

© 2022 Kento Kaneko

AN AUGMENTED BASIS METHOD FOR REDUCED ORDER MODELS OF
TURBULENT FLOW

BY

KENTO KANEKO

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Theoretical and Applied Mechanics
in the Graduate College of the
University of Illinois Urbana-Champaign, 2022

Urbana, Illinois

Doctoral Committee:

Professor Paul Fischer, Chair
Professor Arne Pearlstein
Professor Moshe Matalon
Professor Luke Olson

Abstract

Reduced-order models (ROMs) offer a promising approach for parametric analysis of engineering fluid dynamics applications. The standard procedure consists of using solution snapshots to produce a truncated POD basis, which is in turn used in a Galerkin projection of the governing Navier-Stokes equations (NSE). Unfortunately, the standard POD approach has well-known limitations for high Reynolds number flows that are largely attributable to the lack of fine-scale structure in the low-rank POD bases, which tend to be spatially smooth and therefore unable to generate sufficient small-scale dissipation to stabilize the solution for a small number of modes, N . Even with stabilization, the required value of N is often sufficiently large that these approaches are impaired by the $O(N^3)$ costs associated with evaluation of the third-order advection tensor at each step of the ROM time-advancement. We present a novel non-intrusive stabilization technique in the form of basis augmentation that, in many cases, reduces the total number of modes required to produce a stable and accurate ROM reconstruction for turbulent flows at modest Reynolds numbers. The approach involves augmenting the standard POD modes with divergence-free projections of subsets of POD-expanded terms originating from the advection term. Differing combinations of these basis elements are considered. Bases that include interactions with lifting function and self-interactions have proven to be quite effective for several challenging flow problems with relatively low values of \hat{N} , where \hat{N} is the total number of basis including the augmentation modes. We demonstrate this proposed basis set on several challenging problems and compare its stability properties with alternative stabilization approaches for POD-based ROMs.

To my parents, Keiko and Toru.

Acknowledgments

This work would not have been possible without the support of many people. First, I would like to thank my adviser, Prof. Paul Fischer, who mentored and guided me through my graduate studies and while at the same time, provided freedom to pursue different ideas that culminated in this work. Many thanks to current students of Spectral Element Analysis Lab: Pedro, Thilina, Nick, Ping-Hsuan, Malachi, Yu-Hsiang, and Viral; and former students Li, Ketan, and Som. I would also like thank the committee members, Prof. Arne Pearlstein, Prof. Moshe Matalon, and Prof. Luke Olson who offered guidance and support. And finally, thanks to my parents and friends who always gave their love and support.

Table of contents

Chapter 1 Introduction	1
1.1 Historical Development	5
Chapter 2 Model-Order Reduction for Thermal Fluid Analysis	10
2.1 Galerkin-Based ROM	12
2.1.1 Proper Orthogonal Decomposition	13
2.1.2 Galerkin Formulation	15
2.2 Reproduction Problem	18
2.3 Parametric Model-Order Reduction	24
2.3.1 Axisymmetric Rayleigh-Bénard Problem	25
2.3.2 pMOR Setup	27
2.3.3 Computation of the Nusselt Number	28
2.3.4 pMOR Results	29
2.4 Discussion	31
2.5 Efficient Parallel ROM Generation	32
Chapter 3 ROM Stabilization	36
3.1 Leray Regularization	37
3.2 Constrained Optimization	38
3.3 Results for Stabilized Methods	40
Chapter 4 Augmented Basis Method	45
4.1 Methodology	45
4.1.1 Space of Solutions	46
4.1.2 Time-Evolution of NSE	47
Chapter 5 Applications of ABM-ROM	50
5.1 2D Lid-Driven Cavity (LDC)	51
5.1.1 Comparison of the modes	52
5.1.2 Convergence and Stability Study using Mean and Variance Metric ..	53
5.2 2D Baffle Flow	54
5.3 3D Lid-Driven Cavity	56
5.4 3D Flow over a Hemisphere	59
5.5 3D Flow Past a Cylinder	60
5.6 Forced Convection in Turbulent Pipe Flow	62
5.6.1 FOM Prescription of Constant Flowrate	64
5.6.2 ROM Prescription of Constant Flowrate	64

5.6.3	Enforcing Periodicity in the Thermal Solution	65
5.6.4	ROM Results for Pipe Flow	66
5.7	Forced Convection in Turbulent Channel Flow	69
5.8	Heated Pipe-Flow with Wire Coil Insert	72
5.8.1	FOM Study of Wire Coil and Other Inserts	73
5.8.2	ROM Result for Hydrodynamics of Pipe with Wire-Coil	82
Chapter 6	Investigation of ABM Stabilization Mechanism	84
6.1	Convective Energy-Transfer	84
6.2	Time Evolution of Individual ROM Coefficients	87
6.3	Accuracy Dependency on the Size of Snapshot Set	89
6.3.1	Serial Snapshot Selection Experiment	91
6.4	Towards Higher Reynolds Number	92
Chapter 7	Discussion	97
Chapter 8	Conclusion	101
References	103

Chapter 1

Introduction

Fluid-thermal analysis plays a critical role in understanding and prediction of many important phenomena for engineering and science applications, including engine design, cooling mechanisms in nuclear reactors, fundamentals of stellar convection, and cooling of biological systems. In the early developmental phase of this discipline, rigorous mathematical models, including the Navier-Stokes equations (NSE) were established and used in conjunction with the energy transport equation to accurately represent the thermal states of the flow measured in experimental setups.

A significant challenge with the mathematical formulation is that the governing equations are nonlinear partial differential equations (PDEs) in three space-dimensions and time, which makes their solution difficult to obtain, even in simple geometries. For problems involving complex geometry and/or with variable material properties, it is generally necessary to resort to numerical solution of the governing system of PDEs on a computer using numerical methods. At elevated Reynolds numbers even simple configurations can require a numerical approach if the flow transitions to turbulence. Numerical discretizations, including finite-difference, finite volume, finite element, and spectral methods, have been demonstrated to provide high-fidelity solutions of the NSE with geometry, properties, initial conditions, and boundary conditions that are relevant to engineering applications. This field, known as computational

fluid dynamics (CFD) has become ubiquitous to the point that it is widely accessible by engineers, with a few caveats.

One of these caveats is the limitation in the problem size, which is measured in the number of grid points or degrees-of-freedom (DoFs), \mathcal{N} , required to accurately represent the velocity, pressure, and temperature fields. While CFD was initially limited to small problem sizes that could only describe flows with limited ranges of scales, advances in high-performance computing (HPC) have led to significant increases in the computational capacity available to solve a given CFD problem. These advances are illustrated by the problem of isotropic turbulence in a box over the past four decades. In 1977, Orszag addressed this problem with the *Centicube* code, which could advance the NSE at 20 seconds per timestep on a Cray 1 using a spectral method with $\mathcal{N} = 128^3$ Fourier modes [1, 2]. In 2019, Ravikumar *et al.* [3] could advance the same problem at higher Reynolds numbers using $\mathcal{N} = 18432^3$ Fourier modes in 14 seconds per timestep on the GPU-based Summit platform at Oak Ridge National Laboratory. This latter problem represents a three million-fold increase over the *Centicube* results.

Hero calculations such as those mentioned previously serve multiple roles. First, algorithmic benefits are often manifest in a trickle-down effect since today's supercomputers typically become tomorrow's professional workstations and personal computers. Second, cutting-edge simulations can be used to gain insight into fundamental mechanisms of turbulence, much in the same way as experiments, both of which can lead to turbulence models that might be employed in less expensive (i.e., smaller) thermal-fluid calculations.

Conceivably, HPC could also be used for engineering design with large eddy simulation (LES) or with Reynolds-averaged Navier-Stokes (RANS) turbulence models. Such approaches, however, become challenging for large complicated domains when many design iterations are required. Each expensive simulation incurs as much expense as its predecessor as there is no information sharing between the two calculations. It is in this context where a relatively new approach offers some opportunity. *Reduced-order models* (ROMs) are designed to leverage

existing information (typically coming from a large HPC calculation, known as a *full-order model*, or FOM) by building a low-rank approximation space that is derived from the FOM solution. One then uses a suitable basis for this space in a Galerkin framework to approximate the solutions of the parameterized NSE and energy equations. Possible parameterizations include varying Reynolds or Rayleigh numbers or limited geometric variations.

To be useful as design tools, ROMs for unsteady flows must address two problems. First, they must be able to solve the *reproduction problem*. That is, they must be able to accurately reproduce quantities of interest (QOIs) that are generated by the originating FOM. Second, they should be able to solve the *parametric problem*, in which they are used to evaluate QOIs in the (parametric) neighborhood of the originating problem.

Parametric model-order reduction (pMOR) is a promising approach to leveraging high-performance computing (HPC) for design and analysis in fluid-thermal engineering applications. The governing equations in this context are the time-dependent incompressible Navier-Stokes equations (NSE) and the thermal transport equation,

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}, \quad \nabla \cdot \mathbf{u} = 0, \quad (1.1)$$

$$\partial_t T + \mathbf{u} \cdot \nabla T = \kappa \nabla^2 T. \quad (1.2)$$

where ν and κ parameterize the PDEs and the forcing function \mathbf{f} can be the Boussinesq approximation term, for example.¹ The equations are assumed to hold in a suitable domain Ω with appropriate initial and boundary conditions. The Galerkin statement is

$$\text{Find } (\partial_t \mathbf{u}, p, \partial_t T) \in Y := [\mathbf{H}_0^1 \otimes L^2 \otimes H_0^1] \text{ s.t. } \forall (\mathbf{v}, q, S) \in Y,$$

$$(\mathbf{v}, \partial_t \mathbf{u}) + (\mathbf{v}, \mathbf{u} \cdot \nabla \mathbf{u}) = (\nabla \cdot \mathbf{v}, p) - \nu (\nabla \mathbf{v}, \nabla \mathbf{u}) + (\mathbf{v}, \mathbf{f}), \quad (q, \nabla \cdot \mathbf{u}) = 0, \quad (1.3)$$

$$(S, \partial_t T) + (S, \mathbf{u} \cdot \nabla T) = -\kappa (\nabla S, \nabla T). \quad (1.4)$$

¹These equations are effectively in nondimensional form, which for forced conditions implies that $\nu = Re^{-1}$, the inverse Reynolds number, and $\kappa = Pe^{-1}$, the inverse Peclet number. For buoyancy-driven flows these parameters typically scale with Rayleigh number (Ra) and Prandtl number (Pr), with the precise definition dependent on the chosen scaling.

Here, L^2 is the space of square-integrable functions on Ω ; H^1 is the space of functions in L^2 whose gradient is also in L^2 ; and H_0^1 is the space of functions in H^1 that vanish on subsets of the boundary, $\partial\Omega_D \subset \partial\Omega$, where homogeneous Dirichlet conditions are imposed. \mathbf{H}_0^1 is the vector counterpart to H_0^1 .

To obtain a fully-accurate quantity of interest (QOI) such as friction factor, Nusselt number, or Strouhal number, one formally needs to obtain a full-order model (FOM) solution to the governing equations at discrete points in the parameter space of interest (e.g., spanned by a range of ν and κ , of interest). Typically, the FOM constitutes a high-fidelity spectral- or finite-element solution to the governing equations, which can be expensive to solve, particularly for high Reynolds number cases that are typical of engineering applications. pMOR seeks to develop a sequence of *reduced-order models* (ROMs) that capture the behavior of the FOM and allow for parameter variation. For unsteady flows, the pMOR problem can be broken down into two subproblems: *reproduction*, wherein the ROM captures essential time-transient behavior of the FOM using the same parameter (anchor) point for each, and *parametric variation*, wherein the ROM is run at a *different* parametric point in order to predict the system behavior away from the anchor points at which the FOM simulation was conducted. In this work, we focus primarily on the reproduction problem for challenging unsteady flows. We do, however, also consider pMOR, which we illustrate with an example from Kaneko *et al.* [4] in Chapter 2.

While pMOR is a promising approach for engineering analysis and design, it is well known that even the reproduction problem is challenging for the classical POD-Galerkin approach at high Reynolds numbers after the flow transitions to turbulence. One common issue with this class of problems is that the ROM solution approaches an unphysical attractor. This behavior is attributed to a lack of dissipation, given that the truncated POD space lacks high-wavenumber modes that are capable of dissipating energy. One can induce additional dissipation by including more modes but the cost is high. The convective tensor reduction requires storage of N^3 entries for the advection operator with a corresponding work of $2N^3$

operations per timestep. While $N = 100$, with a cost of a two million operations per step and a million words in memory, may be tolerable, $N = 400$ with a cost of 128 million operations and 64 million words quickly makes pMOR less viable for running on a workstation, which is typically the target for this type of analysis tool.

Existing techniques for addressing the computational cost for the online phase of POD-ROM include the discrete empirical interpolation method (DEIM) [5–7], which can effectively evaluate non-linear terms based on interpolation points, and tensor decomposition, which aims to approximate the convective tensor by a low-rank tensor. We will show in our concluding examples that these methods will not, on their own, address the unphysical ROM dynamics. Stabilization of the ROM is critical and is the primary topic of this work. Several stabilization strategies are introduced in Section 1.1 and a detailed description is given in Section 3. The major contribution of this work is the development of a novel augmented-basis method (ABM) [8], in which we add important modes to the standard POD bases. In many cases, the ABM increases both the stability and accuracy of the ROMs at a cost equivalent to standard POD approaches having the same total number of modes.

1.1 Historical Development

The POD-Galerkin technique in fluid flow emerged from Lumley’s application of POD to identify dominant flow features [9]. Model-reduction using POD modes as basis functions was introduced afterwards, with a comprehensive analysis appearing in a later monograph by Holmes *et al.* [10].

Application of pMOR to more complex PDEs with non-affine parameter dependencies were addressed by Barrault *et al.* [11] using decomposition of the nonlinearity based on the empirical interpolation method (EIM). In this approach, successive interpolation modes are chosen to eliminate the error between the targeted term in FOM and the ROM at “magic points” designated as points in Ω where the error of the current interpolant (i.e.,

the approximant of the next mode) is maximal. This method was further extended by Chaturantabut & Sorensen [5–7] with a POD decomposition of the nonlinear term and the choice of points restricted to discrete points produced by the spatial discretization of the PDE, called discrete empirical interpolation method (DEIM). While these methods enable application of pMOR to nonlinear problems, the issue of insufficient basis and feasible stabilization methods for the NSE persists.

Due to its approach of treating nonlinear terms, DEIM has the potential to address the high-cost issue of including more modes. DEIM replaces the third-order convective tensor with a collocation-like decomposition at the discrete magic points, that can reduce the computational complexity from $O(N^3)$ to $O(N^2)$ if the approximation is sufficiently accurate for a small number of interpolation points. Accounting for the constants, evaluation of advection using DEIM with $N = 200$ modes would be equivalent to using the full tensor with $N = 65$ when fixing the number of interpolation points to N . For the same cost, DEIM thus permits the use of a richer approximation space. In the context of solving the steady Navier-Stokes equations, Elman & Forstall [12] has successfully applied DEIM.

To certify that the error in the ROMs that are produced is smaller than the acceptable tolerance, error indicators have been developed based on the residual of the ROM solutions in the full-order model (FOM) space. This factor in the pMOR methodology is crucial for ensuring the error in the low-dimensional model of the ROM is sufficiently small. Error indicators for coercive elliptic PDEs are described by Veroy *et al.* [13]. In this work, bounds on the QOI from the ROM is established via a posteriori error indicator. The QOI in this case is limited to those produced by a linear functional, but engineering application for which that is the case is numerous (e.g., Nusselt number, velocity component values at particular points, lift/drag, etc.).

An a posteriori error indicator for time-dependent NSE is described by Fick *et al.* [14]. By breaking down the contribution of the error of the ROM in the FOM space at each time-step, they were able to produce an error indicator whose factors are pre-computable in the offline

processing stage and can be evaluated efficiently in the online stage. In the same work by Fick *et al.*, a constrained-optimization approach is also proposed which will be discussed later in this section. We do not consider error indicators further in this dissertation, but they are nonetheless an important component for efficient and accurate pMOR and are discussed by Tsai & Fischer [15].

For addressing the issue of stability, several modifications to the original POD Galerkin approach have been proposed. Iollo *et al.* [16] proposed a modification of the POD mode generation in which the H^1 inner-product is used to produce the Gramian, rather than L^2 inner-product, to emphasize the importance of gradients in the FOM snapshots. Because the velocity gradients are an important factor in the dissipation mechanism of the flow solution, these modes contain more of these high-wavenumber components compared to the L^2 based POD modes.

Wells *et al.* [17] introduced Leray regularization in the context of ROM in which the *advecting* field is smoothed (conveniently, by truncation of the modes in the case of POD-ROM). This regularization enhances the stability property of the dynamical behavior; however, the optimal choice of regularization (e.g., number of modes to truncate or shape of transfer function) is not known a priori. This method has an attractive feature of ease of implementation and also reduction of the cost of evaluating the convective term in the ROM, which is the dominant cost of running the ROM, by a factor of 2 for a sharp cutoff of 50% of the modes.

An alternative stabilization approach, introduced by Fick *et al.* [14], is to replace the discrete ODE system by a constrained minimization problem at each timestep. During the evolution of the system, the basis coefficients are bound by the minimum and maximum coefficient values observed in the snapshot projection onto the truncated POD space (if the constraints are inactive, one recovers the standard Galerkin-based trajectory). With this approach, the ROM tends to stay close to the dynamics of the FOM because each coefficient is forced to evolve between the empirical bounds. A challenge, however, is that this approach requires ad hoc modification of the bounds for parametric values where the FOM snapshots

are unavailable. Applications of several of these stabilization techniques may be found in Kaneko *et al.* [4].

Furthermore, methods that address the stability issue by constructing basis functions that satisfy the energy-balance that closely match the POD basis is introduced by Balajewicz [18]. In this work, existing work on stabilizing LTI systems by finding the optimal combination of snapshots to produce dynamically stable ROM by Amsallem & Farhat [19] is combined with the work by Cazemier *et al.* [20] in which the kinetic energy behavior is stabilized by introducing an empirical turbulence closure term in the ROM. In this combined method by Balajewicz, an a priori non-linear constrained optimization problem is solved that minimizes the difference between the energy captured by the new modes and the energy captured by the POD modes for a given N subject to constraints: the columns of the transformation matrix are orthonormal and the empirical kinetic energy-balance is satisfied. The author found that this approach offers some improvement upon the standard POD-Galerkin approach. Also discovered was the fact that by ensuring the kinetic energy-balance is dissipative to an arbitrary degree, the ROM solution becomes stabilized. Thus, by encapsulating the method by this ROM training stage, the amount of appropriate dissipation to be prescribed in the constrained optimization step can be found to produce a stable ROM with mean TKE behavior close to that of the FOM.

Finally, a basis augmentation approach was introduced by Akkari *et al.* [21, 22] that use combination of L^2 POD modes and H_0^1 POD modes so that a small number of L^2 POD modes capture the dominant features of the flow and the H_0^1 POD modes containing small-scale features will address the lack of dissipation in the former modes. They have found nominal success in using this basis, using qualitative metrics. Although similar to our approach in the sense that they are both an augmentation scheme upon the standard L^2 POD basis, our method can introduce new features that were not present in the snapshots of the FOM solution by considering the dynamical system that drives the solution (i.e., the NSE). In this way, even if there is a lack of snapshots to reproduce the FOM behavior, ABM can help

improve the accuracy as demonstrated in 6.3.

Chapter 2

Model-Order Reduction for Thermal Fluid Analysis

High-performance computing and modern numerical algorithms have made high-fidelity fluid-thermal analysis tractable in geometries of ever increasing complexity. Despite continued advances in these areas, direct numerical (DNS), large-eddy (LES), and even unsteady Reynolds-averaged Navier-Stokes (uRANS) simulations of turbulent thermal transport remain too costly for routine analysis and design of thermal-hydraulic systems, where hundreds of cases must be considered. Reduced-order models (ROMs) offer a promising approach to leverage expensive high-fidelity simulations (referred to as full-order models or FOMs) by extracting from these detailed simulations low-dimensional dynamical systems that capture the principal features of the underlying flow fields [23, 24].

While FOMs for turbulent flows can require $\mathcal{N} = 10^7 - 10^{11}$ degrees-of-freedom, ROMs offer the potential of representing the flow dynamics governing the behavior of quantities of interest (QOIs) with only $N \approx 10-10^3$ basis functions. The amplitudes of these basis functions are typically evolved as a system of nonlinear ordinary differential equations (ODEs) or through other mechanisms, for example, based on machine learning (e.g., [25]), or even through linearized models such as the recently-developed Green's function approach of

Khodkar *et al.* [26]. For reasons of stability, Fick *et al.* [14] replace ODE integration with a sequence of constrained-minimization problems that keep the ROM within observed basin of attraction.

The overall objective of *parametric model-order reduction* (pMOR) is to build ROMs from FOMs that are generated over a small set of *anchor* points $\mathcal{P}_\alpha = \{p_1^*, \dots, p_m^*\}$ in the parameter space and to subsequently use the ROMs to inexpensively explore the parametric dependencies of the QOIs. To realize this goal, pMOR must address two problems: (i) the *reproduction problem*, in which a ROM accurately recovers QOIs for a turbulent flow simulation at a *particular* point in the parameter space (i.e., where the ROM and FOM are evaluated at the same p_j^*); and (ii) the *parametric problem*, in which the ROM provides estimates of QOIs at points $p \notin \mathcal{P}_\alpha$. To make the overall process efficient, the ROM should be equipped with error indicators to assess its fidelity at any given p , which allows one to optimally choose the anchor points p_j^* , thus minimizing the number of expensive FOMs required for effective parametric analysis [14]. We typically choose the anchor points by evaluating the error indicator over a set of *training points*, \mathcal{P}_τ and selecting as the next element of \mathcal{P}_α the point that maximizes the error indicator in \mathcal{P}_τ . In the present study, we do not consider error-indicated selection of \mathcal{P}_τ ; that topic is the focus of future work.

Here, we present some basic steps towards realizing pMOR for buoyancy-driven flows at modest Rayleigh numbers. Initially, we consider the use of proper-orthogonal decomposition (POD) bases in a Galerkin formulation of the momentum and energy equations for relatively low Rayleigh-numbers. We show that ROM reconstruction and pMOR are viable and that pMOR is significantly more accurate than parametric interpolation for nonlinear QOI examples that follow. We will consider the closure problem for high Reynolds number cases, where ROMs typically suffer from inadequate dissipation because of the lack of high-wavenumber functions in the POD approximation space in Chapter 3. There we consider two known stabilization strategies. The first is a Leray-type regularization [17, 27], in which we mollify the nonlinear advection term by regularizing the *advecting field* (only). The second is based

on the constrained-evolution approach of [14], in which the snapshot data (introduced below) is reused to set limits on the basis coefficients in order to keep the dynamics of the system near the observed attractor. We will demonstrate the relative success of these regularization approaches for two-dimensional regularized lid-driven cavity problem which is the subject of Fick *et al.* [14].

Before moving on to these examples, again in Chapter 3, we will focus on successful examples to show the promise of MOR for engineering applications.

2.1 Galerkin-Based ROM

Our point of departure is the Boussinesq approximation for buoyancy-driven flow,

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p = \nu \nabla^2 \mathbf{u} + \eta T \hat{\mathbf{g}}, \quad \nabla \cdot \mathbf{u} = 0, \quad (2.1)$$

$$\partial_t T + (\mathbf{u} \cdot \nabla) T = \kappa \nabla^2 T. \quad (2.2)$$

subject to appropriate Dirichlet or Neumann boundary conditions for the velocity, \mathbf{u} , and temperature, T . Here p is the pressure, $\hat{\mathbf{g}}$ is the unit vector in the (positive) vertical direction, and ν , η , and κ are problem parameters based on the nondimensionalization of the problem.

The ROM for the Boussinesq equations starts with a collection of K snapshots for velocity, $\mathbf{u}^k(\mathbf{x}) := \mathbf{u}(\mathbf{x}, t^k) - \mathbf{u}_b$ and for temperature, $T^k(\mathbf{x}) := T(\mathbf{x}, t^k) - T_b$ corresponding to numerical solutions of the full-order model (FOM) at well-separated timepoints, t^k , minus *base states*, $\mathbf{u}_b(\mathbf{x})$ and $T_b(\mathbf{x})$. The base states (typically time-averaged FOM solutions) satisfy any prescribed inhomogeneous boundary conditions, which are presently assumed to be time-independent.

The FOM is based on the spectral element method (SEM) in the open-source code, Nek5000, and uses the $\mathbb{P}_N - \mathbb{P}_{N-2}$ velocity-pressure coupling where the velocity and pressure are represented in different polynomial bases.¹

¹Nek5000 use two formulations $\mathbb{P}_N - \mathbb{P}_N$ [28–30] and $\mathbb{P}_N - \mathbb{P}_{N-2}$ [31, 32] of which the $\mathbb{P}_N - \mathbb{P}_{N-2}$

For any $\mathbf{u}^k(\mathbf{x})$, we have a corresponding vector of basis coefficients $\underline{\mathbf{u}}^k = [\mathbf{u}_1^k \dots \mathbf{u}_N^k]^T$ such that $\mathbf{u}^k(\mathbf{x}) = \sum_{j=1}^N \mathbf{u}_j^k \phi_j(\mathbf{x})$, with $\phi_j(\mathbf{x})$ the underlying spectral element basis functions spanning the FOM approximation space, $X_0^N \subset Y$. Because the SEM is nodal-based, each \mathbf{u}_j^k represents the three velocity components at gridpoint \mathbf{x}_j in the spectral element mesh at timepoint t^k . Similarly, the temperature is given by $T^k(\mathbf{x}) = \sum_{j=1}^N T_j^k \phi_j(\mathbf{x}) \in X_0^N \subset Y$. Here, $X^N \subset Y$ is the finite dimensional SEM approximation space spanned by $\{\phi_j(\mathbf{x})\}$. We refer to equations ?? and 1.4 for the weak formulation and the definition of spaces.

