{[1..n]} \right)_{\Omega} - \left(\mathbf{w}_i, \mathbf{F} \left(\mathbf{w}_0(\mathbf{x}) + \mathbf{w}^{[1..n]} \right) \right)_{\Omega} = 0, \quad (4)$$

for $i = 1, \dots, n$. This projection yields a set of evolution equations for the mode coefficients a_i

$$\frac{d}{dt} a_i = f_i(\mathbf{a}), \quad (5)$$

where $\mathbf{a} := (a_1, \dots, a_n)^T$ represents the state and $\mathbf{f} := (f_1, \dots, f_n)^T$ its propagator. Given some initial conditions, the evolution equation (5) can be integrated using standard numerical integration techniques. The ROM system (5) is, by construction,

small, and can be integrated forward in time in real or near-real time unlike the high-fidelity CFD model from which it is derived.

2.2. Nonlinear reduction of the compressible Navier–Stokes equations

Consider the 2D compressible Navier–Stokes equations in primitive variables⁴:

$$\zeta_t + u\zeta_x + v\zeta_y - u_x\zeta - v_y\zeta = 0, \quad (6a)$$

$$u_t + uu_x + vv_y + \zeta p_x = \frac{1}{\text{Re}} \zeta \left[\left(\frac{4}{3}u_x - \frac{2}{3}v_y \right)_x + (v_x + u_y)_y \right], \quad (6b)$$

$$v_t + uv_x + vv_y + \zeta p_y = \frac{1}{\text{Re}} \zeta \left[\left(\frac{4}{3}v_y - \frac{2}{3}u_x \right)_y + (v_x + u_y)_x \right], \quad (6c)$$

$$p_t + up_x + vp_y + \gamma p(v_x + u_y) = \frac{\gamma}{\text{Re Pr}} [(p\zeta)_{xx} + (p\zeta)_{yy}] + \frac{1-\gamma}{\text{Re}} \left[u_x \left(\frac{4}{3}u_x - \frac{2}{3}v_y \right) + v_y \left(\frac{4}{3}v_y - \frac{2}{3}u_x \right) + (u_y + v_x)^2 \right]. \quad (6d)$$

Here, $\zeta(\mathbf{x}, t) = 1/\rho(\mathbf{x}, t)$ is the specific volume (the inverse of the density, $\rho(\mathbf{x}, t)$), $u(\mathbf{x}, t)$ and $v(\mathbf{x}, t)$ are the Cartesian components of the flow velocity, $p(\mathbf{x}, t)$ is the pressure, γ is the specific heat ratio, Re is the Reynolds number, Pr is the Prandtl number, and the subscripts denote partial derivatives. A Galerkin projection yields a system of coupled quadratic ODEs whose constant coefficients are calculated off-line and once and for all (see [Appendix A](#) and Iollo et al. [20] for details). This system has the form:

$$\frac{d\mathbf{a}}{dt} = \mathbf{C} + \mathbf{L}\mathbf{a} + [\mathbf{a}^T \mathbf{Q}^{(1)} \mathbf{a} \quad \mathbf{a}^T \mathbf{Q}^{(2)} \mathbf{a} \quad \dots \quad \mathbf{a}^T \mathbf{Q}^{(n)} \mathbf{a}]^T, \quad (7)$$

where $\mathbf{C} \in \mathbb{R}^n$, $\mathbf{L} \in \mathbb{R}^{n \times n}$ and $\mathbf{Q}^{(i)} \in \mathbb{R}^{n \times n}$, $\forall i = 1, \dots, n$.

Remark 1. The compressible Navier–Stokes equations are typically expressed in conservative form. This form is convenient for many applications including CFD. The conservative form contains rational functions of the unknowns and it is therefore not possible to pre-compute ROMs using standard Galerkin projection; to attain any computational speed-up a hyper-reduction step is necessary. Hyper-reduction is not always desirable, as it can destroy energy conservation properties and/or symplectic time-evolution maps [10,21]. On the other hand, if the equations are expressed in primitive variables, hyper-reduction can be avoided because all nonlinearities that appear are polynomial. It is for this reason that, in our approach, we base the ROM on the equations (6). It is possible to extend our proposed approach to the conservative formulation of the compressible Navier–Stokes equations; see [Remark 2](#).

2.3. Construction of optimal reduced-order basis via the POD

As discussed earlier, there exist a number of methods for calculating a reduced basis $\{\mathbf{w}_i(\mathbf{x})\}_{i=1}^n \in H$, e.g., proper orthogonal decomposition (POD) [44,18], Dynamic Mode Decomposition (DMD) [37,41], balanced POD (BPOD) [34,49], balanced truncation [16,26], and the reduced basis method [38,47]. In this paper, attention is restricted to reduced bases constructed using the first of these approaches, namely the POD method. This method is reviewed succinctly below.

Discussed in detail in Lumley [24] and Holmes et al. [18], POD is a mathematical procedure that, given an ensemble of data and an inner product, constructs a basis for the ensemble. The POD basis is optimal in the sense that it describes more energy (on average) of the ensemble in the chosen inner product than any other linear basis of the same dimension n . Let $\mathbf{w}^n \in \mathbb{R}^N$ denote a snapshot vector, computed as the solution of the fully discretized version of Eq. (6), for some instance of its parameters — that is, for some specific time t , some specific value of the set of flow parameters, or some boundary/initial conditions underlying this governing equation. Suppose a total of $K \in \mathbb{N}$ snapshots are collected from a high-fidelity simulation. A snapshot matrix is defined as a matrix $\mathbf{M} \in \mathbb{R}^{N \times K}$ whose columns are individual snapshots. The main focus of this paper is on unsteady flows and on snapshots associated with different time-instances. Hence, $M_{:,i} := \mathbf{w}^i$ for $i = 1, \dots, K$. Mathematically, POD seeks an n -dimensional ($n \ll K$ and $n \ll N$) subspace spanned by the set $\{\mathbf{w}_i(\mathbf{x})\}_{i=1}^n$ such that the difference between the ensemble $\{\mathbf{w}^i\}_{i=1}^K$ and its projection onto the reduced subspace is minimized on average. That is, a POD basis is obtained by solving the following low-rank matrix approximation problem:

For a given snapshot matrix $\mathbf{M} \in \mathbb{R}^{N \times K}$, find a lower rank matrix $\tilde{\mathbf{M}} \in \mathbb{R}^{N \times K}$ that solves the minimization problem

⁴ Presented in two-dimensions for the sake of brevity only. Extension to the three-dimensional equations is straightforward.

$$\min_{\text{rank}(\tilde{\mathbf{M}})=n} \|\mathbf{M} - \tilde{\mathbf{M}}\|_F, \quad (8)$$

where $n \ll N$, and $\|\cdot\|_F$ is the Frobenius norm.

In this problem, the rank constraint can be taken care of by representing the unknown matrix as $\tilde{\mathbf{M}} = \tilde{\mathbf{U}}\tilde{\mathbf{V}}$, where $\tilde{\mathbf{U}} \in \mathbb{R}^{N \times n}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{n \times K}$, so that problem (8) becomes

$$\min_{\tilde{\mathbf{U}} \in \mathbb{R}^{N \times n}, \tilde{\mathbf{V}} \in \mathbb{R}^{n \times K}} \|\mathbf{M} - \tilde{\mathbf{U}}\tilde{\mathbf{V}}\|_F. \quad (9)$$

It is well-known that the solution of the above low-rank approximation problem is given by the Eckart–Young–Mirsky [12,25] theorem via the Singular Value Decomposition (SVD) of \mathbf{M} . Specifically, $\tilde{\mathbf{U}} = U_{:,1:n}$ and $\tilde{\mathbf{V}} = (\Sigma \mathbf{V}^T)_{1:n,:}$ where $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T$. This is the so-called “method of snapshots” for computing a POD basis [44].

2.4. Accounting for modal truncation: eddy-viscosity based closure models

In the Kolmogorov description of the turbulence cascade, the large, energy-carrying flow scales transfer energy to successively smaller scales where finally dissipative forces can dissipate their energy [46,30]. The large, energy-carrying scales are associated with the large singular values of the snapshot matrix \mathbf{M} , while the smaller, energy-dissipative scales of the flow are associated with the smaller singular values. Since low order POD-based ROMs remove modes corresponding to small singular values, these ROMs are, by construction, not endowed with the dissipative dynamics of the flow.

Many of the popular methods for accounting for truncated modes fall in to the family of eddy-viscosity based closure models.⁵ In this family of methods, dynamics of the truncated modes are modeled by modifying the coefficients of the Galerkin model. For example, in the linear eddy-viscosity approach, the linear term of the Galerkin system is modified. Equation (7) is replaced with

$$\frac{d\mathbf{a}}{dt} = \mathbf{C} + (\mathbf{L} + \hat{\mathbf{L}})\mathbf{a} + [\mathbf{a}^T \mathbf{Q}^{(1)} \mathbf{a} \quad \mathbf{a}^T \mathbf{Q}^{(2)} \mathbf{a} \quad \dots \quad \mathbf{a}^T \mathbf{Q}^{(n)} \mathbf{a}]^T \quad (10)$$

Here, $\hat{\mathbf{L}}$ is an additional linear term whose role is to modify the overall eigenvalue distribution of the linear operator $(\mathbf{L} + \hat{\mathbf{L}})$, i.e., decrease the magnitude of this operator's real positive eigenvalues and increase the magnitude of its real negative eigenvalues. This amounts to decreasing energy production and increasing energy dissipation, respectively. In general the appropriate eigenvalue distribution is not known *a priori* and must be identified via a solution matching procedure. For a detailed review of the performance of the various methods based on this approach, the reader is referred to Wang et al. [48], Rempfer and Fasel [33].

