SANDIA REPORT

SAND2018-10824
Unlimited Release
Printed September 2018

Captive Carry Reduced Order Modeling

Jeffrey Fike, Irina Tezaur, Kevin Carlberg, Matthew Barone

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

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories

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

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

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

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

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

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online ordering:	http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2018-10824
Unlimited Release
Printed September 2018

Captive Carry Reduced Order Modeling

Jeffrey Fike
Aerosciences Department
Sandia National Laboratories
P.O. Box 5800, MS 0825
Albuquerque, NM 87185–0825
jafike@sandia.gov

Irina Tezaur (Principal Investigator)
Quantitative Modeling & Analysis Department
Sandia National Laboratories
P.O. Box 969, MS 9159
Livermore, CA 94551–9159
ikalash@sandia.gov

Kevin Carlberg
Extreme Scale Data Science & Analytics Department
Sandia National Laboratories
P.O. Box 969, MS 9159
Livermore, CA 94551–9159
ktcarlb@sandia.gov

Matthew Barone
Aerosciences Department
Sandia National Laboratories
P.O. Box 5800, MS 0825
Albuquerque, NM 87185–0825
mbarone@sandia.gov

Abstract

This report summarizes fiscal year 2018 (FY18) progress towards developing, implementing (within the SPARC finite volume flow solver) and evaluating the viability of projection-based reduced order models (ROMs) for compressible captive-carriage flow problems of interest to Sandia National Laboratories. These ROMs are constructed using the Proper Orthogonal Decomposition (POD)/Least Squares Petrov-Galerkin (LSPG) method to model reduction, a promising approach based on residual minimization that can dramatically reduce the CPU-time requirement for simulations such as these, thereby enabling real-time and multi-query analyses. The key contributions of this report include: (1) the definition of relevant quantities of interest (QoIs) that can be used to evaluate ROM accuracy for this application, (2) a thorough numerical study of the accuracy of POD/LSPG ROMs implemented within SPARC (with and without various "tunings" aimed at improving predictive accuracy, e.g., time step and training window variation) with regard to these quantitative metrics, and (3) the formulation of a preconditioned LSPG ROM, shown to have superior accuracy for long-time predictive runs. We provide compelling evidence that the proposed approach can deliver viable predictive ROMs for the QoIs considered in the context of a two-dimensional (2D) viscous laminar cavity problem simulated using the SPARC code. Moreover, we demonstrate that the aforementioned ROMs can maintain an acceptable level of accuracy with substantial basis truncation through the application of a preconditioner to the discrete LSPG problem. We conclude this report by summarizing the key takeaways from our FY18 findings, and providing some directions for future work.

Acknowledgment

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The project described in this report is aimed at developing Proper Orthogonal Decomposition (POD)/Least Squares Petrov-Galerkin (LSPG) reduced order models (ROMs) of an aircraft weapons bay environment, implementing these models in the Sandia National Laboratories' in-house `SPARC` compressible flow solver [3], and investigating how well these models are able to represent the power spectral densities (PSDs) of the flow field for some representative captive carriage problems (Figure 1.1). The need for model order reduction (MOR) for this application comes from the fact that existing high-fidelity Computational Fluid Dynamics (CFD) codes such as `SPARC` require very fine meshes and long run-times to simulate the environment considered: a single relevant captive-carry simulation can take on the order of *weeks* to complete even when run on a parallel state-of-the-art supercomputer. Such CPU time requirements preclude real-time simulations and multi-query analyses, e.g., uncertainty quantification (UQ) and component design optimization.

Projection-based reduced order modeling is a promising tool for bridging the gap between high-fidelity simulations and real-time/multi-query analyses. Projection-based reduced order models are constructed from one or more high-fidelity simulations, and, if successful, retain the essential physics and dynamics of their corresponding full order models (FOMs) but have a much smaller computational cost. The general procedure for creating these models consists of three main steps: (1) calculation of a low-dimensional subspace through a data compression of a set of snapshots collected from a high-fidelity simulation, (2) projection of the governing partial differential equations (PDEs) or FOM onto this low-dimensional subspace, and (3) hyper-reduction, an approach for efficiently computing the projection of nonlinear terms in the governing PDEs or FOM by effectively evaluating these terms at a small number of carefully-selected points. The result of this procedure is a small set of time-dependent ordinary differential equations (ODEs) for the ROM modal amplitudes that approximately describe the flow dynamics of a FOM for some limited set of flow conditions.

The ROMs considered herein are based on the POD/LSPG approach, a minimal-residual-based method that creates the ROM from a fully-discrete high-fidelity CFD discretization. When combined with a hyper-reduction approach known as gappy POD [11], this method is equivalent to the Gauss-Newton with Approximated Tensors (GNAT) method of Carlberg *et al.* [8]. This methodology and its implementation within `SPARC` is described in detail in [28, 29].

During fiscal year 2017 (FY17), significant progress was made in improving the predictive capabilities of the POD/LSPG ROMs through the discovery of preconditioners that have the effect

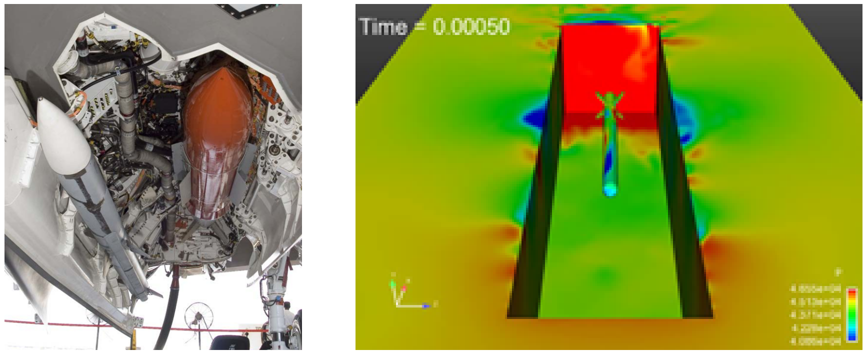

(a) Airplane weapons bay **(b)** Compressible cavity with object

Figure 1.1. Compressible captive-carry problem

of modifying the norm in which the residual is minimized [29]. However, further research was required to gain a complete picture of the viability of the proposed ROM methodology for the targeted application.

This report summarizes fiscal year 2018 (FY18) progress towards gaining a more complete understanding of the viability of the POD/LSPG ROMs implemented in the SPARC flow solver when applied to the captive-carriage problem. We focus our attention to the specific case of a two-dimensional (2D) viscous laminar cavity problem considered earlier [28, 29], targeting the following fundamental research question:

is it possible for a time-predictive LSPG/POD ROM to reproduce the pressure spectrum of its corresponding FOM to a sufficient accuracy?

A "time-predictive" ROM is a ROM that is used to perform a dynamic simulation at the same parameter values as those used to generate the high-fidelity snapshots from which the model was constructed, but run much longer in time than the underlying FOM.

Our first step towards addressing the above research question was the definition of relevant quantities of interest (QoIs) that can be used to evaluate ROM accuracy. For this work, we are primarily interested in pressure PSDs at a point or set of points inside the cavity. Other relevant QoIs included the root-mean-square (RMS) of pressure fluctuations and the global conservation of mass, energy and momentum in the ROM solution. Having defined quantitative metrics focused around these QoIs, we evaluated the performance of a "gold standard" ROM (i.e., an ideally-preconditioned projected solution increment full-basis ROM; see Chapter 2.4). Only after demonstrating that this "ideal" ROM can indeed be predictive for the QoIs considered did we evaluate the performance of more practical and less cost-intensive ROMs (truncated LSPG and preconditioned LSPG ROMs). This study included an investigation of the effect of various "tunings" believed to influence ROM accuracy, such as ROM time-step variation, ROM training window and snapshot sampling frequency, ROM preconditioner, and ROM basis truncation/selection.

The remainder of this report is organized as follows. In Chapter 2, the POD/LSPG approach to nonlinear model reduction is reviewed succinctly to keep this document self-contained. In Chapter

3, we study the viability of the POD/LSPG ROMs in `SPARC` for long-time predictive simulations in the context of a 2D viscous laminar cavity problem with respect to several relevant evaluation criteria defined therein. These results provide good evidence that the proposed model reduction approach is capable of producing a ROM that is predictive for the QoIs considered. It is shown that a critical ingredient to yielding sufficiently accurate time-predictive ROMs is the definition of a preconditioner applied to the discrete system created when constructing the ROM. This preconditioner has the effect of modifying the norm in which the residual is minimized. We summarize the main conclusions of this research effort and outline interesting avenues for future work in Chapter 4. Also included in this report is an Appendix that describes the data contained in a repository created during the course of this project containing relevant data sets, input files and instructions enabling an interested researcher to reproduce the results summarized in this report, and/or continue this research and development (R&D).

Chapter 2

Overview of Proper Orthogonal Decomposition (POD)/Least-Squares Petrov-Galerkin (LSPG) approach to nonlinear model reduction

To keep this report self-contained, we succinctly overview the POD/Least-Squares Petrov-Galerkin approach to model reduction in this chapter. This approach was implemented within the `SPARC` flow solver during fiscal years 2015-2016 (FY15-FY16).

Consider the following system of nonlinear equations

$$\boldsymbol{r}(\boldsymbol{w}) = \boldsymbol{0} \tag{2.1}$$

where $\boldsymbol{w} \in \mathbb{R}^N$ is the state vector and $\boldsymbol{r} : \mathbb{R}^N \to \mathbb{R}^N$ is the nonlinear residual operator. In this case, (2.1) are the compressible Navier-Stokes equations, discretized in space and time, so that (2.1) is the (discrete) full order model (FOM) for which we will build a ROM. Assuming (2.1) is solved using a (globalized) Newton's method, the sequence of solutions generated are

$$\boldsymbol{J}^{(k)} \delta \boldsymbol{w}^{(k)} = -\boldsymbol{r}^{(k)}, \quad k = 1, \ldots, K \tag{2.2}$$

$$\boldsymbol{w}^{(k)} = \boldsymbol{w}^{(k-1)} + \alpha_k \delta \boldsymbol{w}^{(k)}, \tag{2.3}$$

where $\boldsymbol{J}^{(k)} := \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{w}} \left(\boldsymbol{w}^{(k)} \right) \in \mathbb{R}^{N \times N}$, $\boldsymbol{r}^{(k)} := \boldsymbol{r} \left(\boldsymbol{w}^{(k)} \right) \in \mathbb{R}^N$, $\boldsymbol{w}^{(0)}$ is an initial guess for the solution, and $\alpha_k \in \mathbb{R}$ is the step length (often set to one).

2.1 Proper Orthogonal Decomposition (POD)

The first step in the LSPG/POD approach to model reduction is the calculation of a basis of reduced dimension $M << N$ (where N denotes the number of degrees of freedom in the full order model (2.1)) using the POD. The POD is a mathematical procedure that, given an ensemble of data and an inner product, denoted generically by $(\cdot, \cdot)$, constructs a basis for the ensemble that 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 M. The ensemble $\{\mathbf{x}^k : k = 1,\ldots,K\}$ is typically a set of K instantaneous snapshots of a numerical solution field, collected for K values of a parameter of interest, or at K different times. Mathematically, POD seeks an M-dimensional ($M << K$) subspace spanned by the set $\{\boldsymbol{\phi}_i\}$ such that the projection of the difference between the ensemble $\mathbf{x}^k$ and its projection onto the reduced subspace is minimized on average. It is a well-known result [4, 12, 19, 26] that the solution to the POD optimization problem reduces to the eigenvalue problem

$$\mathbf{R}\boldsymbol{\phi} = \lambda\boldsymbol{\phi}, \tag{2.4}$$

where $\mathbf{R}$ is a self-adjoint and positive semi-definite operator with its (i,j) entry given by $R_{ij} = \frac{1}{K}\left(\mathbf{x}^i, \mathbf{x}^j\right)$ for $1 \le i,j \le K$. It can be shown [12, 21] that the set of M eigenfunctions, or POD modes, $\{\boldsymbol{\phi}_i : i = 1,\ldots,M\}$ corresponding to the M largest eigenvalues of $\mathbf{R}$ is precisely the desired basis. This is the so-called "method of snapshots" for computing a POD basis [27].

Once a reduced basis is obtained, we approximate the solution to (2.1) by

$$\tilde{\boldsymbol{w}} = \bar{\boldsymbol{w}} + \boldsymbol{\Phi}_M\hat{\boldsymbol{w}} = \bar{\boldsymbol{w}} + \sum_{i=1}^{M}\boldsymbol{\phi}_i\hat{w}_i \tag{2.5}$$

with $\hat{\boldsymbol{w}} := [\hat{w}_1 \ \cdots \ \hat{w}_M]^T \in \mathbb{R}^M$ denoting the generalized coordinates, and $\bar{\boldsymbol{w}} \in \mathbb{R}^N$ denoting a reference solution, often taken to be the initial condition in the case of an unsteady simulation.

We then substitute the approximation (2.5) into (2.1). This yields

$$\boldsymbol{r}(\bar{\boldsymbol{w}} + \boldsymbol{\Phi}_M\hat{\boldsymbol{w}}) = \mathbf{0}, \tag{2.6}$$

which is a system of N equations in M unknowns $\hat{\boldsymbol{w}}$. As this is an over-determined system, it may not have a solution.

One question that often arises in assessing ROM performance is whether a given POD basis constructed using a set of training data is capable of accurately representing the flow after the end of the training interval. This can be assessed by projecting the snapshots from the FOM run onto the POD basis to create a set of projection coefficients, then using the basis and these projection coefficients to construct an approximation of the flow field. If $\boldsymbol{w}^{(k)}$ is the snapshot of the conservative variables at time instance k from the FOM, then the approximation to the flow field, $\tilde{\boldsymbol{w}}^{(k)}$, is given by

$$\tilde{\boldsymbol{w}}^{(k)} = \bar{\boldsymbol{w}} + \boldsymbol{\Phi}_M\left(\boldsymbol{\Phi}_M^T\left(\boldsymbol{w}^{(k)} - \bar{\boldsymbol{w}}\right)\right), \tag{2.7}$$

where $\boldsymbol{\Phi}_M$ is the POD basis and $\bar{\boldsymbol{w}}$ is some base flow state. This approximated flow state is the best that can be constructed using a given number of modes from the POD basis.

2.2 Least-Squares Petrov-Galerkin (LSPG) projection

In the LSPG approach to model reduction, solving the ROM for (2.1) amounts to solving the following least-squares optimization problem

$$\hat{\boldsymbol{w}}_{\text{PG}} = \arg\min_{\boldsymbol{y}\in\mathbb{R}^M} \|\boldsymbol{r}(\bar{\boldsymbol{w}} + \boldsymbol{\Phi}_M\boldsymbol{y})\|_2^2. \tag{2.8}$$

Here, the approximate solution is $\tilde{\boldsymbol{w}}_{\mathrm{PG}} := \bar{\boldsymbol{w}} + \boldsymbol{\Phi}_M \hat{\boldsymbol{w}}_{\mathrm{PG}}$. The name "LSPG" ROM comes from the observation that solving (2.8) amounts to solving a nonlinear least-squares problem. The two most popular approaches for this are the Gauss–Newton approach and the Levenberg–Marquardt (trust-region) method. Following the work of Carlberg *et al.* [8], we adopt the Gauss–Newton approach[1]. This approach implies solving a sequence of linear least-squares problems of the form

$$\delta \hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = \arg \min_{\boldsymbol{y} \in \mathbb{R}^M} \|\boldsymbol{J}^{(k)} \boldsymbol{\Phi}_M \boldsymbol{y} + \boldsymbol{r}^{(k)}\|_2^2, \quad k = 1, \ldots, K_{\mathrm{PG}} \tag{2.9}$$

$$\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = \hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k-1)} + \alpha_k \delta \hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} \tag{2.10}$$

$$\tilde{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = \bar{\boldsymbol{w}} + \boldsymbol{\Phi}_M \hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k-1)}, \tag{2.11}$$

where K_{PG} is the number of Gauss-Newton iterations. It can be shown that the approximation upon convergence is $\tilde{\boldsymbol{w}}_{\mathrm{PG}} = \tilde{\boldsymbol{w}}_{\mathrm{PG}}^{(K_{\mathrm{PG}})}$ and $\hat{\boldsymbol{w}}_{\mathrm{PG}} = \hat{\boldsymbol{w}}_{\mathrm{PG}}^{(K)}$.[2] Note that the normal-equations form of (2.9) is

$$\boldsymbol{\Phi}_M^T \boldsymbol{J}^{(k)T} \boldsymbol{J}^{(k)} \boldsymbol{\Phi}_M \delta \hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = -\boldsymbol{\Phi}_M^T \boldsymbol{J}^{(k)T} \boldsymbol{r}^{(k)}, \quad k = 1, \ldots, K_{\mathrm{PG}}, \tag{2.12}$$

which can be interpreted as a Petrov–Galerkin process of the Newton iteration with trial basis (in matrix form) $\boldsymbol{\Phi}_M$ and test basis $\boldsymbol{J}^{(k)} \boldsymbol{\Phi}_M$.