2.1.1 Proper Orthogonal Decomposition

For the velocity, we collect the snapshots into the a matrix $\mathbf{U}_K = [\underline{\mathbf{u}}^1 \dots \underline{\mathbf{u}}^K]$. From these, one forms the Gramian, $\mathbf{U} \in \mathbb{R}^{K \times K}$, with $\mathbf{U}_{k,k'} := (\mathbf{u}^k, \mathbf{u}^{k'})_{L^2}$, where $(\mathbf{v}, \mathbf{u})_{L^2} := \int_{\Omega} \mathbf{v} \cdot \mathbf{u} dV$ is the L^2 inner product. One could also consider the H^1 norm and H_0^1 semi-norm that have been considered in [16] and [14].

Following standard POD methodology, the basis functions $\{\underline{\zeta}_n\}$ for the ROM derive from the first N eigenmodes of \mathbf{U} ,

$$\mathbf{U} \underline{z}_k = \lambda_k \underline{z}_k, \quad \underline{z}_k \in \mathbb{R}^K, \quad \lambda_1 \geq \dots \geq \lambda_K \geq 0 \quad (2.3)$$

$$\underline{\zeta}_n := \mathbf{U}_K \underline{z}_n, \quad n = 1, \dots, N < K. \quad (2.4)$$

The continuous functions in the SEM space corresponding to $\underline{\zeta}_n$ are $\boldsymbol{\zeta}_n(\mathbf{x}) := \sum_{j=1}^N (\underline{\zeta}_n)_j \phi_j(\mathbf{x})$.

We perform a similar procedure for temperature, generating $\mathbf{T}_K := [\underline{T}^1 \dots \underline{T}^K]$ and associated Gramian, \mathbf{T} , having entries $\mathbf{T}_{k,k'} := (T^k, T^{k'})_{L^2}$. We then solve the $K \times K$ eigenproblem,

$$\mathbf{T} \underline{\tau}_k = \lambda_k \underline{\tau}_k, \quad \underline{\tau}_k \in \mathbb{R}^K, \quad \lambda_1 \geq \dots \geq \lambda_K \geq 0, \quad (2.5)$$

formulation is used for the ROM operator generation

and define

$$\underline{\theta}_n := \mathbf{T}_K \underline{\tau}_n, \quad n = 1, \dots, N < K. \quad (2.6)$$

The continuous functions in the SEM space corresponding to $\underline{\tau}_n$ are $\theta_n(\mathbf{x}) := \sum_{j=1}^N (\underline{\theta}_n)_j \phi_j(\mathbf{x})$.

Remark 1. We note that, the POD construction provides a set of basis vectors $Z := \{\underline{\zeta}_1, \dots, \underline{\zeta}_N\}$ that minimizes the average distance between Z and U_K in the chosen inner-product (here, $(\cdot, \cdot)_{L^2}$) for any rank- N subset of U_K . Thus, the motivation for POD is its approximation property, which uniformly distributes the error across the snapshot set. That property is not dependent on the subsequent choice of time-evolution of the ROM nor on any particular feature of the Navier-Stokes or Boussinesq equations. A major contribution of this thesis is that we derive, from the POD basis, an augmented basis set that greatly improves stability and convergence of the dynamics as generated by Galerkin ROM.

Remark 2. In the procedures that follow, we treat the velocity and temperature as independent state vectors. One could alternatively work with snapshot sets that combine the velocity and temperature into a single vector and thus form the Gramian from $Q_K = [\underline{q}^1 \dots \underline{q}^K]$, where $\underline{q}^k := [\mathbf{u}^k T^k]$. This approach reduces the online evaluation cost of the ROM by a factor of two because one has one-half the number of variables to track. The principal drawback of using a coupled system is that one then has no mechanism to ensure that the temperature is bounded by its extremal boundary values (where such bounds are meaningful).

2.1.2 Galerkin Formulation

With a given basis sets in hand, the Galerkin formulation follows by inserting the reduced-basis expansions,

$$\tilde{\mathbf{u}}(x, t) = \sum_{n=0}^N \boldsymbol{\zeta}_n(\mathbf{x}) u_n(t) \in Z_b \quad (2.7)$$

$$\tilde{T}(x, t) = \sum_{n=0}^N \theta_n(\mathbf{x}) T_n(t) \in \Theta_b, \quad (2.8)$$

into the weak form of (2.1)–(2.2). In order to set the boundary conditions, we have augmented the trial (approximation) spaces Z_0^N and Θ_0^N with the base states, $\boldsymbol{\zeta}_0 := \mathbf{u}_b$ and $\theta_0 := T_b$. The coefficients for these terms are prescribed: $u_0 \equiv T_0 \equiv 1$. The corresponding test spaces, $Z_0^N := \{\boldsymbol{\zeta}_i\}_{i=1}^N$ and $\Theta_0^N := \{\theta_i\}_{i=1}^N$, satisfy homogeneous boundary conditions, as is standard for Galerkin formulations.

In weak form, the problem can be stated as *Find* $(\tilde{\mathbf{u}}, \tilde{T}) \in Z_b^N \in X^N \in Y$, *s.t.*, $\forall (\mathbf{v}, S) \in Z_0^N \in X_0^N \in Y$,

$$\begin{aligned} & \int_{\Omega} \mathbf{v} \cdot \frac{d\tilde{\mathbf{u}}}{dt} dV + \nu \int_{\Omega} \nabla \mathbf{v} : \nabla \tilde{\mathbf{u}} dV \\ & = - \int_{\Omega} \mathbf{v} \cdot (\tilde{\mathbf{u}} \cdot \nabla) \tilde{\mathbf{u}} dV + \eta \int_{\Omega} \mathbf{v} \cdot (\tilde{T} \hat{\mathbf{g}}) dV, \end{aligned} \quad (2.9)$$

$$\begin{aligned} & \int_{\Omega} S \frac{d\tilde{T}}{dt} dV + \kappa \int_{\Omega} \nabla S \cdot \nabla \tilde{T} dV \\ & = - \int_{\Omega} S (\tilde{\mathbf{u}} \cdot \nabla) \tilde{T} dV + \kappa \int_{\partial\Omega_f} S \nabla \tilde{T} \cdot \hat{\mathbf{n}} dA. \end{aligned} \quad (2.10)$$

First, we remark that $\dim(Z_0^N) \ll \dim(X_0^N)$, meaning these reduced problems are much easier to solve numerically, and we also remark that the presence of the surface integral in (2.10) admits the possibility of an inhomogeneous surface flux $q_s := \kappa \nabla \tilde{T} \cdot \hat{\mathbf{n}}$ on some part of the domain boundary $\partial\Omega_f \subset \partial\Omega$. If $\boldsymbol{\zeta}_0$ and θ_0 satisfy the prescribed Dirichlet conditions, then $\tilde{\mathbf{u}}$ and \tilde{T} will satisfy these conditions as well. Note that the pressure term vanishes from

(2.9) as a result of integration-by-parts because the test functions are weakly divergence-free and they are all homogeneous on the boundary. The continuity equation of the Navier-Stokes equations is also removed since Z_b^N consists solely of weakly divergence-free basis functions.

For temporal discretization of (2.9)–(2.10), we use a semi-implicit scheme with k th-order backward differencing (BDF k) for the time derivative, implicit treatment of the negative-definite dissipation terms, and k th-order extrapolation (EXT k) of the advection and buoyancy terms. We typically use $k = 3$, to ensure that the imaginary eigenvalues associated with skew-symmetric advection operator are within the stability region of the BDF k /EXT k timestepper, as discussed in [33].

When fully discretized in space and time, we get the following systems for the ROM basis coefficients \underline{u}^n and \underline{T}^n at timepoint t^n ,

$$\begin{aligned} \left(\frac{\beta_0}{\Delta t} B_u + \nu A_u \right) \underline{u}^n &= - \sum_{i=1}^k \alpha_i [C_u(\underline{\bar{u}}) \underline{\bar{u}} - \eta B_{uT} \underline{\bar{T}}]^{n-i} \\ &\quad - B_u \sum_{i=1}^k \frac{\beta_i}{\Delta t} \underline{u}^{n-i} - \nu \underline{a}_{u,0}, \end{aligned} \quad (2.11)$$

$$\begin{aligned} \left(\frac{\beta_0}{\Delta t} B_T + \kappa A_T \right) \underline{T}^n &= - \sum_{i=1}^k \alpha_i [C_T(\underline{\bar{u}}) \underline{\bar{T}}]^{n-i} \\ &\quad - B_T \sum_{i=1}^k \frac{\beta_i}{\Delta t} \underline{T}^{n-i} - \kappa \underline{a}_{T,0} + \underline{q}_s. \end{aligned} \quad (2.12)$$

Here, $\underline{\bar{u}}$ and $\underline{\bar{T}}$ are the augmented basis vectors that include the base states for \underline{u} and \underline{T} ; α_i and β_i are the respective BDF k /EXT k coefficients (e.g., for $k=2$, $\alpha_1 = 2$, $\alpha_2 = -1$, $\beta_0 = \frac{3}{2}$, $\beta_1 = -\frac{4}{2}$, and $\beta_2 = \frac{1}{2}$). A , B , and C represent the respective stiffness, mass, and advection

operators, with entries

$$A_{u,ij} = \int_{\Omega} \nabla \zeta_i : \nabla \zeta_j dV \quad (2.13)$$

$$B_{u,ij} = \int_{\Omega} \zeta_i \cdot \zeta_j dV \quad (2.14)$$

$$C_{u,ikj} = \int_{\Omega} \zeta_i \cdot (\zeta_k \cdot \nabla) \zeta_j dV \quad (2.15)$$

$$A_{T,ij} = \int_{\Omega} \nabla \theta_i \cdot \nabla \theta_j dV \quad (2.16)$$

$$B_{T,ij} = \int_{\Omega} \theta_i \theta_j dV \quad (2.17)$$

$$C_{T,ikj} = \int_{\Omega} \theta_i \cdot (\zeta_k \cdot \nabla) \theta_j dV. \quad (2.18)$$

$$B_{uT,ij} = \int_{\Omega} \zeta_i \cdot \hat{\mathbf{g}} \theta_j dV \quad (2.19)$$

$$\underline{q}_{s,i} = \int_{\partial\Omega_f} \theta_i q_s dA \quad (2.20)$$

Remark 3. The computational cost of time-advancing (2.11) and (2.12) is dominated by the application of the third-order advection tensors, C_u and C_T , which requires $O(N^3)$ operations and memory references on each step. The remainder of the terms are $O(N^2)$ or less. Unfortunately, $O(N^3)$ is a very steep cost and prohibits practical consideration of, say, $N > 500$. While not considered in this work, strategies to mitigate this cost are of paramount concern. Instead of work reduction, ABM, to be introduced later, will reduce the number of modes by considering the dynamics (from NSE) in conjunction with the empirical basis from the snapshots.

We have applied the Galerkin-ROM formulation (2.11)–(2.12) for several cases, including the basic test problem of two-dimensional flow past a cylinder at Reynolds number $Re_D = 100$, for which $N \approx 20$ provides a sufficient number of modes for accurate reproduction of the flow field with the drag and lift history to four significant digits. As stated earlier, however, our objective is not reconstruction of the flow field. Rather, we aim to predict QOIs at parameter points p that are not in the set \mathcal{P}_α from which we draw our approximating

bases. In section 2.2, we illustrate use of the Galerkin ROM for both the reproduction for the aforementioned cylinder problem and pMOR problems in the case of a axisymmetric Rayleigh-Bénard convection.

2.2 Reproduction Problem

We first demonstrate reproduction using the POD-Galerkin ROM on the problem of 2D flow past a cylinder. This is a canonical test case for ROMs because of its robust and low-dimensional attractor, which is manifest as a von Karman vortex street for $Re = UD/\nu > 34.37$ [34]. The Reynolds number in our test case is $Re = 100$ and the domain is $\Omega = [-2.5 : 17]D \times [-5 : 5]D$, with the unit-diameter cylinder centered at $[0, 0]$.

We take the drag on the cylinder to be the QOI. Here, we describe only the pressure component, which is more challenging to compute than the viscous component because the pressure solution is not explicitly solved for. Pressure is often solved for in the ROM systems, especially when the QOI includes pressure contributions. Examples of this approach are Stabile *et al.* [35] and Bergmann *et al.* [36]. The challenge for our approach is to compute the contour integral of the pressure without explicitly solving for the pressure field in the ROM. We start from the Incompressible Navier-Stokes Equations:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u}, \quad (2.21)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (2.22)$$

Rearranging the terms, the pressure gradient is

$$\nabla p = -\frac{\partial \mathbf{u}}{\partial t} - (\mathbf{u} \cdot \nabla) \mathbf{u} + \nu \nabla^2 \mathbf{u} \quad (2.23)$$

Now, if we know the pressure gradient along the cylinder surface Γ , we can find the force

exerted by pressure via the following Fourier representation for the pressure

$$p(s) = \sum_{k=1}^{\infty} (A_k \cos(2\pi ks) + B_k \sin(2\pi ks)), \quad s \in [0, 1] \quad (2.24)$$

where s is the parameterization of Γ . Then, the derivative along Γ is:

$$\frac{dp}{ds} = \sum_{k=1}^{\infty} (-2\pi k A_k \sin(2\pi ks) + 2\pi k B_k \cos(2\pi ks)) \quad (2.25)$$

If $\frac{dp}{ds}(s)$ is known, we can find the coefficients by evaluating integrals:

$$A_k = -\frac{1}{k\pi} \int_0^1 \sin(2\pi ks) \frac{dp}{ds}(s) ds, \quad (2.26)$$

$$B_k = \frac{1}{k\pi} \int_0^1 \cos(2\pi ks) \frac{dp}{ds}(s) ds \quad (2.27)$$

Once A_k and B_k are known, we can evaluate the contour integral via the above expression.

Because ∇p can be evaluated from the momentum equation, $\frac{dp}{ds} = \nabla p(s) \cdot \hat{\mathbf{t}}(s)$ can be established, where $\hat{\mathbf{t}}$ is the unit tangent vector along Γ .

$$\tilde{\mathbf{F}}(\nabla p) = - \oint_{\Gamma} p \, d\mathbf{A} \quad (2.28)$$

$$= -L \int_0^1 p(s) \hat{\mathbf{n}}(s) \, ds \quad (2.29)$$

$$= -L \int_0^1 \sum_{k=1}^{\infty} (A_k \cos(2\pi ks) + B_k \sin(2\pi ks)) \hat{\mathbf{n}}(s) \, ds \quad (2.30)$$

$$= -L \sum_{k=1}^{\infty} \left(A_k \int_0^1 \cos(2\pi ks) \hat{\mathbf{n}}(s) \, ds + B_k \int_0^1 \sin(2\pi ks) \hat{\mathbf{n}}(s) \, ds \right) \quad (2.31)$$

$$= -L \sum_{k=1}^{\infty} (A_k \mathbf{I}_{k,1} + B_k \mathbf{I}_{k,2}) \quad (2.32)$$

where L is the arc-length of Γ . $\mathbf{I}_{k,1}$ and $\mathbf{I}_{k,2}$ must be pre-computed once. Therefore, if we find the Fourier coefficients corresponding to ∇p , we find the force on the cylinder due to the pressure. But because we can calculate the pressure gradient field from the velocity field, we

can calculate the pressure drag without knowing the pressure solution i.e.,

$$\mathbf{F} = \tilde{\mathbf{F}}(\nabla p) \approx -L \sum_{k=1}^N (A_k \mathbf{I}_{k,1} + B_k \mathbf{I}_{k,2}), \quad (2.33)$$

$$A_k = -\frac{1}{k\pi} \int_0^1 \sin(2\pi ks) \left(-\frac{\partial \mathbf{u}}{\partial t} - (\mathbf{u} \cdot \nabla) \mathbf{u} + \nu \nabla^2 \mathbf{u} \right) \cdot \hat{\mathbf{t}} ds, \quad (2.34)$$

$$B_k = \frac{1}{k\pi} \int_0^1 \cos(2\pi ks) \left(-\frac{\partial \mathbf{u}}{\partial t} - (\mathbf{u} \cdot \nabla) \mathbf{u} + \nu \nabla^2 \mathbf{u} \right) \cdot \hat{\mathbf{t}} ds \quad (2.35)$$

For the specific case of flow past a cylinder with $\mathbf{u} = \mathbf{0}$ on the cylinder surface Γ , the problem can be further simplified

$$\nabla p|_{\Gamma} = \nu \nabla^2 \mathbf{u}|_{\Gamma} \quad (2.36)$$

$$\Rightarrow \mathbf{F} = -L(A_k \mathbf{I}_{k,1} + B_k \mathbf{I}_{k,2}) \quad (2.37)$$

$$= -2\pi r \langle A_k, B_k \rangle, \quad (2.38)$$

$$A_k = -\frac{\nu}{k\pi} \int_0^1 \sin(2\pi ks) (\nabla^2 \mathbf{u}) \cdot \hat{\mathbf{t}} ds, \quad (2.39)$$

$$B_k = \frac{\nu}{k\pi} \int_0^1 \cos(2\pi ks) (\nabla^2 \mathbf{u}) \cdot \hat{\mathbf{t}} ds \quad (2.40)$$

After application of the above methodology, we obtain an accurate behavior of the pressure component of drag shown in Fig. 2.1, where the projection of the initial condition for the developed FOM was used for the ROM. The error for this ROM reproduction is shown in Fig. 2.2.




Figure 2.1: Circular cylinder pressure drag force comparison between FOM and ROM.




Figure 2.2: Circular cylinder ROM pressure drag force error.

The experimental setup was as follows. Starting with a fully-developed FOM, 200 snapshots (64-bit data) were collected over 200 convective time units (D/U). The tolerance for the pressure solve is 10^{-7} and the tolerance for the velocity solve is 10^{-9} . The number of POD modes in the basis is $N = 20$. The ROM initial condition is based on the POD projection of the initial condition of the FOM. We see that the ROM accurately captures the phase and period of the FOM. Moreover the ROM QOI accurately tracks the pressure drag of the FOM. The relative error of the ROM compared to the FOM in the mean pressure drag value is 4.1×10^{-7} and the relative error in variance is 2.8×10^{-3} . The maximum relative error (based on the mean value) between the two data set is 3.1×10^{-5} .

To demonstrate this methodology on non-trivial cylinder geometry, we will consider

the flow past an elliptic cylinder. Using the same mesh topology, the domain has been stretched in the stream-wise direction by factor 2. Other specifications of the experiments such as the the time-step size, tolerances, number of snapshots, and the snapshot collection time-window are consistent with those of the circular cylinder case. For this problem we see similar success of capturing the pressure drag behavior without solving for the pressure. Figure 2.3 shows the pressure drag history and Fig 2.4 shows the error. For this problem the relative error in the mean drag value is 1.4×10^{-5} and the relative error in the variance is 8.6×10^{-4} . The maximum relative error (based on the mean value) between the two data set is 2.4×10^{-5} . Although these two examples are a trivial cases for MOR, it demonstrate a significant reduction in the degrees of freedom.




Figure 2.3: Elliptic cylinder pressure drag force comparison between FOM and ROM.




Figure 2.4: Elliptic cylinder ROM pressure drag force error.

2.3 Parametric Model-Order Reduction

Here, we demonstrate that the ROM techniques of the preceding section can be effective for the ultimate goal of parametric model-order reduction. Given that the parameters and QOIs for thermal-fluids applications are highly case specific, we introduce the process in the context of a particular model problem. Specifically, we consider the temporal-behavior of the Nusselt number in axisymmetric Rayleigh-Bénard convection, which was studied by Barkley & Tuckerman [37, 38].

2.3.1 Axisymmetric Rayleigh-Bénard Problem

The thermal-fluids problem is the axisymmetric Rayleigh-Bénard configuration depicted in Fig. 2.5(a), which was studied by Tuckerman & Barkley [37, 38]. The problem is parameterized by $\epsilon = \frac{\text{Ra} - \text{Ra}_c}{\text{Ra}_c}$, where $\text{Ra}_c = 1734$ is the critical Rayleigh number. The 2D axisymmetric domain has an aspect ratio of $\Gamma = 5$, shown in Fig. 2.5(a). For $\epsilon > 1.3843$, traveling waves move towards the centerline axis with a period that depends on ϵ . We perform FOM calculations at two anchor points, $\epsilon = 1.6$ and $\epsilon = 2.6$, from which we collect snapshots (full flow/temperature fields). We apply proper orthogonal decomposition (POD) to the snapshot sets from each of the FOMs and use 20 POD modes from each to form a reduced-order subspace Z^N comprised of $N = 40$ basis functions. These modes are used in the weak- (Galerkin-) formulation of the governing equations, where the solution is restricted to $Z^N \subset Y$. The low-dimensional ROM is able to capture short- and long-time behavior as shown in the Nusselt number reproduction traces in Fig. 2.5(b). Moreover, as shown in Fig. 2.5(c), the ROM is able to accurately predict the period of the traveling wave solutions both inside and outside the ϵ range spanned by the anchor points. Note that as $\epsilon \rightarrow 1.3843$, the period goes to infinity and FOM simulations near this limit become intractable. The ROM, however, is able to predict this critical value of ϵ to within a few percent.

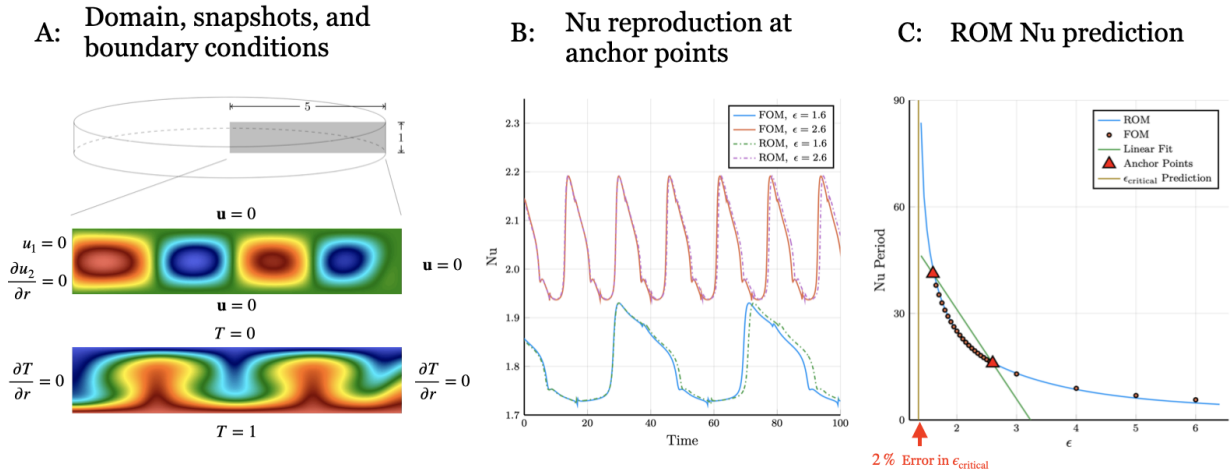


Figure 2.5: ROM application to axisymmetric Rayleigh-Bénard

The spatial domain is a low aspect-ratio cylinder and, following [37], the problem is cast in cylindrical coordinates with $[r, z] \in [0, R] \times [0, H]$ ($R = 5$, $H = 1$), so the flow is effectively two-dimensional. For velocity, we have $\mathbf{u} = 0$ on the upper, lower, and side walls. For temperature, we have $T = 1 - z$ on those surfaces. At $r = 0$, the radial component of the velocity is zero while the vertical component and temperature have homogeneous Neumann conditions. The problem is nondimensionalized with a viscous timescale, resulting in parameters $\nu = \text{Pr}$, $\eta = \text{Ra Pr}$, and $\kappa = 1$, where Ra is the Rayleigh number and Pr is the Prandtl number.

Above a critical Rayleigh number, Ra_c , this flow exhibits an unsteady phenomenon in which vortex rings formed at the outer edge of the cylinder push the interior vortex rings inwards and squeeze out the one at the center. The orientation of the vortices thus periodically flips, with the net effect being that a traveling wave of inward moving rings is observed. For $\text{Ra} > \text{Ra}_c$, the period of this phenomenon increases as $\text{Ra} \rightarrow \text{Ra}_c$. Our goal is to reproduce the FOM results as a function of $\epsilon = (\text{Ra} - \text{Ra}_c)/\text{Ra}_c$ over a range of Rayleigh numbers by performing FOM simulations at only one or two Rayleigh numbers. We take as our parametric domain $\epsilon \in [1.60, 2.60]$ and as our QOI the transition period, indicated by fluctuations in the Nusselt number. By performing this study, we wish to establish the feasibility of pMOR for thermally-driven flows.

Our FOM is based on the axisymmetric formulation in Nek5000 using a 4×1 array ($r \times h$) of 15th-order spectral elements. The FOM is based on 3rd-order semi-implicit timestepping with timestep size $\Delta t = 4 \times 10^{-4}$. The physical parameters are $\text{Ra}_c = 1734$, $\text{Pr} = 10$, domain radius $R = 5$, and height $H = 1$. From the FOM, the Nusselt number period is 41.200 for $\epsilon = 1.60$ and 15.984 for $\epsilon = 2.60$, in accord with the results in [37]. One snapshot from the problem for $\epsilon = 2.60$ is shown in Fig. 2.5 (a).

2.3.2 pMOR Setup

Our pMOR set for the axisymmetric Rayleigh-Bénard problem follows the p -refinement algorithm described in [14] in which we combine basis sets from two points in the ϵ -parameter space to form one large approximation space.

We first start by running FOM simulations using $\epsilon = 1.60$ and $\epsilon = 2.60$ to obtain two sets of snapshots, taken over 200 Time-Units and spaced equally by 0.1 Time-Units. Then, for each set, we form the velocity Gramian matrix by taking the inner product of the velocity snapshots and form the temperature Gramian matrix by taking the inner product of the temperature snapshots. For each set, using these Gramian matrices, we can obtain 20 dominant velocity modes and 20 dominant temperature modes. We use a total of 40 combined velocity modes and 40 combined temperature modes from this procedure as the basis functions for the Galerkin projection. Once the basis functions for the velocity and temperature are generated, the various operators in the systems can be formed as shown in (2.13)–(2.20). (Note that the integrands must be modified to account for the axisymmetric coordinates. This modification is straightforward in Nek5000 because the operator interfaces are independent of the number of space dimensions or type of domain.)