Although the approach described above has been applied successfully to a large number of Galerkin models of complex, high-Reynolds number flows, it has a significant drawback, namely the loss of consistency between the Navier–Stokes equations and the Galerkin system. Since the Galerkin system is modified empirically, the resulting quadratic system of ODEs no longer corresponds to a Galerkin projection of the Navier–Stokes equations.

In the following section, a novel stabilization and enhancement approach that retains consistency is introduced.

3. Stabilization and enhancement of compressible flow ROMs via subspace rotation

In this section the new proposed stabilization and enhancement approach for ROMs is outlined. In this method the projection subspace is “rotated” into a more dissipative regime by modifying the eigenvalue distribution of the linear operator. This new approach may be interpreted as an *a priori* implementation of a traditional eddy-viscosity based closure model.

The modes $\mathbf{w}_i(\mathbf{x})$, $i = 1, 2, \dots, n$ are constructed via linear-superposition of $n + p$ (with $p > 0$) most energetic POD modes. Mathematically this can be expressed as:

$$\tilde{\mathbf{w}}_i = \sum_{j=1}^{n+p} X_{ji} \mathbf{w}_j \quad i = 1, \dots, n, \quad (11)$$

where $\mathbf{X} \in \mathbb{R}^{(n+p) \times n}$ is the orthonormal ($\mathbf{X}^T \mathbf{X} = \mathbf{I}_{n \times n}$) “rotation” matrix. The Galerkin system tensors associated with these new modes are expressed as a function of \mathbf{X} as follows:

⁵ Other ROM stabilization approaches, developed independently from the “modal constant eddy-viscosity” approach and for a broader range of applications than fluid mechanics, e.g., the method of Kalashnikova et al. [22] for stabilizing generic linear ROMs via optimization-based eigenvalue reassignment, give rise to a similar modification to the linear term.

$$\tilde{Q}_{jk}^{(i)} = \sum_{s,q,r=1}^{n+p} X_{si} Q_{qr}^{(s)} X_{qj} X_{rk} \quad i, j, k = 1, \dots, n, \tag{12a}$$

$$\tilde{\mathbf{L}} = \mathbf{X}^T \mathbf{L} \mathbf{X}, \tag{12b}$$

$$\tilde{\mathbf{C}} = \mathbf{X}^T \mathbf{C}^*, \tag{12c}$$

where $\mathbf{C} \in \mathbb{R}^{n+p}$, $\mathbf{L} \in \mathbb{R}^{(n+p) \times (n+p)}$ and $\mathbf{Q}^{(i)} \in \mathbb{R}^{(n+p) \times (n+p)}$, $\forall i = 1, \dots, (n+p)$ are the Galerkin system coefficients corresponding to the first $n+p$ most energetic POD modes. The new Galerkin system is of the form

$$\frac{d\mathbf{a}}{dt} = \tilde{\mathbf{C}} + \tilde{\mathbf{L}}\mathbf{a} + \left[\mathbf{a}^T \tilde{\mathbf{Q}}^{(1)} \mathbf{a} \quad \mathbf{a}^T \tilde{\mathbf{Q}}^{(2)} \mathbf{a} \quad \dots \quad \mathbf{a}^T \tilde{\mathbf{Q}}^{(n)} \mathbf{a} \right]^T, \tag{13}$$

where the matrices $\tilde{\mathbf{Q}}^{(i)}$, $\tilde{\mathbf{L}}$ and $\tilde{\mathbf{C}}$ are given by (12).

The goal of the proposed approach is to find \mathbf{X} such that 1.) the new modes $\tilde{\mathbf{w}}_i$ remain good approximations of the flow, and 2.) the new Galerkin ROM is stable and accurate. To ensure that these properties are satisfied, a constrained optimization problem is formulated for \mathbf{X} . To guarantee that the new modes remain good approximations of the flow, the distance $\|\mathbf{X} - \mathbf{I}_{n+p,n}\|_F$ is minimized, where $\mathbf{I}_{n+p,n}$ are the first n columns of an $n+p$ identity matrix. To ensure that the ROM is stable and accurate, the traditional linear eddy-viscosity closure ansatz is used as a constraint. Specifically, the constraint involves the overall balance between linear energy production and dissipation via the trace of the modified linear operator $\text{tr}(\tilde{\mathbf{L}}) = \text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = \sum_i^n \tilde{\lambda}_i = \eta$ where $\tilde{\lambda}_i$ are eigenvalues of $\mathbf{X}^T \mathbf{L} \mathbf{X}$. Although this constraint alone does not guarantee any particular eigenvalue distribution, the objective function promotes minimal rotations and thus minimal modifications of the eigenvalue distribution. Some alternative candidate constraint and objective functions are outlined in [Appendix B](#).

Mathematically, the constrained optimization problem for \mathbf{X} outlined above reads as follows:

$$\begin{aligned} & \underset{\mathbf{X} \in \mathcal{V}_{(n+p),n}}{\text{minimize}} && -\text{tr}(\mathbf{X}^T \mathbf{I}_{n+p,n}) \\ & \text{subject to} && \text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = \eta \end{aligned} \tag{14}$$

where $\eta \in \mathbb{R}$ and

$$\mathcal{V}_{(n+p),n} \in \{ \mathbf{X} \in \mathbb{R}^{(n+p) \times n} : \mathbf{X}^T \mathbf{X} = \mathbf{I}_n, p > 0 \}. \tag{15}$$

In (15), $\mathcal{V}_{(n+p),n}$ is the Stiefel manifold, defined as the set of $(n+p) \times n$ matrices satisfying the orthonormality condition $\mathbf{X}^T \mathbf{X} = \mathbf{I}_n$, where \mathbf{I}_n is the $n \times n$ identity matrix [32,45]. In Equation (14) the objective function is simplified by utilizing the property that for a real matrix $\|\mathbf{A}\|_F^2 = \text{tr}(\mathbf{A}^T \mathbf{A})$. Thus, minimizing $\|\mathbf{X} - \mathbf{I}_{n+p,n}\|_F$ is equivalent to minimizing $-\text{tr}(\mathbf{X}^T \mathbf{I}_{n+p,n})$.

The appropriate eigenvalue distribution η must be identified using a solution matching procedure. Discussion of an approach for selecting η is deferred until § 3.2.

Remark 2. In this paper, we assume that the ROM advanced forward in time during the online time-integration step of the model reduction is a system of the form (10), which arises when projecting the compressible Navier–Stokes equations in primitive specific volume form (6) onto the reduced basis modes. As suggested in [Remark 1](#), the method described here can be applied in the case the ROM is based on the compressible Navier–Stokes equations in conservative form (with or without hyper-reduction). In this case, the model reduction would proceed as follows:

- Step 1: Run a high-fidelity code to generate snapshots from which the POD basis will be constructed.
- Step 2: Construct from the snapshots collected in Step 1 a POD basis $\{\mathbf{w}_i(\mathbf{x})\}_{i=1}^{n+p}$ for the primitive variables.
- Step 3: Project the compressible Navier–Stokes equations in primitive specific-volume form (6) onto the modes from Step 2 to obtain a system of the form (10).
- Step 4: Use the Galerkin matrices \mathbf{C} , \mathbf{L} and $\mathbf{Q}^{(i)}$ to obtain from the original basis $\{\mathbf{w}_i(\mathbf{x})\}_{i=1}^{n+p}$ a stabilized basis $\tilde{\mathbf{w}}_i = \sum_{j=1}^{n+p} X_{ji} \mathbf{w}_j$, $i = 1, \dots, n$, where \mathbf{X} is the solution to (14).
- Step 5: Transform the stabilized basis $\{\tilde{\mathbf{w}}_i(\mathbf{x})\}_{i=1}^n$ into conservative variables, and use it in a ROM code that projects the compressible Navier–Stokes equations in conservative form (with or without hyper-reduction).

To apply this procedure, two ROM codes are required: a ROM code that projects the compressible Navier–Stokes equations in primitive specific-volume form, and a ROM code that projects the compressible Navier–Stokes equations in conservative form. The former code is only needed to calculate the \mathbf{C} , \mathbf{L} and $\mathbf{Q}^{(i)}$ matrices, which are used for the basis stabilization.

3.1. Solution of constrained optimization problem

A common method for solving constrained optimization problems of the form (14) is the method of Lagrange multipliers [28]. In this method, the Lagrangian of the optimization problem is computed, and its stationary points are sought, yielding necessary optimality conditions for local maxima and minima. The reader can verify that the Lagrangian for Eq. (14) is

$$\mathcal{L}(\mathbf{X}, \mathbf{\Lambda}_1, \mathbf{\Lambda}_2) := -\text{tr}(\mathbf{X}^T \mathbf{I}_{n+p,n}) + \text{tr}(\mathbf{\Lambda}_1 (\mathbf{X}^T \mathbf{L} \mathbf{X} - \frac{\eta}{n} \mathbf{I}_n)) + \text{tr}(\mathbf{\Lambda}_2 (\mathbf{X}^T \mathbf{X} - \mathbf{I}_n)), \quad (16)$$

where $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$ are diagonal matrices of Lagrange multipliers.