The simplest implementation of the Gauss–Newton method for solving (2.8). For details of this approach and its implementation in `SPARC`, the reader is referred to Chapters 2 and 3 of [28], respectively.

2.3 Projected solution increment approach

At the end of Section 2.1, we discussed projecting the snapshots of the flow field created by the FOM onto the POD basis in order to compute projection coefficients. These projection coefficients can then be used to construct an approximation to the flow state at each time instance using (2.7). This provides information on the accuracy to which the flow state could be represented using the POD basis. However, it does not take into account the fact that the true FOM flow states may never be reached by the ROM, and this accuracy may never be realized in practice.

In order to determine the upper limit on the ROM accuracy, we need to take the solution increment that is computed by the FOM at each time step, project it onto the POD basis, and then construct an approximated solution increment. If $\Delta \boldsymbol{w}^{(k)}$ is the solution increment computed by the FOM at time step k, then

$$\Delta \tilde{\boldsymbol{w}}^{(k)} = \boldsymbol{\Phi}_M \left(\boldsymbol{\Phi}_M^T \boldsymbol{\Phi}_M \right)^{-1} \boldsymbol{\Phi}_M^T \Delta \boldsymbol{w}^{(k)} \tag{2.13}$$

is the approximated solution increment. This approximated solution increment is then used to update the flow state, producing a flow state that is realizable by an ideal preconditioned LSPG

[1] The LSPG approach is the basis for the Gauss–Newton with Approximated Tensors (GNAT) method of Carlberg *et al.* [8].

[2] In the event of an unsteady simulation, the initial guess for the generalized coordinates is taken to be the generalized coordinates at the previous time step.

ROM. This is what we call the projected solution increment approach and it is widely used in this project as a means of determining an upper limit on the ROM accuracy.

2.4 Preconditioned LSPG ROMs

One can modify the LSPG projection formulation outlined in Section 2.2 through the addition of a preconditioner.

A preconditioned LSPG ROM can be formed by inserting a preconditioner, $\boldsymbol{M}^{(k)}$, into (2.9) to create

$$\delta\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = \arg\min_{\boldsymbol{y}\in\mathbb{R}^M} \left\| \boldsymbol{M}^{(k)} \left(\boldsymbol{J}^{(k)} \boldsymbol{\Phi}_M \boldsymbol{y} + \boldsymbol{r}^{(k)} \right) \right\|_2^2, \quad k = 1, \ldots, K_{\mathrm{PG}} , \tag{2.14}$$

where the preconditioner, $\boldsymbol{M}^{(k)}$, is often an approximation to the inverse of $\boldsymbol{J}^{(k)}$. The insertion of a preconditioner has the effect of altering the norm we are minimizing in. The corresponding normal-equations form of (2.14) is

$$\boldsymbol{\Phi}_M^T \boldsymbol{J}^{(k)T} \boldsymbol{M}^{(k)T} \boldsymbol{M}^{(k)} \boldsymbol{J}^{(k)} \boldsymbol{\Phi}_M \delta\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = -\boldsymbol{\Phi}_M^T \boldsymbol{J}^{(k)T} \boldsymbol{M}^{(k)T} \boldsymbol{M}^{(k)} \boldsymbol{r}^{(k)}, \quad k = 1, \ldots, K_{\mathrm{PG}} . \tag{2.15}$$

2.4.1 Ideal preconditioned LSPG ROMs

The projected solution increment given by (2.13) is equivalent to what would be produced by a preconditioned LSPG ROM using an exact preconditioner. Using an ideal preconditioner

$$\boldsymbol{M}^{(k)} = \left(\boldsymbol{J}^{(k)} \right)^{-1} \tag{2.16}$$

reduces the normal-equations to

$$\boldsymbol{\Phi}_M^T \boldsymbol{\Phi}_M \delta\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = -\boldsymbol{\Phi}_M^T \left(\boldsymbol{J}^{(k)} \right)^{-1} \boldsymbol{r}^{(k)}, \quad k = 1, \ldots, K_{\mathrm{PG}} . \tag{2.17}$$

Substituting in the true FOM solution increment,

$$\Delta\boldsymbol{w}^{(k)} = \left(\boldsymbol{J}^{(k)} \right)^{-1} \boldsymbol{r}^{(k)} , \tag{2.18}$$

produces

$$\boldsymbol{\Phi}_M^T \boldsymbol{\Phi}_M \delta\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = -\boldsymbol{\Phi}_M^T \Delta\boldsymbol{w}^{(k)}, \quad k = 1, \ldots, K_{\mathrm{PG}} . \tag{2.19}$$

The right hand side is the projection of the true FOM solution increment. Solving for $\delta\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)}$ and constructing the approximate solution increment produces (2.13).

This use of this ideal preconditioner will lead to an optimal ℓ^2 projection of the solution increment $\Delta\boldsymbol{w}$ onto the basis $\boldsymbol{\Phi}_M$,

$$\delta\hat{\boldsymbol{w}}_{\mathrm{PG}}^{(k)} = \arg\min_{\mathbf{x}} \left\| \boldsymbol{\Phi}_M \mathbf{x} - \Delta\boldsymbol{w}^{(k)} \right\|_2^2 . \tag{2.20}$$

2.5 Hyper-reduction

The LSPG projection and preconditioned LSPG projection approaches described in Sections 2.2 and 2.4 respectively are not efficient for nonlinear problems. This is because the solution of the ROM system requires algebraic operations that scale with the dimension of the original full-order model N. This problem can be circumvented through the use of hyper-reduction. A number of hyper-reduction approaches have been proposed, including the discrete empirical interpolation method (DEIM) [9], “best points” interpolation [23, 24], collocation [20] and gappy POD [11]. Implementation of the latter two approaches has been started within `SPARC`, but is not complete at the present time, and hence a detailed discussion of this methodology goes beyond the scope of this report. The basic idea behind these approaches is to compute the residual at some small number of points q with $q << N$, encapsulated in a “sampling matrix” $\boldsymbol{Z}$. This set of q points is typically referred to as the “sample mesh”. The “sample mesh” is computed by solving an optimization problem offline; see Section 2.3 of [28] and [8] for details. The LSPG projection approach combined with gappy POD hyper-reduction is equivalent to the GNAT method [8].

2.6 Implementation in `SPARC`

The POD/LSPG method described above has been implemented in the `SPARC` compressible flow solver. For details of this implementation, the reader is referred to Chapter 3 of [28].

Several improvements were made to the implementation this FY, including

- Improvements to the projected solution increment approach,
- Implementation of preconditioned LSPG ROMs,
- Improved restart capabilities throughout `SPARC`,
- Modification of POD calculation routine to handle larger numbers of snapshots.

Of particular importance is the implementation of preconditioned LSPG ROMs. There are several preconditioners that can be utilized. There is a standalone diagonal (Jacobi) preconditioner. This was simple to implement and is likely the easiest to adapt for hyper-reduction. `SPARC` can make use of the preconditioners defined in the `Ifpack2` package of `Trilinos` [25]. Preconditioned LSPG ROMs have been run using Jacobi and ILUT-1 preconditioners. There is also an exact preconditioner, formed by computing the inverse of the Jacobian. This produces results equivalent to the projected solution increment implementation, however it is computationally expensive and is only feasible for small test cases.

Chapter 3

Investigation of ROM viability on 2D viscous laminar cavity

3.1 2D viscous laminar cavity problem description

The LSPG/POD ROMs implemented in SPARC are evaluated on a test case involving a 2D viscous laminar flow around an open cavity geometry, described in this subsection.

The computational domain for this test case is composed of two rectangular regions: a cavity region, $\Omega_{cavity} = [0.0m, 0.0917136m] \times [0.0m, -0.0458568m]$, and an outer flow region, $\Omega_{flow} = [-4.58568m, 4.58568m] \times [0.0m, 6.87852m]$. This 2D domain is made into a 3D domain for SPARC by making the mesh 1 cell thick and imposing symmetry (inviscid, slip wall) boundary conditions on the faces parallel to the plane. The nominal mesh used for these simulations contains 104,500 hexahedral cells and is shown in Figure 3.1. A preliminary mesh convergence study for this geometry was performed during FY17[1]. The large extent of the outer flow domain is intended to minimize the effects of any pressure waves reflecting off the boundaries. Reflections off the boundaries were seen to have a significant impact on the accuracy and stability of ROMs in previous work [14]. Previous runs using other codes, such as Sigma CFD, utilized a sponge region to eliminate reflected pressure waves. As SPARC does not currently have a sponge boundary condition implemented, the outer domain was made very large and the cells stretched in the far field in order to minimize any pressure wave reflections.

The flow conditions for this test case are chosen to produce approximately Mach 0.6 and a Reynolds number of approximately 3000. The exact parameters specified in the SPARC input file are given in Table 3.1. The nominal full-order model runs used BDF2 time stepping with a fixed time step corresponding to a CFL number of under 50.0. The time step was chosen in order to limit the number of snapshots that are produced during the ROM training interval to a reasonable number.

Viscous, no-slip boundary conditions are imposed on the left, right, and bottom surfaces of the cavity domain, Ω_{cavity}. Far-field boundary conditions are imposed on the left, right, and top surfaces of the outer flow domain, Ω_{flow}. A combination of inviscid slip wall and viscous no-slip wall boundary conditions are imposed on the lower boundary of Ω_{flow}. The regions immediately before and after the cavity have no-slip walls, but the regions closer to the inflow and outflow surfaces

[1] See Chapter 4 of [29] for some SPARC mesh convergence studies in the context of a simplified box geometry.

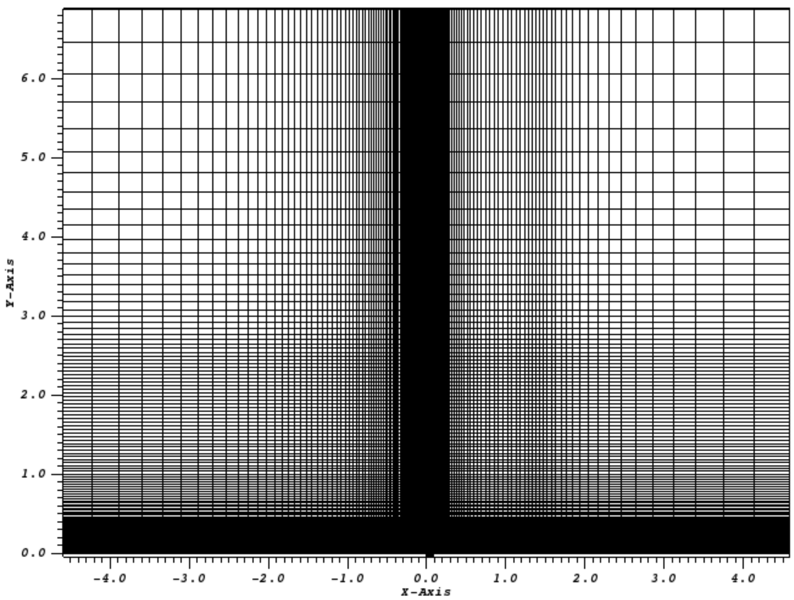

Figure 3.1. The computational mesh used for 2D viscous laminar cavity simulations in SPARC.

Parameter	Dimensional Value	Non-Dimensional Value
Free-stream Velocity, u_1	$208.7816\,m/s$	0.601247049354
Temperature, T	$300\,K$	1.0
Density, ρ	$2.9026155498083859 \times 10^{-4}\,kg/m^3$	1.4
Pressure, p	$25.0\,Pa$	1.0
Viscosity, μ	$8.46 \times 10^{-7}\,kg/(m\,s)$	$1.17508590713 \times 10^{-5}$
Specific Gas Constant, R	$287.097384766765\,m^2/(s^2\,K)$	0.714285714286
γ	1.4	1.4
Prandtl Number	0.72	0.72
Time Step	$2.0 \times 10^{-6}\,s$	$6.9449521698 \times 10^{-4}$

Table 3.1. Parameters used for the 2D viscous laminar cavity test case in SPARC.

have inviscid slip walls. This strategy allows constant far-field inflow conditions to be specified without having to impose a boundary layer profile. The boundary layer begins to grow at the upstream transition from inviscid wall to no-slip wall. The extent of the no-slip wall was chosen to allow the boundary layer to attain the desired thickness at the beginning of the cavity.

Reduced-order models are created and run in `SPARC` in several stages. First, the full-order model in `SPARC` was run for 270,000 time steps. During this initial run restart files were saved every 5000 time steps. Once the run was completed, we identified regions where we wanted to generate training data. The full-order model is then restarted at the desired times, and snapshots of the flow field are saved every time step. This is done in 1000 time step chunks to aid in forming arbitrary sets of training data. The sets of 1000 snapshots are then formed into larger sets of various sizes. POD is then performed on these larger set of snapshots to create a modal basis. This modal basis is then used when reduced-order models are run in `SPARC`.

3.2 Evaluation criteria

Reduced-order models typically trade reduced accuracy for reduced computational cost. It is therefore important to be able to assess the accuracy of the ROMs and determine if the reduced accuracy is acceptable. This can be done by running both the full-order model and reduced-order models for the same flow and comparing the results.

There are several ways to compare the results from the full-order model and reduced-order models. No method of comparison is perfect; each method has its strengths and weaknesses. A good comparison between the full-order model and the reduced-order model will likely require a combination of methods.

One way to compare two simulations is to use a flow visualization package to visualize the flow field as the solution evolves in time. This enables us to see if the overall general flow features are accurately captured by the reduced-order model. Flow field animations also allow us to see if there are discrepancies in frequency or phase between the FOM and the ROM. Discrepancies in the amplitude are a little more difficult to discern. Flow field visualizations can provide a qualitative ("eyeball norm") comparison between simulations but do not provide a quantitative measure of the error.