Although the physics is fully coupled, semi-implicit time-stepping of the ROM allows the velocity and temperature fields to be advanced without implicit coupling. Thus, the third-order tensors C_T and C_u are created with basis functions for each field; resulting in 40^3 entries² in each of the advection tensors. In short, although we are using 80 total velocity and temperature basis functions, the total cost of evaluating the convection for both equations is 2×40^3 multiplication operations, instead of 80^3 multiplications.

With this setup, we adjust the parameters in (2.11)–(2.12) to run the ROM for $\epsilon \in \{1.6, 1.65, \dots, 2.6\}$. The endpoint data, $\epsilon = 1.6$ and 2.6 , are used to generate the parametric model while the interior points are used for verification of the pMOR procedure. Presently,

²The full C_T contains $40 \times 40 \times 41$ elements, with the extra elements arising from the inhomogeneous boundary term

we use the same timestep size for the ROM as for the FOM, but this choice is not strictly necessary. For each ϵ , we obtain Nusselt number vs. time behavior after first having let the flow settle to a steady-periodic state for 100 time units. We extract the period from the subsequent 100-unit time interval.

2.3.3 Computation of the Nusselt Number

The Nusselt number for this problem is defined as:

$$\text{Nu} = \frac{2H}{R^2 \Delta T} \int_0^R T_z(r, 0, t) r dr, \quad (2.41)$$

where $\Delta T = 1$ is the temperature differential between the upper and lower surfaces. Calculation of Nu by reconstructing the FOM temperature field from the basis coefficients will result in a heavy computational cost of $O(\mathcal{N} \times N)$ where \mathcal{N} is the total number of degrees of freedom for the FOM.

Fortunately, because Nu is a linear functional in T , it can be calculated indirectly by computing the contribution from each basis function to the sum in the offline (FOM) phase without reconstruction of the solution at each timestep in the online (ROM) phase. Specifically,

$$\text{Nu}(t) = \frac{2H}{R^2 \Delta T} \int_0^R \sum_{i=0}^N T_i(t) \theta_{i,z}(r, 0) r dr \quad (2.42)$$

$$= \frac{2H}{R^2 \Delta T} \sum_{i=0}^N T_i(t) \int_0^R \theta_{i,z}(r, 0) r dr \quad (2.43)$$

$$= \frac{2H}{R^2 \Delta T} \sum_{i=0}^N T_i(t) I_i \quad (2.44)$$

We calculate the I_i in the offline phase so we can determine the Nusselt number in the online phase without any knowledge of the FOM (i.e., SEM) basis. Thus, the computational cost of evaluating the Nusselt number at each timestep in the ROM is only $O(N)$.

To calculate the period of the Nusselt number, we subtract the mean value from the data and measure the average time between successive downward zero-crossings over the final 100 time units. We note that the period is *not* a linear functional of the solution but is nonetheless accessible via an accurate ROM.

2.3.4 pMOR Results

Figures 2.6-2.8 show the Nusselt number histories for the FOMs and ROMs at $\epsilon = 1.60, 2.35,$ and 2.60 . The 1.60 and 2.60 cases correspond to anchor points in the pMOR study—the ROM bases derive directly from FOMs at these parameter values. The 2.35 case derives from projection onto the 40×2 basis vectors and is thus truly the result of pMOR-driven analysis. We remark that, even at the anchor points (e.g., $\epsilon = 1.60$), the histories are not perfect. This discrepancy arises because the time-dependent ROM is not a perfect reconstruction of the FOM. Despite the variation in period, the ROM is able to reproduce both short and long timescale variations in Nu with reasonably high fidelity, as evidenced by the small kinks in the Nu history curves in Fig. 2.5 (b) and Fig. 2.9.

In Fig. 2.9, we see good agreement between the FOM data and the ROM data. The largest error is found at $\epsilon = 1.60$, with a difference of 1.08 (i.e., $\approx 2.6\%$ relative error). Thus, for this problem, instead of running a FOM simulation for all 21 points, we can obtain fairly accurate results by running 2 FOM simulations and running pMOR on all the points, resulting in an accurate representation of the parameter dependence that is shown in Fig. 2.5 (c). Furthermore, if a period value for an additional ϵ within the range is desired, the ROM, without any modification, can obtain that value fairly accurately.

We note that the potential value of the pMOR procedure is clearly evident in this example. If one only had knowledge of the period at the points $\epsilon = 1.6$ and 2.6 , there would be little recourse but to estimate the periods at interior points through linear interpolation, which leads to $\approx 20\%$ overestimate of the period at $\epsilon \approx 2$. The low-rank dynamical system of the ROM does a remarkably good job of improving upon this estimate.




Figure 2.6: Nusselt number behavior for $\epsilon = 1.60$




Figure 2.7: Nusselt number behavior for $\epsilon = 2.35$




Figure 2.8: Nusselt number behavior for $\epsilon = 2.60$




Figure 2.9: Period comparison of FOM and ROM using two anchor points at $\epsilon = 1.60$ and 2.60

2.4 Discussion

In this section, we have presented several technical components towards the development of a reduced-order modeling (ROM) capability that aims to harness the output of high-fidelity simulations on leadership computing facilities for thermal-hydraulic design studies.

For the *parametric problem*, we have presented results for a two-anchorpoint model-order reduction procedure applied to the problem of buoyancy-driven flows. We have demonstrated that the technique is capable of accurately predicting a challenging QOI—the period of roll-reversal—as a function of Rayleigh number for the highly nonlinear problem of Rayleigh–Bénard convection in a low aspect-ratio cylindrical domain.

Despite the success of the methods shown here, several developments remain to make pMOR viable for engineering design. The first of these is a need to incorporate more (or better) modal information. Presently, the $O(N^3)$ costs in the ROM reconstruction phase prohibit consideration of N much larger than 200. There are several possible mitigation strategies. One idea is the discrete empirical interpolation method [5–7], which leads to low-rank approximations of the advection operator. For example, one can split the velocity as $\mathbf{u} = \mathbf{U} + \mathbf{u}'$ and then account for $\mathbf{U} \cdot \nabla \mathbf{U}$, $\mathbf{U} \cdot \nabla \mathbf{u}'$, and $\mathbf{u}' \cdot \nabla \mathbf{U}$, but model the $\mathbf{u}' \cdot \nabla \mathbf{u}'$ term (again giving rise to the classic turbulence closure problem). In our discrete model, this separation would lead to modeling the convective term by a interpolation function on the

velocity coefficients in equation (2.11). (In some sense, sharp-cut-off Leray regularization is an example of this approach.) Another approach is to develop ROM and pMOR strategies that are optimized for the quantities of interest. In this work, however, we will introduce ABM in section 4 that reduced the N required for ROM.

As indicated in the introduction, error indicators are of paramount importance for efficient pMOR design tools as they allow one to optimally select anchor points. We have initiated development of error indicators following the approach described in [14]. From the results of Section 5, it seems clear that error indicators that focus on the QOIs will be important for overall efficiency. Another concern is that the p -greedy pMOR strategy leads to an increase in N . In [14], an h -greedy strategy is proposed in which one *changes* the approximation space, rather than *augmenting* the space, as the parameter p moves through the range of anchor space, \mathcal{P}_α .

2.5 Efficient Parallel ROM Generation

The ROM/pMOR procedure requires a significant amount of software above and beyond that of a standard code used for the full-order models. Many of the offline steps, such as forming the Gramian, require manipulation of the FOM vectors of length $\mathcal{N} \gg N$, which are generally not practical on a laptop. On the other hand, it is inconvenient if one is forced to use a supercomputer to perform all of the requisite offline tasks. It is therefore desirable to have a set of memory-efficient offline tools that could be run on a modest-sized workstation.

To get a sense of the offline-online procedure in which the FOM snapshots are produced and processed (offline) and ROM advancement (online), we have gathered the timings for a pipe flow ($\text{Re} = 5300$) FOM-ROM in the style of Fick *et al.* [14] in Table 2.1. Here, the FOM snapshot generation (64 compute-nodes) and post-processing ROM generation (16 compute-nodes) were performed on the OLCF Summit supercomputer (CPU-only), and the ROM was run on a workstation with an Intel Xeon E5-2620 CPU with 2 threads with ROM

simulation time being 10 times longer than the FOM.

DNS + snapshot acquisition ($\Delta t = 5 \times 10^{-4}$, $K = 1000$)	6658.46s
POD space construction ($N = 200$, $K = 1000$)	340.60s
Construction of Galerkin structures ($N = 200$)	1550.48s
POD-Gal: online evaluation ($N = 60$)	200.57s

Table 2.1: FOM & ROM Timings for pipe flow (Re = 5300)

Although the process for generating the ROM operator is straightforward when run on the same machine that produced the snapshots, it lacks the portability property that is attractive for pre-existing snapshots sets. Therefore, we propose that the data be processed as they are being read in. This way, the snapshots and the POD basis functions need not be fully stored in the same instance. In order to produce the correct inner-products, we break them up into subdomains.

$$G_{ij}^A = \sum_{e=1}^{N_e} \int_{\Omega^e} \nabla \cdot \mathbf{u}(\mathbf{x}, \tau^i) : \nabla \cdot \mathbf{u}(\mathbf{x}, \tau^j) dV \quad (2.45)$$

$$G_{ij}^B = \sum_{e=1}^{N_e} \int_{\Omega^e} \mathbf{u}(\mathbf{x}, \tau^i) \cdot \mathbf{u}(\mathbf{x}, \tau^j) dV \quad (2.46)$$

$$G_{ijk}^C = \sum_{e=1}^{N_e} \int_{\Omega^e} \mathbf{u}(\mathbf{x}, \tau^i) \cdot (\mathbf{u}(\mathbf{x}, \tau^k) \cdot \nabla) \mathbf{u}(\mathbf{x}, \tau^j) dV \quad (2.47)$$

Then, these inner-products can be evaluated by accumulating the integrals amongst the subdomains. In practice the subdomains correspond to the spectral elements in the domain. Figure 2.10 shows the reading process in more detail: we read a full-width block of the snapshot matrix and perform a transpose so that each mpi-rank has access to a subdomain data of all the snapshots. Thus, the inner-products can be performed independently in each rank until the very end of the reading loop at which the inner-products are summed together to produce the full G^A , G^B , and G^C . The notation used to illustrate the process are:

$$N_s := \text{Number of Snapshots}, \quad (2.48)$$

$$N_e := \text{Number of Total Elements}, \quad (2.49)$$

$$N'_e := \text{Number of Elements to be Read Incrementally}, \quad (2.50)$$

$$N_p := \text{Number of MPI Ranks}. \quad (2.51)$$

Normally, in Nek5000, the elements in the domain are distributed evenly throughout the mpi-ranks. To evaluate the inner-products, Nek5000 will load the snapshot fields and every mpi-rank will have direct access to all the snapshot solutions in it's own subdomain. By integrating the inner-products locally in the subdomain and summing the local contributions, the total integral value can be evaluated. This method is efficient in the use of computational resources, but requires that the information of all the snapshots ($N_e \times N_s$) be stored in memory at the same time.

In our proposed approach, a subdomain's worth of data ($N'_e \times N_s$) is read that can be stored to the memory of a workstation, where each mpi-ranks can load a chunk of data $N'_e \times (N_s/N_p)$ which is contiguous for each snapshot. Then, we transpose the data so each mpi-rank has a sub-subdomain's worth of field data for every snapshot. This way, the inner-products between different snapshots can be evaluated and accumulated internally in the mpi-ranks then at the end of the process, the global summation is performed. This does not significantly increase the total amount of computational work and IO (because the read-time of the snapshots dominate the transpose operation by crystal router in gslib [39, 40]) of all mpi-ranks. Figure 2.10 shows how this process of evaluating the inner-products can be performed by limiting the communication between the MPI ranks.




Figure 2.10: Demonstration of data transpose

Once the inner-products are evaluated, we can generate the ROM operators by applying the W matrix, whose columns are the eigenvectors of G^B to the inner-products:

$$A = Z^T A_{\mathcal{N}} Z = W^T (U^T A_{\mathcal{N}} U) W = W^T G^A W \quad (2.52)$$

Likewise,

$$B = W^T G^B W \quad (2.53)$$

$$C = W^T G^C(W) W \quad (2.54)$$

Following a similar procedure, QOIs that are linear, bilinear, or trilinear forms of the fields can be evaluated as a function of the coefficients of the POD basis.

Overall, this efficient process (in terms of memory footprint) allows a user to apply the POD-ROM methodology to existing high-fidelity snapshot data without needing access to HPC centers as long as the total computational work can be feasibly be performed on a local workstation or other accessible systems. This process significantly reduces the memory footprint of the ROM operator generation step while only slightly increasing the total computational cost and inter-rank IO (not disk IO).

Chapter 3

ROM Stabilization

A significant issue with the standard Galerkin-based ROM at higher Reynolds numbers is that (2.11)–(2.12) can deviate substantially from the projection of the FOM behavior on to the reduced basis due to the presence of potentially spurious attractors. This unstable behavior is an artifact of modal truncation (often manifest as energy pile-up in the high wavenumber modal coefficients) and relates to the well-known closure problem in LES and other turbulence-modeling scenarios [14, 41]. The smoothness of the approximating modes $\{\zeta_n\}$ prevents generation of adequate dissipation (through $\nu(\zeta_j, \zeta_j)_A$) unless the corresponding coefficients have large amplitude. One can potentially resolve this issue by increasing N to the point where the reduced bases include modes capable of dissipating the energy. For turbulent flows, however, the number of required modes can be quite high, particularly in light of the $O(N^3)$ costs associated with the advection operators C_u and C_T . Hence, for engineering applications, some sort of closure model is typically required to dissipate energy or otherwise stabilize the ROM.

In this chapter we consider two stabilization approaches introduced by other authors. The first is a relatively simple Leray-type regularization introduced in the ROM context by Wells *et al.* [17]. The second is the constraint-based ROM of Fick and co-workers [14]. We first introduce both methodologies, then show an application of these stabilization methods

on a ROM example. Further examples will be presented in Chapter 5 in comparison with augmented basis method introduced in Chapter 4.

3.1 Leray Regularization

For the Navier-Stokes equations, Leray regularization amounts to filtering the *advecting* field, such that the momentum equation reads

$$\frac{\partial \mathbf{u}}{\partial t} + (\bar{\mathbf{u}} \cdot \nabla) \mathbf{u} + \nabla p = \nu \nabla^2 \mathbf{u} + \eta T \hat{\mathbf{g}}, \quad (3.1)$$

where $\bar{\mathbf{u}} = F(\mathbf{u})$ is the result of some sort of smoothing or regularization of the current velocity field. As noted by Guermond *et al.* [42], even just a small amount of regularization suffices to make gains in proving existence and uniqueness of the solution in the continuous case, a task that remains insurmountable for the full Navier-Stokes equations. Thus, Leray regularization is of interest both from a numerical (and physical) stabilization viewpoint and from a theoretical perspective.

Guermond *et al.* [27] discuss a Helmholtz filter of the form

$$\bar{\mathbf{u}} = (1 - \delta^2 \nabla^2)^{-1} \mathbf{u} \quad (3.2)$$

which is a 2nd-order filter that suppresses modes of wavenumber $k > 1/\delta$ at a rate $\sim O(k^2)$ and which can be implemented in general geometries (e.g., [43]). For Fourier-based discretizations, filtering is particularly simple as one can simply scale the wavenumber coefficients, $\hat{\bar{\mathbf{u}}}_k = \sigma_k \hat{\mathbf{u}}_k$, where $\sigma_k \rightarrow 1$ as $k \rightarrow 0$ and $\sigma_k \rightarrow 0$ as $k \rightarrow \infty$.

For the Galerkin-based ROM, the modification equivalent to (3.1) would be to replace $\tilde{\mathbf{u}}$ by $\bar{\mathbf{u}}$ in the energy equation (2.10) and to similarly replace the advector in the momentum

equation (2.9) to yield

$$\int_{\Omega} \mathbf{v} \cdot \frac{d\tilde{\mathbf{u}}}{dt} dV + \nu \int_{\Omega} \nabla \mathbf{v} : \nabla \tilde{\mathbf{u}} dV = - \int_{\Omega} \mathbf{v} \cdot (\tilde{\mathbf{u}} \cdot \nabla) \tilde{\mathbf{u}} dV + \eta \int_{\Omega} \mathbf{v} \cdot (\tilde{T}\hat{\mathbf{g}}) dV. \quad (3.3)$$

Wells *et al.* [17] consider several possible filters to map $\tilde{\mathbf{u}}$ to $\bar{\mathbf{u}}$. The simplest one, suggested by one of the authors, is to treat the POD modes as one would do with Fourier bases (personal communication, 2019). That is, one can simply damp out the amplitudes of the higher modes when constructing $\bar{\mathbf{u}}$ in (3.3). In the following Leray-ROM examples, we do precisely that. We use a sharp cut-off filter in which we set a few of the highest-numbered basis coefficients to zero when constructing the advector, $\bar{\mathbf{u}}$. Notice that this approach, while blunt, also has the advantage of reducing the leading-order $O(N^3)$ cost of evaluating the advection terms. As seen in the examples at the end of this section, simple Leray regularization is remarkably effective and merits further development. One improvement is to use the differential filter (3.2), which is readily implemented in the POD-ROM by using (2.13) and (2.14). This approach was shown in [17] to be superior to simple modal damping for the example of three-dimensional flow past a cylinder at $Re = 1000$.

3.2 Constrained Optimization

Fick *et al.* [14] proposed a novel stabilization procedure that ensures that the trajectory of the ROM stays bounded within the range of values observed in the snapshot space of the FOM. In this way, if there is an instability due to a lack of sufficient dissipation in the ROM, this constraint keeps the ROM trajectory closer to the FOM trajectory.

To introduce the method, we initially consider only the ROM momentum update (2.12), which can be written as

$$H\underline{\mathbf{u}}^n = \underline{\mathbf{b}}^n, \quad (3.4)$$

where $H = \frac{\beta_0}{\Delta t} B_u + \nu A_u$ and \underline{b}^n is the right-hand side of (2.11). Rather than solving (3.4), Fick *et al.* propose to solve the constrained minimization problem,

$$\underline{u}^n = \underset{v \in \mathbb{R}^N}{\operatorname{argmin}} \|H\underline{u} - \underline{b}\|_2 \quad \text{s.t. } m_i < u_i < M_i \quad (3.5)$$

where the constraint bounds are $m_i = \min_{\mathbf{u}_j \in U_K} (\zeta_i, \mathbf{u}_j)$ and $M_i = \max_{\mathbf{u}_j \in U_K} (\zeta_i, \mathbf{u}_j)$. These bounds amount to the ranges observed in the snapshot space. Hence, (3.5) is making further use of the snapshot information.

The minimization problem must be solved at each timestep, which can be done, for example, using or standard optimization techniques such as BFGS with barriers to enforce the inequality constraints. While the resulting ROM is still much faster than the equivalent FOM, the cost of this approach is relatively high compared to the standard POD-Galerkin ROM as it requires multiple solutions in the SPD H -matrix for each BFGS iteration. (BFGS has updates because of the nonlinear barrier functions.)

The alternative procedure described in Kaneko *et al.* [4] seeks to transform the coordinate system such that the evaluation of $f(\underline{u})$ becomes less costly. That is, we transform the system such that the forward operator is diagonal, reducing the evaluation cost from $O(N^2)$ to $O(N)$ for each iteration. While this is much more efficient, especially for a large N , the non-linear iterative solver has undesirable characteristics such as variable number of iterations, higher cost than a standard solve.

A further simplification of the above transformation can be made by understanding the diagonalization as re-orientation of the principal axis of the energy bowl aligned with the coordinates of the ROM system. With this understanding, we can simply solve the discrete system in a transformed coordinates than apply the constraints derived in the new coordinates. This can be done because the energy isocontours are tangent to the bounding box on the principal axis. With this new approach, we can further reduce the cost to $O(N^2)$ per timestep, applying the transformation to and from the old coordinates to the new coordinates, which

is less expensive than a single solve of the original system dominated by $O(N^3)$ convective evaluation. With the BFGS algorithm for the constrained optimization, the cost per iteration of the nonlinear optimization is $O(N^2)$ so this reduction in cost is significant as well as reducing the need to set tolerances and other parameters for the optimization step.

3.3 Results for Stabilized Methods

Here, we present results for two examples in which we compare stabilized and unstabilized methods. First, we reproduce work done by Fick *et al.* [14] with the Lid-Driven Cavity Problem. Then we add the advection-diffusion equation to the NSE to show the effect of the constrained optimization on a relevant QOI which is the Nusselt number. We also apply Leray regularization to compare the effectiveness of the two stabilization techniques.

Following [14], we consider the following unsteady lid-driven cavity problem using parameter $\nu = \frac{1}{\text{Re}}$ and $\eta = 0$ subject to steady boundary conditions.

$$\begin{cases} \mathbf{u} = \mathbf{u}_b & \text{on } \Gamma_{top} \times \mathcal{R}_+, \\ \mathbf{u} = 0 & \text{on } \partial\Omega \setminus \Gamma_{top} \times \mathcal{R}_+, \\ \mathbf{u} = 0 & \text{on } \Omega \times \{0\}, \end{cases} \quad (3.6)$$

where $\mathbf{u} : \Omega \times \mathcal{R}_+ \rightarrow \mathcal{R}^2$ is a two-dimensional vector field, $\Omega = [-1, 1]^2$, $\Gamma_{top} = \{\mathbf{x} \in \bar{\Omega} : x_2 = 1\}$, the Dirichlet datum is given by

$$\mathbf{u}_b(\mathbf{x}) = \begin{bmatrix} (1 - x_1^2)^2 \\ 0 \end{bmatrix}. \quad (3.7)$$

The FOM results were produced using Nek5000 [44] with 256 spectral elements of order $N = 7$ at $\text{Re} = 30,000$. In order to generate the ROM solution, we collect 500 snapshots at the sampling times $\{t_s^k = 500 + \Delta t_s k\}_{k=1}^K$ with $\Delta t_s = 1$ and $K = 500$, where the data are




Figure 3.1: Velocity magnitude comparison at $t = 1740$ between the FOM, and the Galerkin, Leray, and constrained ROMs with $N = 20$.

taken from the statistically steady state solution. For the lid-driven cavity problem, we also apply the Leray regularization for the ROM. The velocity magnitude at time $t = 1740$ is shown in Fig. 3.1 for the FOM, and for the Galerkin, Leray, and constrained ROMs, each using $N = 20$.

Figure 3.2 shows the sample mean and variance for the FOM, Galerkin ROM, Leray ROM, and constrained ROM. We observe that the results for the Leray regularization and the constrained ROM are superior to those of the Galerkin ROM.

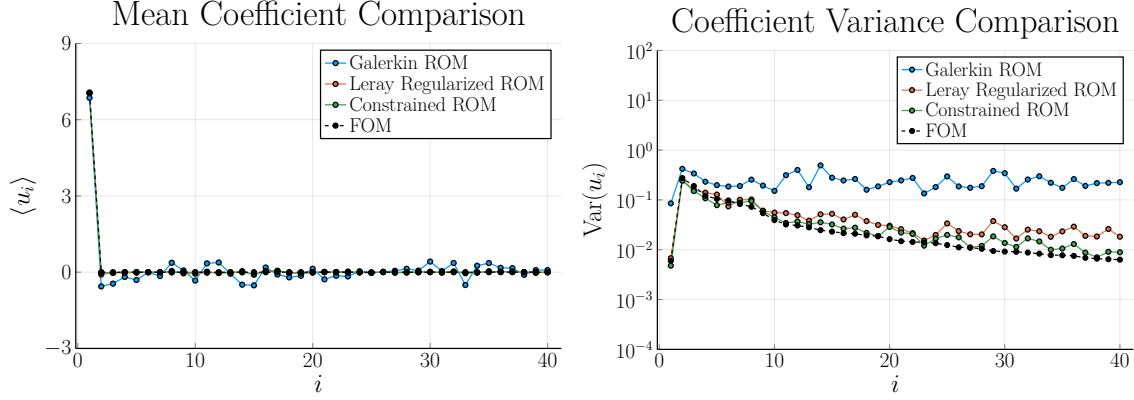


Figure 3.2: Behavior of the sample mean and variance of the coefficients $\{u_n^j\}_j$ for the lid-driven cavity with $N = 40$.

Figure 3.3 shows the estimated velocity variance,

$$\text{Var}_u(t) = \frac{1}{2} \int_{\Omega} \|\mathbf{u}(x, t) - \langle \mathbf{u} \rangle(x)\|_2^2, \quad (3.8)$$

where $\|\cdot\|_2$ is the Euclidean norm. The Galerkin, Leray, and constrained ROMs using $N = 40$ are compared with the FOM. For the Leray regularization, the last 50% of the modes are set to zero. Note that the percentage of modes to be filtered is not known beforehand. For this problem, we have experimented with filtering out 10%, 20%, 30%, 40% and 50% of the modes, and found that 50% filtering produced the most accurate results. The question of optimal filter choice depends on the number of POD modes and on the problem and will be a topic of future study.




Figure 3.3: Velocity variance behavior comparison for FOM, Galerkin ROM, Leray ROM and constrained ROM with $N = 40$.

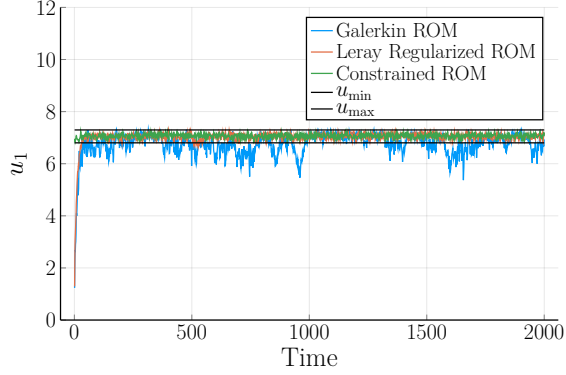


Figure 3.4: Comparison of the first mode behavior for various stabilization techniques with $N = 40$.