Suppose that \mathbf{X} is a local maximizer of problem (14). Then \mathbf{X} satisfies the first-order optimality condition $\mathcal{L}_{\mathbf{X}} = -\mathbf{I}_{n+p,n} + \mathbf{\Lambda}_1 (\mathbf{L} + \mathbf{L}^T) \mathbf{X} + 2\mathbf{\Lambda}_2 \mathbf{X} = \mathbf{0}$, $\text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X} - \frac{\eta}{n} \mathbf{I}_n) = \mathbf{0}$, and $\mathbf{X}^T \mathbf{X} - \mathbf{I}_n = \mathbf{0}$. Solving Eq. (14) using Lagrange multipliers is possible; however, it is inefficient. Significant speed-ups are possible by satisfying the orthonormality constraint directly via optimization on the Stiefel matrix manifold. In this method, with the help of the augmented Lagrange method, the constrained optimization problem is reduced to an unconstrained optimization problem on the Stiefel manifold as follows:

$$\underset{\mathbf{X} \in \mathcal{V}_{(n+p),n}}{\text{minimize}} \quad -\text{tr}(\mathbf{X}^T \mathbf{I}_{n+p,n}) + \frac{\mu_k}{2} \text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X} - \frac{\eta}{n} \mathbf{I}_n)^2 - \lambda_{\mathcal{L}} \text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X} - \frac{\eta}{n} \mathbf{I}_n), \quad (17)$$

where μ_k is increased until the constraint is satisfied to some desired precision. The variable $\lambda_{\mathcal{L}}$ is an estimate of the Lagrange multiplier and is updated according to the rule

$$\lambda_{\mathcal{L}} \leftarrow \lambda_{\mathcal{L}} - \mu_k \text{tr}(\mathbf{X}^{(k)T} \mathbf{L} \mathbf{X}^{(k)} - \frac{\eta}{n} \mathbf{I}_n), \quad (18)$$

where $\mathbf{X}^{(k)}$ is the solution of the unconstrained problem at the k th step. In this work, the `Manopt` MATLAB toolbox [7] is used to solve (17). The algorithm is initialized with $\mathbf{X}^{(0)} = \mathbf{I}_{n+p,n}$ that corresponds to the standard Galerkin ROM. All derivatives in the optimization algorithm are calculated analytically.

3.2. Solution matching procedure for η

In this section a solution matching procedure for the appropriate eigenvalue distribution, η is outlined. The relative error of the solution delivered by a ROM is defined as

$$e(\eta) = \frac{\langle E_{CFD} \rangle_T - \langle E_{ROM} \rangle_T}{\langle E_{CFD} \rangle_T} \times 100 \quad (19)$$

where $E(t) = \sum_i^n a_i^2(t)$ is the solution “energy” and $\langle \cdot \rangle_T$ is the mean value (temporal average). The solution matching procedure consists of identifying the root $e(\eta) = 0$ using the bisection method as summarized in Algorithm 1. For all numerical experiments the endpoint values are set to $\eta_b = \text{tr}((\mathbf{X}^{(0)})^T \mathbf{L} (\mathbf{X}^{(0)})) = \text{tr}(\mathbf{L}_{1:n,1:n})$, $\eta_a = \eta_b - 5|\eta_b|$, where η_b corresponds to the standard n -order Galerkin ROM. The convergence tolerance, TOL and the maximum number of iterations, $NMAX$ are set to 0.1% and 100, respectively.

Remark 3. Here, we provide some general guidelines and remarks pertaining to the proposed stabilization and fine-tuning algorithm.

- Our numerical experiments suggest $n = p$ provides best performance; however the optimal choice of p remains an open question.
- The proposed algorithm requires that $p \geq 1$. For $p = 0$, the rotation matrix \mathbf{X} is square and therefore corresponds to a coordinate transformation that, by definition, can have no effect on the dynamics of the system.
- Uniqueness results for the solution to (14) are not provided in this work. We have found the results to be insensitive to the initial conditions, and the solutions to (14) to be unique for a large number of random initial conditions.
- For some choices of n , p and η the optimization problem (17) does not have feasible solutions. Consider the constraint $\text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = \eta$. Since the trace is invariant under cyclic permutations, we have that $\text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = \text{tr}(\mathbf{L} \mathbf{X} \mathbf{X}^T)$. The rectangular matrix \mathbf{X} is orthonormal $\mathbf{X}^T \mathbf{X} = \mathbf{I}_n$ so the product $\mathbf{X} \mathbf{X}^T$ is positive semidefinite. It was proved in [39] that for an arbitrary matrix \mathbf{A} , and positive semidefinite matrix \mathbf{B}

$$\text{tr}(\mathbf{A} \mathbf{B}) \leq |\text{tr}(\mathbf{A} \mathbf{B})| \leq \|\mathbf{A}\|_2 \text{tr}(\mathbf{B}) \quad (20)$$

where $\|\mathbf{A}\|_2 = \sigma_{\max}(\mathbf{A})$ is the spectral norm, i.e., the largest singular value of \mathbf{A} . Setting $\mathbf{A} = \mathbf{L}$ and $\mathbf{B} = \mathbf{X} \mathbf{X}^T$, and utilizing the fact that the trace of a projection matrix is equal to its rank, $|\frac{\eta}{n}| \leq \sigma_{\max}(\mathbf{L})$ is a necessary condition for the existence of a solution to (14).

Algorithm 1: Stabilization algorithm for compressible Navier–Stokes equations.

```

input : endpoint values  $[\eta_a, \eta_b]$  (where  $\eta_a < \eta_b$ ), convergence tolerance  $TOL$ , maximum iterations,  $NMAX$ , ROM size  $n$  and  $p \geq 1$ , and Galerkin
system matrices associated with the first  $n + p$  most energetic POD modes,  $\mathbf{C} \in \mathbb{R}^{n+p}$ ,  $\mathbf{L} \in \mathbb{R}^{(n+p) \times (n+p)}$  and
 $\mathbf{Q}^{(i)} \in \mathbb{R}^{(n+p) \times (n+p)}$ ,  $i = 1, \dots, n + p$ 
output: stabilizing rotation matrix  $\mathbf{X}$ , and eigenvalue distribution  $\eta$ 

1 while  $N \leq NMAX$  do
2   set new midpoint to  $\eta_c = (\eta_a + \eta_b)/2$ 
3   solve constrained optimization problem on Stiefel manifold:

       minimize    $-\text{tr}(\mathbf{X}^T \mathbf{I}_{n+p,n})$ 
        $\mathbf{X} \in \mathcal{V}_{(n+p),n}$ 
       subject to  $\text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = \eta_c$ 

4   construct new Galerkin matrices using (12)
5   integrate numerically new Galerkin system
6   calculate relative ROM error  $e(\eta_c)$ 
7   if  $|e(\eta_c)| < TOL$  then
8      $\eta := \eta_c$ 
9     output  $\mathbf{X}$ , and  $\eta$ 
10    terminate the algorithm
11  end
12  if  $\text{sign}(e(\eta_c)) = \text{sign}(e(\eta_a))$  then
13     $\eta_a = \eta_c$ 
14  else
15     $\eta_b = \eta_c$ 
16  end
17 end
18 output (“Method failed, maximum number of iterations exceeded”)

```

4. Numerical experiments

In this section, we evaluate the performance of [Algorithm 1](#) for stabilizing compressible flow ROMs on several 2D problems: a problem involving a laminar flow around an inclined airfoil, and a channel-driven cavity problem at two Reynolds numbers. In all cases, the flow is governed by the full compressible Navier–Stokes equations with constant viscosity. Direct Numerical Simulations (DNS) are performed and POD basis functions are derived from snapshots collected during these simulations. ROMs are derived by projecting the fully compressible Navier–Stokes equations in specific volume (ζ -) form onto the first n most energetic basis modes. The projection is performed off-line, and once and for all, resulting in a system of n coupled quadratic ODEs in the form of (7). From this point forward, such ROMs are referred to as “standard POD–Galerkin ROMs”, where “standard” refers to the fact that no model is used to account for the dynamics of the truncated modes, $n + 1, n + 2, \dots, \infty$. ROMs derived using the stabilization approach proposed in this paper are referred to as “stabilized POD–Galerkin ROMs”.

4.1. High angle of attack laminar airfoil

The first test case involves the 2D flow around an inclined NACA0012 airfoil at Mach 0.7, and $Re = 500$ at 25 degrees angle of attack. The Reynolds number is based on the chord of the airfoil, c . At this Reynolds number and angle of attack, the flow is separated and the solution corresponds to a stable limit cycle. High-fidelity simulation snapshots are generated using a second-order, finite-difference, embedded boundary (EB) solver. The no-slip adiabatic boundary conditions are satisfied using a first-order ghost fluid method. For a detailed description of the scheme the reader is referred to Balajewicz and Farhat [4]. The snapshots correspond to a DNS of the compressible Navier–Stokes equations. The viscosity of the fluid inside the domain is assumed constant. The computational domain extends $20c$ in all directions and a sponge zone of thickness $2c$ is used to help absorb waves entering and leaving the domain. The domain is discretized using a non-uniform 300×300 cartesian grid. The flow is initialized by setting the solution at all grid points to the free stream values. Time integration is performed using the second-order BDF scheme and a constant time step corresponding to a $CFL = 1$ is used. Snapshot collection begins after 5000 time steps to ensure the solution has reached the limit cycle. A total of $K = 500$ snapshots are collected every 5 simulation time steps. The first four basis functions capture approximately 86% of the snapshot energy.