The error between two simulations can be computed by taking the norm of the difference between the two data sets at every point in the domain. The relative error between two flow solutions can be computed as

$$\mathscr{E} = \frac{\|q_2 - q_1\|_2}{\|q_1\|_2}, \tag{3.1}$$

where q is some quantity of interest. The subscript 1 is for the full-order model or true solution, and the subscript 2 is for the reduced-order model solution or whatever solution we are comparing to the true solution. Using a constant time step for all simulations guarantees that this difference can be performed at every time step. The quantity of interest, q in (3.1), could be a scalar quantity such as the pressure, or it could be a vector quantity such as the conservative solution variables. We often compute the total error between the conservative variables as

$$\mathscr{E}_{total} = \frac{\left(\sum_{cells} (\rho_1 - \rho_2)^2 + ((\rho u)_1 - (\rho u)_2)^2 + ((\rho v)_1 - (\rho v)_2)^2 + ((\rho w)_1 - (\rho w)_2)^2 + ((\rho e)_1 - (\rho e)_2)^2\right)^{\frac{1}{2}}}{\left(\sum_{cells} (\rho_1)^2 + (\rho u)_1^2 + (\rho v)_1^2 + (\rho w)_1^2 + (\rho e)_1^2\right)^{\frac{1}{2}}} . \tag{3.2}$$

Here, ρ denotes the fluid density; u, ,v and w are the x–, y– and z– components of the fluid velocity respectively; e is the fluid energy. The sum in (3.2) is over the cells in the mesh. As written, this expression is sensitive to the relative magnitudes of the conservative variables and will be dominated by the variable with the largest magnitude. This sensitivity is reduced when the variables are non-dimensional and of the same magnitude.

It is often not feasible to record snapshots of the entire flow field at every time step. Therefore, flow visualizations or flow field errors have lower temporal resolution than the simulations. An alternative is to output data every time step but at only a few specified locations. The use of probes or traces of this type increases the temporal resolution of the data, but there is much less spatial information. It is therefore important to place these probes in locations where the variations in the quantities of interest are desired. For the cavity flow, we have data at many points in the cavity but we focus on a point midway up the downstream wall of the cavity. Other locations of interest may be on the floor of the cavity or in the shear layer at the top of the cavity.

In `SPARC`, probes output the conservative variables at each time step. For the cavity flow, the primary quantity of interest are the pressure fluctuations, so the conservative variables are post-processed to produce a pressure time history. Plotting the pressure time histories from both the FOM and the ROM enables a comparison of the amplitude, phase, and frequency of oscillations.

During this FY, we also explored other quantities of interest besides the pressure time histories. We briefly explored the use of averaged or integrated quantities such as the pressure or drag along portions of the bottom wall of the cavity. It was hypothesized that these types of quantities might not show as much variation in time, and thus might prove to be a less volatile means of comparison. However, the variations of these averaged or integrated quantities showed the same spectral content as the point pressure histories.

We also explored quantities aimed at assessing whether the ROM simulations satisfied the physical constraints of the full-order model. The full-order model solver the RANS equations, meaning that mass, momentum, and energy are conserved within each control volume. On the other hand, reduced-order models may not exactly satisfy these physical constraints unless they are specifically enforced [7, 22]. Figure 3.2 shows a comparison of the violation of the mass conservation constraint for the full-order model and a standard LSPG ROM using 100 modes from a 2000 snapshot training interval. While the amount of constraint violation is higher for the ROM than for the FOM, the magnitude of the violation is relatively small. For the LSPG ROM the amount of constraint violation jumps at the end of the 2000 time step training interval, but still remains fairly small.

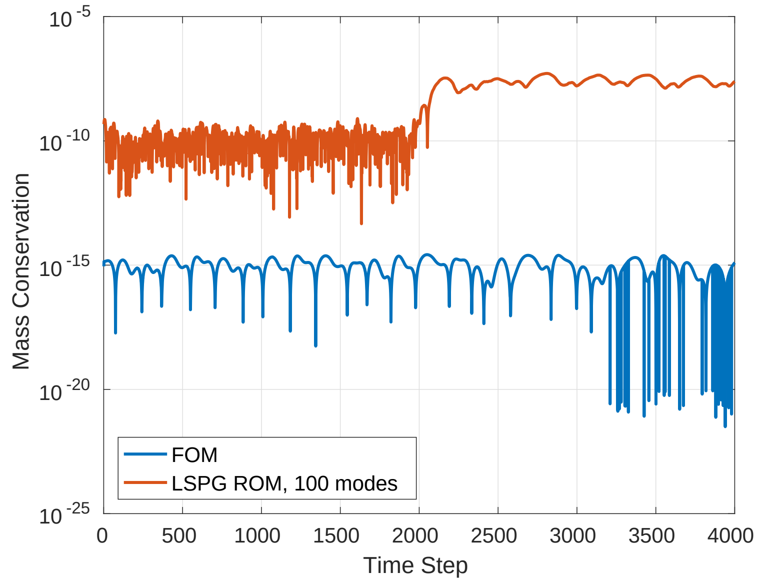

Figure 3.2. A comparison of the violation of the mass conservation constraint for the full-order model and a standard LSPG ROM using 100 modes from a 2000 snapshot training interval.

The pressure time history, and quantities computed from it, remains the primary QoI for this study. It is not anticipated that the ROM will be able to exactly reproduce the pressure time history from the FOM. Small differences in the flow state can produce substantially different time histories as the simulations are run for long time periods. However, the statistics of these time histories, such as the root-mean-square (RMS) value or the power spectral density (PSD), should be similar.

The power spectral density is computed by taking the magnitude of the Fourier transform of the time signal. This provides information on the power of the oscillations at each frequency in the time history. The area under the PSD curve is the RMS value of the oscillations, and provides information on the average amplitude of the oscillations.

Examining the PSD curves produced by the FOM and ROMs provides a good means of comparison. However, it is difficult to define a quantitative measure of error for PSD curves. Taking the difference is simple enough, but small shifts in frequency can lead to large differences. The different parts of the frequency range are also likely to have different error characteristics. The dominant low frequency peaks are likely of more importance than the high frequency content where the magnitudes are very small. Several quantitative measures that act on a frequency range of interest were explored, however none were completely satisfactory.

A difficulty with comparing PSD curves is that a very long time signal is required to produce a converged PSD. Given the finite time signal generated during the simulations, the resulting PSDs are expected to have a large statistical sampling error. In standard PSD calculations, the full time history is broken up into several overlapping segments whose length is determined by the desired frequency resolution of the PSD. Essentially, a PSD is computed for each segment of the time history and then averaged to produce the final PSD. The PSDs for each segment can be plotted to show how the frequency content of the signal changes over time. Viewing these waterfall plots or spectrograms is useful to discern if the training interval contains enough frequency content to represent the long time signal, or if frequencies gain importance later in the time history.

Averaging the individual segment PSDs helps remove noise from the final PSD, but very many segments are required for the mean PSD to fully converge. The mean PSD can be thought of as the most likely realization of the true PSD given the finite time signal. However, curves within the spread of the PSDs of the individual segments are also likely realizations. An effective way of presenting this spread of values is to plot the mean PSD as a solid line, and plot the range of the individual segment PSDs as a shaded band. This allows for a meaningful comparison between two simulations by showing how well the shaded bands and mean match between two simulations, rather than artificially constraining the accuracy to be within some fixed dB range. Visualizing the PSD data in this way not only allows for a comparison between two mean PSD curves, but shows how close the mean PSD of one simulation is to the spread of the segment PSDs for the other. However, this is still a qualitative comparison and effectively quantifying this comparison has proved elusive.

None of the comparison methods discussed above are perfect. A combination of the methods should be used to provide a sufficient comparison between the ROM and the FOM. Using probes we can obtain pressure time histories at one or more points. PSDs can then be computed at these points. Plots comparing the pressure time histories and PSDs from different simulations provide a good means of comparing the two simulations. However, these comparisons are localized to the probe locations. Flow visualizations of the entire flow field provide a more global comparison.

Plotting time histories or PSD curves and visualizing the entire flow field provides a qualitative comparison between two simulations. A quantitative measure of the error can be computed using (3.1) on either the full flow field, the pressure time histories, or the PSDs. However, there can be situations where the solutions are qualitatively similar but have a large computed error. Small shifts in frequency or phase can produce large errors between flow fields or time histories. Small shifts in the frequencies of the peaks on the PSD can also produce large error values.

Quantitative measures of the error are useful in assessing how much of an effect changes to the ROMs, such as the number of modes, have on the accuracy of the simulations. It is important to emphasize that these quantitative measures should be combined with qualitative comparisons to provide a complete assessment of the difference between two simulations. The primary means of comparisons used in this report are pressure time histories and the PSD plots showing the mean and range of support. In terms of quantitative comparisons, the primary metric is the value of the root-mean-square (RMS) overall-sound-pressure level (OASPL).

3.3 Accuracy assessments of LSPG ROMs in `SPARC`

The accuracy of a reduced-order model can be affected by many factors. Some of these are standard concerns that arise in most reduced-order modeling situations. These standard choices include:

- The type of projection employed (Galerkin, LSPG, etc.).
- The number of modes used for the ROM.
- The time step used for the ROM.
- Hyper-reduction choices.
- The type of reduced basis employed (POD, PGD, DMD, etc.).

Here we study the first three of these five choices, restricting our attention to ROMs with no hyper-reduction and employing a POD reduced basis.

Some other choices are specific to our chosen application and the goal of training over a small time interval but then predicting for a much longer time interval. The more application-specific choices include:

- The choice of how many snapshots to collect for training.
- The choice of when to collect snapshots for training.
- Whether the training data are composed of contiguous or disjoint sets of snapshots.

In this section, we examine the effect of these choices on the accuracy of the predictions made by the resulting reduced-order model. Since the primary goal of this project is to determine whether reduced-order models produce sufficiently accurate predictions of long time statistical behavior for the cavity flow case, the focus has been on accuracy not computational cost. Thus, the effect of hyper-reduction has not yet been explored, and is to be studied in future work.

3.3.1 Projection type

The following types of reduced-order models have been implemented in `SPARC`.

- Galerkin.
- Least-Squares Petrov-Galerkin (LSPG) (see Section 2.2).
- Projected Solution Increment or Ideally-Preconditioned LSPG (see Section 2.3 and Section 2.4.1).
- Preconditioned LSPG (see Section 2.4 and 2.6).

Although attention is restricted here to reduced bases computed using the POD approach, the ROM implementation within `SPARC` is in no way limited to this type of reduced basis.

The mathematical details of the various projection approaches considered here and itemized above are discussed in Section 2 of this report. Galerkin ROMs were briefly explored during the efforts in

FY16, but are not expected to work well for this application, as demonstrated in some of our past work [15, 4, 13, 16, 17, 18, 1]. This is largely due to the fact that Galerkin projection in the L^2 inner product is not guaranteed to yield an energy-stable formulation for the equations of compressible flow. The current work therefore focuses on the various flavors of LSPG ROMs.

Standard LSPG ROMs are expected to work well as a nonlinear model reduction method [8, 6]; however, they can behave poorly for the particular application we consider here [29]. This led to the introduction of preconditioned LSPG ROMs and the development of the projected solution increment approach. The projected solution increment or ideally-preconditioned LSPG ROM is expected to be the most accurate approach considered here for reproducing the full-order model behavior, as discussed in Sections 2.3 and 2.4. Preconditioned LSPG ROMs, using one of several varieties of preconditioners, are expected to approach the accuracy of the projected solution increment but be more practical to use.

Figure 3.3 shows a comparison between the pressure time histories for the full-order model and the projected solution increment approach using 327 modes. These results are for a basis constructed from snapshots collected during training interval 5, as will be described in Section 3.3.4. The blue box indicates the extent of the training data used to create the POD basis. In this case, snapshots of the flow field were collected at every time step for a total of 5000 snapshots. These 5000 snapshots were used to create a POD basis containing 5000 modes. The projected solution increment run started at the end of the training interval and ran for 40,000 time steps. The projected solution increment run used 327 modes out of the 5000 available. This level of truncation corresponds to retaining 99.9% of the energy captured by the snapshots, as determined from the magnitudes of the singular values computed when forming the basis. The choice to use 327 modes was made out of consideration for the long runtimes required for the LSPG ROMs without hyper-reduction.

The pressure time histories shown in Figure 3.3 do not match on a point-by-point comparison, but we did not expect it to do so. The more important comparison is whether the statistics computed from these time histories show fairly good agreement with the underlying FOM statistics. Figure 3.4 depicts a comparison of the power spectral densities computed from the relevant time histories. The solid lines show the mean PSDs for the FOM and projected solution increment data. These curves demonstrate that the time history created by the projected solution increment approach has nearly the same spectral content as the full-order model. Generally speaking, the peaks of the projected solution increment data match the peaks of the FOM over a range of frequencies of interest, up to a non-dimensional frequency of 20. The agreement breaks down at high frequencies, but the magnitude indicates that these frequencies are unlikely to be of great significance.

Figure 3.5 shows a comparison between the pressure time histories for the full-order model and the standard LSPG ROM using 327 modes. For this case, the standard LSPG ROM produced a time history with much larger magnitude oscillations than the FOM or projected solution increment approach. The magnitude of the oscillations is generally growing, indicating a possible instability in the ROM. The LSPG ROM eventually produces a non-physical quantity, such as a negative pressure or density, and the solution terminates. It is possible that using more modes may improve the behavior of the standard LSPG ROM, but without hyper-reduction the runtime is prohibitively high.

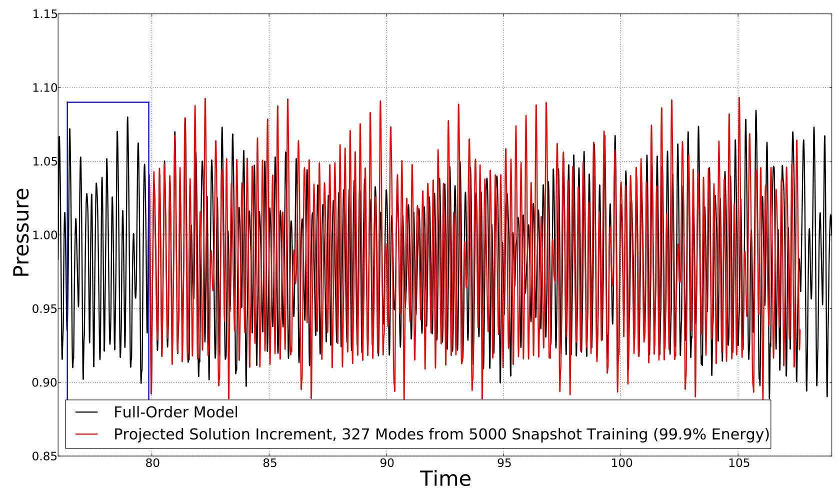

Figure 3.3. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using 327 modes. The blue box indicates the extent of the training data used to create the POD basis.