To further demonstrate the effects of stabilization, we plot the coefficient behavior versus time. Figure 3.4 clearly shows that, without regularization, the coefficient for the first POD basis function ventures out of the range observed in the data. Other spurious phenomena include effects such as false stable steady flows.

As thermal transport is one of our principal motivators, we consider an energy transport problem using parameter $\kappa = \frac{1}{\text{Re}}$ with the solution to the lid-driven cavity problem as the advecting field. The thermal boundary conditions are:

$$\left\{ \begin{array}{l} T(x, -1, t) = 0, \\ T(x, 1, t) = 1, \\ T(x, -1, t) = T(x, 1, t) = 1. \end{array} \right. \quad (3.9)$$




Figure 3.5: Heated lid-driven cavity Nusselt number vs. time for Galerkin ROM with various N .




Figure 3.6: Heated lid-driven cavity Nusselt number vs. time for ROM with Leray regularization for various N

The QOI of choice is the Nusselt number, which is the mean temperature gradient on the lid

$$Nu = \frac{\int_{\Gamma_{top}} \nabla T \cdot \hat{n} dS}{\int_{\Gamma_{top}} dS}. \quad (3.10)$$

Figure 3.5 shows the Nu history for the FOM and Galerkin ROM for different values of N . We observe that, even with the unconstrained Galerkin approach, the accuracy of the QOI can be recovered as the number of modes increases from $N = 10$ to 80. Figure 3.6, however, illustrates that the stabilized Leray and constrained approaches are able to produce accurate Nu with much lower N (as low as 10) than the standard Galerkin approach.

Chapter 4

Augmented Basis Method

In prior sections, we have outlined some mitigation strategies to ensure proper dynamics in the ROM evolution. Although effective in appropriate settings, these methods do not possess qualities that are desirable for general use: non-intrusive, highly effective, and amenable to pMOR. Non-intrusive methods such as the use of H_0^1 POD basis [16] follow the Galerkin statement without any modeling or modifications, but offers modest improvement. The constrained optimization approach is highly effective for the reproduction problem, but in the pMOR setting, its guarantee in the coefficient behavior is only accurate at the anchor point from which the ROM was produced. Therefore, we will introduce a method that have all three characteristics: the augmented basis method (ABM) [8].

4.1 Methodology

To motivate the ABM, we consider the Leray-projected form of the NSE, in which the velocity field evolution is described as:

$$\partial_t \mathbf{u} = \mathbb{P}[-\mathbf{u} \cdot \nabla \mathbf{u} + \nu \nabla^2 \mathbf{u}]. \quad (4.1)$$

Here, the pressure has been formally eliminated and its effects are represented by an abstract operator, \mathbb{P} , sometimes called the Leray projector, which will project the operand onto a space of divergence-free fields. While the Leray projector is a projection using the H_0^1 inner-product, we will use the L^2 inner-product for our definition. For the discretized system, particularly with the $\mathbb{P}_N - \mathbb{P}_{N-2}$ spectral element discretization, this operator is well-defined [31].

4.1.1 Space of Solutions

For the full-order model, we look for the solution in the finite-dimensional space $X^{\mathcal{N}}$ spanned by the spectral element method basis functions. This space consists of P th-order tensor-product polynomial bases mapped from a reference unit cube to each of E spectral elements, for a total of $\mathcal{N} \approx EP^3$ degrees-of-freedom per field (in 3D). For the POD-Galerkin approach we restrict our attention to solutions $Z^N \subset X^{\mathcal{N}}$, where the basis is generally formed from a proper orthogonal decomposition of a sequence of SEM solution snapshots, using the method of snapshots introduced by Sirovich [45].

The method of snapshots forms a basis from a linear combination of FOM solution fields (each involving $O(\mathcal{N})$ spectral element basis coefficients). One forms the Gramian matrix, whose first N eigenvectors (ranked by eigenvalues from largest to smallest) are used to determine the linear combination of the snapshots that forms the N -dimensional basis for the ROM approximation space, Z^N . Because the snapshots are (weakly) divergence-free, so are all elements of Z^N , which means that pressure term drops out of the equation that governs the ROM. Recall, the velocity POD modes are denoted as ζ_i , and the thermal POD modes are denoted as θ_i . For both of these collections of modes, the $i = 0$ modes correspond to a lifting function that satisfies the boundary conditions and is always associated with a coefficient value of $u_0 = 1$ and $T_0 = 1$. The choice of the lifting function may be a solution to the Stokes problem, the Poisson equation, or the time-averaged solution. For the examples in Section 5, the lifting function is based on time-averaged FOM solutions. For the POD-ROM, the hierarchy of the spaces of interest is $Z^N \subset X^{\mathcal{N}} \subset Y$. We assume that the FOM discretization

is well-resolved such that the projection error from Y to X^N is minimal. We next show how the Z^N space derived by the classical POD-Galerkin method can be augmented such that the time-evolution of the solution in the extended space better approximates the time-evolution of the solution in X^N .

4.1.2 Time-Evolution of NSE

Assuming that the solution to (4.1) exists near t^* , we can describe the local temporal behavior through a Taylor-series expansion involving a linear combination of all time-derivatives.

$$\mathbf{u}(\mathbf{x}, t^* + \epsilon) = \mathbf{u}(\mathbf{x}, t^*) + \epsilon \partial_t \mathbf{u}(\mathbf{x}, t^*) + \dots \quad (4.2)$$

$$= \underbrace{\mathbf{u}(\mathbf{x}, t^*)}_{\text{Snapshot}} + \epsilon \mathbb{P} \{ -\mathbf{u} \cdot \nabla \mathbf{u} + \nu \nabla^2 \mathbf{u} \} + \dots \quad (4.3)$$

Therefore, in addition to capturing the dominant modes of the snapshots, we propose to augment the POD basis set Z^N with modes that can accurately represent the order ϵ terms on the right-hand side of (4.3) in order to construct $\mathbf{u}(\mathbf{x}, t^* + \epsilon)$. The consequence of not representing the $O(\epsilon)$ term is deviation in the trajectory of the physical solution and the projected (Galerkin) solution.

Consider a solution \mathbf{u} that lives in Z^N , meaning $\mathbf{u} = \sum_{i=0}^N u_i \zeta_i$. Using (4.1), we can describe the time-derivative of the solution as

$$\partial_t \mathbf{u} = \mathbb{P}[-\mathbf{u} \cdot \nabla \mathbf{u} + \nu \nabla^2 \mathbf{u}] = - \sum_{i,j=0}^N u_i u_j \mathbb{P}[\zeta_i \cdot \nabla \zeta_j] + \sum_{i=0}^N u_i \nu \mathbb{P}[\nabla^2 \zeta_i] \quad (4.4)$$

Thus, we can accurately describe the time-derivative with $(N + 1)^2$ terms for the nonlinear term and $(N + 1)$ terms for the viscous term.

We consider an example with Fourier basis to highlight this issue. When the solution is band-limited with the highest wavenumber k , the convection term would produce a solution at the next timestep of highest wavenumber $2k$, which does not live in the original space. Thus,

the wavenumber $2k$ behavior is never observed in the evolution of the projected Galerkin system. With an augmentation of the basis with the high-wavenumber modes, we will face the same issue through lack of $4k$ mode representation. This issue is of course recursive. We are helped, however, by the fact that the higher wavenumber modes have higher rates of dissipation. Continuation of this process will therefore eventually yield only marginal returns in improved solution fidelity. We shall see, however, that addition of just a few modes can have a significant impact on the overall ROM performance.

Because of nonlinear interaction through advection, the solution will evolve outside the N -dimensional span of Z^N . We note that as the basis includes more fine-scale components, the convective contribution becomes small relative to the diffusive contribution; thus, the solution becomes closed as the minimal grid-size approaches 0, as is the case in FOM solvers (i.e., the exact solution is band-limited). In the POD-ROM, however, the basis is typically far from completing the relevant approximation space and the addition of the modes $\mathbb{P}[\zeta_i \cdot \nabla \zeta_j]$ and $\mathbb{P}[\nabla^2 \zeta_i]$ can provide an important first-order correction to Z^N .

For advection dominated problems, we can focus on the nonlinear contributions,

$$\partial_t \mathbf{u} \approx \mathbb{P}[\mathbf{u} \cdot \nabla \mathbf{u}] = \sum_{j,k=0}^N u_j u_k \mathbb{P}[\zeta_k \cdot \nabla \zeta_j] \quad (4.5)$$

Whenever we evolve the solution in the space Y , where the current solution lies in the truncated POD space Z^N , we see that the time derivative be reasonably represented with an additional $(N + 1)^2$ basis functions of the form $\phi_{l=j+k(N+1)} = \mathbb{P}[\zeta_k \cdot \nabla \zeta_j]$. Obviously, this process is not closed, since more basis functions are required in the next timestep. Worse still, even starting with $\mathbf{u} \in Z^N$, the required number of additional basis functions will be $O(N^2)$ if we include all terms in (4.5), which comes with an $O(N^6)$ computational cost that is untenable, even for a small number of POD modes, N . We therefore seek to augment the original POD basis with subsets of these evolution basis that are most relevant to the dynamics.

The first subset captures the interaction between the lifting function, ζ_0 , and all other modes. This choice ensures that both the Taylor dispersion induced by the lifting function and the transport of the mean momentum by the POD modes are accurately captured. This choice is also rationalized by the fact that the lifting function is ever-present in the solution so its convective interaction is important in accurate reproduction of the time-evolution by the ROM. Thus, we add the modes $\mathbb{P}[\zeta_0 \cdot \nabla \zeta_j + \zeta_j \cdot \nabla \zeta_0]$. Note that the two interactive terms can be combined because it is linear in each POD basis, ζ_j , meaning we only add $N + 1$ modes, which is still an $O(N)$ augmentation.

Next, we extract the diagonal entries, $\mathbb{P}[\zeta_j \cdot \nabla \zeta_j]$. This choice is justified by the fact that for each mode, ζ_j , the mode that is the most correlated with it is itself (i.e., when other modes might have a phase-shift, or different temporal frequencies associated with it, the auto-correlation dominates other interactions). So we consider addition of these N modes with the total additional modes being $2N + 1$.

For a thermal system with an advection-diffusion equation to describe its state, we can follow the same procedure as above for the lifting function interaction in the form of $\zeta_0 \cdot \nabla \theta_j$; however, the auto-interaction modes are not obvious. For this work, we will choose $\zeta_j \cdot \nabla \theta_j$, but there is no one-to-one correspondence between the dominant thermal modes and dominant velocity modes. One may come up with a more coherent substitute, but this choice remains an open question.

In summary, the ABM starts with N standard POD modes in Z^N and adds $2N + 1$ modes corresponding to advection by the lifting function, $\mathbb{P}[\zeta_0 \cdot \nabla \zeta_j + \zeta_j \cdot \nabla \zeta_0]$ and auto-advection, $\mathbb{P}[\zeta_j \cdot \nabla \zeta_j + \zeta_j \cdot \nabla \zeta_j]$, resulting in a total of $\hat{N} = 3N + 1$ basis functions, which are used in a standard Galerkin formulation. We will use \hat{N} for the comparison against other (classic or stabilized) methods so that we have a fair cost comparison. The standard POD Galerkin ROM and ABM differ only in the choice of the underlying basis set.

Chapter 5

Applications of ABM-ROM

Here, we demonstrate the effectiveness of the proposed augmentation method on several examples, including flow in a 2D Lid-Driven cavity ($\text{Re} = 30,000$), 2D flow past baffles ($\text{Re} = 800$), 3D lid-driven cavity flow ($\text{Re} = 10,000$), flow over a hemisphere ($\text{Re} = 2,000$), turbulent pipe flow with forced convection for $\text{Re} = 4000 - 10,000$, channel flow ($\text{Re} = 10,000$), and turbulent pipe flow with wire-coil insert.

We denote the FOM solution with $(\tilde{\mathbf{u}}, \tilde{T})$ and the various ROM solutions with (\mathbf{u}, T) . Time averages are defined as $\langle \cdot \rangle = \frac{1}{\tau} \int_{t_0}^{t_0+\tau} \cdot dt$ with integration times, τ , prescribed on a case-by-case basis. In these cases, we compare different ROM strategies: Standard POD Galerkin (L^2 -Gln), Galerkin with H_0^1 POD basis (H_0^1 -Gln), Leray-filtered (L^2 -Lry), Constrained Optimization (L^2 -Cst), ABM with lifting function interaction (Aug0), ABM with auto-interaction (AugD), and ABM with both interactions (AugC).

For each case, we track the convergence by increasing the number of modes, \hat{N} , that are used in the ROM. (This count excludes the 0-mode, which is the lifting function.) For the standard cases, $\hat{N} = N$, corresponding to the number of active POD modes. For the ABM, \hat{N} is total number of modes including the base POD modes and their augmentation counterparts. For example, if we are using $N = 3$ with both 0- and diagonal-interaction modes, we would have $\hat{N} = 10$.




Figure 5.1: FOM snapshot (velocity magnitude) of 2D LDC ($\text{Re} = 30,000$), $|\tilde{\mathbf{u}}(\mathbf{x}, 1000)|$, used to produce the initial condition for all ROMs.

5.1 2D Lid-Driven Cavity (LDC)

We begin with the regularized lid-driven cavity problem. This problem is identical to what was used by Fick *et al.* in Chapter 3, save that we have increase the number of elements from uniform 16×16 array to 32×32 array with a Chebyshev distribution. The spectral element polynomial order is 7 with time-step size $\Delta t = 2 \times 10^{-3}$. The boundary condition is kept the same with: $u(\Gamma) = (1 - x^2)^2 \delta(1 - y)$. We ran the simulation starting with the Stokes solution to the problem with the same boundary condition for 2000 CTUs and will use 1000 snapshots collected over the last 1000 CTUs. A velocity snapshot at time 1000 with the spectral element discretization is shown in Fig. 5.1.

Aside from the constrained case, all the ROMs are advanced in an identical manner: Galerkin method with BDF/EXT3 time-discretization using the projection of $\tilde{\mathbf{u}}(\mathbf{x}, 1000)$ to each of the solution space as the initial condition. Furthermore, they all use $\langle \mathbf{u}(\mathbf{x}, t) \rangle$, the mean solution from time 1000 to 2000, as the lifting function. Each ROM is advanced from 1000 CTUs and the time-average of the coefficients and the outer-product of the coefficients are collected so we calculate the mean and variance for each case.

5.1.1 Comparison of the modes

An example is illustrated in Figs. 5.2–5.4, which shows the first 14 L^2 –AugC modes for the case of a lid-driven cavity at $\text{Re} = 30,000$ (based on the cavity width, equivalent to $\text{Re} = 15,000$ in [14]). For $j = 0, \dots, 4$, the first five POD modes, $\zeta_j \in Z^N$, are in the top row; the first five 0-modes, $\mathbb{P}\{\mathbf{u}_0 \cdot \nabla \mathbf{u}_j + \mathbf{u}_j \cdot \nabla \mathbf{u}_0\}$, are in the center row; and the first five diagonal-modes, $\mathbb{P}\{\mathbf{u}_j \cdot \nabla \mathbf{u}_j\}$, are in the lower row.




Figure 5.2: Magnitude of first 5 L^2 POD modes, ζ_i .




Figure 5.3: Magnitude of first 5 augmentation of $\mathbb{P}[\zeta_0 \cdot \nabla \zeta_i + \zeta_i \cdot \nabla \zeta_0]$.




Figure 5.4: Magnitude of first 5 augmentation of $\mathbb{P}[\zeta_i \cdot \nabla \zeta_i]$.

We note that the diagonal augmentation in Fig. 5.4 in particular have high wavenumber content compared to the original POD basis. This will contribute to additional dissipation and we can expect a more stable solution. This predicted stable behavior is shown in the mean and variance results.

5.1.2 Convergence and Stability Study using Mean and Variance Metric

Now we look at how the various approaches, including the ABM, compare in a dynamic context. We begin with the error in the mean field prediction shown in Fig. 5.5. We see that the augmented approach performs consistently better than all other cases, hovering around 2×10^{-2} or lower. The other methods converge as \hat{N} is increased, particularly the H_0^1 basis, however, ABM is the clearly better.

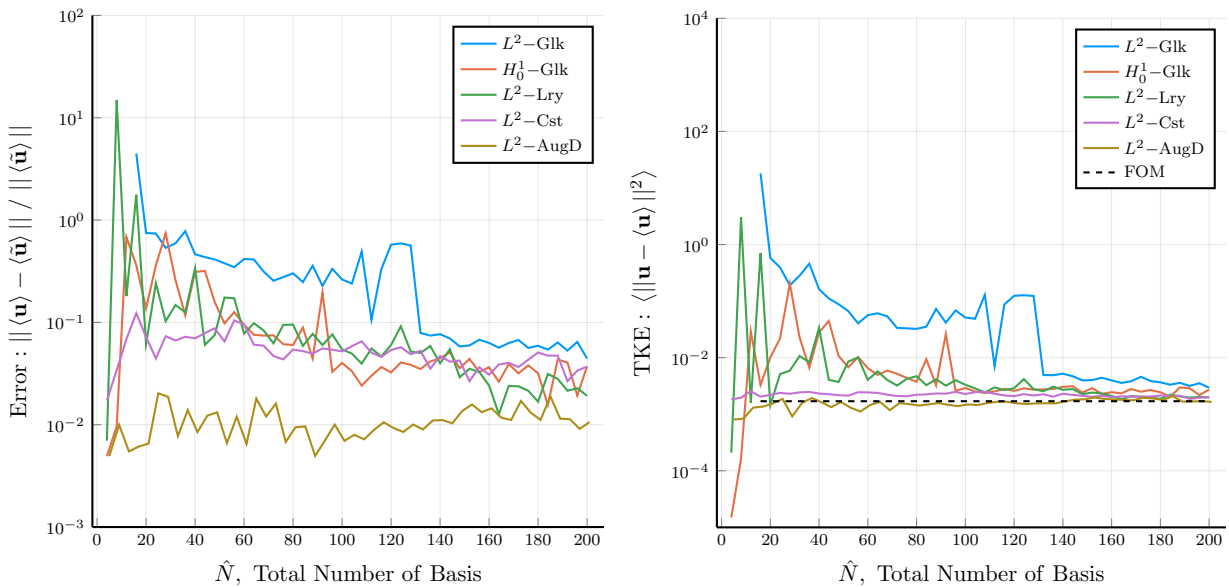


Figure 5.5: 2D LDC ($Re = 30,000$): Comparison of ROM velocity field accuracy w.r.t. \hat{N} .

For the velocity variance (labeled as TKE in Fig. 5.5, right), we see remarkably consistent agreement between the ABM and FOM starting with as few as $N = 9$ modes. As with the mean error behavior, all the other approaches converge for the TKE metric, but the augmented approach is again clearly better. Thus, we see that simply capturing the most amount of energy in the snapshots per given number of modes, as is done by POD, is not the optimal choice of basis for capturing the highly nonlinear dynamics described in the Navier–Stokes equations. Furthermore, we emphasize the fact that the constrained optimization shows fairly good result, but once the Reynolds number is moved away from 30,000, for




Figure 5.6: Velocity magnitude plot of FOM Snapshot of 2D baffle flow ($Re = 800$).

pMOR application, we expect that the empirical bounds will be less relevant.

5.2 2D Baffle Flow

As a surrogate for turbulent channel flow, we consider the challenging problem of 2D flow in a channel with baffles, as illustrated in Fig. 5.6. This flow features periodic shedding and transport of large-scale vortices in the wake of each baffle whose snapshot is shown in Fig. 5.6. 1000 snapshots are taken from a 2D baffle problem with $Re = 800$ over 100 CTUs after steady statistical state is reached. The flow is then reconstructed over the same time.

All the means converge to the same magnitude with $N = 80$, shown in Fig. 5.7; however, they reach a non-zero value of ≈ 0.1 . This may indicate they all got stuck near the same unphysical attractor or simply the chaotic nature of the flow configuration requires a longer integration time for the mean to fully converge. However, the variance converges to the correct value for all three cases with the augmented basis showing much better stability properties with low values of N , shown in Fig. 5.8.

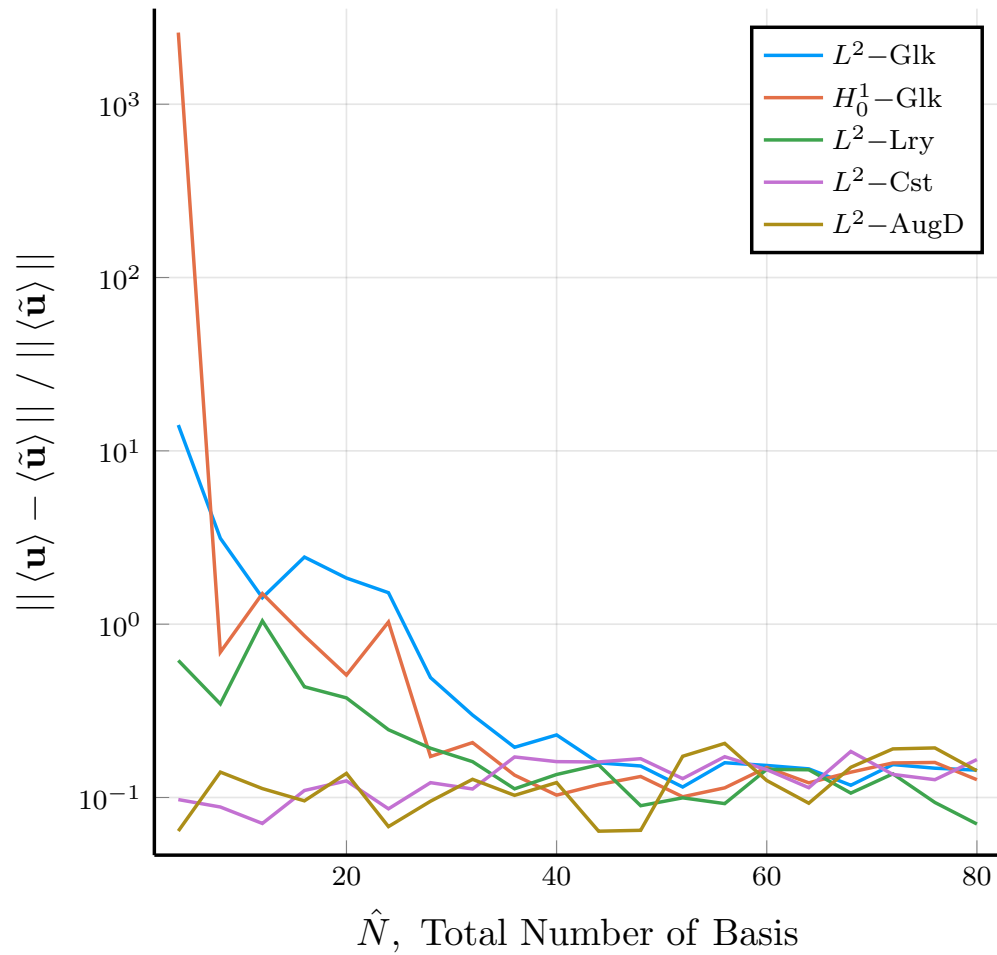


Figure 5.7: 2D Baffle flow (Re = 800): Mean error comparison.




Figure 5.8: 2D Baffle flow (Re = 800): Variance (TKE) comparison.

5.3 3D Lid-Driven Cavity

We next consider a 3D problem with moderate Reynolds number. We choose the non-regularized 3D lid-driven cavity problem of Reynolds number 10,000. The FOM mesh consists of a tensor-product array of $E = 16^3$ elements with a Chebyshev distribution with polynomial order 7 (2 million DoFs). We take 2000 snapshots over the span of 1000 CTUs (with initial condition being the Stokes solution). We use the latter half of the snapshots from $t_0 = 500$ to $t_f = 1000$ to generate the POD modes and run the ROM from in $[t_f, t_f + 500]$. One of these snapshots is shown in Fig. 5.9.

As is the case in the 2D problems, we see that the ABM outperforms standard POD bases

for the mean field and TKE quantities shown in Fig. 5.10 and Fig. 5.11. This is a promising result which establishes the merits of the proposed approach on this class of problems.




Figure 5.9: FOM snapshot (velocity magnitude), $\tilde{\mathbf{u}}(\mathbf{x}, 500)$, for the 3D LDC problem ($\text{Re} = 10,000$).

3D LDC: mean comparison between FOM and ROM




Figure 5.10: 3D LDC: Mean error comparison.

3D LDC: variance (TKE) comparison between FOM and ROM




Figure 5.11: 3D LDC: Variance (TKE) comparison.




Figure 5.12: FOM Velocity magnitude snapshot of flow over hemisphere ($\text{Re} = 2,000$) with overlaid λ_2 contour.

5.4 3D Flow over a Hemisphere

Figure 5.12 shows a snapshot of flow past a wall-mounted hemisphere of height $D/2$ at $\text{Re}_D = DU/\nu = 2000$. This problem was studied experimentally by Arcalar and Smith [46] to understand the role of hairpin vortices in transition to turbulence. In these numerical simulations, the FOM uses a Blasius profile with boundary-layer thickness $\delta_{99} = 0.6D$, prescribed at the inlet that is $3.2D$ units upstream of the hemisphere center. Periodic boundary conditions are prescribed at $\pm 3.2D$ units in the spanwise direction and a stress-free condition is applied on the top surface, $3.2D$ units above the wall. Under these conditions, the flow exhibits periodic shedding of hairpin vortices, evidenced by the velocity distribution and λ_2 contours [47] in the hemisphere wake. The FOM, based on a spectral element mesh with $\mathcal{N} \approx 2$ million gridpoints was run for 100 convective time units ($1 \text{ CTU} = D/U$) and 1000 snapshots we collected to form the ROM POD bases. The mean-velocity error as a function of \hat{N} is shown in Fig. 5.13 (left) for the five different ROMs. The unstable L^2 and H^1 Galerkin results have several drop-outs for conditions that did not converge for this relatively high-Reynolds number application. Given enough basis functions, however, all cases converge, with the constrained case (L^2 -Cst) being the best performer for $\hat{N} < 120$.