[Fig. 1](#) illustrates the performance of a stabilized $n, p = 4$ ROM of the laminar airfoil using $\eta = -3.513 \times 10^{-1}$. The stabilization algorithm terminated after 19 iterations with a relative error $e(\eta) = 0.017\%$. In [Fig. 1\(a\)](#), the global energy of the high-fidelity CFD model, and standard (i.e., unstabilized) and stabilized ROMs are illustrated. The stabilized ROM is shown to track very accurately the mean of the fluctuating energy of the CFD solution while the standard ROM overpredicts the mean by an order of magnitude. To investigate the long term stability of the stabilized ROM, the system was numerically integrated $100 \times$ the duration of the original snapshots. No change or drift in trajectory was observed during this long integration period. In [Fig. 1\(b\)](#) the trajectories of the first and second temporal coefficient, $a_1(t)$, and $a_2(t)$ respectively,

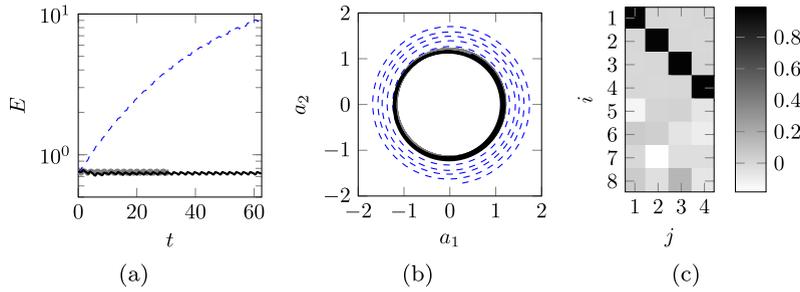


Fig. 1. Nonlinear model reduction of the laminar airfoil. Evolution of modal energy (a), and phase plot of the first and second temporal basis, $a_1(t)$ and $a_2(t)$ (b); CFD (thick gray line), standard $n = 4$ ROM (dashed blue line), stabilized $n, p = 4$ ROM (solid black line). Stabilizing rotation matrix, \mathbf{X} (c). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

are illustrated. The stabilized ROM accurately reproduces the closed orbit of the stable limit cycle while the standard ROM predicts an unstable spiral. Figs. 1(a) and 1(b) demonstrate how the stabilized ROM reproduces reliably both the global mean and fluctuating components of the CFD solution. The stabilizing transformation matrix \mathbf{X} for this problem is illustrated in Fig. 1(c). As expected the rotation of the projection subspace is small as demonstrated by the fact that $\mathbf{X} \approx \mathbf{I}_{n+p,n}$. For this configuration, the normalized error defined as $\|\mathbf{X} - \mathbf{I}_{n+p,n}\|_F/n$ is 0.083.

Finally, the predicted velocity magnitude at the final snapshot is illustrated in Fig. 2. The stabilized ROM (Fig. 2(a)) reproduces the velocity contours of the original high-fidelity CFD solution (Fig. 2(c)) remarkably well, in contrast to the standard ROM (Fig. 2(b)). This demonstrates the effectiveness of the proposed model reduction approach.

4.2. Channel driven laminar cavity

For the results presented in this section, the high-fidelity fluid simulation data are generated using a Sandia National Laboratories' in-house finite volume flow solver known as SIGMA CFD. This code is derived from LESLIE3D [40], a Large Eddy Simulations (LES) flow solver originally developed in the Computational Combustion Laboratory at the Georgia Institute of Technology. The code has LES as well as DNS capabilities. For the channel-driven laminar cavity problem considered here, the code was run in DNS mode. For a detailed description of the schemes and models implemented within LESLIE3D, the reader is referred to [14,15].

ROMs for the channel-driven laminar cavity problem are constructed using a Sandia in-house parallel C++ model reduction code known as Spirit, which constructs ROMs for compressible flow problems using the POD and continuous projection method. This code, detailed in [21], reads in the snapshot and mesh data written by a high-fidelity flow solver, creates a finite element representation of the snapshots and computes the numerical quadrature necessary for evaluation of the inner products arising in the Galerkin projection step of the model reduction. All calculations are performed in parallel using distributed matrix and vector data structures and parallel eigensolvers from the Trilinos project [17], which allows for large data sets and a relatively large number of POD modes. The libmesh finite element library [23] is used to compute the element quadratures.

In the discussion that follows, two variants of the 2D channel-driven laminar cavity problem are considered: a low Reynolds number variant ($Re \approx 1500$) and a moderate Reynolds number variant ($Re \approx 5500$). Both tests cases involve a Mach 0.6 viscous laminar flow over a cavity in a T-shaped domain (Fig. 3). The flow conditions for both tests are similar to case L2 in [35]. The free stream pressure is 25 Pa, the free stream temperature is 300 K, and the free stream velocity is 208.8 m/s. The viscosity μ is spatially constant and calculated such that the above Reynolds numbers are achieved. The thermal conductivity κ is also constant, calculated such that the Prandtl number is $Pr = 0.72$. At the inflow boundary, a value of the velocity and temperature that is above the free stream values is specified. The flow at the cavity walls is assumed to be adiabatic and to satisfy a no-slip condition. The remaining outflow boundaries are open, and a far-field boundary condition that suppresses the reflection of waves into the computational domain is implemented here. The high-fidelity simulation is initialized by setting the flow in the cavity to have a zero velocity, free stream pressure, and temperature. The region above the cavity is initialized to free stream conditions and the flow is allowed to evolve. As both SIGMA CFD and Spirit are 3D codes, a 2D mesh of the domain Ω is converted to a 3D mesh by extruding the 2D mesh in the z -direction by one element. Finally, it is noted that SIGMA CFD and Spirit work with different meshes. The former code requires a structured hexahedral discretization, whereas the latter assumes a tetrahedral discretization. To overcome this difference, the hexahedral high-fidelity meshes associated with the snapshots are converted to tetrahedral meshes prior to constructing the ROMs. This is accomplished by breaking up each hexahedral element into six tetrahedral elements.

4.2.1. Low Reynolds number ($Re \approx 1500$)

For the first, low Reynolds number variant of the channel-driven laminar cavity problem, the free-stream viscosity is set to $\mu_\infty = 3.17 \times 10^{-6}$ kg/(m·s), so that $Re = 1453.9$. The discretized domain, illustrated in Fig. 3, consists of 98,408 nodes, cast as 292,500 tetrahedral finite elements within the ROM code, Spirit. The 2D extent of the domain is: $[(-6.42, 10) \times (-1, 10)] \setminus [(-6.42, 10) \times (-1, 0) \cup (2, 10) \times (-1, 0)]$ m. The reader can observe that the mesh is structured but non-uniform.

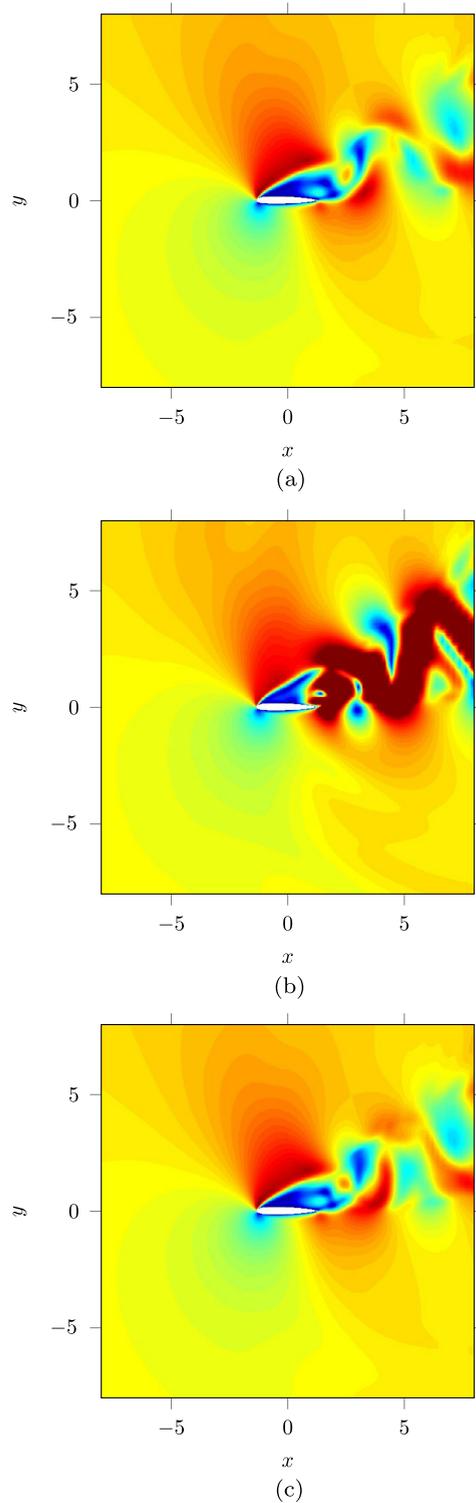


Fig. 2. Snapshot of high angle of attack airfoil at final snapshot; contours of velocity magnitude. CFD (a), standard $n=4$ ROM (b), and stabilized $n, p=4$ ROM (c).