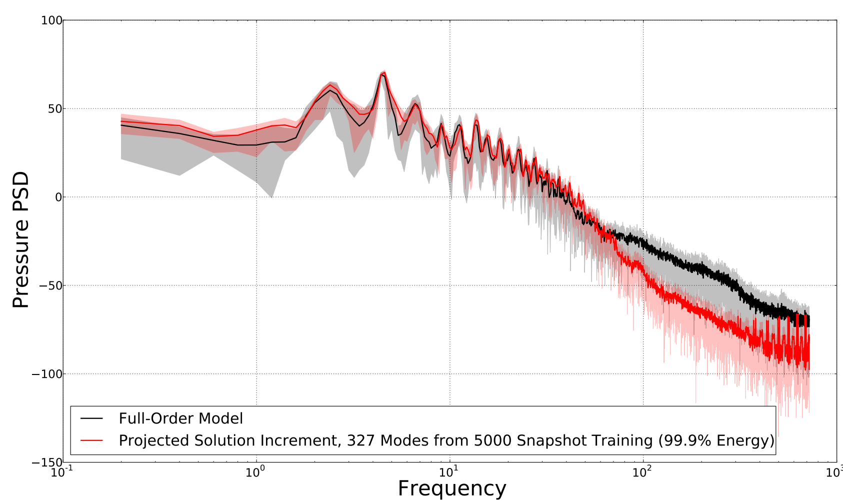

Figure 3.4. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using 327 modes. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

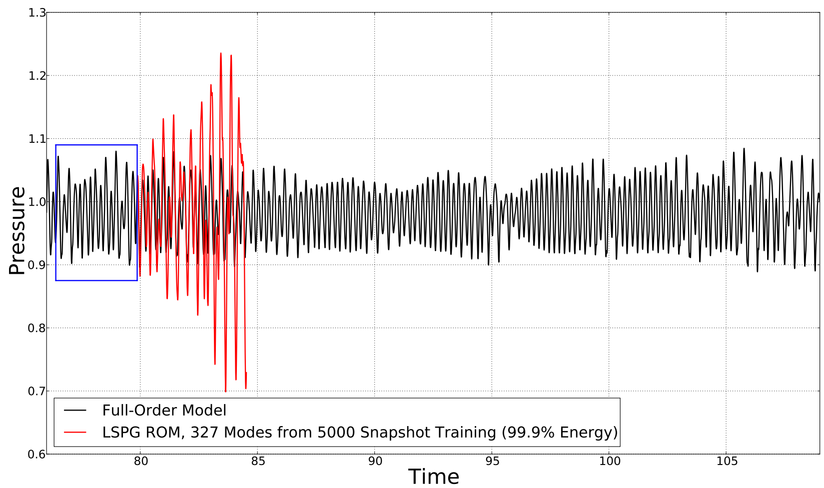

Figure 3.5. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and a standard Least-Squares Petrov-Galerkin ROM using 327 modes. The blue box indicates the extent of the training data used to create the POD basis.

While the standard LSPG ROM shown in Figure 3.5 did not run successfully, the projected solution increment approach shows much more promise, as demonstrated by Figures 3.3 and 3.4. Given that the projected solution increment approach is equivalent to an ideally-preconditioned LSPG ROM, it is reasonable to suppose that a preconditioned LSPG ROM might be more successful than the standard LSPG ROM without a preconditioner. Figure 3.6 shows a comparison between the pressure time histories for the full-order model and a preconditioned LSPG ROM using 327 modes. For this case, the preconditioner was the hand-coded diagonal preconditioner that is equivalent to a Jacobi preconditioner. The action of the preconditioner allows the LSPG ROM to run successfully. Figure 3.7 shows a comparison between the PSDs the full-order model and the preconditioned LSPG ROM. Overall, the mean PSD for the preconditioned LSPG ROM matches the PSD for the full-order model fairly well (especially at higher frequencies), although generally not as well as the projected solution increment approach (in particular at low frequencies).

Figure 3.8 shows a comparison between the pressure time histories for the full-order model and a preconditioned LSPG ROM using 327 modes, this time using the ILUT-1 preconditioner from The Trilinos package `Ifpack2` [25]. In this case, the preconditioned LSPG ROM shows promise, but the solution has not yet finished despite the fact that we only used 327 modes to try to keep the LSPG runtime reasonable.

In addition to the qualitative comparisons provided by the plots of pressure time histories and power spectral densities, it is useful to consider a more quantitative comparison. Table 3.2 shows the values of the root mean square overall sound pressure level (RMS OASPL) for the pressure time histories. The value for the projected solution increment approach is slightly better than the preconditioned LSPG ROM, and both are within 2 dB of the full-order model.

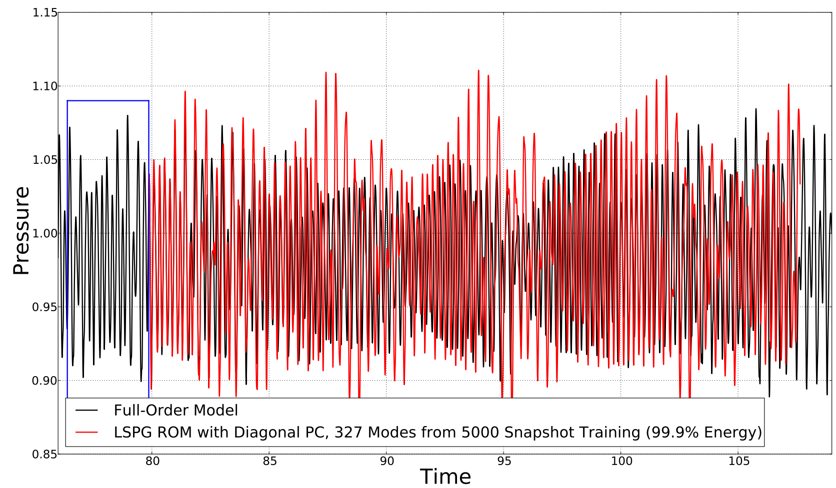

Figure 3.6. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and a preconditioned LSPG ROM using the diagonal(Jacobi) preconditioner and 327 modes. The blue box indicates the extent of the training data used to create the POD basis.

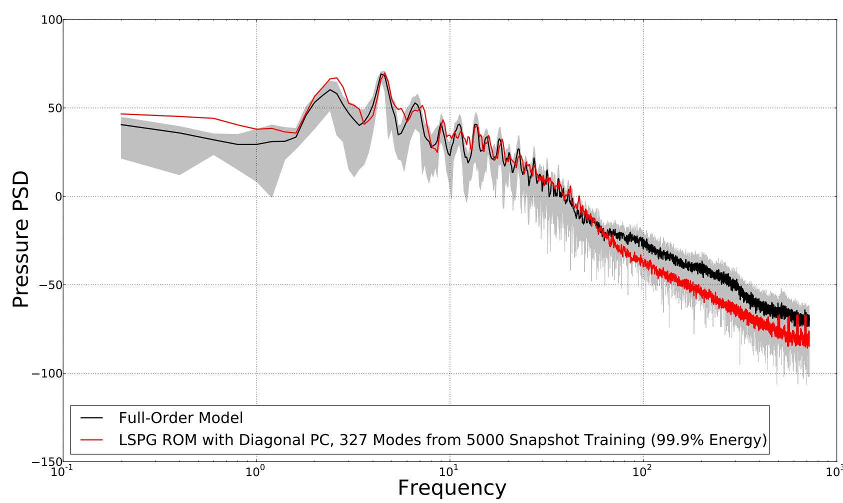

Figure 3.7. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and a preconditioned LSPG ROM using the diagonal(Jacobi) preconditioner and 327 modes. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

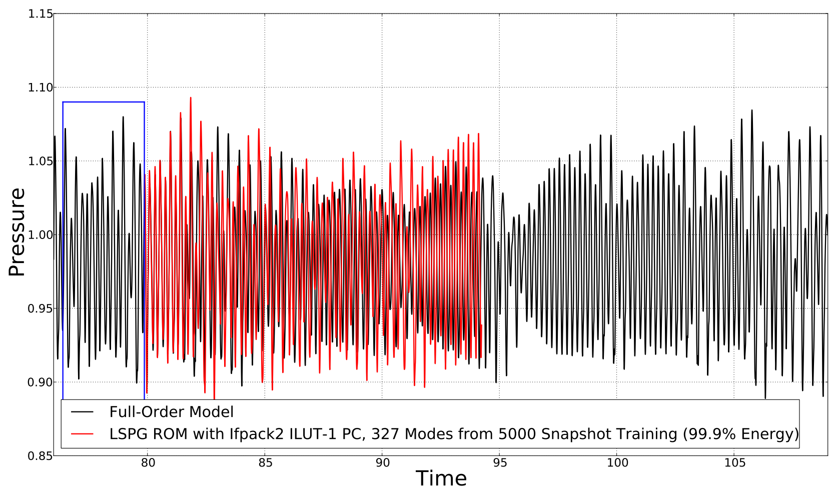

Figure 3.8. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and a preconditioned LSPG ROM using the ILUT-1 preconditioner from Ifpack2 and 327 modes. The blue box indicates the extent of the training data used to create the POD basis. Note that this simulation has not finished running.

Method	RMS OASPL in dB	Difference from FOM in dB
Full-Order Model	66.176	–
Proj. Sol. Incr.	67.552	1.376
LSPG w/o PC	N/A	N/A
LSPG w/ diagonal PC	68.033	1.857
LSPG w/ ILUT-1 PC	N/A	N/A

Table 3.2. A comparison of the root mean square overall sound pressure level for the full-order model and ROMs using 327 modes.

The following conclusions can be made based on this section. The projected solution increment is expected to provide an upper bound on the accuracy of an LSPG ROM. Standard LSPG ROMs without preconditioning can experience difficulties for the cavity flow case we are considering. Preconditioned LSPG ROMs can perform better than standard LSPG ROMs without preconditioning. The choice of preconditioner is likely to have an impact on the results, with more accurate preconditioners improving the accuracy but at a higher computational cost.

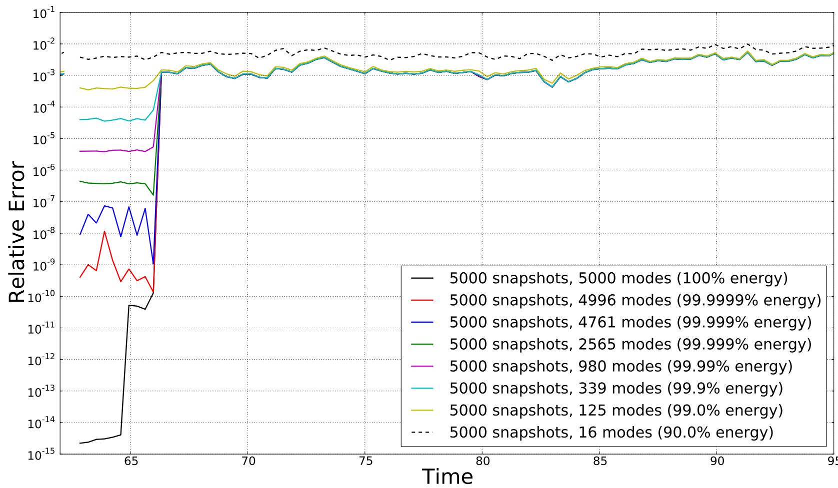

Figure 3.9. The projection error as the number of modes is varied. The sharp increase in error occurs at the end of the training interval used to construct the basis.

3.3.2 Number of modes

The number of modes used by a ROM is an important factor affecting both accuracy and computational cost. Using a large fraction of the total modes is likely to result in a more accurate ROM but will be computationally more expensive. Using a smaller fraction of the total modes is likely to reduce the accuracy, but also reduce the computational cost.

One method of assessing the impact of varying the number of modes (in our case POD modes) is to look at the projection error (see discussion at the end of Section 2.1). Figure 3.9 shows the projection error as the number of modes used for the projection is varied. These results are for a basis constructed from snapshots collected during training interval 4, as will be described in Section 3.3.4. The sharp increase in error occurs at the end of the training interval used to construct the basis. Within the training interval the projection error increases as fewer modes are used for the projection. However, outside of the training interval the projection error appears insensitive to the number of modes, provided that over 339 modes (99.9% of the energy) are used.

Figure 3.10 shows the PSDs for the full-order model and the projected solution increment approach as the number of modes used for the projection is varied. The PSDs for the 980 mode case (99.99% of the energy) and higher are largely indistinguishable from one another except at high frequencies. The 339 mode case (99.9% of the energy) exhibits some differences from the higher mode cases, but overall it yields a PSD that matches the corresponding FOM PSD fairly well. The discrepancy is larger for the 125 mode case (99.0% of the energy). The 16 mode case (90.0% of the energy) failed to run successfully, eventually producing a non-physical quantity before terminating. This behavior for ROMs with small basis sizes has been observed in other work, e.g., [1]. These results reinforce the conclusion reached based on the projection error, that when considering data outside

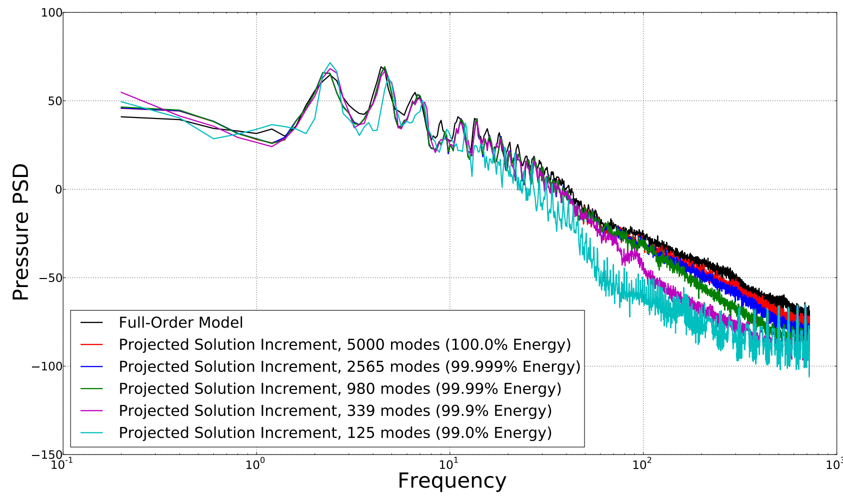

Figure 3.10. The mean PSDs for the full-order model and projected solution increment approach as the number of modes is varied.

Modes	% Energy	RMS OASPL in dB	Difference from FOM in dB
FOM	–	66.72782	–
5000	100.0	66.71432	0.01350
2565	99.999	66.71764	0.01018
980	99.99	66.66956	0.05824
339	99.9	66.78537	0.05755
125	99.0	67.17222	0.44440

Table 3.3. A comparison of the root mean square overall sound pressure level for the full-order model and projected solution increment runs as the number of modes is varied.

the training interval the results are largely insensitive to the number of modes used, provided that we are above the 99.9% threshold.

Table 3.3 provides a more quantitative comparison by showing the discrepancy in the RMS OASPL. This quantitative measure also reinforces the conclusion that the overall results are largely insensitive to the number of modes used by the ROM. For cases with enough modes to capture over 99.9% of the energy the difference between the projected solution increment results and the FOM is less than 0.06 dB.

The overall conclusion from this section is that results are fairly insensitive to the number of modes used, as long as enough modes are used to capture at least 99.9% of the energy.

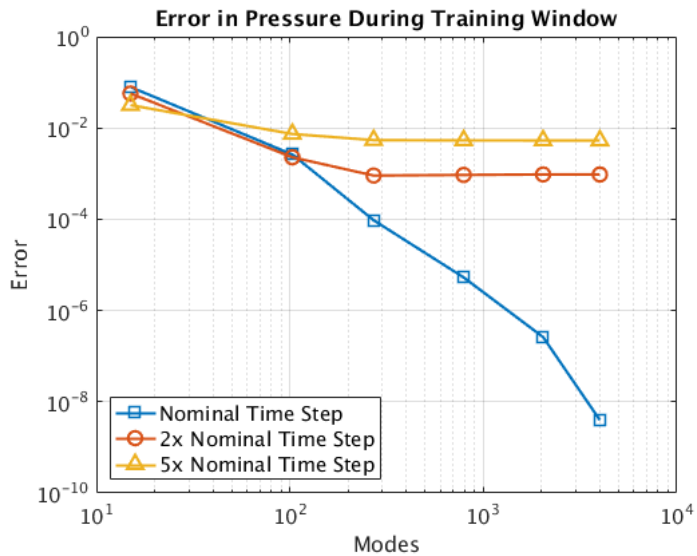

Figure 3.11. The pressure time history error during the training interval for the projected solution increment approach as the time step and number of modes are varied.

3.3.3 Varying the time step used by the ROM

The time step used by the ROM is expected to have an impact on the accuracy of LSPG ROMs. Error bounds were derived that showed a dependence on the time step were derived in [6]. It was found that there is an optimal time step which minimizes the error of the LSPG ROMs, and that error increases if the time step is too small. This was verified computationally for a cavity flow using the LSPG ROM implementation in the AERO-F code. The optimal time step was found to be larger than that used for the full-order model.