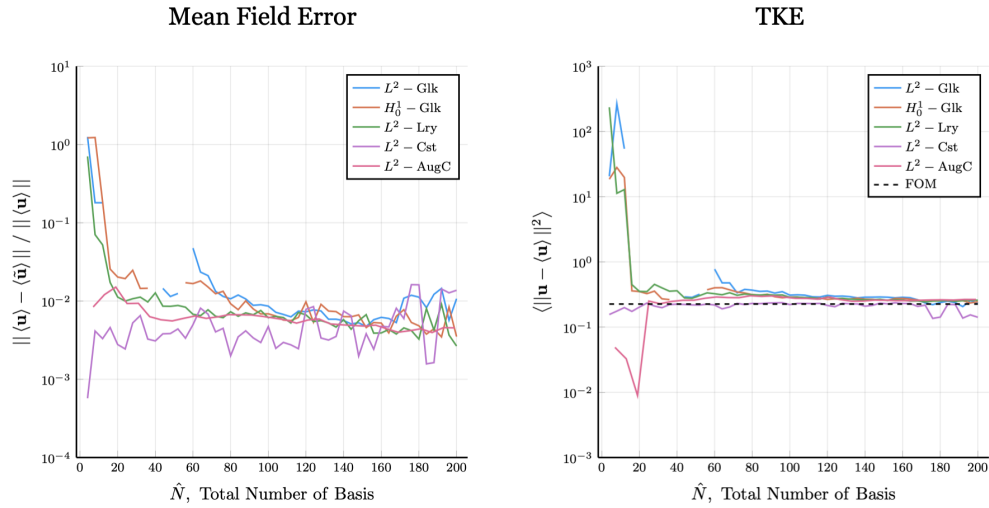


Figure 5.13: Flow over hemisphere ($\text{Re} = 2,000$): Mean error (left) and TKE (right) comparison, curve broken by blowup solution.

Both L^2 -Cst and L^2 -AugC yield mean-field errors $< .01$ for the majority of the cases, with L^2 -Cst generally being the best performer. Similar conclusions hold for the TKE results, shown in the right panel of 5.13.

5.5 3D Flow Past a Cylinder

We follow the procedure by Wells *et al.* [17] and compute 3D flow past a cylinder for Reynolds number 1000. The mesh has 4.3 million DoFs (544 elements in the x-y plane and 8 elements in the spanwise direction, each element of order 9). The domain size is 19.5 by 10 by 4 where the cylinder axis is placed 2.5 diameters from the inlet and centered in the x-z plane.

The results are shown in Fig. 5.15. Again, broken curves correspond to cases did not converge. In this case, constrained ROM (L^2 -Cst) has the best performance. The ABM also performs well, particularly for the TKE. In this case, however, the ABM convergence is slower than is observed for other cases. One hypothesis to be verified is that this slow convergence is due to the lack of resolution in the mesh used for the FOM.




Figure 5.14: λ_2 contour plotted over a slice of the velocity magnitude plot for a 3D cylinder flow $Re = 1,000$ snapshot.




Figure 5.15: Mean velocity error (left) and TKE (right) for the 3D cylinder problem ($Re = 1000$, broken curves indicate solution blowup).

5.6 Forced Convection in Turbulent Pipe Flow

The next example is that of forced convection in turbulent pipe flow with Reynolds number $Re = 4000, 5300, \text{ and } 10000$ (based on pipe diameter), and Prandtl number $Pr = 1$. All the cases use the same spectral element distribution with differing polynomial orders. The mesh consists of 12.5 million grid points for $Re = 4000$ and 5300 , and 24.5 million points for $Re = 10,000$. The periodic domain length is $L = 4D$, which is generally inadequate for a full DNS of turbulence but deemed sufficient for the numerical tests in this study. For $Re = 4000, 5300, \text{ and } 10000$, the respective FOM Nusselt numbers are $Nu = 16.38, 21.42, \text{ and } 36.14$, which is in good agreement with the Dittus-Boelter relationship, $Nu = 0.023 Re^{4/5} Pr^{2/5}$. For all cases, the FOM is run until the solution is relaxed to a statistically steady state prior to gathering statistics or snapshot data. For each case, 1000 snapshots are collected over 50 CTUs to form the Gramian, from which the POD basis is generated. Figure 5.16 shows typical snapshots of velocity magnitude and temperature that reveal the variation in range of scales for the different cases.




Figure 5.16: FOM velocity (top) and temperature (bottom) snapshots of 3D pipe flow.

The governing equations for the FOM are the incompressible Navier–Stokes equations and the thermal advection-diffusion equation. Because of the constant-flow rate and periodic restriction on the solutions, we provide a brief discussion of modifications to the standard equations for the FOM and their effect on the ROM formulation. We start with the Navier–Stokes equations:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}(\mathbf{u}), \quad \nabla \cdot \mathbf{u} = 0 \quad (5.1)$$

Here, $\mathbf{f}(\mathbf{u})$ is a uniform forcing vector field function in the streamwise direction, $\hat{\mathbf{z}}$, that enforces a time-constant flow-rate. In the time-discrete problem, the forcing term effectively adds an impulse-response streamwise velocity field with boundary layer thickness proportional to $\sqrt{\nu \Delta t}$. This impulse response is scaled appropriately at each time step to ensure that the mean velocity at each timestep conforms to the prescribed flow-rate [48, 49].

5.6.1 FOM Prescription of Constant Flowrate

The semi-implicit Navier-Stokes update has the form

$$\begin{bmatrix} H & -D^T \\ -D & 0 \end{bmatrix} \begin{pmatrix} \tilde{\underline{u}} \\ \tilde{\underline{p}} \end{pmatrix} = \begin{pmatrix} \underline{f} \\ \underline{0} \end{pmatrix}, \quad (5.2)$$

where the right-hand side vector $\begin{pmatrix} \underline{f} \\ \underline{0} \end{pmatrix}$ has contributions from the convection term and boundary conditions. We note that (5.2) is *linear* and superposition therefore applies. When solved, the flowrate of the solution, \tilde{Q} , may not satisfy the prescribed flow-rate, Q . To fix the discrepancy, we can solve an auxiliary system,

$$\begin{bmatrix} H & -D^T \\ -D & 0 \end{bmatrix} \begin{pmatrix} \underline{u}' \\ \underline{p}' \end{pmatrix} = \begin{pmatrix} \underline{g} \\ \underline{0} \end{pmatrix}, \quad (5.3)$$

where \underline{g} is the vector field filled with a unit-vector in the intended stream-wise direction. We can find the volumetric flow rate of the solution, Q' , to this system, scale it by the difference between the flow rate and the intended flow rate then add it to the solution

$$\underline{u} = \tilde{\underline{u}} + \frac{Q - \tilde{Q}}{Q'} \underline{u}' \quad (5.4)$$

Using this method, the solution \underline{u} can be guaranteed to satisfy the prescribed flowrate at each timestep, meaning all the snapshots satisfy the prescribed flowrate.

5.6.2 ROM Prescription of Constant Flowrate

In the case of the ROM, the lifting function has the prescribed flow-rate and the remaining POD basis functions have zero flow-rate, meaning that the test-space Z_0^N only contains

members with zero flow-rate. In the weak-form, the ROM forcing term therefore becomes

$$(\mathbf{v}, \mathbf{f}) = \int_{\Omega} \mathbf{v} \cdot \mathbf{f} dV = f_z \left(\int_{\Omega} v_z dV \right) = 0. \quad (5.5)$$

Thus, the forcing term in the ROM formulation is zero, meaning there is no additional forcing term that needs to be added to the existing ROM formulation.

5.6.3 Enforcing Periodicity in the Thermal Solution

Because the computational domain is finite and the fluid problem is solved with periodic boundary condition, the thermal solution should also be periodic. To impose periodicity, we follow [48, 49]. Because constant flux is applied through the outer surface and there is net transport of heat from the $z = 0$ plane to the $z = L$ plane, the physical temperature cannot be at $z = 0$ and $z = L$. Therefore, we introduce another variable, T_{Lift} so that T satisfy the periodic boundary condition and $T + T_{\text{Lift}}$ satisfy conservation of energy. For this purpose we define

$$T_{\text{Lift}} = \gamma z \quad (5.6)$$

where γ is a constant that is proportional to the total flux through the domain boundary. By substituting the expression into the original equation, we get

$$\frac{\partial T}{\partial t} + (\mathbf{u} \cdot \nabla)T = \kappa \nabla^2 T - u_z \gamma \quad (5.7)$$

where u_z is the z-component of the velocity vector field. In effect, by applying a periodic boundary condition on this arbitrary variable, an additional sink term is introduced. To ensure that thermal energy is conserved in the domain we must set the net energy transported by the fluid equal to the energy injected through thermal flux. Therefore,

$$\int_{\partial\Omega_{\text{out}}} u_z(T + \gamma(L)) dS - \int_{\partial\Omega_{\text{in}}} u_z(T + \gamma(0)) dS = \int_{\partial\Omega_{\text{wall}}} \kappa \nabla T \cdot \hat{n} dS \quad (5.8)$$

$$\gamma L \int_{\partial\Omega_{\text{out}}} u_z = \int_{\partial\Omega_{\text{wall}}} (1) dS \quad (5.9)$$

$$\gamma = \frac{P}{Q} \quad (5.10)$$

where P is the perimeter of the z cross-section of the domain, where the flux is applied. Thus, for a circular pipe-flow of unit diameter, we have $\gamma = \frac{P}{Q} = 4$.

5.6.4 ROM Results for Pipe Flow

The results for the ROMs are presented in Figs. 5.17 – 5.19, which show the error and variance for the velocity and temperature as well as the Nusselt number behavior as a function of the total number of modes. The mean Nusselt number definition is based on the time-averaged streamwise velocity and temperature,

$$\overline{\text{Nu}} = \frac{1}{\alpha(\overline{T_s} - \overline{T_b})}, \quad \overline{T_s} = \int_{\partial\Omega} \langle T \rangle dS, \quad \overline{T_b} = \frac{\int_{\Omega} \langle T \rangle \langle u_z \rangle dV}{\int_{\Omega} \langle u_z \rangle dV}. \quad (5.11)$$

A common observation for Figs. 5.17 – 5.19 is nominal convergence for the $\text{Re} = 4000$ case for the L^2 , H_0^1 , and Leray regularization methods, albeit to relatively large asymptotic values. As the Reynolds number increases, more modes are required for the L^2 and H^1 formulations to converge, with the required number of modes apparently exceeding $\hat{N} = 200$ for $\text{Re} = 10,000$. Clearly, Leray outperforms standard L^2 and H_0^1 , but is inferior to L^2 -Cst and L^2 -AugC, with the latter two having mean velocity error of just a few percent at $\text{Re} = 10,000$ (Fig. 5.17, top right). The thermal behavior is similar, save that the mean-field error (Fig. 5.18, top) is above 10% with the exception of L^2 -AugC for $\text{Re} = 10,000$. Remarkably, this same case exhibits too little thermal variance, as seen in the lower right frame of Fig. 5.18. (We explore this anomalous behavior in the next section.) On the other hand, the error in Nu for L^2 -AugC at $\text{Re} = 10,000$ is uniformly less than 1% (Fig. 5.19).

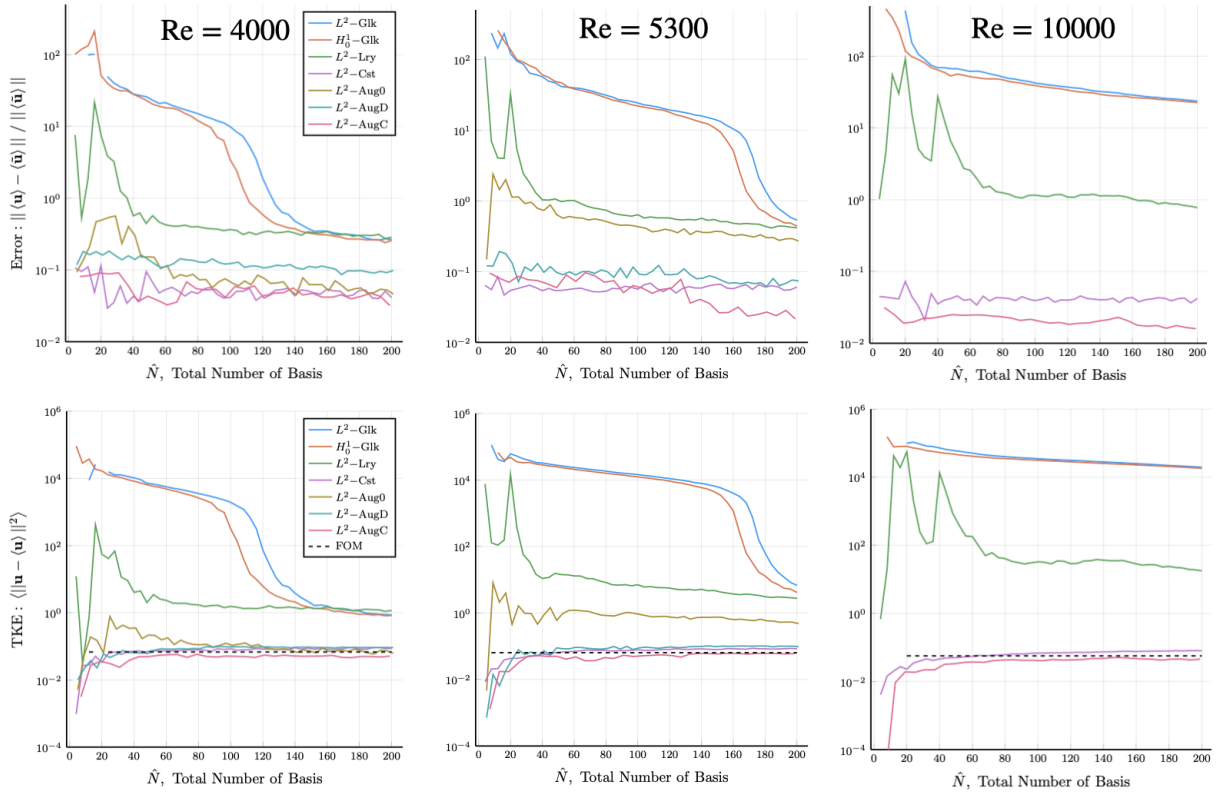


Figure 5.17: Pipe flow mean velocity error (top) and TKE results (bottom), for $Re = 4000$, 5300 , and 10000 . (FOM: $\tau = 50$, ROM: $\tau = 500$).

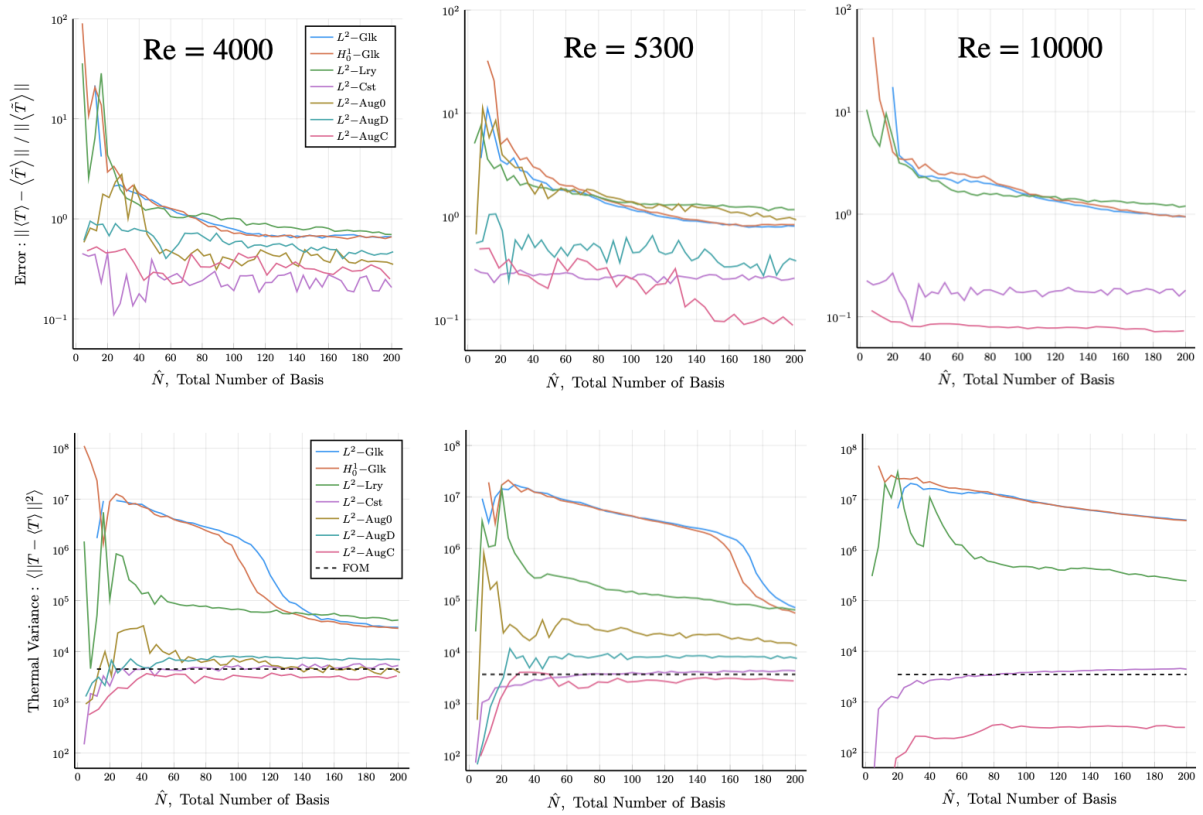


Figure 5.18: Pipe flow mean temperature error (top) and thermal variance results (bottom), for $Re = 4000, 5300, \text{ and } 10000$. (FOM: $\tau = 50$, ROM: $\tau = 500$).

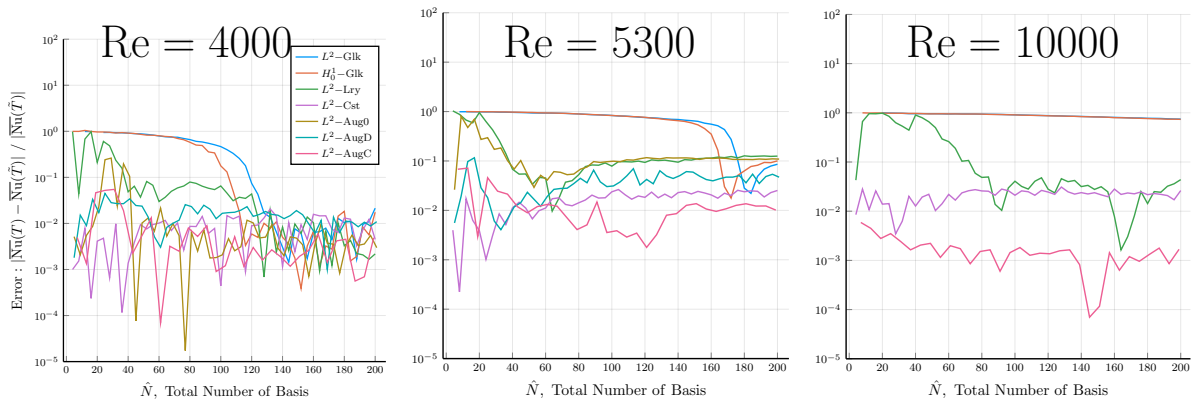


Figure 5.19: Pipe flow error in Nusselt number for $Re = 4000, 5300, \text{ and } 10000$.

5.7 Forced Convection in Turbulent Channel Flow

The next example is that of forced convection in turbulent channel flow with Reynolds number $Re = 10000$, (based on channel height), and Prandtl number $Pr = 1$. The physical domain is $L = 8H$, $W = 4H$, where

$$H := \text{Channel Height}, \quad (5.12)$$

$$L := \text{Streamwise Length}, \quad (5.13)$$

$$W := \text{Spanwise Width}. \quad (5.14)$$

The spectral elements are in a $[88 \times 48 \times 22]$ regular mesh with a uniform distribution in stream-wise and span-wise directions and Chebyshev distribution in wall-normal direction with a total of $\mathcal{N} = 32$ million. For this problem a periodic boundary condition is applied in the stream-wise and span-wise directions with constant-flow rate prescription of mean velocity of 1. For this problem, $\gamma = \frac{P}{Q} = 2$. The FOM is run until the solution is relaxed to a statistically steady state prior to gathering statistics or snapshot data. 2000 snapshots are collected over 50 CTUs to form the Gramian, from which the POD basis is generated. Figure 5.20 shows typical snapshots of velocity magnitude and temperature that reveal the range of scales for this flow and thermal fields.




Figure 5.20: FOM velocity (top) and temperature (bottom) snapshots of 3D channel flow ($Re = 10,000$).

The results for the ROMs are presented in Figs. 5.21 – 5.23, which show the error and variance for the velocity and temperature as well as the Nusselt number behavior as a function of the total number of modes. The mean Nusselt number definition is based on the time-averaged streamwise velocity and temperature.

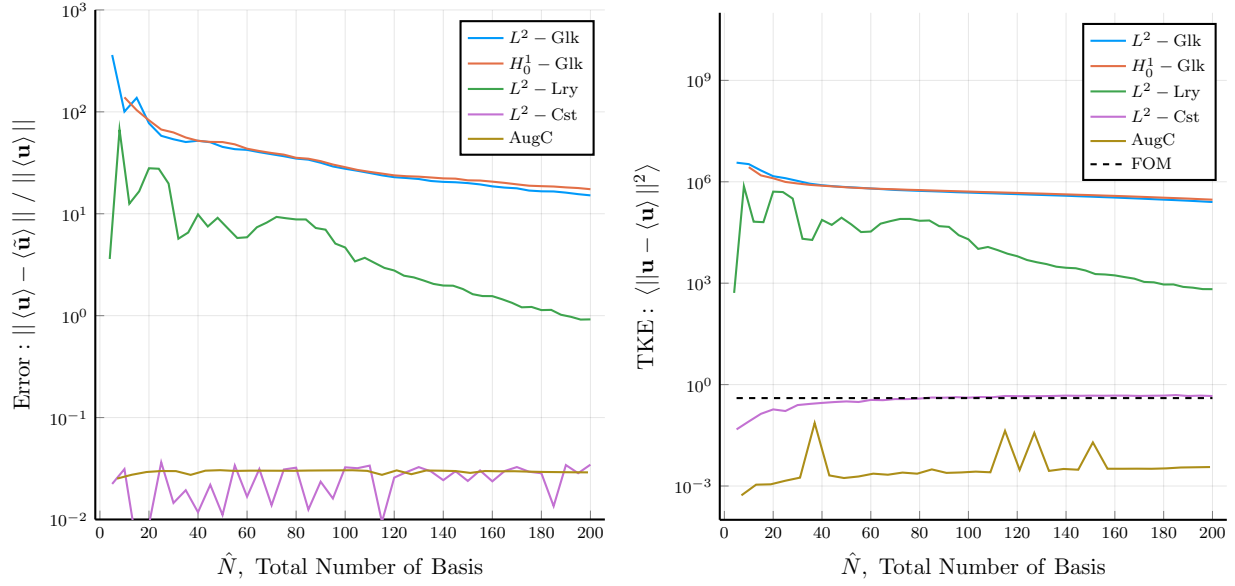


Figure 5.21: Mean velocity error (left) and TKE results (right), for channel flow, $\text{Re} = 10,000$. (FOM: $\tau = 50$, ROM: $\tau = 500$).

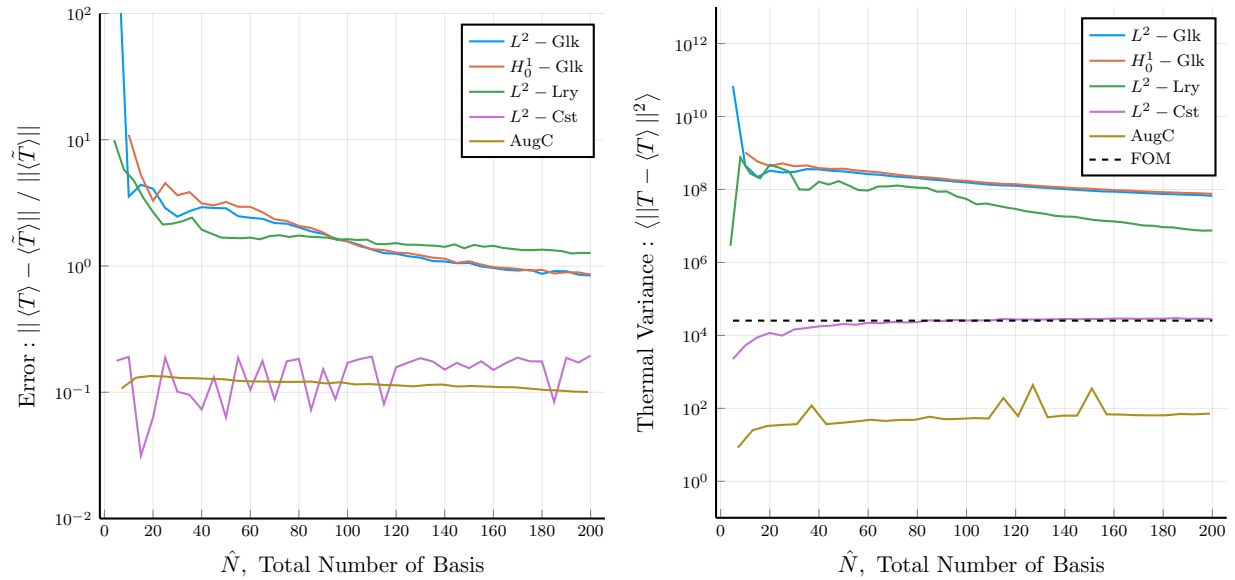


Figure 5.22: Mean thermal error (left) and thermal variance results (right), for channel flow, $\text{Re} = 10,000$. (FOM: $\tau = 50$, ROM: $\tau = 500$).




Figure 5.23: Mean Nusselt number error for channel flow, $\text{Re} = 10,000$. (FOM: $\tau = 50$, ROM: $\tau = 500$).