The high-fidelity solver, SIGMA CFD, is initiated with the conditions described above and allowed to run until a statistically stationary flow regime is reached. At this point, a total of $K=500$ snapshots are collected from SIGMA CFD, taken every $\Delta t_{snap} = 1 \times 10^{-4}$ seconds. The snapshots are used to construct a POD basis of size 4 modes in the L^2 inner product. This basis captured about 91% of the snapshot energy. For more details on this test case, the reader is referred to [21].

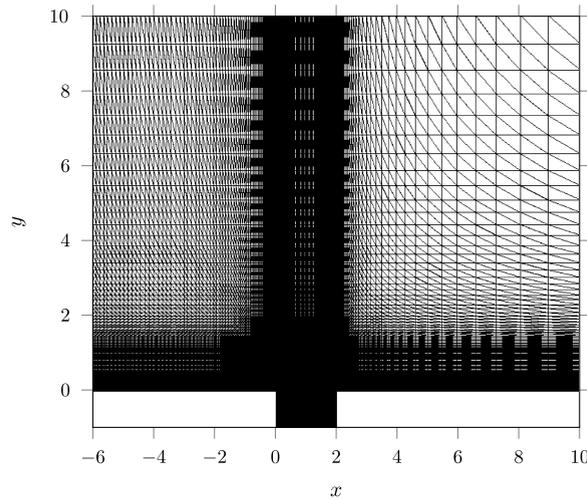


Fig. 3. Domain and mesh for viscous channel-driven cavity problem.

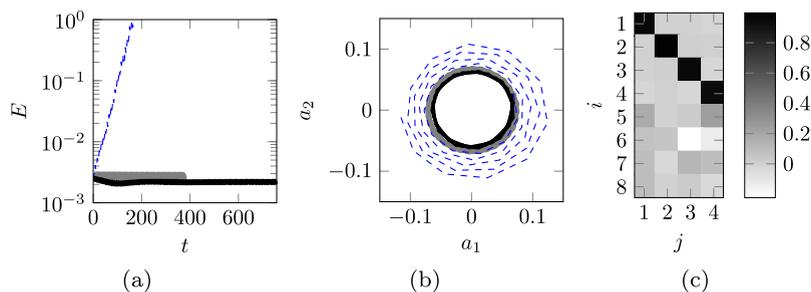


Fig. 4. Nonlinear model reduction of channel-driven cavity at $Re \approx 1500$. Evolution of modal energy (a) and phase plot of the first and second temporal basis, $a_1(t)$ and $a_2(t)$ (b); CFD (thick gray line), standard $n = 4$ ROM (dashed blue line), stabilized $n, p = 4$ ROM (solid black line). Stabilizing rotation matrix, \mathbf{X} (c). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

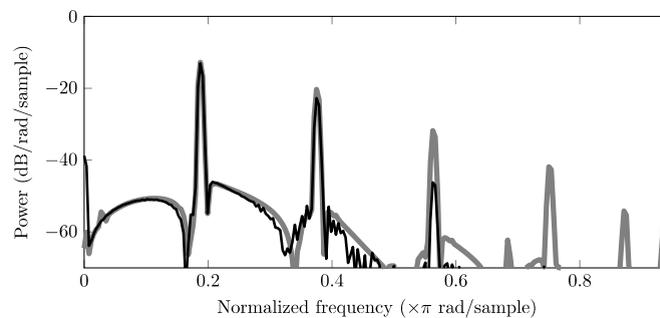


Fig. 5. PSD of $p(\mathbf{x}, t)$ where $\mathbf{x} = (2, -1)$ of channel-driven cavity $Re \approx 1500$. CFD (thick gray line), stabilized $n, p = 4$ ROM (black line).

Fig. 4 illustrates the performance of a stabilized $n, p = 4$ ROM of the low Reynolds number cavity using $\eta = -3.330 \times 10^{-3}$. The stabilization algorithm terminated after 18 iterations with a relative error $e(\eta) = -0.012\%$. In Fig. 4(b), the modal energy of the CFD, standard, and stabilized ROMs are illustrated. The stabilized ROM is shown to track very accurately the energy of the original CFD solution while the standard ROM is unable to reproduce this trajectory. The long term stability of the stabilized ROM was validated by numerically integrating the system $100\times$ the duration of the original snapshots. No change or drift in trajectory was observed during this long integration period. In Fig. 4(b) the trajectories of the first and second temporal coefficient, $a_1(t)$, and $a_2(t)$ respectively, are illustrated. The stabilized ROM predicts correctly the closed orbit of the stable limit cycle while the standard ROM predicts an unstable spiral. The stabilizing transformation matrix \mathbf{X} for this problem is illustrated in Fig. 4(c). As before, the rotation of the projection subspace is small as demonstrated by the fact that $\mathbf{X} \approx \mathbf{I}_{n+p,n}$. For this configuration, the normalized error defined as $\|\mathbf{X} - \mathbf{I}_{n+p,n}\|_F/n$ is 0.1182.

Fig. 5 shows the Power Spectral Density (PSD) of the predicted pressure fluctuations at the bottom right corner of the cavity, $\mathbf{x} = (2, -1)$. Both the fundamental and first harmonic of the response is accurately predicted by the stabilized

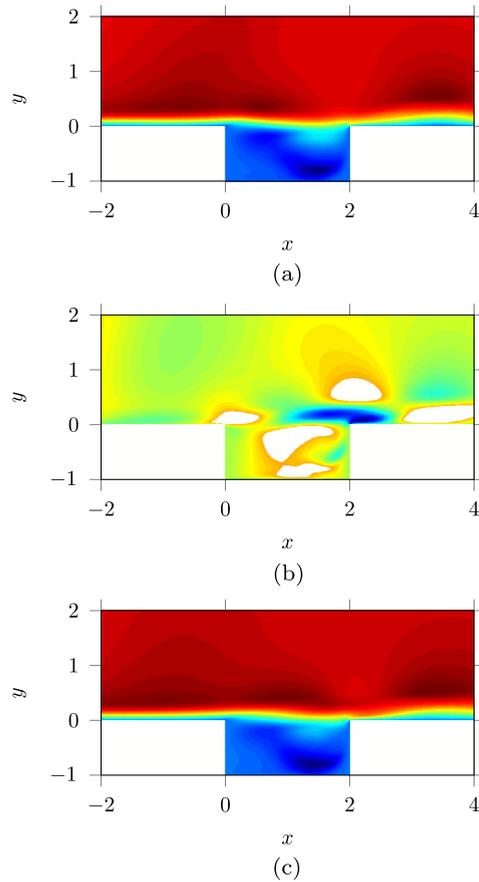


Fig. 6. Snapshot of channel-driven cavity $Re \approx 1500$; contours of u -velocity magnitude at the final snapshot. CFD (a), standard $n = 4$ ROM (b) and stabilized $n, p = 4$ ROM (c).

$n, p = 4$ ROM. The PSD of the CFD signal was computed using all available snapshots from $t = 0$ to $t = 380$ where t is non-dimensional. On the other hand, the PSD of the stabilized ROM was computed from the signal $100\times$ past the duration of the original snapshots; i.e. $t = (38000 - 380) = 37620$ to $t = 38000$.

Finally, a snapshot of the predicted velocity magnitude at the final snapshot is illustrated in Fig. 6. The stabilized ROM (Fig. 6(a)) reproduces the velocity contours of the original high-fidelity CFD simulation (Fig. 6(c)) remarkably well. In contrast, the standard ROM (Fig. 6(b)) is unstable and inaccurate.

4.2.2. Moderate Reynolds number ($Re \approx 5500$)

The next test case considered is also a channel-driven laminar cavity problem, but at a higher Reynolds number. The only parameter that is different is the free-stream viscosity, now set to $\mu_\infty = 8.46 \times 10^{-7}$ kg/(m·s), so that $Re = 5452.1$. Also changed is the size of the geometry extent, which has a larger sponge region near the outflow regions. This is needed to suppress adequately the reflection of waves into the computational domain for this problem. Toward this effect, the 2D extent of the domain is: $[(-6.42, 30) \times (-1, 30)] \setminus [(-6.42, 30) \times (-1, 0) \cup (2, 30) \times (-1, 0)]$ m. The geometry is discretized by 117,328 nodes, cast as 345,900 tetrahedral elements in *Spirit*. As before, the mesh is structured but non-uniform. The flow is significantly more chaotic than the $Re \approx 1500$ case considered in Section 4.2.1.

A total of $K = 500$ snapshots are collected from SIGMA CFD at increments $\Delta t_{snap} = 1 \times 10^{-5}$ seconds. As before, snapshots collection does not begin until a statistically stationary flow regime has been reached. From these snapshots, a POD basis of size 20 modes is constructed in the L^2 inner product. This basis captures about 72% of the snapshot energy. Typically, n would be selected such that the POD basis captures a greater percentage of the snapshot ensemble energy (e.g., $\approx 90\%$ or more). We choose a basis that captures less energy of the snapshot set to highlight the effectiveness of our approach for low-dimensional POD expansions.