A study of the effect of the ROM time step was begun during FY17 but was incomplete. The results of this study are summarized in [29]. This effort continued during the first part of FY18.

Figure 3.11 shows the error in the pressure time history during the training window as the number of modes and time step are varied. The data in this plot show that there is a benefit to using a larger time step when the number of modes is small, around 100 or less for this case. However, when more modes are used increasing the time step has a negative effect on the error. In the earlier observations about the effect of the number of modes on the PSDs of the projected solution increment runs, it was observed that the accuracy of the results was insensitive to the number of modes as long as there were enough modes to capture 99.9% of the energy or more. This brings into question the utility of increasing the time step for this cavity flow case.

The above results only looked at the behavior of the ROMs when used to reproduce the training interval. The projected solution increment cases were also run beyond the training interval to produce pressure time histories over a long time duration (40,000 time steps as before). The pressure time histories were then used to compute the PSDs of the signal. Given the large number

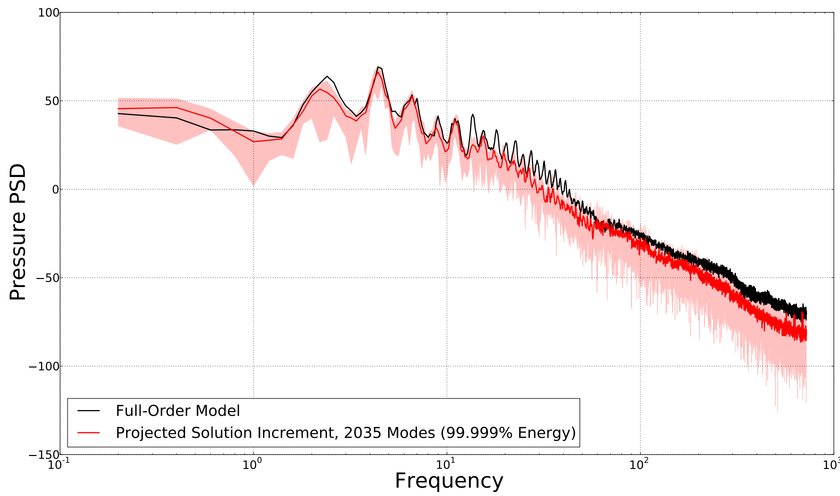

Figure 3.12. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and the projected solution increment approach using 2035 modes and run at the nominal time step. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

of combinations of time step and number of modes, only a few representative plots are shown here.

Figure 3.12 shows a comparison of the PSDs for the full-order model and the projected solution increment approach using 2035 modes. For this case, the projected solution increment case was run at the same nominal time step as that used by the full-order model. Figure 3.13 shows the PSD for the projected solution increment approach when run at twice the nominal time step. This PSD is similar to that for the nominal time step case with a few exceptions. The discrepancies between the magnitudes of the peaks are generally larger for this case than for the nominal time step case. The cutoff frequency is also reduced due to the larger time step, so there is not as much high frequency information.

Figure 3.14 shows the PSD for the projected solution increment approach when run at 5 times the nominal time step. The discrepancies between the PSD for this case and the full-order model are larger than those for the smaller time step cases. These qualitative comparisons would seem to indicate that increasing the time step over the nominal case results in a decrease in accuracy of the resulting PSDs.

We have explored several ways of computing a quantitative measure of the error between two PSDs. However, none of the methods we looked at are completely satisfactory. Nevertheless, these quantitative measures applied to this time step study generally reinforce the qualitative observations and indicate that the accuracy of the PSDs decreases as the time step is increased.

Table 3.4 shows a comparison of the RMS OASPL for the full-order model and projected solution

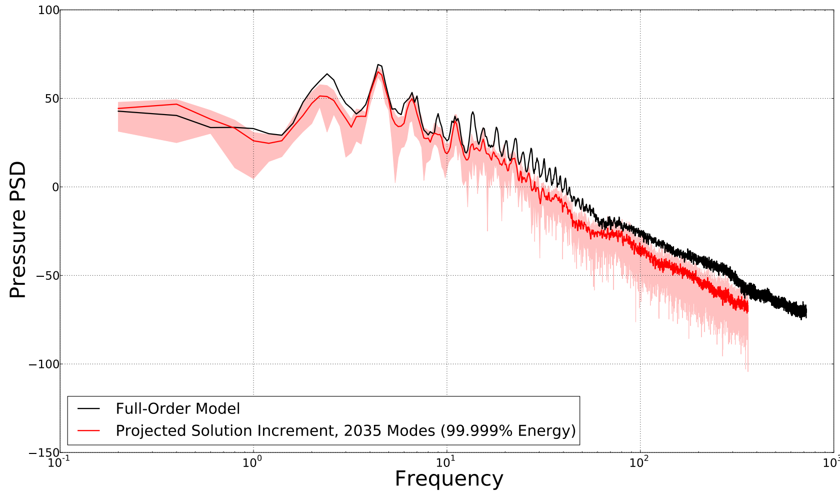

Figure 3.13. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and the projected solution increment approach using 2035 modes and run at twice the nominal time step. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

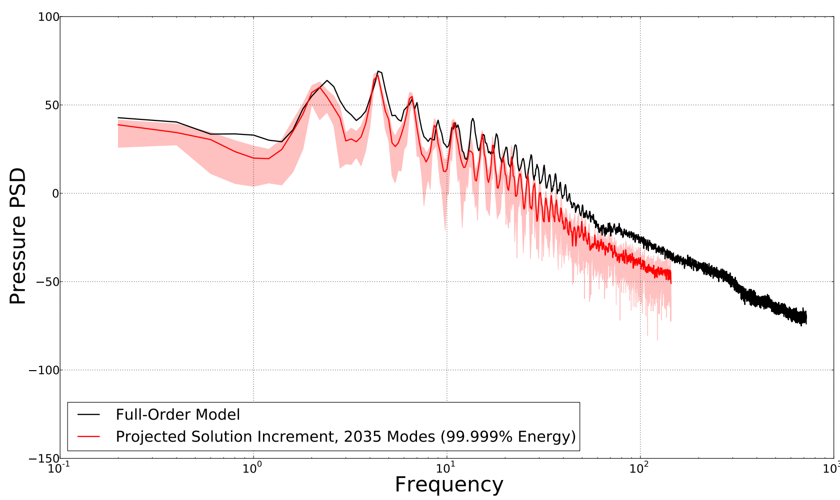

Figure 3.14. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and the projected solution increment approach using 2035 modes and run at 5 times the nominal time step. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

Modes	% Energy	1× dt	2× dt	5× dt
FOM	–	66.67 dB	–	–
4000	100.0	62.71 dB	62.28 dB	63.37 dB
2035	99.999	62.96 dB	61.82 dB	63.26 dB
791	99.99	62.07 dB	62.54 dB	62.87 dB
271	99.9	62.16 dB	62.68 dB	63.82 dB
103	99.0	60.23 dB	60.31 dB	60.35 dB
15	90.0	76.84 dB	74.67 dB	72.38 dB

Table 3.4. A comparison of the root mean square overall sound pressure level for the full-order model and projected solution increment runs as the time step and number of modes are varied.

increment runs as the time step and number of modes are varied. In general, the RMS values for the projected solution increment runs are 3 to 4 dB lower than the full-order model. As the time step is increased the RMS values also generally increase, making them closer to the full-order model. These data seem to suggest that there may be a benefit to increasing the time step, in contrast to the PSD comparisons.

To summarize, our studies suggest that there is mixed evidence about the effect of increasing the time step used by the ROMs. During the training interval, our results clearly show that increasing the time step lowers the error if the number of modes is very low; otherwise, it has a negative impact on the error. When considering the time history beyond the training interval the results are not as clear. A comparison of the PSDs seems to indicate that increasing the time step has a negative impact on the accuracy of the ROM results, but a comparison of the RMS values shows the opposite trend. Given the mixed evidence, there is no specific conclusion that can be reached based on the available data at the present time.

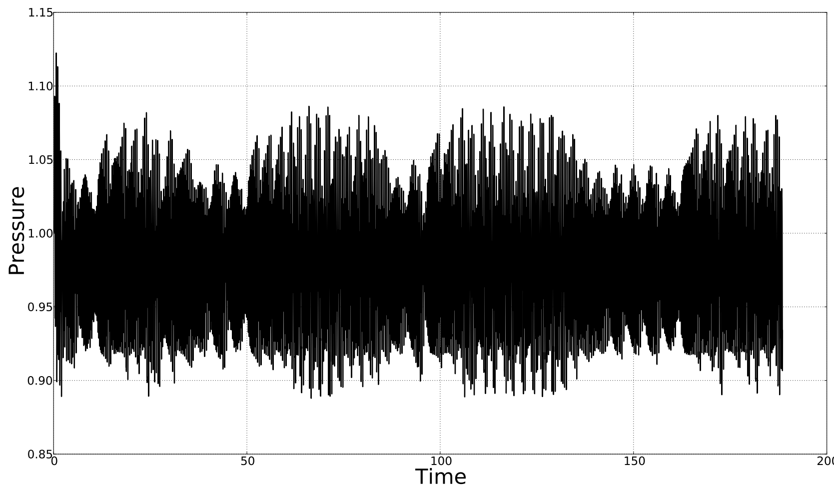

Figure 3.15. The pressure time history from the full-order model for a point midway up the downstream wall of the cavity.

3.3.4 Snapshot collection strategies

In this project, the overall goal was to train a ROM on a short training interval but then use the ROM to run for a much longer time. The cavity flow case that is under consideration exhibits quasi-periodic behavior. In other words, the physics of the flow follow a more or less periodic behavior, but it is not truly periodic and there are variations in the flow on large time scales. Figure 3.15 shows the pressure time history at a point midway up the downstream wall of the cavity when the full-order model is run for 270,000 time steps. This plot shows that the character of the pressure oscillations changes over fairly long time scales.

As discussed, the goal of this project is to train over a short training interval, but then use the ROM over a much longer time interval. The decision of when to collect snapshots for training is important. Figure 3.16 shows the pressure time history and five separate training intervals. This time history is of shorter duration than that in Figure 3.15 to emphasize the range over which the ROMs in this work are trained and tested. The blue boxes in Figure 3.16 indicate five separate training intervals where snapshots were collected and a POD basis was computed. Throughout this report we refer to these as training intervals 1 through 5, with training interval 1 being the earliest and 5 being the latest. As shown, training interval 1 consists of 4000 time steps, while training intervals 2 through 5 consist of 8000 time steps. These are standard durations; however, throughout this project, the actual durations were varied.

In general, it is desirable to start the training interval as early in the simulation as possible in order to minimize the computational effort of running the full-order model. Training interval 1 was chosen based on this consideration. As shown in Figure 3.16, training interval 1 consists of 4000 time steps. However, training intervals of shorter and longer duration were also utilized. During FY16 and FY17, training intervals of 2000 time steps were also used. During FY18 the primary

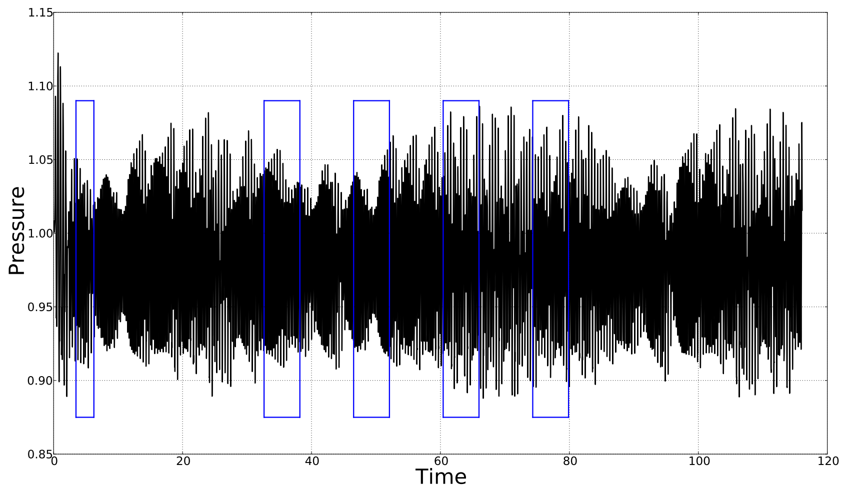

Figure 3.16. The pressure time history from the full-order model for a point midway up the downstream wall of the cavity. The blue boxes indicate five separate training intervals where snapshots were collected and a POD basis was computed.

duration was 4000 time steps, although 5000 to 7000 time step cases were also run. Results from training interval 1 were shown in Section 3.3.3.

While having the training interval early in the simulation is desirable, this can cause some problems. One concern with an early training interval is that the flow may still be undergoing some initial transient behavior. This is indeed an issue with training interval 1. At the end of training interval 1, the initial pressure disturbance created by the cavity has not fully propagated upstream to the inlet. This means that the training snapshots and resulting basis vectors cannot represent pressure wave propagation beyond this time. This would have a negative impact on the overall accuracy of any longer time simulations, in particular far-field results. However, it is uncertain how much of an impact this has on the flow near the cavity.

A possible remedy to this issue is to use a basis refinement technique [5]. A version of this technique was implemented for the projected solution increment approach, where the solution increment was modified as usual in regions near the cavity but the far-field regions were left to evolve as in the standard full-order model. Unfortunately, the runs using this approach did not complete successfully, eventually producing non-physical quantities and terminating. This may not indicate a problem with the approach, but may have to do the details of the implementation. Further effort could enable this approach to work, however it was decided that it was simpler to just change to a later training interval.

Training intervals 2 through 5 were chosen to try to ensure that the full-order model simulation was beyond any initial transient behavior. As shown in Figure 3.16, training intervals 2 through 5 consist of 8000 time steps. However, durations as short as 2000 snapshots were also considered. We were limited to a maximum of around 8000 snapshots due to (currently unresolved) issues in our basis calculation code that occurred when attempting to operate on 9000 or more snapshots, likely due to excessive memory requirements. Results for 5000 snapshots from training interval 4 were shown in Section 3.3.2, while results for 5000 snapshots from training interval 5 were shown in Section 3.3.1. For these later training intervals, the point at which the ROMs become predictive was fixed at the end of the training interval and all snapshots were collected prior to that point.

In general, as might be expected, collecting as many snapshots as possible leads to the most accurate ROMs. Typically, we used 8000 snapshots as the standard, but as noted there were situations where we used 5000 snapshots. It is therefore useful to know what effect the number of snapshots has on the resulting ROM accuracy. Projected solution increment runs were carried out using untruncated bases constructed from 5000, 6000, 7000, and 8000 snapshots from training interval 5. The PSDs for these runs (not shown) are all fairly similar. It is possible that more significant effects would be observed for smaller numbers of snapshots. Given more time, it would be interesting to determine what is the minimum number of snapshots required to produce a sufficiently accurate ROM.

Training intervals 2 through 5 capture different flow behavior as shown in the pressure time history of Figure 3.16. The question then becomes how well do the resulting ROMs perform during the relevant long time simulations. Figure 3.17 shows the time history for the full-order model and the projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 2. Figure 3.18 shows the same time histories, but is zoomed in to better show the details. The time history of the projected solution increment approach seems to enter a periodic state with oscillations of limited amplitude. Figure 3.19 shows a comparison of the PSDs for these time histories. As can be observed, the PSD for the projected solution increment approach fails to capture a significant number of the peaks. The reason for this poor performance is unclear. The current hypothesis is that during the training interval the magnitudes of the oscillations are generally decreasing, so it is possible that the resulting basis is not capable of supporting the growing oscillations observed for the full-order model. In other words, the thinking is that the flow features contained in the training interval are different than those we are trying to predict in the longer time simulations.