A common observation for Figs. 5.21 – 5.23 is non-convergent behavior for the standard (L^2) and H_0^1 POD basis. The Leray regularization approach shows nominal convergence, but not adequate even at $\hat{N} = 200$. The constrained optimization and ABM appear to produce accurate mean velocity and thermal solutions, but the TKE and thermal variance behavior appears too dissipative for ABM. There are possible explanations for this behavior including one discussed in Section 6.3. Another factor may be the size of the domain compared to the characteristic length (the channel height) of this problem since for the turbulent pipe-flow, this issue is not prevalent in the result. This possible dependence on the domain size (for spatially periodic problems) will be investigated in future work.

5.8 Heated Pipe-Flow with Wire Coil Insert

In [50], Collins and co-authors measured the effects of wire coil inserts on heat transfer in low-Reynolds number pipe flow. Because of reduced fouling concerns, wire coils were preferred

to, say, copper-mesh inserts, for heat transfer enhancement in the target application of cooling Joule-heated copper blocks that are part of Argonne’s advanced photon source (APS). The authors explored a wide range of wire diameters and pitches to maximize the heat transfer coefficient for a *fixed flow rate*, which is appropriate for any system where flow resistance in the critical region (e.g., the copper blocks) is small compared to the overall resistance in the coolant-supply loop. For the medium (optimal) wire shape, they found a three- to five-fold increase in Nusselt number over a range of Reynolds numbers $Re_D = DU/\nu = 2000\text{--}50,000$.

5.8.1 FOM Study of Wire Coil and Other Inserts

In this section, we numerically revisit several of the flow configurations considered in [50] in an effort to understand the mechanisms of heat transfer enhancement. We consider straight pipe, optimal and off-optimal wire-coil configurations, and configurations that do not induce a mean swirl such as a “meandering” wire and rings.

We first validate this approach by reproducing several of the data points from [51] for both wire-coil inserts and straight pipe results. We then investigate variations on the experiment including replacing the copper coils with glass, with copper rings, and with meandering wires.

Domain and Governing Equations




Figure 5.24: Wire-coil domain: (a) copper pipe with wire-coil insert used in [50]; (b) snapshot of temperature distribution in spectral element simulation for medium (optimal) wire pitch at $Re = 5300$; (c) spectral element surface-mesh mesh distribution for the same simulation.

The flow domain consists of internal pipe flow with a wire-coil insert, as illustrated in Fig. 5.24. In this study, we consider conjugate heat transfer with water as a coolant, finite-thickness copper tubing on the outside, and either glass or copper inserts. Figure 5.24 (a) shows a typical insert configuration for the experiments in [50], (b) shows the temperature for a simulation at $Re_D = UD/\nu = 5300$, and (c) shows the boundary of the fluid domain for the spectral element mesh used for (b). For all cases the outer diameter of the tube is $D_o = 1.33D$, where D is the internal diameter. The wire diameter is $d = 0.2507D$, and the wire pitches are $p = .4454D$ (short), $.5867D$ (medium), and $2.667D$ (long). The results for other configurations, including wire rings and a meandering wire, are included in Table 5.2 and Fig. 5.27.

By including the copper tubing, we can determine the average temperature at the interior wall in a manner consistent with the experiments. Specifically, we integrate the temperature at the outer surface of the pipe, where we apply a uniform heat flux, and compute the inner wall temperature as the solution to the solid conduction problem that would be obtained

in the axisymmetric case without the wire-coil insert. For a uniform thermal flux \tilde{q}'' on the outer surface, $\partial\Omega$, we define the mean inner-wall temperature,

$$T_w = \frac{\int_{\partial\Omega} \langle T \rangle dS}{\int_{\partial\Omega} dS} - \frac{\tilde{q}'' D_o \ln(D_o/D)}{2k_{\text{copper}}}, \quad (5.15)$$

where $\langle T \rangle$ indicates a time average of the field over tens of convective time units and k_{copper} is the thermal conductivity of copper. This approach gives a consistent definition of T_w across all configurations and is relevant to the copper-block cooling problem as it provides a direct measure of the temperature in the device.

For the computations, we denote the domain as $\Omega = \Omega_f \cup \Omega_s$, where Ω_f is the coolant passageway and Ω_s represents the conducting solid including the tubing and the wire insert. In all cases, Ω is streamwise periodic, meaning that the velocity satisfies

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}(\mathbf{x} + L_z \hat{\mathbf{z}}, t), \quad (5.16)$$

where $\hat{\mathbf{z}}$ is the unit vector in the streamwise (z) direction and L_z is the domain length. As in the pipe flow cases described in Section 5.6, periodic boundary condition in velocity and lifted temperature solution is applied. For this problem, since the outer diameter, where the thermal flux is applied, is $\frac{4}{3}$, $\gamma = \frac{16}{3}$. Because constant flow rate is prescribed for all problems instead of mean streamwise velocity, the definition of Re is well-defined regardless of the amount of channel blockage caused by the wire coils or the rings. This definition of Re is consistent with that of the experiments. The flow rate Q is fixed to be a time-independent constant by a Green's function approach described in Section 5.6.

Nondimensionalization of the energy equation is less straightforward than that of the Navier-Stokes equations because of the conjugate heat-transfer formulation, so we review the

material	k	ρ	C_p	$\hat{\alpha}$	$\hat{\sigma}$
Water	0.58	1000	.004180	1.0	1.0
Copper	400	8900	.000385	690	0.82
Glass	1.2	2800	.000840	2.06	0.56

Table 5.1: Material properties for conjugate heat-transfer simulations. MKS units are: $[k] = W/mK$, $[\rho] = kg/m^3$, $[C_p] = J/kg^\circ K$.

steps here to clarify notation. The dimensional energy equation in $\Omega \times [0, t_f]$ is

$$[\rho C_p] \left(\frac{\partial T}{\partial \tilde{t}} + \mathbf{u} \cdot \tilde{\nabla} T \right) = \nabla \cdot k(\tilde{\mathbf{x}}) \tilde{\nabla} T, \quad (5.17)$$

where $\tilde{\mathbf{x}}$ represents the dimensional position vector and $\tilde{\nabla}$ represents the dimensional gradient.

The constant-flux boundary condition on the outer surface of the pipe is

$$k \tilde{\nabla} T \cdot \hat{\mathbf{n}} \Big|_{\partial\Omega} = \tilde{q}'', \quad (5.18)$$

where $\hat{\mathbf{n}}$ is the outward point normal on the domain surface, $\partial\Omega$. The volumetric heat capacity, ρC_p , and conductivity, k , are functions of space because of the material differences between the fluid and solid domains. Within a given material they are taken to be constant with values listed in Table 5.1.

To nondimensionalize (5.17), we first define

$$\hat{\sigma}(\tilde{\mathbf{x}}) = \frac{[\rho C_p]}{[\rho C_p]_{H_2O}} \quad (5.19)$$

$$\hat{\alpha}(\tilde{\mathbf{x}}) = \frac{k}{k_{H_2O}}, \quad (5.20)$$

which are unity in Ω_f and constant in Ω_s . (For the glass wire cases, $\hat{\sigma}$ and $\hat{\alpha}$ will differ between the wire and the outer pipe.) Dividing (5.17) by $[\rho C_p]_{H_2O}$ yields

$$\hat{\sigma} \left(\frac{\partial T}{\partial \tilde{t}} + \mathbf{u} \cdot \tilde{\nabla} T \right) = \alpha_{H_2O} \nabla \cdot \hat{\alpha} \nabla T, \quad (5.21)$$

where α_{H_2O} is the molecular diffusivity of water. We next rescale with respect to length (D) and time (D/U) scales to arrive at the nondimensional form,

$$\hat{\sigma} \left(\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T \right) = \frac{1}{Pe} \nabla \cdot \hat{\alpha}(\mathbf{x}) \nabla T, \quad (5.22)$$

where $\mathbf{x} = \tilde{\mathbf{x}}/D$ and ∇ is the gradient with respect to \mathbf{x} . The Peclet number is defined with respect to the fluid properties,

$$Pe = \frac{UD}{\alpha_{H_2O}} = RePr. \quad (5.23)$$

We take the Prandtl number for water to be $Pr = 5.858$ in all cases. Equation (5.22) is homogeneous and we can therefore choose a thermal scale corresponding to a unit surface flux, $q'' \equiv 1$, which is prescribed on $\partial\Omega$,

$$q'' = \frac{\hat{\alpha}}{Pe} \nabla T \cdot \hat{\mathbf{n}} \Big|_{\partial\Omega}. \quad (5.24)$$

The physical flux is recovered as $\tilde{q}'' = [\rho C_p]_{H_2O} U q''$. Note that $\mathbf{u} \equiv \mathbf{0}$ in Ω_s . No further modification of (5.22) is required to accommodate the conjugate heat transfer problem.

Our quantity of interest is the Nusselt number, Nu , which is the the nondimensional heat-transfer coefficient

$$Nu = \frac{Dh}{k_{H_2O}}, \quad (5.25)$$

where h is the heat-transfer coefficient defined by the relationship

$$\tilde{q}'' = h(T_w - T_{bulk}). \quad (5.26)$$

Here, T_{bulk} is the bulk-mean temperature,

$$T_{bulk} = \frac{\int_{\Omega} \langle u_z T \rangle dV}{\int_{\Omega} \langle u_z \rangle dV}, \quad (5.27)$$

where $\langle \cdot \rangle$ indicates a temporal average of the integrand over long time, as specified in Table 5.2. The wall temperature, T_w is meant to represent the average temperature at the interior wall of the pipe. Because of the presence of the wire coil, this temperature is difficult to quantify directly and we therefore estimate this value from the solution of the conduction problem in the copper tubing. The nondimensional equivalent of (5.28) is

$$\theta_w = \frac{\int_{\partial\Omega} \langle T \rangle dS}{\int_{\partial\Omega} dS} - \frac{q'' Pe}{2\hat{\alpha}_{copper}} \log(D_o/D). \quad (5.28)$$

Figure 5.25 shows the scale of the flow for the long helix case with $Re = 10,000$.




Figure 5.25: Velocity magnitude (a), temperature (b), and λ_2 contour (c) for pipe flow with wire-coil insert ($Re = 10,000$).

Geometry

Collins *et al.* has experimentally studied the effects of varying the pitch of the wire-coil inserts in a tube on the Nusselt number. Experiments have been performed for non-dimensional pitches $p/D = (2.67, 2.00, 1.33, 1.00, 0.667, 0.587, 0.533, 0.445)$, and have found that $p/D = 0.587$ maximizes the average Nusselt number of the flow. For our numerical simulations, we have chosen to study the pitches $p/D = (2.67, 0.587, 0.445)$ corresponding to the longest,

medium (optimal), and shortest pitches for $\text{Re} = (5300, 10000)$. We have also performed a study of pipe-flow with an array of rings with spacing lengths equivalent to the pitches of the wire-coil geometries. Finally, we performed a study of a meandering coil-insert: the coil volume is determined by sweeping a ball through a path with sinusoidally $\frac{d\theta}{dz}$ behavior.

Results


Table 5.2 contains the summary of the DoFs, domain length, transient time and integration time. For the averaging process, all of the invariance of the flow in the domain was used for each of the cases.

Table 5.2: Run information for all configurations, $\text{Re}=5,300$.


Type	DoFs	Length (D)	Pitch (D)	VF	τ_{trans} (CTU)	τ_{integ} (CTU)
Straight Pipe	3.7×10^7	15	N/A	1.00	100+	40
Straight Wire	2.3×10^8	15	N/A		50	100
Short Helix	2.6×10^7	0.891	0.445	0.66	200+	50
Medium Helix	1.1×10^7	1.17	0.587	0.74	40	50
Long Helix	8.8×10^7	2.67	2.67	0.92	200+	100
Meandering Helix	2.5×10^7	6.3	6.3	0.92	200+	100
Short Ring	8.2×10^6	0.891	0.446	0.67	200+	50
Medium Ring	2.0×10^7	1.17	0.585	0.74	200+	200+
Long Ring	2.9×10^7	2.67	2.67	0.94	150	50



(a) Short-pitch case ($p/D = 0.445$).



(b) Medium-pitch case ($p/D = 0.587$).



(c) Long-pitch case ($p/D = 2.67$).

Figure 5.26: λ_2 contours for pipe flow wire-coil ($Re = 10,000$) for three pitches.

The following is an alternative view of the previous plots shown in Fig. 5.27. As shown, the numerically experiments is in good agreement with physical experimental results where it exists.




Figure 5.27: Nusselt number vs. Reynolds number comparison between numerical result and experiment.

We observe that ring inserts spaced at equivalent pitches as the wire-coils is more performant than the wire-coil devices.

For the helical cases studied by Collins *et al.* [50], we see fairly good agreement between our numerical result and the experimental result. We have performed an additional set of numerical simulations of flow through uniformly spaced rings in a pipe. The rings show a significant increase in the Nusselt number compared to their wire-coil counterparts. This increase seems to suggest that the recirculation created by the cavity-like geometry is a major factor in the heat-transfer efficiency of the wire-coil geometry. In contrast to the ring cases, the meandering wire-coil case did not show significant increase in the Nusselt number, indicating that the swirling motion that is prevalent in this case does not contribute significantly to the high Nusselt number. In other words, the Taylor dispersion from the

shearing motion is likely not the dominant heat transfer mechanism, thereby showing that the thermal transport from the recirculation is the likely candidate for the heat transfer efficiency of these wire-coil cases compared to the standard pipe flow.

5.8.2 ROM Result for Hydrodynamics of Pipe with Wire-Coil

In the first step of applying the ABM-ROM methodology to this engineering problem, we look to reproduce the hydrodynamics of the long-pitch wire-coil case with Reynolds number 5300. 2000 snapshots were collected over 50 CTUs. The computational domain consists of 44.3 million DoFs with wire-coil pitched at $p/D = 2.67$. One snapshot obtained from the FOM calculation is shown in Fig. 5.28.




Figure 5.28: Velocity magnitude of a pipe flow with wire-coil insert ($Re = 5,300$) snapshot.

We produced both the standard Galerkin POD and ABM ROMs and ran them for 500 CTUs with the projection of the FOM initial condition on to the ROM basis. The result

shown in Fig. 5.29 display accurate solutions for both the ABM-ROM and POD-ROM with constrained optimization, a result consistent with previous examples. From all the examples that we have shown with ABM-ROM, it was successful in reproducing stable behavior that is consistent with the FOM with regard to the mean and TKE metrics. With the adoption of ABM-ROM for challenging flows, we anticipate success in many more problems than is outlined in this document.

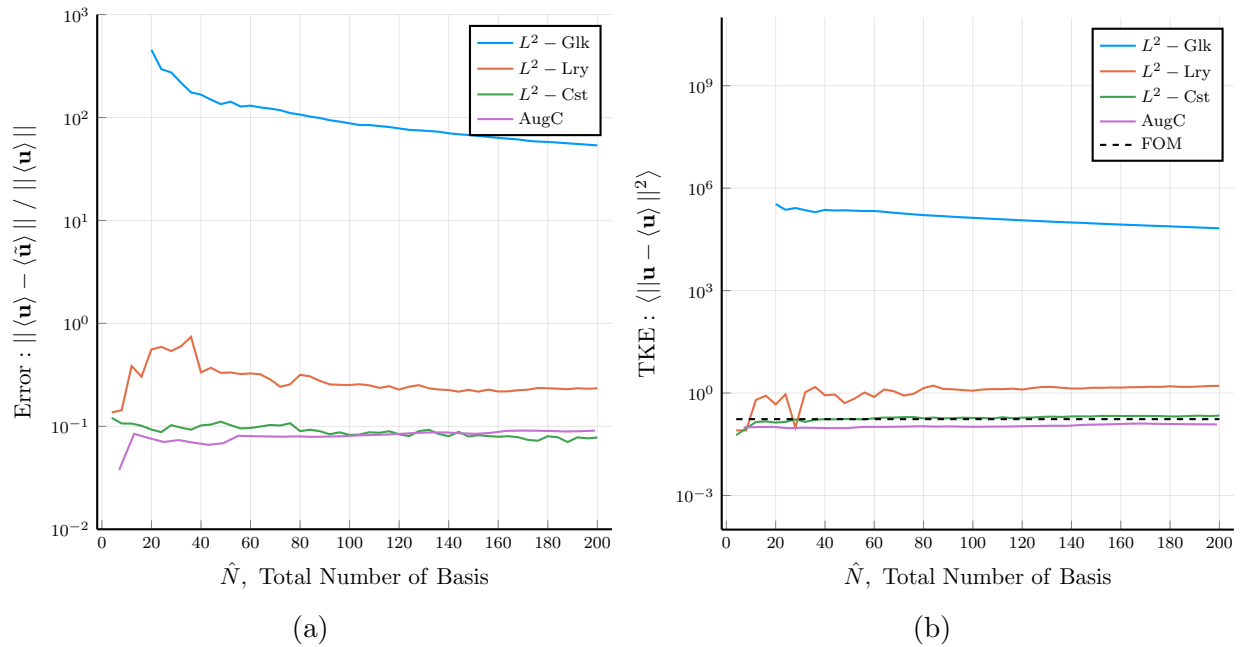


Figure 5.29: ROM mean velocity error (a) and TKE (b) for Wire-coil flow ($Re = 5,300$).

Chapter 6

Investigation of ABM Stabilization Mechanism

In this section, we explore and discuss potential mechanisms for the success of the ABM-ROM and apply these findings to a higher Reynolds number case of pipe flow at $Re = 15,000$.

6.1 Convective Energy-Transfer

In many chaotic flows, the POD truncation can result in issues where the ROM solution is attracted to an unobserved attractor. We suspect a contributing factor of this phenomenon stems from the inability for the existing modes mimic the energy-transfer mechanism of the analytical flow solution. One remedy to this problem is to add more and more POD modes, but this may be costly since there may be many small-scale modes that need to be represented in order to fully diffuse the large-scale energy. Therefore, we consider the energy transfer between the resolved and unresolved modes and show how they are related to the ABM modes.

Previous analysis on convective energy-transfer of POD modes have been performed by Couplet *et al.* [52]. In this work, they have studied the POD modes of a flow past a backward-facing and the convective transfer of energy between these modes. In a similar

manner, we will consider the transfer of energy between the resolved scales (originally the POD modes) and the unresolved scales via a variational multi-scale (VMS) like decomposition.

First, we write the weak form of the equation with known solution $\tilde{\mathbf{u}}(\mathbf{x}, t)$, $\mathbf{x} \in \Omega$, $t \in [t_0, t_f]$

$$(\tilde{\mathbf{v}}, \partial_t \tilde{\mathbf{u}}) = -c(\tilde{\mathbf{v}}, \tilde{\mathbf{u}}, \tilde{\mathbf{u}}) - \nu a(\tilde{\mathbf{v}}, \tilde{\mathbf{u}}), \quad \forall \tilde{\mathbf{v}} \in \tilde{V} = \{\langle \alpha(t) \tilde{\mathbf{u}}(\mathbf{x}, t) \rangle \mid \alpha \in L^2\} \quad (6.1)$$

We can now introduce the truncated POD series to form a space and its complement.

$V = \{\zeta_i\}$, $V' = \tilde{V} \setminus V$. We also introduce projection operators

$$\Pi(\cdot) = \sum_{i=1}^N \frac{(\cdot, \zeta_i)_V}{(\zeta_i, \zeta_i)_V} \zeta_i \quad (6.2)$$

$$\Pi'(\cdot) = (\cdot) - \Pi(\cdot) \quad (6.3)$$

s.t. $V = \Pi \tilde{V}$ and $V' = \Pi' \tilde{V}$. Now, we can restrict the space to get additional equations that describe the evolution of the POD resolved modes and the unresolved modes:

$$(\mathbf{v}, \partial_t \mathbf{u}) = -c(\mathbf{v}, \tilde{\mathbf{u}}, \mathbf{u}) - c(\mathbf{v}, \tilde{\mathbf{u}}, \mathbf{u}') - \nu a(\mathbf{v}, \mathbf{u}) - \nu a(\mathbf{v}, \mathbf{u}'), \quad \forall \mathbf{v} \in V \quad (6.4)$$

$$(\mathbf{v}', \partial_t \mathbf{u}') = -c(\mathbf{v}', \tilde{\mathbf{u}}, \mathbf{u}') - c(\mathbf{v}', \tilde{\mathbf{u}}, \mathbf{u}) - \nu a(\mathbf{v}', \mathbf{u}) - \nu a(\mathbf{v}', \mathbf{u}'), \quad \forall \mathbf{v}' \in V' \quad (6.5)$$

The next step is to consider the energy evolution in each space

$$\partial_t E = -c(\mathbf{u}, \tilde{\mathbf{u}}, \mathbf{u}) - c(\mathbf{u}, \tilde{\mathbf{u}}, \mathbf{u}') - \nu a(\mathbf{u}, \mathbf{u}) - \nu a(\mathbf{u}, \mathbf{u}') \quad (6.6)$$

$$\partial_t E' = -c(\mathbf{u}', \tilde{\mathbf{u}}, \mathbf{u}') - c(\mathbf{u}', \tilde{\mathbf{u}}, \mathbf{u}) - \nu a(\mathbf{u}', \mathbf{u}) - \nu a(\mathbf{u}', \mathbf{u}') \quad (6.7)$$

where

$$E = \frac{1}{2}(\mathbf{u}, \mathbf{u}) \quad (6.8)$$

$$E' = \frac{1}{2}(\mathbf{u}', \mathbf{u}') \quad (6.9)$$

so that we can track the energy evolution of the resolved part of the solution and the unresolved part of the solution independently.

The terms in red in equations (6.6–6.7) evaluate to 0 given $\nabla \cdot \tilde{\mathbf{u}} = 0$ and $\int_{\partial\Omega} \tilde{\mathbf{v}} \cdot \nabla \tilde{\mathbf{u}} \cdot d\mathbf{A} = 0$, $\forall \tilde{\mathbf{v}} \in \tilde{V}$. Thus, the terms that quantify the convective energy transfer between the resolved and unresolved modes are

$$c(\mathbf{u}, \tilde{\mathbf{u}}, \mathbf{u}') \text{ and } c(\mathbf{u}', \tilde{\mathbf{u}}, \mathbf{u}) \quad (6.10)$$

which consequently sum to 0 since $c(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = 0$ i.e., these terms do not represent a net source or sink of energy, but rather a transfer of energy between the resolved scales and the unresolved scales.

One observation is that when the resolved solution space becomes larger and consequently the unresolved space becomes smaller, the magnitude of these terms become smaller. This reduction in energy-transfer between the two spaces can be understood as the resolved solution space becoming self-contained, indicating that the dynamics of the autonomous evolution of the resolved scale should approach the dynamics of the projection of the real solution on to the resolved scales, at least in the global energy sense.

Therefore, when considering augmenting the current basis, we would like to choose modes such that the receiving part of the energy-transfer is well-represented by these additional modes i.e., we want to force the unresolved scales to be more orthogonal to $\tilde{\mathbf{u}} \cdot \nabla \mathbf{u}$. In the case of ABM, our hypothesis is that the additional basis functions contain significant parts of this term such that the unresolved scales become more orthogonal to this term. In this way, ABM can be understood to approximately augmenting the space by selecting pockets of

the energy-cascade that are the receiving end of the energy transfer of the POD modes.

In this framework, ABM can be described as a turbulence-closure model of sorts. Instead of introducing additional transport equations to be solved, as is the case in the $k - \epsilon$ or $k - \omega$ models, we model the action of unresolved scales by an low-dimensional auxiliary system which we can be described as additional modes in the velocity basis so that this auxiliary system can be solved concurrently with the original POD-ROM system.

A more careful study on this subject may reveal more insight into the ABM mechanism and perhaps more optimal basis augmentation that can be employed in the ROM of turbulent flows.

6.2 Time Evolution of Individual ROM Coefficients

Our initial hypothesis for the development of the ABM was that it would provide a mechanism to capture the quadratic interactions of the advection term in the NSE. A similar approach was taken by Akkari and co-workers [21, 22], who develop an alternative to POD basis-extraction that nonetheless uses linear combinations of the FOM snapshots to generate H_0^1 POD modes which is more capable of absorbing (and dissipating) the energy arising from the quadratic interactions. They show nominal success in applying this method to a 3D turbulence problem and track some of the coefficient behavior.