Fig. 7 illustrates the performance of a stabilized $n, p = 20$ ROM of the higher Reynolds number cavity problem using $\eta = -5.679 \times 10^{-1}$. The stabilization algorithm terminated after 25 iterations with a relative error $e(\eta) = -0.065\%$. In Fig. 7(a), the modal energy of the CFD, standard, and stabilized ROMs are illustrated. The standard ROM is shown to overpredict the energy of the original CFD solution by an order of magnitude. The predictive power of the stabilized ROM is demonstrated by numerically integrating the ROM $10\times$ the duration of the original snapshots. The stabilizing transformation matrix \mathbf{X} for

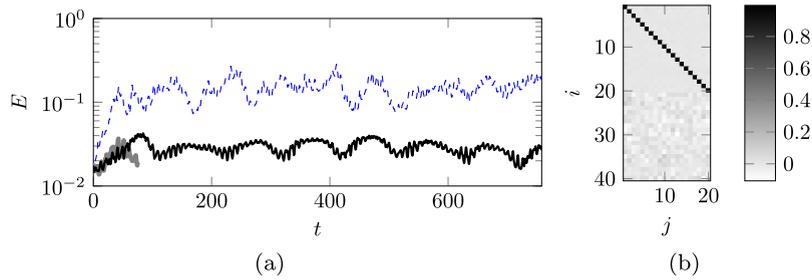


Fig. 7. Nonlinear model reduction of channel-driven cavity at $Re \approx 5500$. Evolution of modal energy (a); CFD (thick gray line), standard $n = 20$ ROM (dashed blue line), stabilized $n, p = 20$ ROM (solid black line). Stabilizing rotation matrix, X (b). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

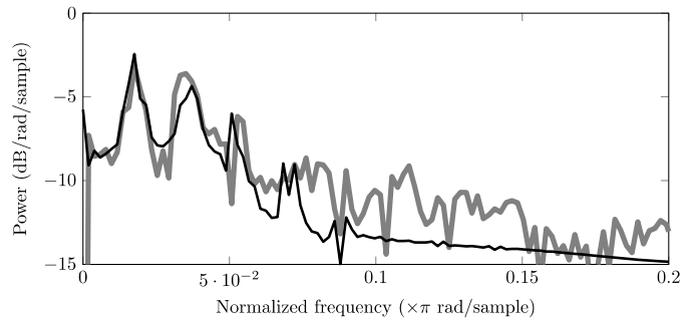


Fig. 8. PSD of $p(x_1, t)$ where $x_1 = (2, -0.5)$ of channel-driven cavity at $Re \approx 5500$. CFD (thick gray line), stabilized $n, p = 20$ ROM (black line).

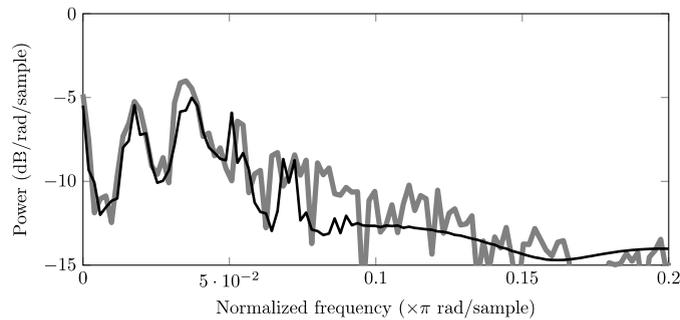


Fig. 9. PSD of $p(x_2, t)$ where $x_2 = (0, -0.5)$ of channel-driven cavity at $Re \approx 5500$. CFD (thick gray line), stabilized $n, p = 20$ ROM (black line).

this problem is illustrated in Fig. 7(b). As before, the rotation of the projection subspace is small as demonstrated by the fact that $X \approx I_{n+p,n}$. For this configuration, the normalized error defined as $\|X - I_{n+p,n}\|_F/n$ is 0.0384.

Figs. 8 and 9 show the PSDs of the predicted pressure fluctuations at locations $x_1 = (2, -0.5)$ and $x_2 = (2, 0.5)$, respectively. The PSD of the CFD signal was computed using all available snapshots from $t = 0$ to $t = 67$ where t is non-dimensional. On the other hand, the PSD of the stabilized ROM was computed from the signal $10\times$ past the duration of the original snapshots; i.e. $t = (670 - 67) = 603$ to $t = 670$. The stabilized ROM accurately predicts the chaotic pressure fluctuations at both locations. Fig. 10 illustrates the Cross Power Spectral Density (CPSD) for pressure fluctuations at x_1 and x_2 . Both the power and phase lag at the fundamental frequency, and the first two super harmonics (normalized frequency ($\times \pi$ rad/sample) $\approx 0.18, 0.35$, and 0.53) are predicted accurately using the stabilized ROM. The phase lag at these three frequencies in Fig. 10 as predicted by the CFD and the stabilized ROM is identified by red squares and blue triangles, respectively. As expected, the low-dimensional ROM is unable to reproduce the phase lag of low-amplitude frequencies or higher-order super harmonics.

Finally, a snapshot of the predicted velocity and pressure magnitudes at the final snapshot are illustrated in Fig. 11 and 12. Since the flow at this higher Reynolds number is chaotic, the low-dimensional model can not be expected to track the original snapshots exactly. However, the snapshot demonstrates that the stabilized ROMs faithfully reproduce the large features of the flow. The same cannot be said of the standard ROMs.

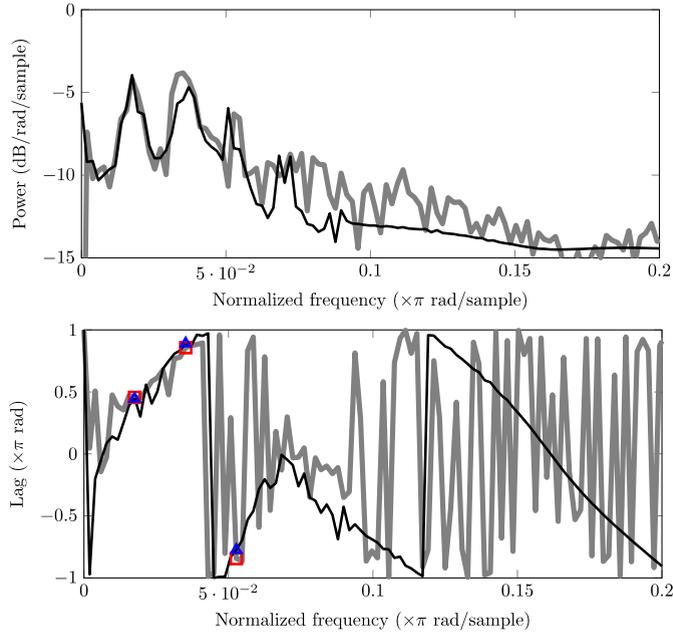


Fig. 10. CPSD of $p(\mathbf{x}_1, t)$ and $p(\mathbf{x}_2, t)$ where $\mathbf{x}_1 = (2, -0.5)$ and $\mathbf{x}_2 = (0, -0.5)$ of channel-driven cavity at $Re \approx 5500$. CFD (thick gray line), stabilized $n, p = 20$ ROM (black line). (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

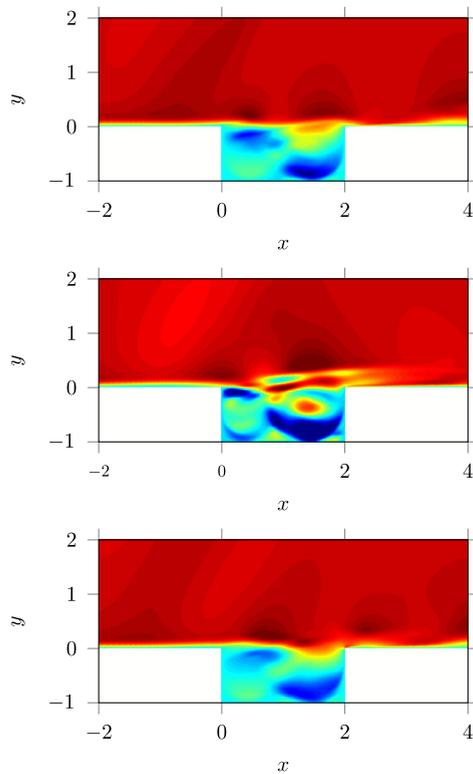


Fig. 11. Snapshot of channel-driven cavity $Re \approx 5500$; contours of u -velocity magnitude at the final snapshot. CFD (top), standard $n = 20$ ROM (middle), and stabilized $n, p = 20$ ROM (bottom).