Figures 3.20, 3.21 and 3.22 repeat the above comparisons for training interval 3. Figures 3.23, 3.24 and 3.25 repeat the same comparisons for training interval 4. And Figures 3.26, 3.27 and 3.28 repeat the comparisons for training interval 5. For these three cases, the projected solution increment approach produces quite accurate results. The reason for this excellent agreement likely has to do with the training interval containing features similar to those that we are trying to predict in the longer time simulations.

The discussion in this section has served to illustrate the impact of choosing when to collect the training data. However, we have the benefit of running the full-order model for the entire time range of interest. We can then select a training interval that is likely to capture the flow features that can represent the behavior we are trying to predict in the time range of interest. In practice, the full-order model would only be run to collect training data, which would not provide information on later behavior. This makes it difficult to know *a priori* if the ROM will be successful.

Figure 3.29 shows the projection error for the untruncated bases constructed for training intervals

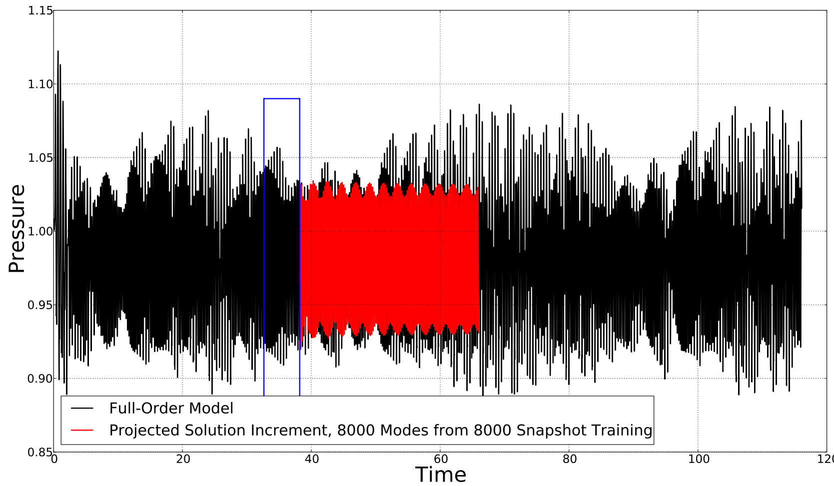

Figure 3.17. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 2. The blue box indicates the extent of the training data used to create the POD basis.

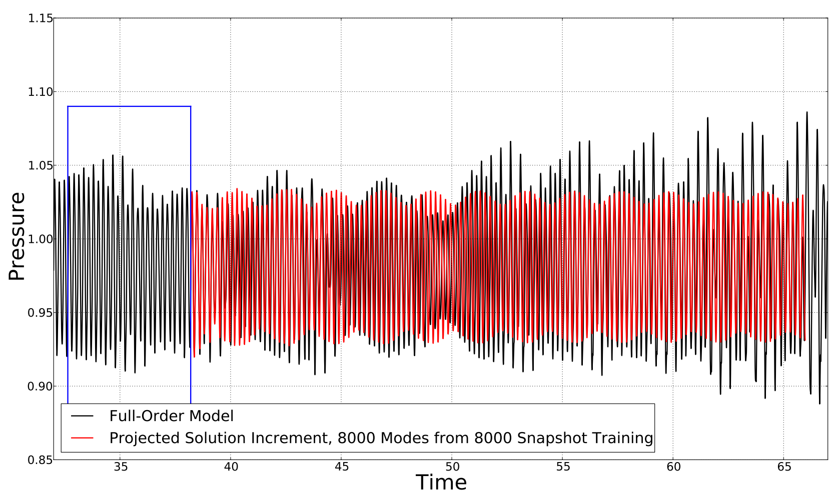

Figure 3.18. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 2. The blue box indicates the extent of the training data used to create the POD basis. This figure is zoomed in to better show the details of the time history comparison.

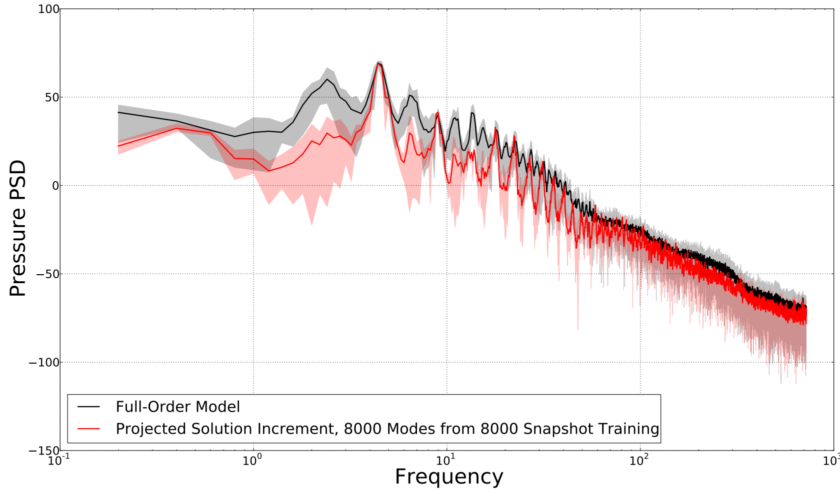

Figure 3.19. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 2. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

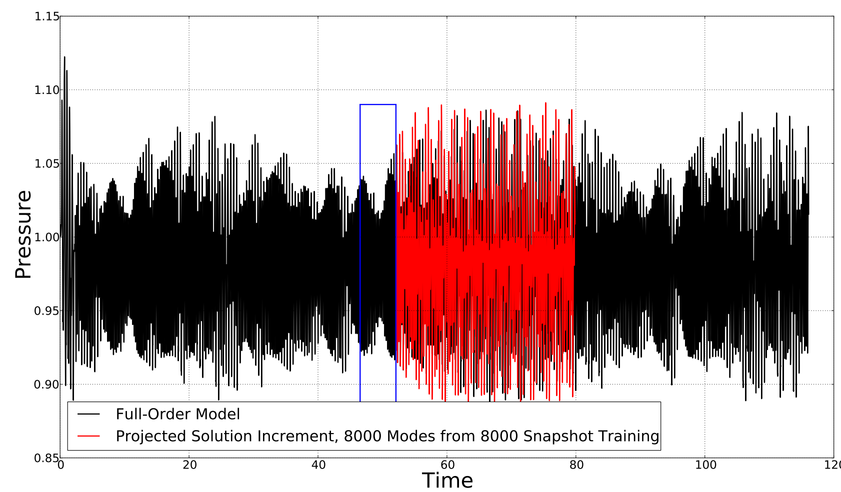

Figure 3.20. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 3. The blue box indicates the extent of the training data used to create the POD basis.

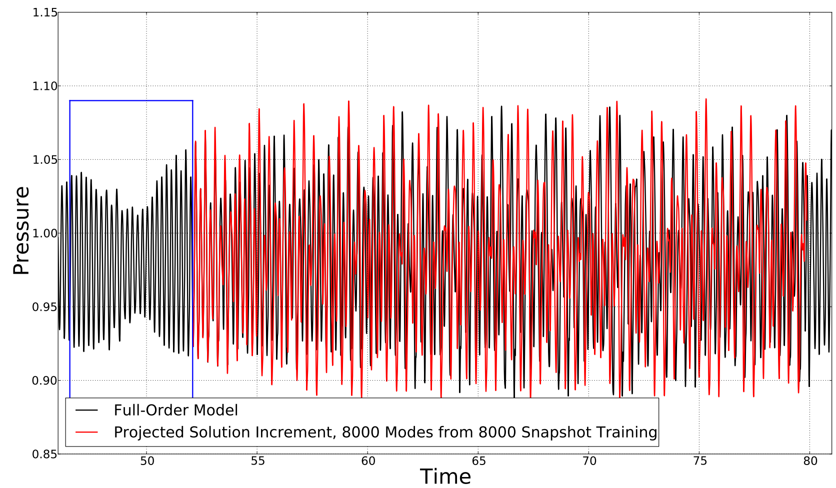

Figure 3.21. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 3. The blue box indicates the extent of the training data used to create the POD basis. This figure is zoomed in to better show the details of the time history comparison.

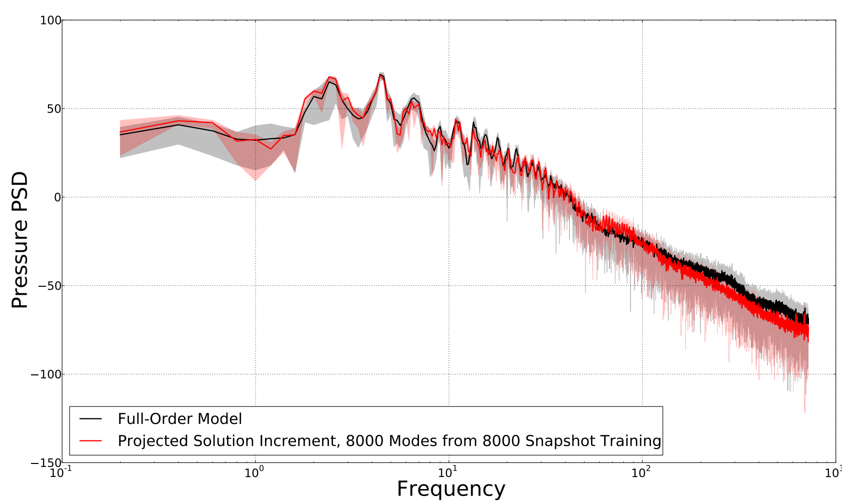

Figure 3.22. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 3. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

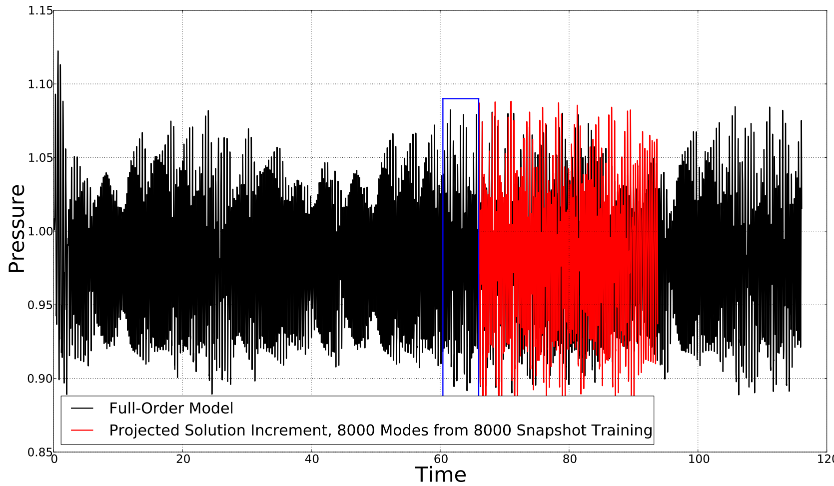

Figure 3.23. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 4. The blue box indicates the extent of the training data used to create the POD basis.

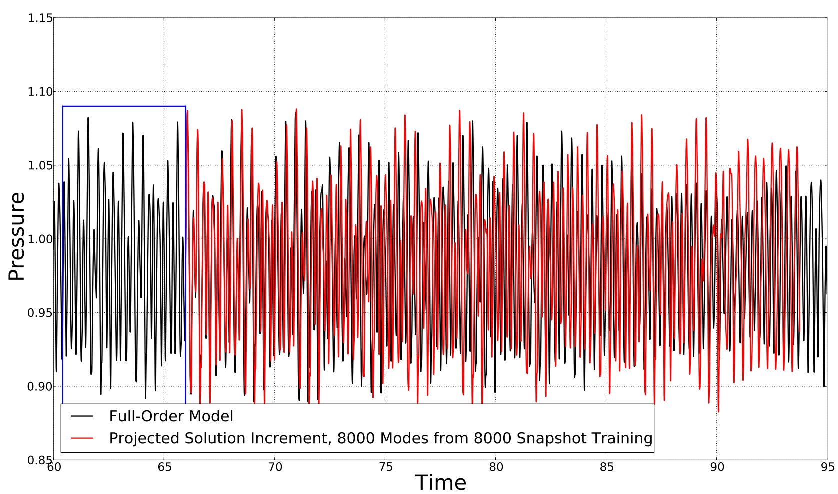

Figure 3.24. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 4. The blue box indicates the extent of the training data used to create the POD basis. This figure is zoomed in to better show the details of the time history comparison.

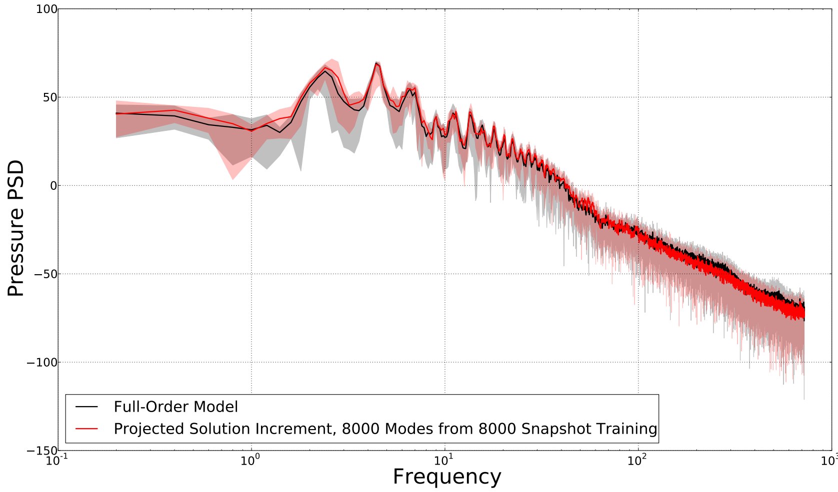

Figure 3.25. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 4. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

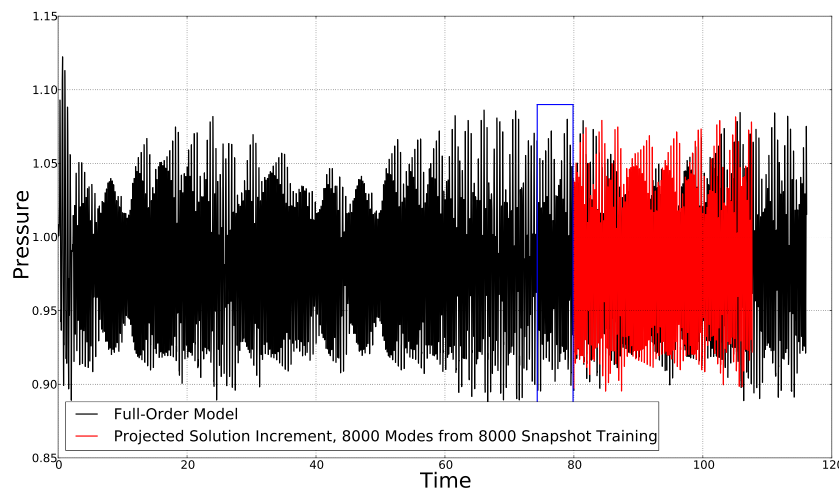

Figure 3.26. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 5. The blue box indicates the extent of the training data used to create the POD basis.