The energy-transfer stabilization hypothesis is supported by the graphs of Fig. 6.1, which show the amplitudes of the basis coefficients for POD and ABM-ROM solutions to pipe flow at $\text{Re} = 5300$ as a function of time and mode number. The coefficient evolutions are shown over three time windows, $[0, 10]$, $[0, 100]$, and $[0, 500]$, which reveal the growth and saturation of the amplitudes. We see that for the ABM (in the lowest row), the coefficients are all smaller than unity, save for the ζ_0 coefficient, which is unity. (All modes have unit L^2 -norm, so the coefficients represent the true amplitude of each scaled mode.) The ABM results also show that most of the energy is in the POD bases, corresponding roughly to

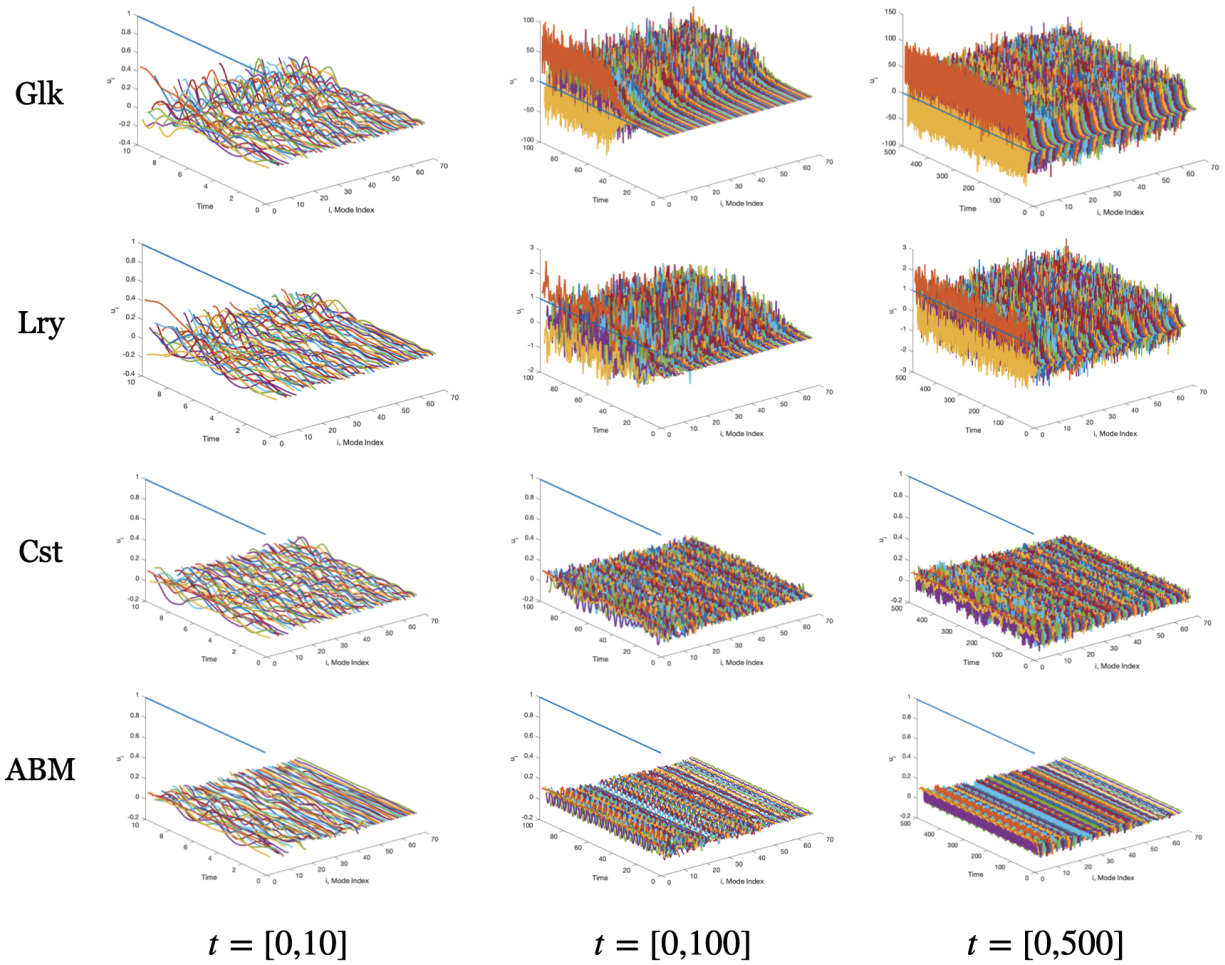


Figure 6.1: Temporal behavior of basis coefficients for standard POD (Glk), Leray-filtered (Lry), Constrained (Cst), and ABM-ROM solutions of turbulent pipe flow at $Re = 5300$ over (convective) time intervals $t \in [0, 10]$, $[0, 100]$, and $[0, 500]$.

the lower third of the mode indices. By contrast, the coefficients for the standard POD Galerkin modes (top row) quickly saturate to amplitudes in excess of 100, and all modes are excited. As is well known from under-resolved Navier-Stokes simulations, there is an energy pile up—manifest as high amplitude modal coefficients—when the representation lacks high wavenumber bases capable of dissipating energy. The Leray-regularized coefficients (second row) exhibit a behavior similar to the standard Galerkin approach, save that the coefficients are much more controlled, which peak amplitudes much closer to unity. The constrained approach (third row) also exhibits chaotic coefficient behavior but at much more controlled amplitudes than either the standard or Leray cases. Remarkably, the evolution on the $t = [0, 100]$ window indicates that the ABM coefficient behavior is nearly time periodic. This indicates a playback of a solution evolution loop or rather a stable attractor that is close to that of the FOM.

6.3 Accuracy Dependency on the Size of Snapshot Set

For the application of the ABM, we have used a fine sampling time of the snapshots to produce the POD modes. In particular, for the pipe-flow case for varying Reynolds number, we have used the same number of snapshots within the sampling interval of the same time-length. We suspect this may be a contributing reason why the high-Reynolds number cases suffer from over-dissipation. That is, the sampling frequency required for the low Reynolds number cases and high Reynolds number cases are different. To test this hypothesis, we have applied subsampling to the set of 1000 snapshots for the $Re = 5, 300$ case to see how the ROMs perform with fewer snapshots taken in the same time-interval. The subsampling comparison was performed for both the standard POD procedure and the AugC procedure. An experiment on the relationship between snapshot sampling interval and POD convergence was performed by Brenner [53]. This only considers convergence in a static sense: projecting the snapshot on to the POD basis generated by different number of snapshots. Unsurprisingly, more snapshots

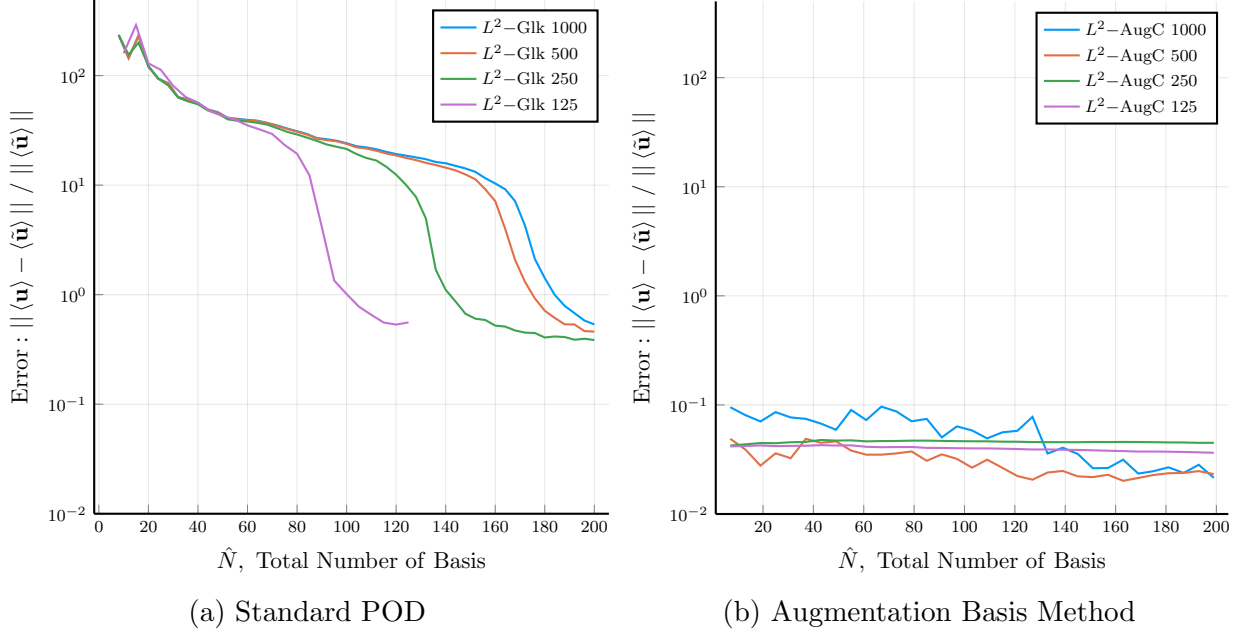


Figure 6.2: Mean error result from subsampled snapshot set, pipe flow (Re = 5300).

resulted in reduction of projection error. Here, we consider dynamical evolution of the POD modes vs. number of snapshots and in particular for the L^2 POD-Galerkin system, we found the opposite is true for the dynamics of these modes.

The results of the subsampling study are shown in Figs. 6.2–6.3. We focus initially on 6.2 (a) which shows the mean-flow error for the standard POD-ROM case as a function of the number of modes $\hat{N} = N$. The modes are drawn from a set of POD bases functions based on K snapshots, where $K = 125, 250, 500, \text{ or } 1000$. Whenever $N = K$ it's clear that Z^N is equivalent to the snapshot space, which implies that the modes contain all the high frequency content present in the snapshots of a turbulent flow solution. We see that these cases have a lower error than cases where the number of modes is a relatively small fraction of the number of snapshots. The same trends are indicated in the TKE plots for the POD-ROM in Fig. 6.3 (a). By contrast, the ABM-ROM needs very few total modes to yield a better estimate of the mean flow Fig. 6.2 (b) and the best TKE predictions are obtained when the snapshot set is large (e.g., $K \geq 500$ in the 6.3 (b)). If we have too few modes in the snapshot space, along with the nonlinear augmentation modes, the ABM-ROM appears to be overly dissipative.

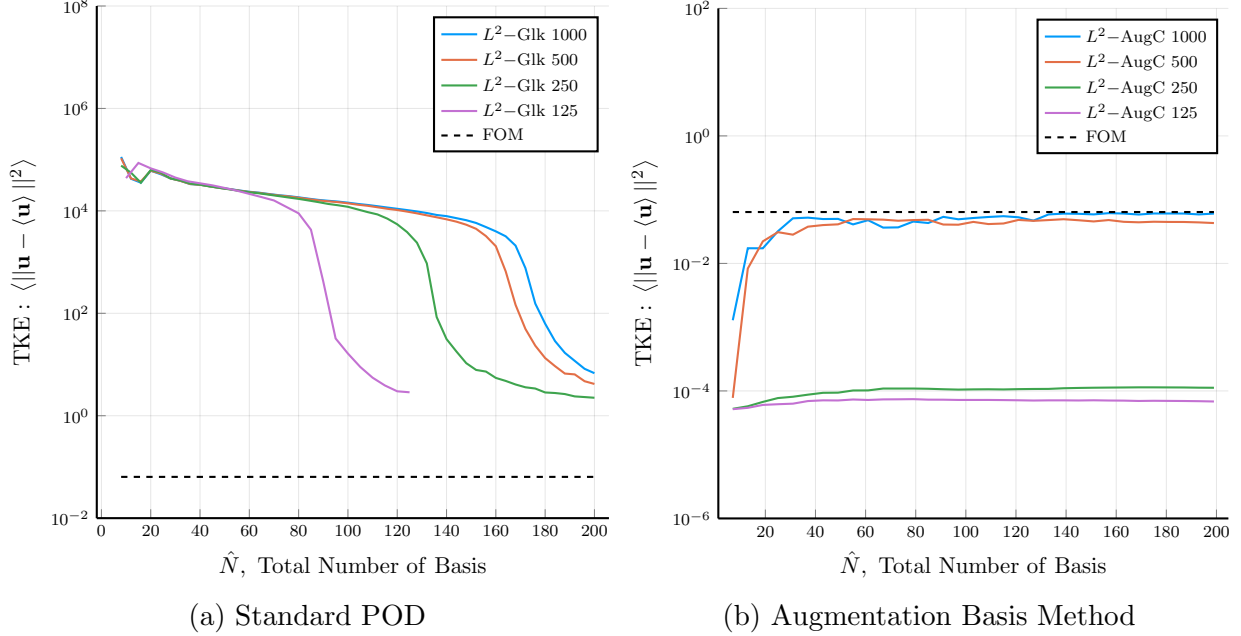


Figure 6.3: TKE result from subsampled snapshot set, pipe flow ($Re = 5300$).

Therefore, we suspect that using more snapshots to produce a more accurate POD series may ensure an accurate ROM reproduction for high Reynolds number pipe flow cases.

6.3.1 Serial Snapshot Selection Experiment

We also show the same experiment as described above except the snapshots are selected serially for the smaller snapshot size cases e.g., for the 125-snapshot case, the first 125 snapshots are used out of the 1000 total to produce the POD modes. Figures 6.4 and 6.5 show the result for the mean velocity error and TKE comparison, respectively.

For the standard L^2 POD-ROM result, we see a similar behavior as before, except there appears to be a higher floor for the lower snapshot count result, although the 500 snapshot case appears to start a convergent behavior earlier on, at about $\hat{N} = 80$. We also note that the 125 snapshot case produced a blowup solution throughout the range of \hat{N} considered. This may suggest that accurately capturing the turbulent flow behavior for a short time (6.25 CTUs for the 125 snapshot case) is not a sufficient condition for long-time stability.

For the ABM approach, we see inconclusive results for the mean metric, but for the

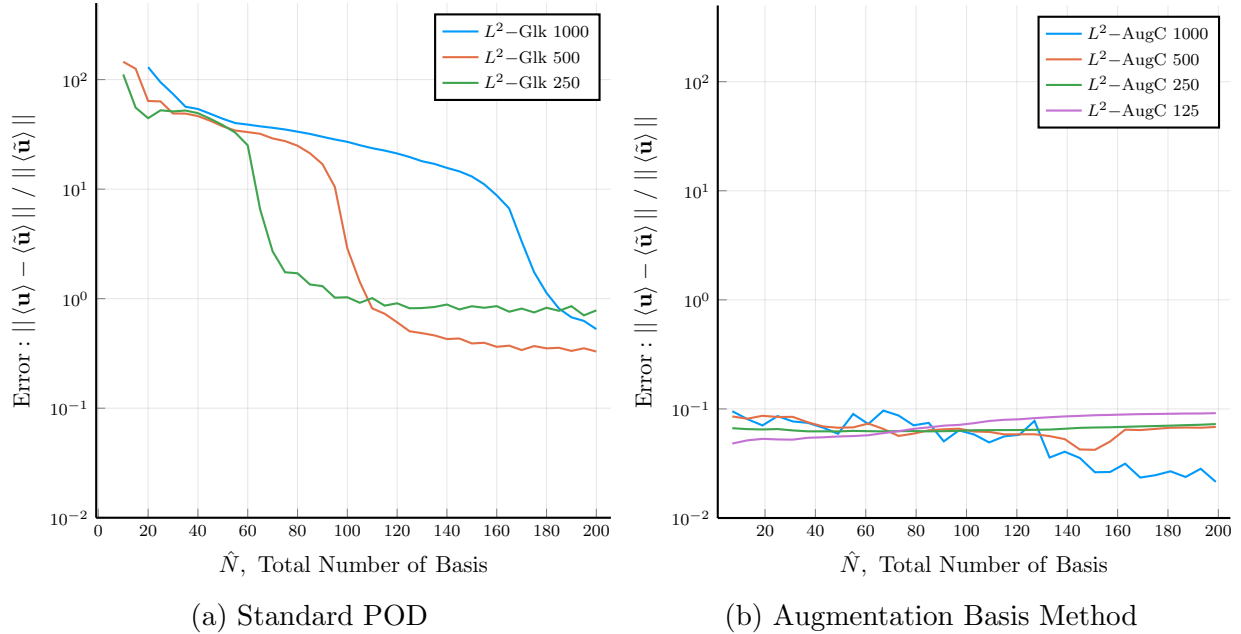


Figure 6.4: Mean error result from serially selected snapshot set, pipe flow ($\text{Re} = 5300$).

TKE comparison, we see that 500 evenly-spaced snapshots perform slightly better compared to 500 serial snapshots. For lower number of snapshots, 250 and 125, we see less of the overly-dissipative behavior.

6.4 Towards Higher Reynolds Number

We have seen promising success in the three Reynolds number that we have considered for the heated pipe flow configuration in Section 5.6, especially for the flow solution. The obvious question is: Are there any limitations or shortcomings to this methodology? One possible limitation is the range of Reynolds number where this method applies. We predict that for a very high Reynolds number solution which have a large range of scales represented, it would still require a large number of base POD modes to capture the large-scale motion which would make the cost prohibitive, even with ABM. Another possibility is that ABM is only effective in the moderate or lower Reynolds number regime. To understand the limitation, we extend the pipe flow case to $\text{Re} = 15,000$. We show the velocity magnitude of the flow field and temperature field of one snapshot of this new FOM case in Fig. 6.6.

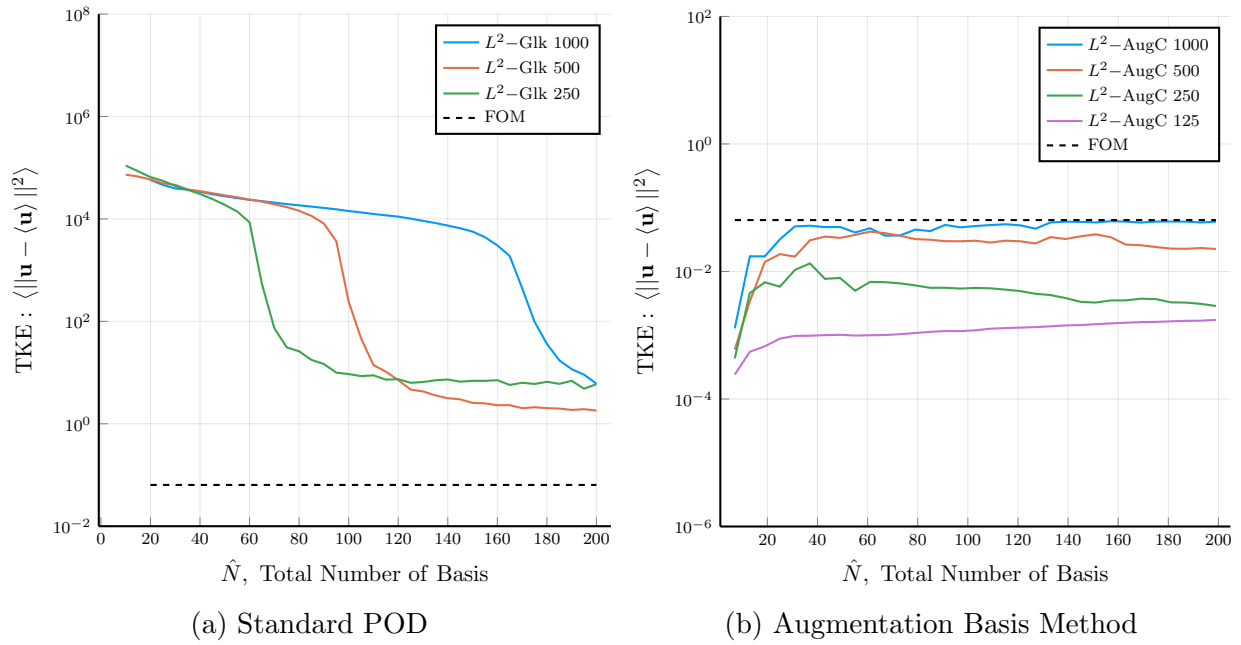


Figure 6.5: TKE result from serially selected snapshot set, pipe flow ($Re = 5300$).




Figure 6.6: Velocity (left) and temperature (right) snapshots for pipe flow $Re = 15,000$.

We start with the same setup as the other three cases: time-stepping up to 50 CTUs and taking 1000 snapshots to be post-processed to generate the ROMs. The mesh consists of 130,048 spectral elements with polynomial order 7 resulting in 62 million DoFs. Figure 6.7 shows the result of the ROM run.

Figure 6.7 shows over-dissipative behavior both for the velocity field and thermal field at Reynolds number 15,000. One hypothesis for the increased error for the velocity evolution is that the higher Reynolds number flows have solution components that have a smaller timescale associated with it than the smallest timescales in the low Reynolds number flows. Thus, if we increased the number of snapshots in the same time-span, we may be able to capture that fast timescale events in the POD modes. Another possibility is that the fidelity of the dominant POD scales are important for the ROM dynamics to match the FOM dynamics. Both of these hypotheses can be addressed by a set of experiments laid out in Section 6.3. These experiments showed that ABM-ROM is successful when there is a large number of snapshots that are processed to produce the POD modes. Thus, we suspect that increasing the number of snapshots for this higher Reynolds number pipe flow case may alleviate the issue. Therefore, we re-ran the simulation and gathered 2000 total snapshots in the same 50 CTU timespan. The resultant ABM-ROM result are shown in Fig. 6.8. Note that the results other than the AugC curves are carried over from the previous plot i.e., they were produced by ROM generated from 1000 snapshots.

We see a significant improvement in the ABM accuracy, especially in the TKE metric. Here, we have demonstrated the snapshot dependency of ABM and this phenomenon must be fully investigated in a future study in order to establish firm conditions for a successful ABM.

With this result in mind, we anticipate that higher Reynolds number flows require a larger number of snapshots i.e., for $Re = 30,000$ flow, more than 2000 snapshots will be required for a successful ABM-ROM to be generated. We do note, however, that this increase did not result in improvement of accuracy for the thermal field evolution.

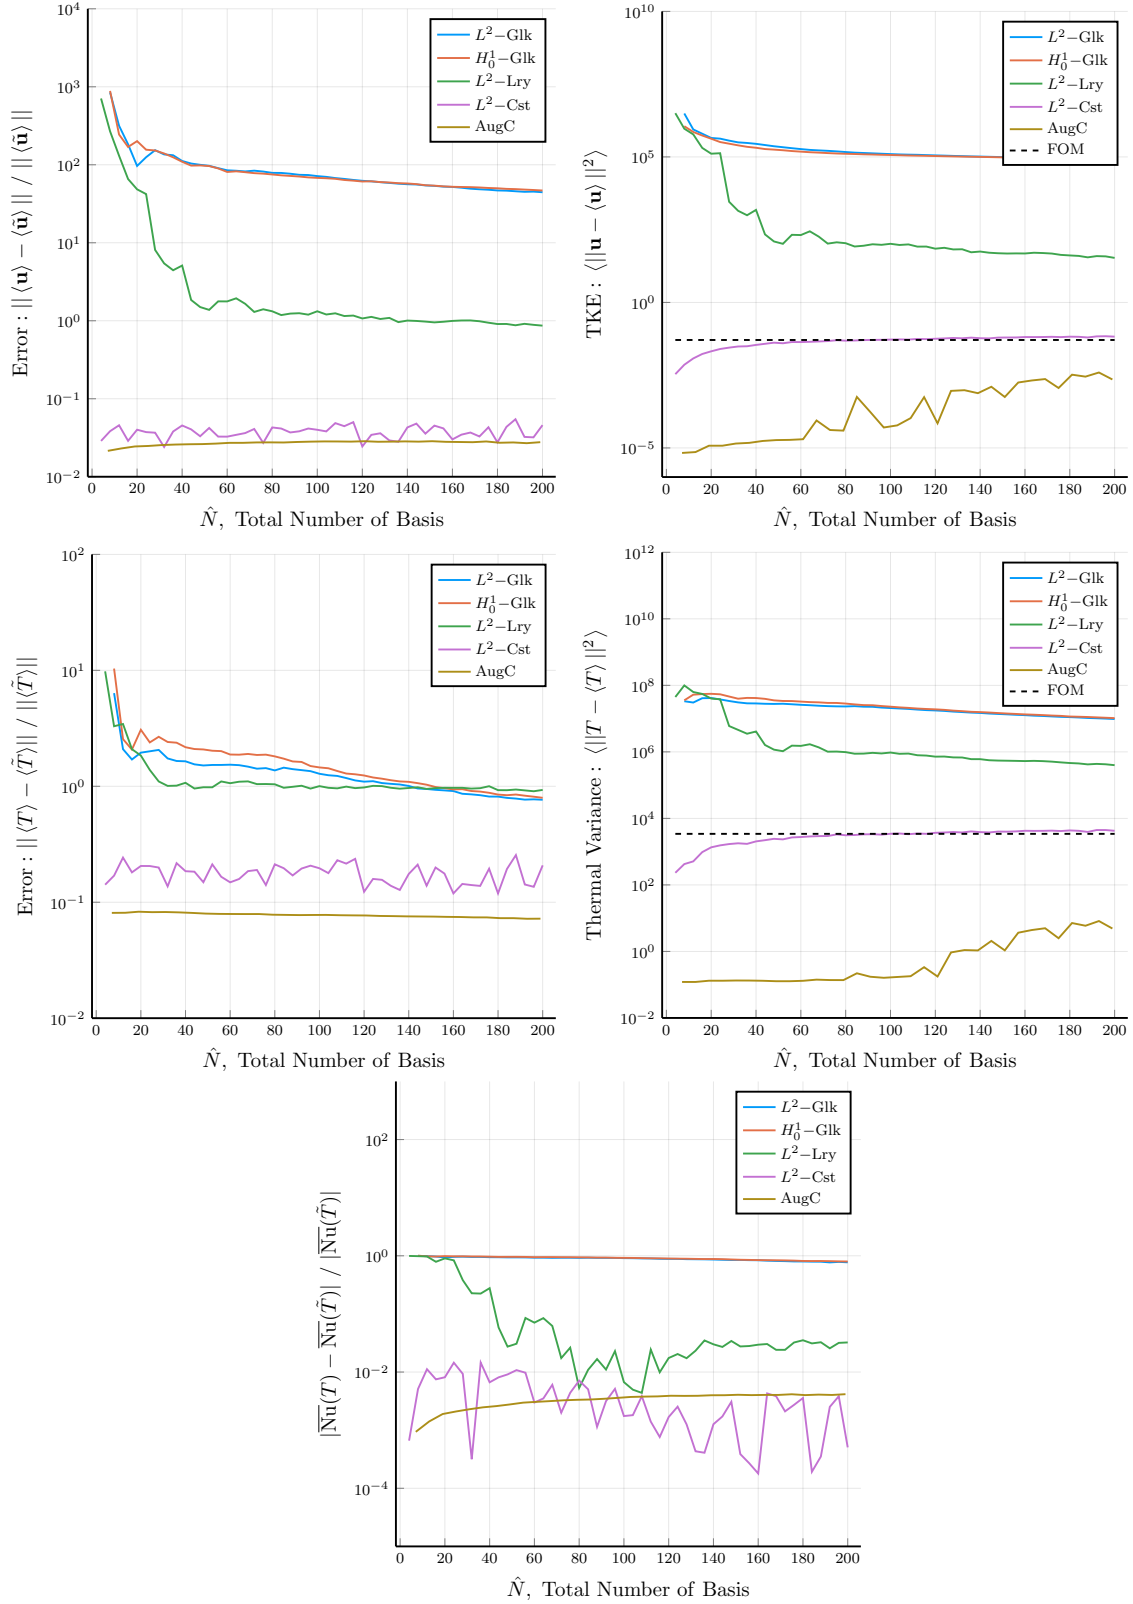


Figure 6.7: From left to right, top to bottom: mean velocity error, TKE, mean temperature error, thermal variance, and Nusselt number result with the new ABM-ROM constructed using 2000 snapshots, pipe flow ($\text{Re} = 15,000$).