4.3. Computational speed-up

For each problem considered, the speed-up factor delivered by its ROM for the online computations is reported in Table 1. All ROMs are solved in MATLAB using ODE45s. For more details on this algorithm, the reader is referred to Shampine and

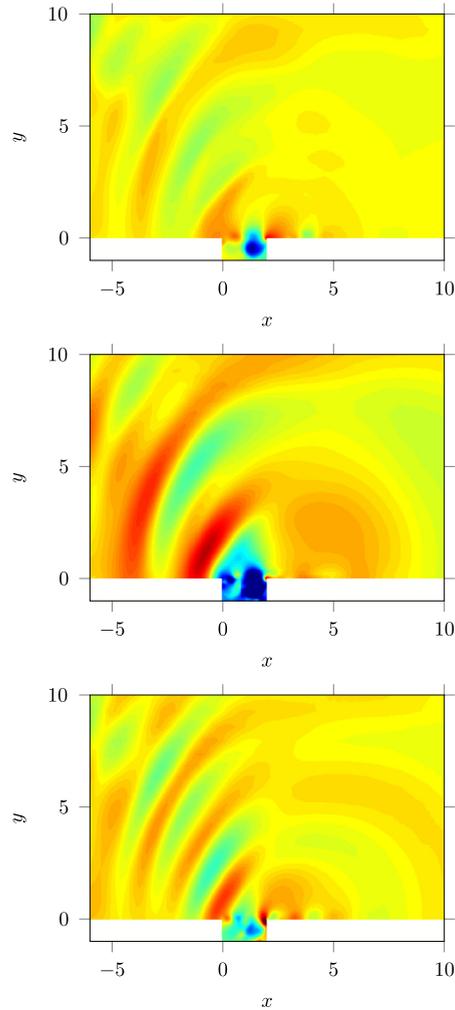


Fig. 12. Snapshot of channel-driven cavity $Re \approx 5500$; contours of pressure at the final snapshot. CFD (top), standard $n = 20$ ROM (middle), and stabilized $n, p = 20$ ROM (bottom).

Table 1

CPU times for off-line and on-line computations.

| Procedure | Numerical experiment | | |
|--|----------------------|-------------------|------------------------|
| | Airfoil | Cavity, low-Re | Cavity, moderate-Re |
| CFD # of DOF | 360,000 | 288,250 | 243,750 |
| Time-integration of CFD | 7.8 hrs | 72 hrs | 179 hrs |
| Basis construction (size $n + p$ ROM) | 0.16 hrs | 0.88 hrs | 3.44 hrs |
| Galerkin projection (size $n + p$ ROM) | 0.74 hrs | 5.44 hrs | 14.8 hrs |
| Stabilization | 28 sec | 14 sec | 170 sec |
| ROM # of DOF | 4 | 4 | 20 |
| Time-integration of ROM | 0.31 sec | 0.16 sec | 0.83 sec |
| Online computational speed-up | 9.1×10^4 | 1.6×10^6 | 7.8×10^5 |

Reichelt [43]. All online time-integration CPU times were measured using the `tic-toc` function on a single computational thread via the `-singleCompThread` start-up option. The CFD time-integration, basis construction and Galerkin projection times given in the table are reported in CPU-hours, calculated as the product of the number of processors used in the computation and the mean CPU time over all processors. The number of processors employed varied between 1 and 128. The online speed-up is calculated by evaluating the ratio between the time-integration of the CFD and the time-integration of the ROM. The reader can observe that the ROM online speedup is on the order of at least 10^4 for all three problems considered. Moreover, the stabilization step takes very little time (on the order of seconds/minutes).

5. Conclusions

In this paper, an approach for stabilizing and enhancing projection-based fluid ROMs of the compressible Navier–Stokes equations is developed. Unlike traditional approaches, no empirical turbulence modeling terms are required, and consistency between the ROM and the CFD model from which the ROM is derived is maintained. Mathematically, the approach is formulated as a trace minimization on the Stiefel manifold. The method is shown to yield both stable and accurate low-dimensional models of several representative compressible flow problems. In particular, the method is demonstrated on flows at higher Reynolds number where the dynamics are chaotic. Future work will include the extension of the proposed approach to problems with generic non-linearities, where the ROM involves some form of hyper-reduction (e.g., DEIM, gappy POD) following the procedure described in Remark 2, as well as to predictive applications with varying Reynolds number and geometry.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. NSF-CMMI-14-35474. The authors would like to thank Dr. Srinivasan Arunajatesan and Dr. Matthew Barone at Sandia National Laboratories for some useful discussions on model reduction for compressible flows. The authors would also like to thank Dr. Srinivasan Arunajatesan for generating the snapshots from which the channel-driven cavity reduced order models examined in the “Numerical experiments” section of this paper were constructed, and Dr. Jeffrey Fike for implementing the nonlinear compressible Navier–Stokes equations in specific volume code in the *Spirit* code.

Appendix A. Construction of Galerkin matrices for the compressible Navier–Stokes equations

In this section the construction of Galerkin matrices for the compressible Navier–Stokes equations (6) is outlined. Consider the standard orthonormal POD vectors u^i , v^i , ζ^i , and p^i , $i = 0, \dots, n$ where $i = 0$ identifies the constant mean flow. The following products are generated for $j, k = 0, \dots, n$:

$$\phi_1^{(j,k)} = u^j \odot \zeta_x^k + v^j \odot \zeta_y^k - u_x \odot \zeta^k - v_y^j \odot \zeta^k, \quad (\text{A.1a})$$

$$\phi_2^{(j,k)} = u^j \odot u_x^k + v^j \odot u_y^k + \zeta^j \odot p_x^k - \frac{1}{\text{Re}} \zeta^j \odot \left[\left(\frac{4}{3} u_x^k - \frac{2}{3} v_y^k \right)_x + (v_x^k + u_y^k)_y \right], \quad (\text{A.1b})$$

$$\phi_3^{(j,k)} = u^j \odot v_x^k + v^j \odot v_y^k + \zeta^j \odot p_y^k - \frac{1}{\text{Re}} \zeta^j \odot \left[\left(\frac{4}{3} v_y^k - \frac{2}{3} u_x^k \right)_y + (v_x^k + u_y^k)_x \right], \quad (\text{A.1c})$$

$$\begin{aligned} \phi_4^{(j,k)} &= u^j \odot p_x^k + v^j \odot p_y^k + \gamma p^j \odot (v_x^k + u_y^k) - \frac{\gamma}{\text{Re Pr}} \left[(p^j \odot \zeta_k)_{xx} + (p^j \odot \zeta_k)_{yy} \right] \\ &+ \frac{1-\gamma}{\text{Re}} \left[u_x^j \odot \left(\frac{4}{3} u_x^k - \frac{2}{3} v_y^k \right) + v_y^j \odot \left(\frac{4}{3} v_y^k - \frac{2}{3} u_x^k \right) + (u_y^j + v_x^j) \odot (u_y^k + v_x^k) \right], \end{aligned} \quad (\text{A.1d})$$

where subscripts identify partial spatial derivatives, and \odot is the Hadamard element-by-element product. A standard L^2 Galerkin projection for $i, j, k = 0, \dots, n$ is performed as follows:

$$b^{(i)}(j, k) = \sum_{m=0}^N \left[h \odot \left(\zeta^i \odot \phi_1^{(j,k)} + u^i \odot \phi_2^{(j,k)} + v^i \odot \phi_3^{(j,k)} + p^i \odot \phi_4^{(j,k)} \right) \right]_m, \quad (\text{A.2})$$

where h is a vector of element volumes. Finally, the standard Galerkin matrices for $i, j, k = 1, \dots, n$ with $a_0 = 1$, are given by

$$C_i = b^{(i)}(0, 0), \quad (\text{A.3a})$$

$$L_{i,j} = b^{(i)}(j, 0) + b^{(i)}(0, j), \quad (\text{A.3b})$$

$$Q_{j,k}^{(i)} = b^{(i)}(j, k). \quad (\text{A.3c})$$

Appendix B. Alternative formulations of the constrained optimization problem

In this section some alternative formulations of the proposed stabilization algorithm are laid out. In the current implementation, i.e. Eq. (14), the eddy-viscosity closure ansatz appears in the form of the trace of the modified linear operator. This alone does not place any guarantees on the eigenvalues of the linear operator. Indeed there are an infinite number of matrices, and thus eigenvalue distributions, with the same trace. It is straight forward to modify our approach to include a constrain for each individual eigenvalue

$$\tilde{\lambda}_i = \alpha_i, i = 1, \dots, n \quad (\text{B.1})$$

where $\tilde{\lambda}_i$ are eigenvalues $\tilde{\mathbf{L}} = \mathbf{X}^T \mathbf{L} \mathbf{X}$. Since the subspace rotation modifies all terms of the Galerkin system, the most general constraint would include the complete Galerkin system

$$\mathbf{X}^T \mathbf{C} - \Phi_C = \mathbf{0} \quad (\text{B.2a})$$

$$\mathbf{X}^T \mathbf{L} \mathbf{X} - \Phi_L = \mathbf{0} \quad (\text{B.2b})$$

$$\sum_{s,q,r=1}^{n+p} X_{si} Q_{qr}^{(s)} X_{qj} X_{rk} - \Phi_Q^{(s)} = \mathbf{0}, s = 1, \dots, n \quad (\text{B.2c})$$

where Φ_C , Φ_L , and $\Phi_Q^{(s)}$ are of appropriate dimension. Such fine grained control would in principle yield better performing ROMs. Unfortunately since α_i , Φ_C , Φ_L , and $\Phi_Q^{(s)}$ are not known *a priori*, the solution matching algorithm would be very high dimensional.

There are also opportunities for exploring alternative objective functions. For example, one may be interested in minimizing the rotation based on some weighted norm

$$\|\mathbf{W}(\mathbf{X} - \mathbf{I}_{n+p,n})\|_F. \quad (\text{B.3})$$

One natural choice for a weighted might be the POD eigenvalues, Σ but in principle arbitrary weightings can be implemented.