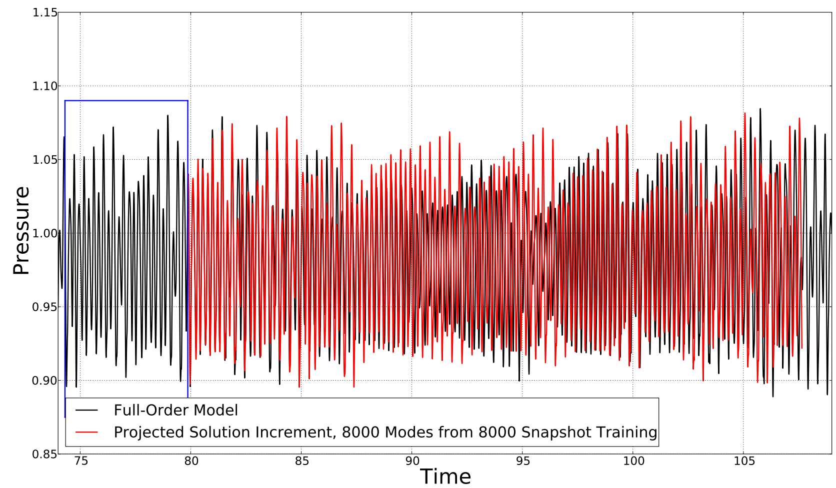

Figure 3.27. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 5. The blue box indicates the extent of the training data used to create the POD basis. This figure is zoomed in to better show the details of the time history comparison.

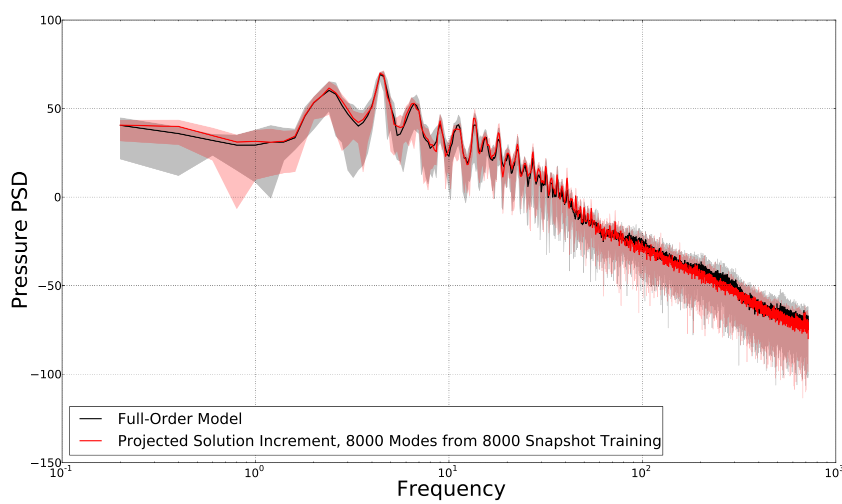

Figure 3.28. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from 8000 snapshots recorded during training interval 5. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

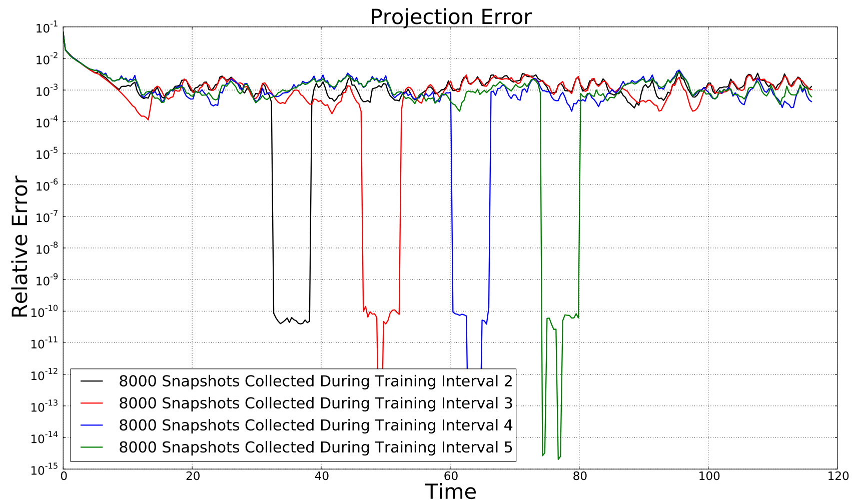

Figure 3.29. The projection error for the bases constructed from 8000 snapshots in training intervals 2 through 5. The sharp changes in error occur at the beginning and end of the individual training intervals.

2 through 5. The projection error serves to indicate how well a given basis can represent the flow field from the full-order model. However, as discussed above, this relies on running the full-order model for the entire time range of interest. From the plots of the projection error, it can be observed that, in general, training intervals 2 and 3 behave similarly as do training intervals 4 and 5. There are regions where intervals 2 and 3 have lower error, and other regions where 4 and 5 have lower error. This leads to the conjecture that combining data from the different training intervals might result in a basis with improved accuracy.

There are several different ways of combining data from different training intervals. We considered 8 different ways of combining data, called groups A through H. Groups A through E are composed of 8 disjoint sets of 1000 snapshots. Groups F through H are composed of 4 disjoint sets of 2000 snapshots. Each group therefore has 8000 total snapshots and is used to construct an 8000 mode basis. Figure 3.30 shows the projection error for group A compared to the 4 contiguous training intervals. The projection error for group A is generally in the middle, at any given point the error is generally lower than the worst performing cases but generally not as low as the best performing cases. The projection errors for the other groups are similar, and the plots are omitted for the sake of brevity.

In the interest of time, we only focused on 3 of the groups, namely A, E, and F. For each group, the start of the projected solution increment run is taken as the end of training interval 5. We can therefore compare these results with those of training interval 5, as presented above. Figure 3.31 shows the pressure time histories for the full-order model and the projected solution increment approach using the basis constructed from the snapshots contained in group A. Figure 3.32 shows the corresponding PSDs. Figures 3.31 and 3.32 show the same comparisons for group E. While Figures 3.31 and 3.32 show the comparisons for group F.

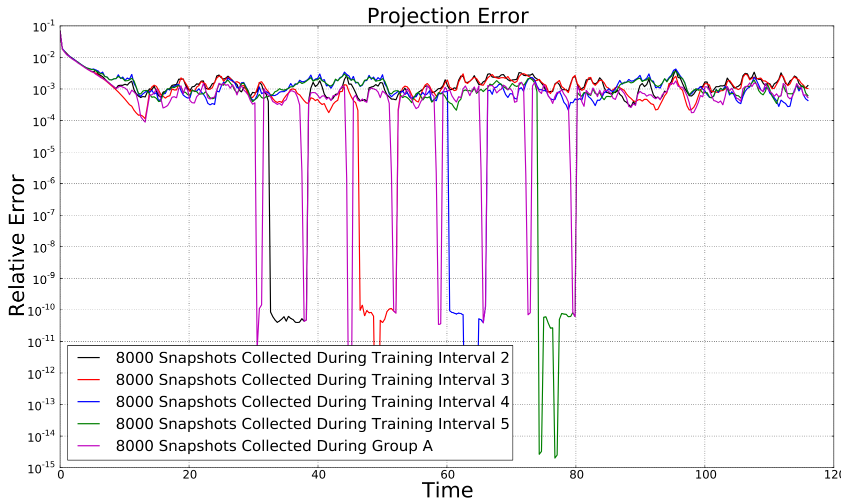

Figure 3.30. The projection error for the bases constructed from 8000 snapshots in group A. The sharp changes in error occur at the beginning and end of the individual training intervals.

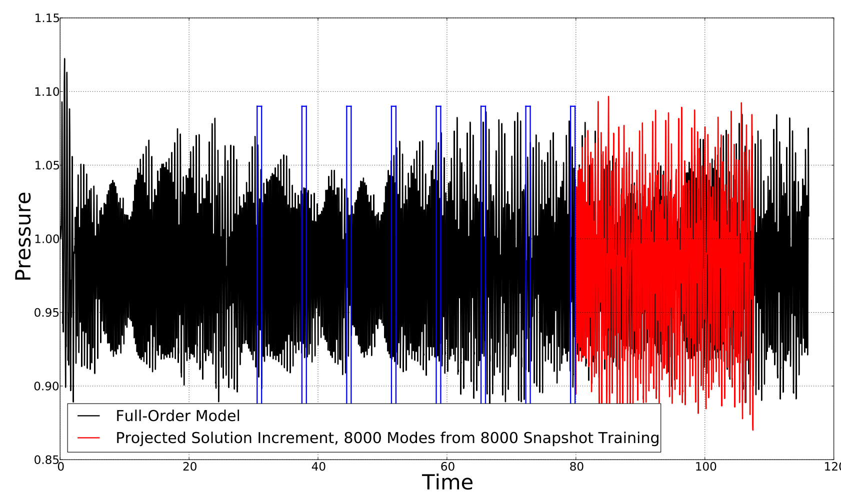

Figure 3.31. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from the 8000 snapshots in group A. The blue boxes indicate the ranges of the training data used to create the POD basis.

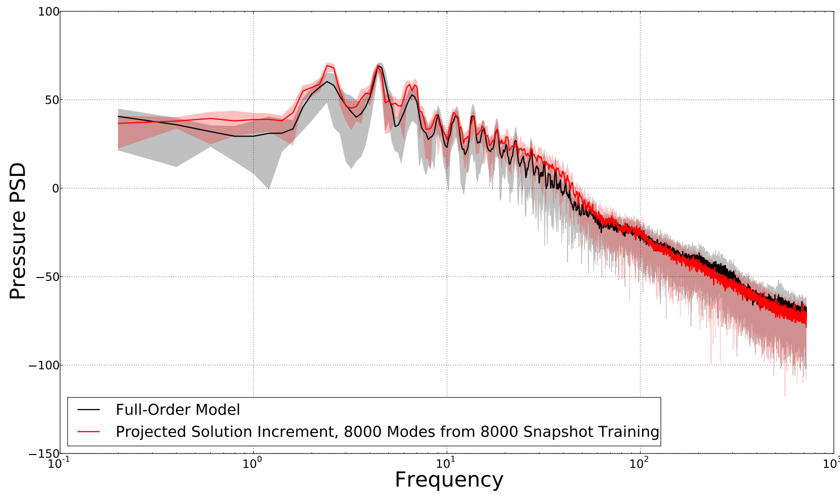

Figure 3.32. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from the 8000 snapshots in group A. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

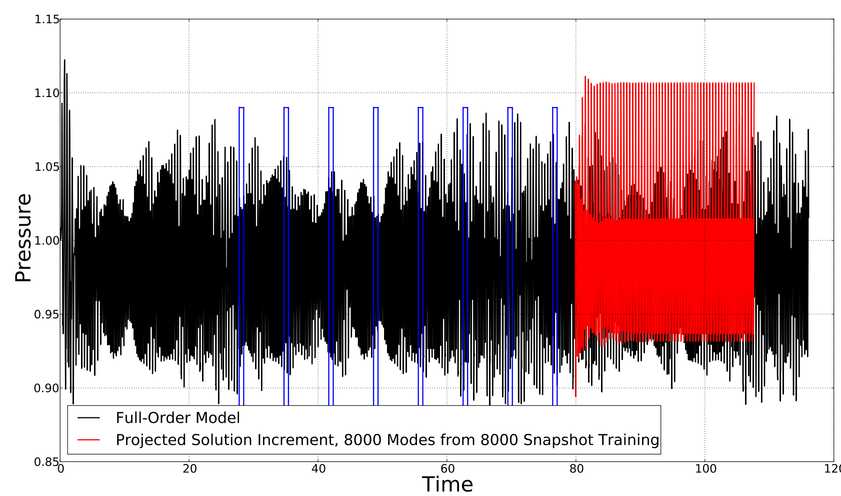

Figure 3.33. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from the 8000 snapshots in group E. The blue boxes indicate the ranges of the training data used to create the POD basis.

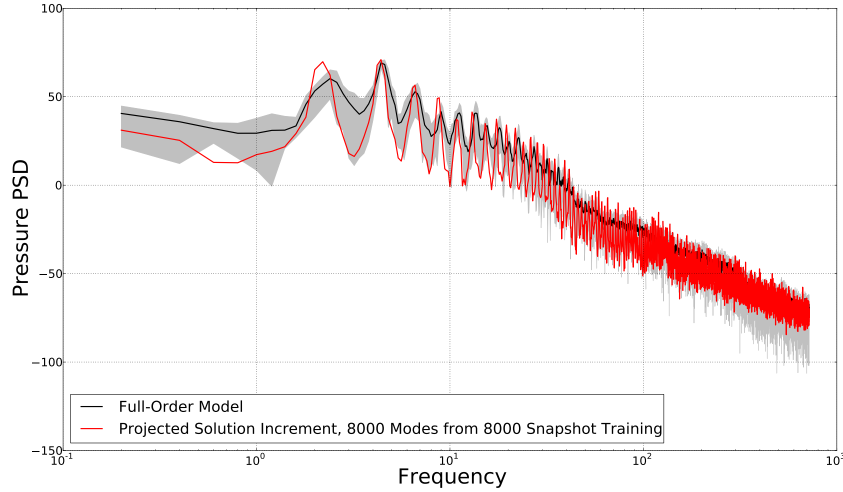

Figure 3.34. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from the 8000 snapshots in group E. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

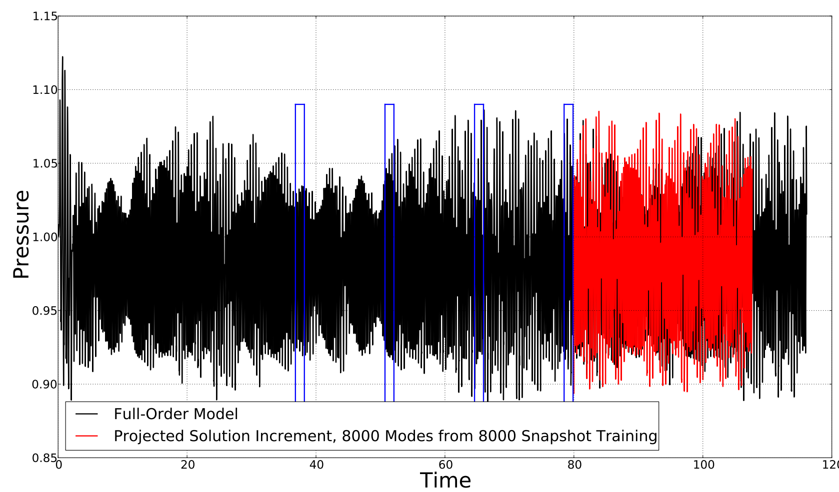

Figure 3.35. The pressure time history for a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from the 8000 snapshots in group F. The blue boxes indicate the ranges of the training data used to create the POD basis.

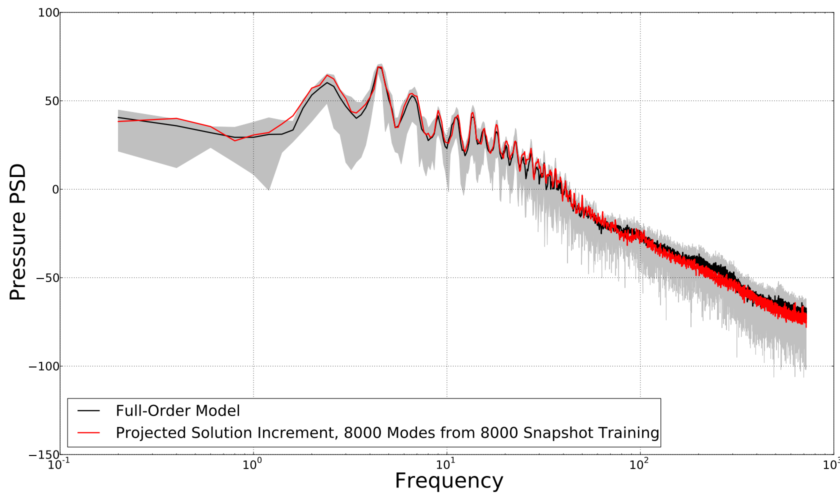

Figure 3.36. The power spectral density of the pressure at a point midway up the downstream wall of the cavity. Shown are results for the full-order model and projected solution increment approach using an untruncated basis constructed from the 8000 snapshots in group F. The solid line indicates the mean PSD. The shaded regions indicate the range of values used to construct the mean.