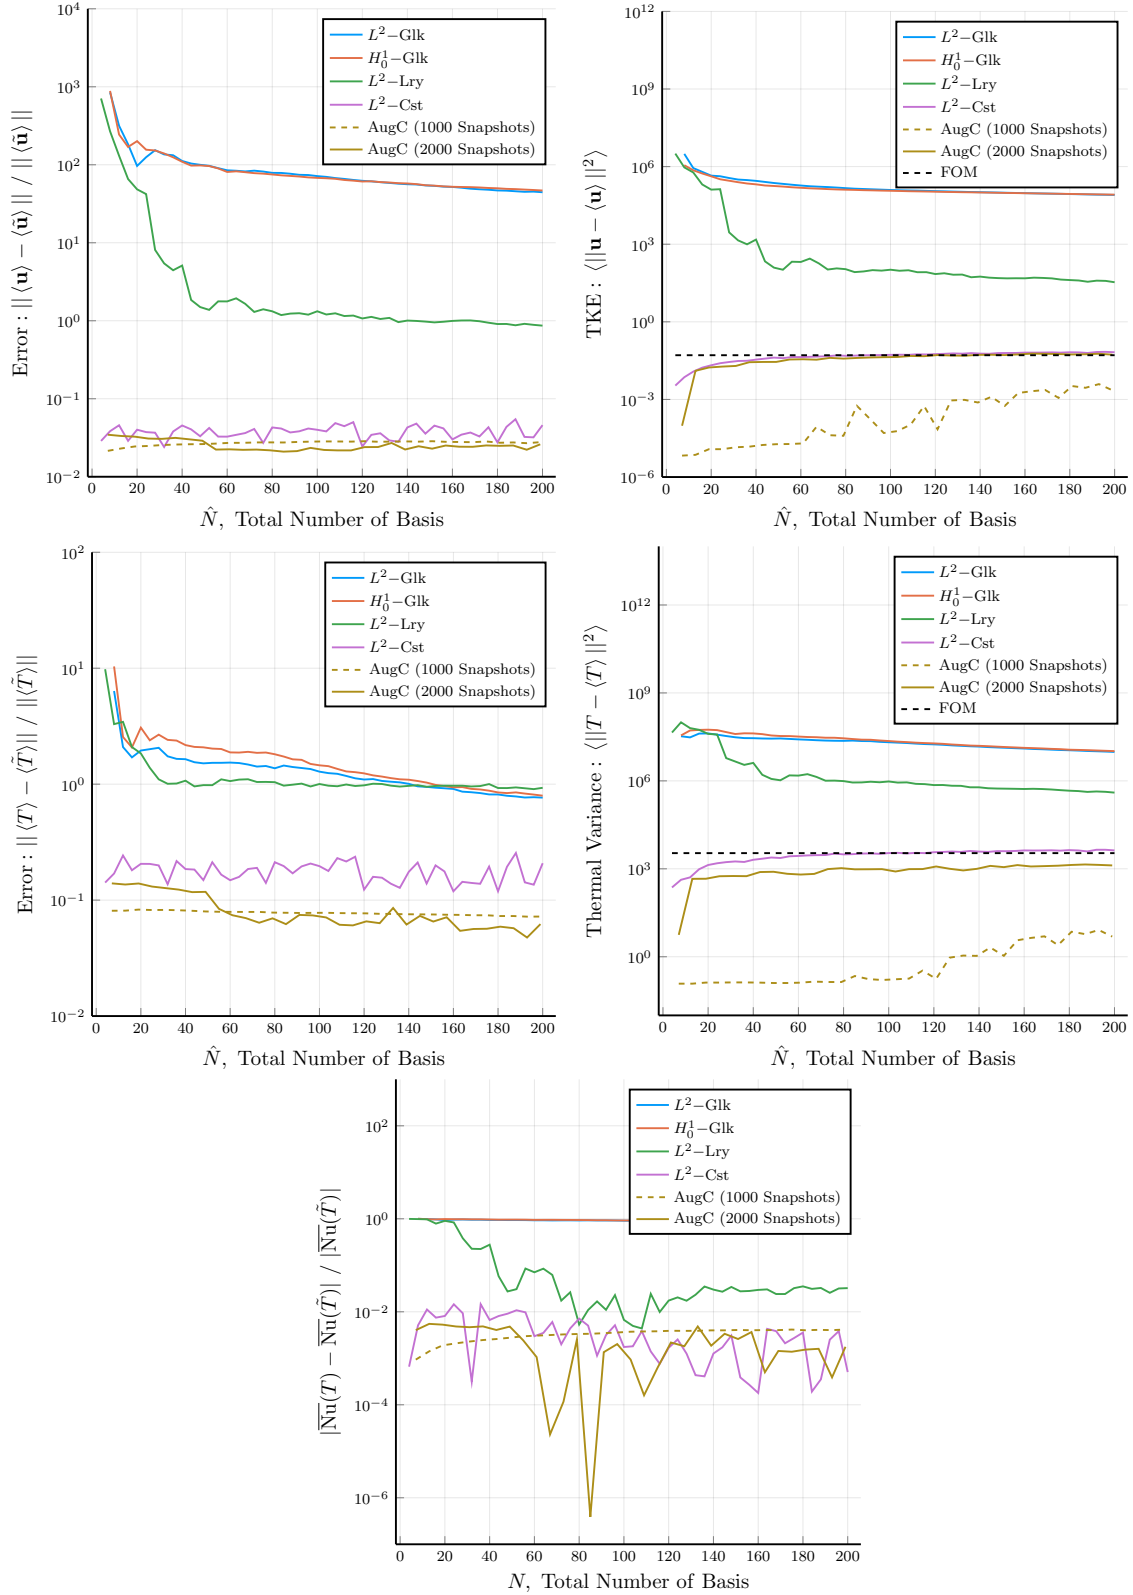


Figure 6.8: From left to right, top to bottom: mean velocity error, TKE, mean temperature error, thermal variance, and Nusselt number result with the new ABM-ROM constructed using 2000 snapshots, pipe flow ($\text{Re} = 15,000$)

Chapter 7

Discussion

ABM has been remarkably successful in advancing our ability to apply ROMs to high-Reynolds number flows. Several observations point to the stabilization properties of the ABM, rather than its approximation quality, as the principal driver for its success. Inspection of the modes for several cases indicate that the augmenting modes in the ABM have high wavenumber content that is localized in Ω to regions of active flow dynamics. We again note that An example is illustrated in Figs. (5.4–5.4), which shows the first 14 L^2 –AugC modes for the case of a lid-driven cavity at $\text{Re} = 30,000$. For $j = 0, \dots, 4$, the first five POD modes, $\zeta_j \in Z^N$, are in the top row; the first five 0-modes, $\mathbb{P}\{\mathbf{u}_0 \cdot \nabla \mathbf{u}_j + \mathbf{u}_j \cdot \nabla \mathbf{u}_0\}$, are in the center row; and the first five diagonal-modes, $\mathbb{P}\{\mathbf{u}_j \cdot \nabla \mathbf{u}_j\}$, are in the lower row. (The 0–0 mode is of course not used twice when forming the augmented basis.) We see that the auto-interaction modes in particular feature high wavenumber content in regions of Ω where the POD modes have significant amplitude. Although it is not shown here, the augmented bases develop high wavenumber content at a much faster rate (i.e., lower mode number) than their high mode-number POD counterparts, which explains why it takes so long for the standard POD Galerkin method to stabilize in the $\text{Re} = 4000, 5300$, and 10000 pipe flow cases of the preceding section. In this sense, the augmenting modes are more wavelet-like than Fourier-like and therefore quite efficient in providing a localized dissipation mechanism

for quadratic interactions. Using these bases thus makes some sacrifice on approximation properties (because we use fewer POD modes, which are optimal in generating low-rank approximations to the snapshot space in the same spirit as low-rank SVD-based matrix decompositions) in favor of better stabilization. Despite this trade-off, the ABM generally yields a much better overall approximation of the dynamics than even its stabilized POD counterparts, as is evident in the turbulent pipe flow case.

This stabilization hypothesis is supported by the graphs of Fig. 6.1, which shows the amplitudes of the basis coefficients for POD and ABM Galerkin ROM solutions to pipe flow at $Re = 5300$ as a function of time and mode number. The coefficient evolutions are shown over three time windows, $[0, 10]$, $[0, 100]$, and $[0, 500]$, which reveal the growth and saturation of the amplitudes. We see that for the ABM (in the lowest row), the coefficients are all smaller than unity, save for the ζ_0 coefficient, which is unity. (All modes have unit 2-norm, so the coefficients represent the true amplitude of each scaled mode.) We reiterate that the ABM results also show that most of the energy is in the POD bases, corresponding roughly to the lower third of the mode indices. By contrast, the coefficients for the standard POD Galerkin modes (top row) quickly saturate to amplitudes in excess of 100, and all modes are excited. As is well known from under-resolved Navier-Stokes simulations, there is an energy pile up—manifest as high amplitude modal coefficients—when the representation lacks high wavenumber bases capable of dissipating energy. The Leray-regularized coefficients (second row) exhibit a behavior similar to the standard Galerkin approach, save that the coefficients are much more controlled, which peak amplitudes much closer to unity. The constrained approach (third row) also exhibits chaotic coefficient behavior but at much more controlled amplitudes than either the standard or Leray cases. Remarkably, the evolution on the $t = [0, 100]$ window indicates that the ABM coefficient behavior is nearly *time periodic*.

We have seen over-dissipative effects that hinder the accuracy of the thermal solution for the pipe-flow cases. When the dissipative effect in the velocity field for the $Re = 15,000$ case has been addressed by a large snapshot set, but the thermal issue remained. One contributing

factor may be the ad hoc nature of the combination of thermal POD modes and velocity POD modes that were used to produce the augmented basis. For the momentum equation, the velocity POD modes are all correlated with itself in the self-interaction modes, but that is not the case for the thermal problem. Since there is no obvious one-to-one correlation, between the two POD basis, it is difficult to establish an appropriate augmented basis. One such idea maybe to use the coordinates associated with each snapshot to create a discrete trajectory of each POD-ROM coefficient for each field. Then find the linear combination of the discrete thermal coefficient evolution to find the best match.

For the present work, we have shown that there exists scenarios in which the ABM-ROM system is too dissipative. We can consider applying the constrained optimization approach on top of ABM, but we anticipate that there will not fully resolve the problem because the optimal mean value of individual coefficients is 0 if the lifting function was the time-averaged solution. Thus, the maximum coefficient value should be a positive number and the minimum value should be a negative number. Therefore, if the dynamical system was maximally dissipative, meaning the coefficients were all 0, the constraints will activate. In this sense the constrained approach may fail.

As an alternative, we can look at the range of TKE of the time-dependent FOM solution and use that as a new constraint metric. Because the TKE in the ROM system can be computed by a sum of the squares of the coefficient values,

$$\text{TKE} = \frac{1}{2} \int_{\Omega} \mathbf{u}' \cdot \mathbf{u}' dV = \frac{1}{2} \sum_{i=1}^N u_i^2 \quad (7.1)$$

we can apply a spherical shell constraint on the ROM solution at each time-step. This way, we can ensure that the TKE observed in the ROM is within the observed range. Although this may mitigate the short-comings of the standard constrained optimization procedure for the NSE-ROM, again, this method will not apply in a p-greedy algorithm.

We close with a remark about an alternative strategy to addressing high Reynolds number

flows, namely, approximation approaches that enable one to use large numbers of modes, N . Based on the results of Figs. 5.17–5.19, we can see that the naive brute-force approach of increasing the number of POD basis functions is doomed to fail: the number of modes required increases rapidly with increasing Reynolds number and even in the converged limit the accuracy is low. It seems unlikely that these shortcomings could be averted even if one could reach, say, $N = 500$ in an economic way.

We close with a remark about other reduction algorithm (e.g., DEIM) as a possible alternative to ABM. Approximation methods such as DEIM [5–7] and tensor-decomposition address the $O(N^3)$ costs (for the advection operator, and other high-order costs for other operators), but do not directly address the slow convergence of POD-Galerkin to the dynamics of the NSE. Although DEIM allows for larger number of modes in the ROM for a given cost, its accuracy will not surpass that of the underlying ROM formulation on which it is based. So, for a classic L^2 - or H^1 -based formulation, DEIM will not yield an acceptable reconstruction result even at $N = 200$, whereas the constrained and ABM formulations realize convergence at much lower values of \hat{N} and much lower costs.

Chapter 8

Conclusion

We introduced a novel stabilization method, ABM, for ROM-based simulations of incompressible turbulent flows that augments the standard POD basis with approximate temporal derivatives. For a space of POD basis functions, $Z^N = \{\zeta_i\}$, $i = 0, \dots, N$, we include and additional $2N + 1$ functions that are the Leray (divergence-free) projections of the nonlinear interactions with the lifting mode, $\{\zeta_0 \cdot \nabla \zeta_i + \zeta_i \cdot \nabla \zeta_0\}$, and nonlinear auto-interactions, $\{\zeta_i \cdot \nabla \zeta_i\}$. With these basis functions, the ROM proceeds in the standard Galerkin fashion and is seen to dramatically outperform standard L^2 - and H_0^1 -POD Galerkin ROM approaches as well as Leray-stabilized methods introduced by Wells *et al.* [4, 17]. ABM performs comparably to the constraint-based stabilization approach of [14], but the latter is restricted to the ROM reproduction problem (i.e., running at the same parameter points as the originating FOM) because, in a pMOR setting, the correct basis-coefficient limits are not known at training points other than the anchor points.

We showed that the auxiliary modes of the ABM have high wavenumber content that is localized to regions in Ω where flow gradients are large and thus provide efficient dissipation mechanisms that are lacking in standard POD bases. We further demonstrated that, for standard POD methods, having a more complete POD space (i.e., incorporating $N \approx K$ modes from a relatively small snapshot space of rank K) yields lower errors than having

$N' > N$ POD modes from a larger snapshot space of rank $K' > N'$. The reasoning is the same—the more complete space includes high wavenumber content in the ROM basis set that provides dissipation and hence stability. Analysis of the ROM coefficient time-traces for turbulent pipe flow at $\text{Re} = 5300$, illustrated that the the amplitudes of all the modes for non-stabilized POD-ROM are orders of magnitude larger than their stabilized counterparts. While Leray-based stabilization mitigates this behavior, it still yields coefficient amplitudes that are roughly a factor of ten greater than observed in either the constrained or ABM-based formulations.

The ABM was also shown to be effective for predicting thermal QOIs such as Nusselt numbers. It was, however, a slightly over-dissipative at $\text{Re} = 10,000$. The study of the interplay between N and K indicates that this dissipation can be controlled with these two parameters and one might therefore use these parameters to gain insight to the root cause of the over-dissipation. Future work will include application of the ABM to even higher Reynolds number flows and to more complex domains.

References

- [1] D. Gottlieb and S. Orszag, *Numerical Analysis of Spectral Methods: Theory and Applications*. Philadelphia: SIAM-CBMS, 1977.
- [2] S. Orszag, “Advanced stability theory analyses for laminar flow control,” NASA Langley Research Center, Hampton, VA, United States, Tech. Rep. NASA Contractor Report 3233, 1980.
- [3] K. Ravikumar, D. Appelhans, and P. Yeung, “Acceleration of extreme scale pseudo-spectral simulations of turbulence using asynchronism,” in *Proc. of the Int. Conf. for High Perf. Comp., Net., Storage and Analysis*, 2019, 8 (1–22). [Online]. Available: <https://doi.org/10.1145/3295500.3356209>.
- [4] K. Kaneko, P.-H. Tsai, and P. Fischer, “Towards model order reduction for fluid-thermal analysis,” *Nuclear Engineering and Design*, vol. 370, p. 110 866, 2020.
- [5] S. Chaturantabut and D. C. Sorensen, “Nonlinear model reduction via discrete empirical interpolation,” *SIAM J. Sci. Comput.*, vol. 32, pp. 2737–2764, 5 2010.
- [6] —, “Discrete empirical interpolation for nonlinear model reduction,” *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pp. 4316–4321, 2009.
- [7] —, “Application of POD and DEIM on dimension reduction of non-linear miscible viscous fingering in porous media,” *Mathematical and Computer Modelling of Dynamical Systems*, vol. 17, no. 4, pp. 337–353, 2011.

- [8] K. Kaneko and P. Fischer, “Augmented reduced order models for turbulence,” *Frontiers in Physics*, 2022.
- [9] J. Lumley, “The structure of inhomogeneous turbulent flows. in: Yaglom, a.m. and tartarsky, v.i., eds. atmospheric turbulence and radio wave propagation,” pp. 166–177, 1967.
- [10] P. Holmes, J. L. Lumley, and G. Berkooz, *Turbulence, Coherent Structure, Dynamical Systems and Symmetry*. Cambridge University Press, 1996.
- [11] M. Barrault, Y. Maday, N. Nguyen, and A. Patera, “An ‘empirical interpolation’ method: Application to efficient reduced-basis discretization of partial differential equations,” *Comptes Rendus Mathematique*, vol. 339, pp. 667–672, 9 2004.
- [12] H. C. Elman and V. Forstall, “Numerical solution of the parameterized steady-state Navier-Stokes equations using empirical interpolation methods,” *Computer Methods in Applied Mechanics and Engineering*, vol. 317, pp. 380–399, 2017.
- [13] K. Veroy, D. V. Rovas, and A. Patera, “A posteriori error estimation for reduced-basis approximation of parametrized elliptic coercive partial differential equations: “convex inverse” bound conditioners,” *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 8, pp. 1007–1028, 2002.
- [14] L. H. Fick, Y. Maday, A. T. Patera, and T. Taddei, “A stabilized pod model for turbulent flows over a range of reynolds numbers: Optimal parameter sampling and constrained projection,” *Journal of Computational Physics*, vol. 371, pp. 214–243, 2018.
- [15] P.-H. Tsai and P. Fischer, “Parametric model-order-reduction development for unsteady convection,” *Frontiers in Physics*, 2022.
- [16] A. Iollo, S. Lanteri, and J. A. Desideri, “Stability properties of pod–galerkin approximations for the compressible navier–stokes equations,” *Theoret. Comput. Fluid Dynamics*, vol. 13, pp. 377–396, 2000.

- [17] D. Wells, Z. Wang, X. Xie, and T. Iliescu, “An evolve-then-filter regularized reduced order model for convection-dominated flows,” *International Journal for Numerical Methods in Fluids*, vol. 84, pp. 598–615, 10 2017.
- [18] M. Balajewicz, “A new approach to model order reduction of the navier-stokes equations,” Ph.D. dissertation, Duke University, 2012.
- [19] D. Amsallem and C. Farhat, “Stabilization of projection-based reduced-order models,” *International Journal for Numerical Methods in Engineering*, vol. 91, no. 4, pp. 358–377, Jun. 2012.
- [20] W. Cazemier, R. W. C. P. Verstappen, and A. E. P. Veldman, “Proper orthogonal decomposition and low-dimensional models for driven cavity flows,” *Physics of Fluids*, vol. 10, no. 7, pp. 1685–1699, 1998.
- [21] N. Akkari, R. Mercier, G. Lartigue, and V. Moureau, “Stable POD-Galerkin Reduced Order Models for unsteady turbulent incompressible flows,” in *55th AIAA Aerospace Sciences Meeting*, Grapevine, United States, 2017.
- [22] N. Akkari, F. Casenave, and V. Moureau, “Time stable reduced order modeling by an enhanced reduced order basis of the turbulent and incompressible 3d navier–stokes equations,” *Mathematical and Computational Applications*, vol. 24, no. 2, 2019.
- [23] E. Merzari, H. Ninokata, A. Mahmood, and M. Rohde, “Proper orthogonal decomposition of the flow in geometries containing a narrow gap,” *Theoretical and Computational Fluid Dynamics*, vol. 23, no. 5, pp. 333–351, 2009.
- [24] E. Merzari, W. D. Pointer, and P. Fischer, “A pod-based solver for the advection-diffusion equation,” in *ASME-JSME-KSME 2011 Joint Fluids Engineering Conference*, American Society of Mechanical Engineers Digital Collection, pp. 1139–1147.
- [25] H. F. S. Lui and W. R. Wolf, “Construction of reduced-order models for fluid flows using deep feedforward neural networks,” *JFM*, vol. 872, pp. 963–994, 2019.

- [26] M. Khodkar, P. Hassanzadeh, S. Nabi, and P. Grover, “Reduced-order modeling of fully turbulent buoyancy-driven flows using the green’s function method,” *Physical Review Fluids*, vol. 4, no. 1, p. 013 801, 2019.
- [27] J.-L. Guermond, S. Prudhomme, and J. T. Oden, “An interpretation of the Navier-Stokes-alpha model as a frame-indifferent leray regularization,” *Physica D*, vol. 177, pp. 23–30, 2003.
- [28] S. Orszag, M. Israeli, and M. Deville, “Boundary conditions for incompressible flows.,” *J. Sci. Comp.*, vol. 1, pp. 75–111, 1986.
- [29] A. Tomboulides, M. Israeli, and G. Karniadakis, “Efficient removal of boundary-divergence errors in time-splitting methods,” *J. Sci. Comput.*, vol. 4, pp. 291–308, 1989.
- [30] A. Tomboulides, J. Lee, and S. Orszag, “Numerical simulation of low Mach number reactive flows,” *J. of Scientific Computing*, vol. 12, pp. 139–167, 1997.
- [31] P. Fischer, “An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations,” *J. Comp. Phys.*, vol. 133, pp. 84–101, 1997.
- [32] A. T. Patera, “A spectral element method for fluid dynamics: Laminar flow in a channel expansion,” *Journal of Computational Physics*, vol. 54, pp. 468–488, 1983.
- [33] P. Fischer, M. Schmitt, and A. Tomboulides, “Recent developments in spectral element simulations of moving-domain problems,” in *Recent Progress and Modern Challenges in Applied Mathematics, Modeling and Computational Science*, R. Melnik, Ed. Fields Institute Communications, Springer, 2017, vol. 79, pp. 213–244.
- [34] K. Ding, “Free response of a freely-rotatable eccentric linearly-sprung circular cylinder in or absent a cross-flow,” Ph.D. dissertation, University of Illinois Urbana-Champaign, 2021.

- [35] G. Stabile, S. Hijazi, A. Mola, S. Lorenzi, and G. Rozza, “Pod-galerkin reduced order methods for cfd using finite volume discretisation: Vortex shedding around a circular cylinder,” *Communications in Applied and Industrial Mathematics*, vol. 8, no. 1, pp. 210–236, 2017.
- [36] M. Bergmann, C.-H. Bruneau, and A. Iollo, “Enablers for robust pod models,” *Journal of Computational Physics*, vol. 228, no. 2, pp. 516–538, 2009.
- [37] D. Barkley and L. Tuckerman, “Traveling waves in axisymmetric convection: The role of sidewall conductivity,” *Physica D.*, vol. 37, pp. 288–294, 1989.
- [38] L. Tuckerman and D. Barkley, “Global bifurcation to traveling waves in axisymmetric convection,” *Phys. Rev. Lett.*, vol. 61, pp. 408–411, 4 1988.
- [39] G. Fox, M. Johnson, G. Lyzenga, S. Otto, J. Salmon, D. Walker, and R. L. White, “Solving problems on concurrent processors vol. 1: General techniques and regular problems,” *Computers in Physics*, vol. 3, no. 1, pp. 83–84, 1989.
- [40] P. F. Fischer, J. W. Lottes, D. W. Pointer, and A. R. Siegel, “Petascale algorithms for reactor hydrodynamics,” 2008.
- [41] D. Rempfer, “On low-dimensional Galerkin models for fluid flow,” *Theor. and Comput. Fluid Dyn.*, vol. 14(2), pp. 75–88, 2000.
- [42] J.-L. Guermond, J. T. Oden, and S. Prudhomme, “Mathematical perspectives on large eddy simulation models for turbulent flows,” *J. Math. Fluid Mech.*, vol. 6, pp. 194–248, 2004.
- [43] J. Mullen, “Development of a parallel spectral element based large eddy simulation model for the flow of incompressible fluids in complex geometries,” Division of Engineering, Ph.D. dissertation, Brown University, 1999.
- [44] P. Fischer, J. Kruse, J. Mullen, H. Tufo, J. Lottes, and S. Kerkemeier, “Nek5000: Open source spectral element cfd solver,” *Argonne National Laboratory, Mathematics and Computer Science Division, Argonne, IL*, see <https://nek5000.mcs.anl.gov>, 2008.

- [45] L. Sirovich, “Turbulence and the dynamics of coherent structures. part 1 : Coherent structures,” *Quarterly of Applied Mathematics*, vol. 45, pp. 561–571, 1987.
- [46] M. Acalar and C. Smith, “A study of hairpin vortices in a laminar boundary layer: Part 1, hairpin vortices generated by a hemisphere protuberance,” *J. Fluid Mech.*, vol. 175, pp. 1–41, 1987.
- [47] J. Jeong and F. Hussain, “On the identification of a vortex,” *Journal of Fluid Mechanics*, vol. 285, pp. 69–94, 1995.
- [48] N. Ghaddar, G. Karniadakis, and A. Patera, “A conservative isoparametric spectral element method for forced convection: Application to fully developed flow in periodic geometries,” *Numer. Heat Transfer*, vol. 9, pp. 277–300, 1986.
- [49] S. Patankar, C. Liu, and E. Sparrow, “Fully developed flow and heat transfer in ducts having streamwise-periodic variations of cross-sectional area,” *J. Heat Transfer*, vol. 99, pp. 180–186, 1977.
- [50] J. T. Collins, C. M. Conley, J. M. Attig, and M. M. Baehl, “Enhanced heat transfer using wire-coil inserts for high-heat-load applications,” *Int. Workshop on Mechanical Engineering Design of Synchrotron Radiation Equipment and Instrumentation*, pp. 409–419, 2002.
- [51] J. T. Collins, private communication, 2020.
- [52] M. Couplet, P. Sagaut, and C. Basdevant, “Intermodal energy transfers in a proper orthogonal decomposition–galerkin representation of a turbulent separated flow,” *Journal of Fluid Mechanics*, vol. 491, pp. 275–284, 2003.
- [53] T. Brenner, “Practical aspects of the implementation of reduced-order models based on proper orthogonal decomposition,” Ph.D. dissertation, Texas A&M University, 2011.