References

- [1] D. Amsallem, C. Farhat, Stabilization of projection-based reduced-order models, *Int. J. Numer. Methods Eng.* 91 (4) (2012) 358–377.
- [2] N. Aubry, P. Holmes, J.L. Lumley, E. Stone, The dynamics of coherent structures in the wall region of a turbulent boundary layer, *J. Fluid Mech.* 192 (1988) 115–173.
- [3] M. Balajewicz, E.H. Dowell, Stabilization of projection-based reduced order models of the Navier–Stokes, *Nonlinear Dyn.* 70 (2) (2012) 1619–1632.
- [4] M. Balajewicz, C. Farhat, Reduction of nonlinear embedded boundary models for problems with evolving interfaces, *J. Comput. Phys.* 274 (2014) 489–504.
- [5] M.J. Balajewicz, E.H. Dowell, B.R. Noack, Low-dimensional modelling of high-Reynolds-number shear flows incorporating constraints from the Navier–Stokes equation, *J. Fluid Mech.* 729 (2013) 285–308.
- [6] M.F. Barone, I. Kalashnikova, D.J. Segalman, H.K. Thornquist, Stable Galerkin reduced order models for linearized compressible flow, *J. Comput. Phys.* 228 (6) (2009) 1932–1946.
- [7] N. Boumal, B. Mishra, P.A. Absil, R. Sepulchre, Manopt, a Matlab toolbox for optimization on manifolds, *J. Mach. Learn. Res.* 15 (2014) 1455–1459.
- [8] T. Bui-Thanh, K. Willcox, O. Ghattas, B. van Bloemen Waanders, Goal-oriented, model-constrained optimization for reduction of large-scale systems, *J. Comput. Phys.* 224 (2) (2007) 880–896.
- [9] K. Carlberg, C. Farhat, J. Cortial, D. Amsallem, The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows, *J. Comput. Phys.* 242 (2013) 623–647.
- [10] K. Carlberg, R. Tuminaro, P. Boggs, Preserving Lagrangian structure in nonlinear model reduction with application to structural dynamics, *SIAM J. Sci. Comput.* 37 (2) (2015) B153–B184.
- [11] W. Cazemier, R.W.C.P. Verstappen, A.E.P. Veldman, Proper orthogonal decomposition and low-dimensional models for driven cavity flows, *Phys. Fluids* 10 (7) (1998) 1685–1699.
- [12] C. Eckart, G. Young, The approximation of one matrix by another of lower rank, *Psychometrika* 1 (3) (1936) 211–218.
- [13] B. Galletti, C.H. Bruneau, L. Zannetti, A. Iollo, Low-order modelling of laminar flow regimes past a confined square cylinder, *J. Fluid Mech.* 503 (2004) 161–170.
- [14] F. Génin, S. Menon, Dynamics of sonic jet injection into supersonic crossflow, *J. Turbul.* 11 (4) (2010) 1–30.
- [15] F. Génin, S. Menon, Studies of shock/turbulent shear layer interaction using large-eddy simulation, *Comput. Fluids* 39 (5) (2010) 800–819.
- [16] S. Gugercin, A.C. Antoulas, A survey of model reduction by balanced truncation and some new results, *Int. J. Control* 77 (8) (2004) 748–766.
- [17] M.A. Heroux, R.A. Bartlett, V.E. Howle, R.J. Hoekstra, J.J. Hu, T.G. Kolda, R.B. Lehoucq, K.R. Long, R.P. Pawlowski, E.T. Phipps, A.G. Salinger, H.K. Thornquist, R.S. Tuminaro, J.M. Willenbring, A. Williams, K.S. Stanley, An overview of the Trilinos project, *ACM Trans. Math. Softw.* 31 (3) (2005) 397–423.
- [18] P. Holmes, J.L. Lumley, G. Berkooz, C.W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edition, Cambridge University Press, 2012.
- [19] T. Iliescu, Z. Wang, Variational multiscale proper orthogonal decomposition: Navier–Stokes equations, *Numer. Methods Partial Differ. Equ.* 30 (2) (2014) 641–663.
- [20] A. Iollo, S. Lanteri, J. Désidéri, Stability properties of POD–Galerkin approximations for the compressible Navier–Stokes equations, *Theor. Comput. Fluid Dyn.* 13 (6) (2000) 377–396.
- [21] I. Kalashnikova, S. Arunajatesan, M.F. Barone, B.G. van Bloemen Waanders, J.A. Fike, Reduced order modeling for prediction and control of large-scale systems, Tech. rep., Sandia National Laboratories, Albuquerque, NM, 2014.
- [22] I. Kalashnikova, B. van Bloemen Waanders, S. Arunajatesan, M. Barone, Stabilization of projection-based reduced order models for linear time-invariant systems via optimization-based eigenvalue reassignment, *Comput. Methods Appl. Mech. Eng.* 272 (2014) 251–270.
- [23] B.S. Kirk, J.W. Peterson, R.H. Stogner, G.F. Carey, libMesh: a c++ library for parallel adaptive mesh refinement/coarsening simulations, *Eng. Comput.* 22 (3–4) (2006) 237–254.
- [24] J.L. Lumley, *Stochastic Tools in Turbulence*, Academic Press, New York, 1971.
- [25] L. Mirsky, Symmetric gauge functions and unitarily invariant norms, *Q. J. Math.* 11 (1960) 50–59.
- [26] B. Moore, Principal component analysis in linear systems: controllability, observability, and model reduction, *IEEE Trans. Autom. Control* 26 (1) (1981) 17–32.
- [27] B.R. Noack, M. Schlegel, B. Ahlborn, G. Mutschke, M. Morzyński, P. Comte, A finite-time thermodynamics of unsteady fluid flows, *J. Non-Equilib. Thermodyn.* 33 (2) (2008) 103–148.

- [28] J. Nocedal, Numerical Optimization, Springer, 1999.
- [29] J. Östth, B.R. Noack, S. Krajnović, D. Barros, J. Borée, On the need for a nonlinear subscale turbulence term in POD models as exemplified for a high-Reynolds-number flow over an Ahmed body, *J. Fluid Mech.* 747 (2014) 518–544.
- [30] S.B. Pope, Turbulent Flows, 6th edition, Cambridge University Press, 2000.
- [31] B. Protas, B.R. Noack, J. Östth, Optimal nonlinear eddy viscosity in Galerkin models of turbulent flows, arXiv preprint, arXiv:1406.1912, 2014.
- [32] T. Rapsák, On minimization on Stiefel manifolds, *Eur. J. Oper. Res.* 143 (2) (2002) 365–376.
- [33] D. Rempfer, H.F. Fasel, Dynamics of three-dimensional coherent structures in a flat-plate boundary layer, *J. Fluid Mech.* 275 (1994) 257–283.
- [34] C.W. Rowley, Model reduction for fluids, using balanced proper orthogonal decomposition, *Int. J. Bifurc. Chaos* 15 (3) (2005) 997–1013.
- [35] C.W. Rowley, T. Colonius, A.J. Basu, On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities, *J. Fluid Mech.* 455 (2002) 315–346.
- [36] C.W. Rowley, T. Colonius, R.M. Murray, Model reduction for compressible flows using pod and Galerkin projection, *Physica D* 189 (1) (2004) 115–129.
- [37] C.W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, D.S. Henningson, Spectral analysis of nonlinear flows, *J. Fluid Mech.* 641 (2009) 115–127.
- [38] G. Rozza, Reduced basis approximation and error bounds for potential flows in parametrized geometries, *Commun. Comput. Phys.* 9 (2011) 1–48.
- [39] J.M. Saniuk, I.B. Rhodes, A matrix inequality associated with bounds on solutions of algebraic Riccati and Lyapunov equations, *IEEE Trans. Autom. Control* 32 (8) (1987) 739–740.
- [40] V. Sankaran, S. Menon, LES of scalar mixing in supersonic mixing layers, *Proc. Combust. Inst.* 30 (2) (2005) 2835–2842.
- [41] P.J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* 656 (2010) 5–28.
- [42] G. Serre, P. Lafon, X. Gloerfelt, C. Bailly, Reliable reduced-order models for time-dependent linearized Euler equations, *J. Comput. Phys.* 231 (15) (2012) 5176–5194.
- [43] L.F. Shampine, M.W. Reichelt, The MATLAB ODE suite, *SIAM J. Sci. Comput.* 18 (1) (1997) 1–22.
- [44] L. Sirovich, Turbulence and the dynamics of coherent structures. Part III: dynamics and scaling, *Q. Appl. Math.* 45 (3) (1987) 583–590.
- [45] E. Stiefel, Richtungsfelder und fernparallelismus in n -dimensionalen mannigfaltigkeiten, *Comment. Math. Helv.* 8 (1) (1935) 305–353.
- [46] H. Tennekes, J.L. Lumley, A First Course in Turbulence, MIT Press, 1972.
- [47] K. Veroy, A.T. Patera, Certified real-time solution of the parametrized steady incompressible Navier–Stokes equations: rigorous reduced-basis a posteriori error bounds, *Int. J. Numer. Methods Fluids* 47 (8–9) (2005) 773–788.
- [48] Z. Wang, I. Akhtar, J. Borggaard, T. Iliescu, Proper orthogonal decomposition closure models for turbulent flows: a numerical comparison, *Comput. Methods Appl. Mech. Eng.* 237 (2012) 10–26.
- [49] K. Willcox, J. Peraire, Balanced model reduction via the proper orthogonal decomposition, *AIAA J.* 40 (11) (2002) 2323–2330.