Group E performs much worse than group A, despite both cases being constructed from 8 equally spaced sets of 1000 snapshots. The primary difference is that group A contains training data immediately prior to the start of the predictive time range. Based on a comparison of the PSDs, group F outperforms group A but is not as accurate as the contiguous training interval 5 results shown in Figure 3.28. These results appear to support the conclusion that the best solution collection strategy is to collect the snapshots immediately prior to the start of the predictive time range. However, based on the projection error shown in Figure 3.30 there should be a benefit to including some information from other time intervals that exhibit different flow features.

Chapter 4

Conclusions and future work

4.1 Conclusions

This report assessed the impact of several decisions on the accuracy of ROMs for the cavity flow test case under consideration. Based on the results presented in this report several conclusions can be made about what choices yield the most accurate ROMs. These conclusions are not listed in order of importance, they are merely in the same order as the results presented in Section 3.3 of this report.

An important decision is the type of projection to use. Here, we considered standard LSPG, preconditioned LSPG, and the projected solution increment approach (ideally-preconditioned LSPG). The projected solution increment approach produces the most accurate results, but relies on solving the full-order model and is therefore computationally expensive and not amenable to hyper-reduction. Standard LSPG without preconditioning can have issues when applied to the cavity flow test case and may fail to successfully complete the simulation (e.g., due to the computation of non-physical quantities such as negative densities or pressures). Preconditioned LSPG resolves some of the issues with standard LSPG and can produce accurate results. The choice of the preconditioner affects the accuracy, with more accurate preconditioners leading to more accurate ROMs but at a greater computational cost.

A somewhat unexpected result is that the ROM accuracy is fairly insensitive to the number of modes used, provided that there are enough modes to capture 99.9% of the energy of the training set. We can therefore truncate the basis significantly without having a large impact on the accuracy. The number of modes does have an impact when reproducing the training interval, but once the ROM is in the predictive time range the accuracy is insensitive to the number of modes.

Based on the work summarized herein, it is unclear if the time step used for the ROM has a significant impact on ROM accuracy. When reproducing the training data there is a benefit to increasing the time step when the number of modes used is very small. However, when considering the PSDs computed from the predictive time range the evidence is less clear. A comparison of the PSDs seems to indicate that increasing the time step negatively impacts the accuracy, while a comparison of the RMS values may provide evidence of a benefit to increasing the time step.

The choice of training interval is an important consideration given that the goal of this project is to train a ROM on a short time interval and then predict for a much longer time interval. Making the training interval as large as possible will likely result in a more accurate ROM. We observed little difference between using 5000 snapshots and using 8000 snapshots, although using fewer snapshots would likely result in a less accurate ROM. We looked at several different sets of 8000 snapshots: 4 contiguous 8000 snapshot sets, 3 combinations of 4 disjoint sets of 2000 snapshots, and 5 combinations of 8 disjoint sets of 1000 snapshots. The results seem to indicate that the optimal choice is to collect the snapshots in a contiguous set immediately prior to start of the predictive time range. However, there may be a benefit to including some information from other time intervals that exhibit different flow features.

The overall conclusion of this work is that it is possible to train a ROM on a fairly short time interval and accurately reproduce the long time statistical quantities of the full-order model. As evidence of this conclusion, please refer to the plots of the PSDs for training interval 5 shown in Figure 3.28. However, it is still possible for things not to work well even when using a contiguous 8000 snapshot training set, as evidenced by the pressure time histories and PSDs for training interval 2 shown in Figures 3.17 and 3.19. This behavior leads to the conclusion that while we can expect a ROM to work well, there may be situations where it does not work as well as anticipated. It is therefore difficult to know *a priori* whether a ROM will yield accurate results.

4.2 Future work

While the R&D described in this report helped to answer many important questions about the viability of projection-based ROMs for the captive-carriage application, further theoretical and implementational work is required to make these models "production ready" (i.e., computationally tractable, reliable predictive tools). We briefly describe some worthwhile avenues for future R&D below.

- *Development of goal-oriented snapshot sampling strategies.* It was shown in this report that the accuracy of a ROM can depend on the set of snapshots employed in creating the reduced basis. Using a contiguous set of snapshots is not always better than employing several sets of noncontiguous snapshots. It is not clear at the present time why certain snapshot sets yield more accurate ROMs. A worthwhile future research endeavor would be to perform a more rigorous study of the effect of the snapshot sampling strategy on ROM accuracy, towards developing an "optimal" (error minimizing) goal-oriented algorithm/criterion for snapshot selection.

- *Addition of physics-based constraints to LSPG ROM formulation.* It was shown in several past works, e.g., [29, 7, 22], that ROMs can benefit from "structure preservation": the introduction of physics-based constraints to ensure that the ROM formulation respects certain properties of the underlying flow physics and/or full order discretization. For compressible flow problems, relevant physics properties include global conservation and entropy-stability. During the course of this project, we began a numerical study of ROM global conservation properties, which revealed the presence of small conservation violations in the ROM solution in the predictive regime, as shown in Figure 3.2. A more thorough study of the extent of these violations and their effect on ROM accuracy is worthwhile. If conservation or entropy-stability violations are deemed to be a significant source of error, a future development task should be the extension of the LSPG formulation in `SPARC` to include such constraints.

- *Space-time projection.* Another way to potentially improve the long-time behavior of the ROM is to adopt a recently-developed space-time LSPG projection [10]. This approach has several advantages over typical spatial-projection-based ROMs such as traditional Galerkin and LSPG projection. First, because the full space–time problem is solved via projection, errors do not accumulate in time, which leads to error bounds that grow sub-quadratically (rather than exponentially) in time. Furthermore, these bounds are *a priori* in nature, and are formulated in terms of the *best approximation error* of the solution over the space–time trial subspace. Second, because this method is equipped with space–time hyper-reduction, the complexity of the resulting ROM is independent of both the original spatial and temporal dimensions. The main drawback of the method is the implementation (the ROM requires accessing the residual at multiple time instances simultaneously) and storage requirements, although the latter are mitigated by using tensor-decomposition methods (e.g., sequentially truncated high-order SVD) to define the space–time trial basis.

- *Development/integration of basis enrichment and/or refinement strategies.* As shown in this report, moderately large basis sizes (capturing 99.9% of the snapshot energy) are required to achieve a time-predictive ROM that is sufficiently accurate for the captive-carry application. This can lead to a prohibitively large number of modes, especially when considering high Reynolds number problems. To mitigate this problem, several basis enrichment [2, 1] and refinement [5] strategies have been proposed in recent years. The integration of such approaches into the ROMs implemented within `SPARC` can potentially yield more computationally efficient and accurate models.

- *Computational improvements.* Finally, as noted earlier in this report, the focus of the R&D described herein has been on ROM accuracy, not ROM efficiency. Having determined that it *is* possible to obtain time-predictive ROMs that are sufficiently accurate at representing our key QoI, the pressure PSD, with respect to a full-order model, the next step is to implement approaches to reduce the computational complexity of these ROMs. The key ingredient missing from the current implementation is hyper-reduction. Following the implementation of hyper-reduction (Section 2.5) in `SPARC`, the numerical studies detailed herein should be repeated to gauge the amount of error introduced by the hyper-reduction. Other computational improvements to the ROM implementation include the introduction of TSQR ("tall-skinny QR") for solving the optimization problem (2.9). The current implementation solves the normal equations, an approach that is in general less stable and more cost-intensive than TSQR.

Chapter 5

Appendix: Repository of data sets

An archive of the data and other files used to create the results in this report has been created and is accessible internally to Sandia. A subset of this archive is intended to be made available externally as well. This archive is meant to serve as a reference for future ROM efforts.

The internal archive includes:

- Data files containing the snapshots from training intervals 2 to 5 and groups A through H.
- The POD basis files generated from these snapshots.
- Selected results from the projected solution increment runs.
- Selected results for standard LSPG and preconditioned LSPG ROM runs.
- Input files for the selected runs.
- The version of `SPARC` and `Trilinos` used to generate these results.
- Post-processing tools used to extract QoIs and plot results.

The files in the internal archive will remain static. However, `SPARC` is currently under active development. Given the chaotic nature of the cavity flow test case, there is no guarantee that newer versions of `SPARC` will exactly reproduce the archived results. The input files in the archive will be static, however, versions of the input files are also contained in the `SPARC` verification repository. The versions in the verification repository will be updated periodically to account for changes in `SPARC`, such as input file syntax changes. However, there is likely to be very limited testing to see if the archived results can be replicated. The archive is therefore meant to serve as a reference for future ROM efforts at Sandia but may not represent the current state of the ROM capabilities in `SPARC`.

References

[1] M. Balajewicz, I. Tezaur, and E. Dowell. Minimal subspace rotation on the stiefel manifold for stabilization and enhancement of projection-based reduced order models for the compressible navier-stokes equations. *J. Comput. Phys.*, 321:224–241, 2016.

[2] M.J. Balajewicz, E.H. Dowell, and B.R. Noack. Low-dimensional modelling of high-Reynolds-number shear flows incorporating constraints from the Navier-Stokes equation. *J. Fluid Mech.*, 729:285–308, 2013.

[3] D. Barnette, R. Barrett, S. Hammond, J. Jayaraj, and J. Laros III. Using miniapplications in a mantevo framework for optimizing sandias sparc cfd code on multi-core, many-core, and gpu-accelerated compute platforms. Proceedings of 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition (AIAA 2013–1126), 2013.

[4] M.F. Barone, I. Kalashnikova, D.J. Segalman, and H. Thornquist. Stable Galerkin reduced order models for linearized compressible flow. *J. Comput. Phys.*, 288:1932–1946, 2009.

[5] K. Carlberg. Adaptive *h*–refinement for reduced-order models. *Int. J. Numer. Meth. Engng.*, 102:1192–1210, 2014.

[6] K. Carlberg, M. Barone, and H. Antil. Galerkin v. least-squares Petrov–Galerkin projection in nonlinear model reduction. *J. Comput. Phys.*, 330:693–734, 2017.

[7] K. Carlberg, Y. Choi, and S. Sargsyan. Conservative model reduction for finite-volume models. *J. Comput. Phys.*, (371):280–314, 2018.

[8] K. Carlberg, C. Farhat, J. Cortial, and D. Amsallem. The GNAT method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows. *J. Comput. Phys.*, 24(2):623–647, 2013.

[9] S. Chaturantabut and D.C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM J. Sci. Comput.*, 32:2737–2764, 2010.

[10] Y. Choi and K. Carlberg. Space-time least-squares petrov-galerkin projection for nonlinear model reduction. *arXiv e-Print*, (1703.04560), 2017.

[11] R. Everson and L. Sirovich. Karhunen-Loeve procedure for gappy data. *J. Optical Society of America A*, pages 1657–1664, 1995.

[12] P. Holmes, J.L. Lumley, and G. Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, 1996.

[13] I. Kalashnikova and S. Arunajatesan. A stable Galerkin reduced order model for compressible flow. WCCM-2012-19407, 10th World Congress on Computational Mechanics (WCCM), Sao Paulo, Brazil, 2012.

[14] I. Kalashnikova, S. Arunajatesan, M.F. Barone, B.G. van Bloemen Waanders, and J.A. Fike. Reduced order modeling for prediction and control of large-scale systems. Sandia National Laboratories Report, SAND No. 2014-4693. Sandia National Laboratories, Albuquerque, NM, 2014.

[15] I. Kalashnikova and M.F. Barone. On the stability and convergence of a Galerkin reduced order model (ROM) of compressible flow with solid wall and far-field boundary treatment. *J. Comput. Phys.Int. J. Numer. Meth. Engng.*, 83:1345–1375, 2010.

[16] I. Kalashnikova and M.F. Barone. Stable and efficient Galerkin reduced order models for non-linear fluid flow. AIAA-2011-3110, 6th AIAA Theoretical Fluid Mechanics Conference, Honolulu, HI, 2011.

[17] I. Kalashnikova, M.F. Barone, S. Arunajatesan, and B.G. van Bloemen Waanders. Construction of energy-stable Galerkin reduced order models. Sandia National Laboratories Report, SAND No. 2013-4063, Sandia National Laboratories, Albuquerque, NM, 2013.

[18] I. Kalashnikova, M.F. Barone, S. Arunajatesan, and B.G. van Bloemen Waanders. Construction of energy-stable projection-based reduced order models. *Appl. Math. Computat.*, pages 569–596, 2014.

[19] K. Kunisch and S. Volkwein. Galerkin proper orthogonal decomposition for a general equation in fluid dynamics. *SIAM J. Num. Anal.*, 40(2):492–515, 2002.

[20] P. LeGresley. Application of proper orthogonal decomposition (pod) to design decomposition methods. Ph.D. thesis, Stanford University, 2006.

[21] J.L. Lumley. *Stochastic tools in turbulence*. Academic Press: New York, 1971.

[22] D. Maddix, I. Tezaur, and K. Carlberg. Methods for structure preserving reduced order models in fluid dynamics. *In Preparation*, 2017.

[23] N.C. Nguyen, A.T. Patera, and J. Peraire. A 'best points' interpolation method for efficient approximation of parametrized functions. *Int. J. Numer. Meth. Engng.*, 73:521–543, 2008.

[24] N.C. Nguyen and J. Peraire. An efficient reduced-order modeling approach for non-linear parametrized partial differential equations. *Int. J. Numer. Meth. Engng.*, 76:27–55, 2008.

[25] A. Prokopenko, C. Siefert, J. Hu, M. Hoemmen, and A. Klinvex. Ifpack2 User's Guide 1.0 (Trilinos version 12.6). Sandia Report, SAND2016-5338, 2016.

[26] M. Rathinam and L.R. Petzold. A new look at proper orthogonal decomposition. *SIAM J. Num. Anal.*, 41(5):1893–1925, 2003.

[27] L. Sirovich. Turbulence and the dynamics of coherent structures, part iii: dynamics and scaling. *Q. Appl. Math.*, 45(3):583–590, 1987.

[28] I. Tezaur, J. Fike, K. Carlberg, M. Balajewicz, M. Barone, and E. Mussoni. Model reduction for compressible cavity simulations towards uncertainty quantification of structural loading. Sandia National Laboratories Report, Sand No. 2016-9463. Sandia National Laboratories, Albuquerque, NM, 2016.

[29] I. Tezaur, J. Fike, K. Carlberg, M. Barone, D. Maddix, E. Mussoni, and M. Balajewicz. Advanced fluid reduced order models for compressible flow. Sandia National Laboratories Report, Sand No. 2017-10335. Sandia National Laboratories, Albuquerque, NM, 2016.

DISTRIBUTION:

1	MS 0825	Srini Arunajatesan, 01515
1	MS 0825	Matt Barone, 01515
1	MS 9159	Patrick Blonigan, 08759
1	MS 9159	Kevin Carlberg, 08759
1	MS 0828	Paul Crozier, 01541
1	MS 0825	Jeff Fike, 01515
1	MS 0828	Brian Freno, 01544
1	MS 9159	Chi Hoang, 08759
1	MS 0825	Micah Howard, 01515
1	MS 9159	Kookjin Lee, 08759
1	MS 9159	Jerry McNeish, 08754
1	MS 9159	Eric Parish, 08759
1	MS 9152	Francesco Rizzi, 08759
1	MS 0825	Justin Smith, 01515
1	MS 0899	Technical Library, 9536 (electronic copy)

Sandia National Laboratories