SANDIA REPORT
SAND2022-12742
Printed Sept 2022

Basis Adaptation for Steady Reduced-Order Models

David S. Ching, Francesco Rizzi, Jaideep Ray

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

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

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

Available to DOE and DOE contractors from

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

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

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.gov
Online order:	https://classic.ntis.gov/help/order-methods

ABSTRACT

Projection based reduced-order models generally have poor accuracy in problems with shocks or near-discontinuities due to the inability of a linear basis to resolve shocks at different locations in the domain. Here we develop two novel basis adaptation methods to improve reduced-order model accuracy for steady problems. The first method uses gradients to adapt the basis and the second solves a optimization problem to improve the basis at selected points with high error. The basis adaptation methods are tested on two steady problems: the one dimensional Burgers equation with a spatially varying source and two dimensional hypersonic flow. In both cases, basis adaptation decreases errors, particularly the heat flux error. This is a preliminary study to examine accuracy of the proposed methods, and much work remains to be done for an efficient and robust implementation.

ACKNOWLEDGMENT

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

This project builds on previous work done at Sandia by a number of researchers, including Patrick Blonigan, Elizabeth Krath, Eric Parish, and Marco Arienti. This project would not have been possible without their contributions. We are also grateful to the SPARC team for supporting the Pressio interface and providing advice on implementation.

We would also like to thank Elizabeth Krath for her assistance setting up scripts and her review of our work.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

In many applications, high-fidelity simulations are computationally expensive, and running a suite of high-fidelity simulations for many-query problems becomes computationally intractable, necessitating the need for computationally inexpensive surrogate models. Projection-based reduced order models (pROMs) are surrogate models that solve the governing equations on a reduced basis, thereby achieving high accuracy at relatively low cost. Projection-based reduced order models have shown success in a variety of applications such as heat transfer, low speed fluid flow, and solid mechanics, but have also been demonstrated to have poor accuracy in applications with shocks or near-discontinuities. One example of this is hypersonic aerodynamics, where the shock position can change with Mach number or angle-of-attack and a linear basis is insufficient to represent the shock at different locations in the domain.

A variety of methods have been developed to improve the accuracy for shock driven problems. Methods that replace the linear basis with a convolutional autoencoder have been demonstrated to have higher accuracy than a linear basis [10], but the cost of training an autoencoder makes the method prohibitive for large problems such as hypersonic aerodynamics where the state space dimension can be well over a million. Another method is online basis vector sieving, which successively splits basis vectors, but this method is inefficient at resolving shocks [6]. Registration-based methods move the location of the shock in the domain using a PDE or rules so that the shock location is at nearly fixed location [12, 11, 19, 7, 16, 4]. One particular shock moving method of interest here is grid-tailoring, in which the location of grid cells are adjusted during the high-fidelity simulation so that the shock is a fixed number of cells from the inlet [4]. In Ref [4], it was shown that standard linear pROMs have poor accuracy in hypersonic flows, while grid-tailored pROMs generally have much higher accuracy. However, it was also found that significant errors will occur if the grid-tailoring algorithm does not closely align the grid with the shock. Grid-tailored pROMs also cannot deal with secondary shocks. We believe that registration-based methods are needed to obtain accurate pROMs for shock-dominated problems, but further improvements are still needed.

One method which has shown success is basis adaptivity methods, developed by Ref. [15] for transient transport-dominated flows. Basis adaptivity methods solve the full order model (FOM) governing equations at a limited number of points in the domain and use those solutions to improve the basis during the online stage [15, 14, 9, 1]. These methods have been shown to improve accuracy significantly but have a only been developed for transient flows, whereas in hypersonic aerodynamics the simulations are simulated as steady problems due to the extremely short timescales of the problem. In this work we develop new basis adaptivity methods targeting steady problems. These methods are intended to be applied in combination with grid-tailoring, where the shock locations in different solutions are at nearly the same location in the grid.

2. PROJECTION-BASED REDUCED-ORDER MODELS (PROMS)

2.1. Steady-state Projection-based Reduced-Order Models

Projection-based Reduced-Order Models (pROMs) use the full-order model (FOM) governing equations with solutions constrained to a basis computed from FOM solutions. For time-varying systems, the governing equation is defined as:

$$\frac{\partial \boldsymbol{x}}{\partial t} = \boldsymbol{f}(\boldsymbol{x}, t, \boldsymbol{\mu}) \tag{2.1}$$

where $\boldsymbol{x} \in \mathbb{R}^N$ is the state vector and $\boldsymbol{\mu}$ is the parameter vector. For steady-state problems, this reduces to

$$0 = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{\mu}) \tag{2.2}$$

The state $\boldsymbol{x}$ is approximated as

$$\tilde{\boldsymbol{x}} = \boldsymbol{\Phi}\hat{\boldsymbol{x}} + \boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}) \tag{2.3}$$

where $\boldsymbol{\Phi} \in \mathbb{R}^{N \times p}$ is the basis and $\hat{\boldsymbol{x}} \in \mathbb{R}^p$ is the ROM state. Generally, $p << N$, so the linear basis assumption allows the pROM to operate in a much lower dimensional space.

In Least-Squares Petrov-Galerkin (LSPG) pROMs, the governing equation is not solved exactly; rather the 2-norm of the residual is minimized:

$$\hat{\boldsymbol{x}} = \underset{\hat{\boldsymbol{z}} \in \mathbb{R}^p}{\arg\min} \left\| \boldsymbol{f}(\boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{z}}; \boldsymbol{\mu}) \right\|_2^2 . \tag{2.4}$$

We note that this formulation uses conservative variables and no weighting matrix for the inner product; see Ref. [4] for details on primitive variable pROMs and Ref. [13] for pROMs with entropy-based inner products.

The first order optimality conditions can be determined analytically as:

$$\boldsymbol{\Psi}^T \boldsymbol{f}(\boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{x}}; \boldsymbol{\mu}) = \boldsymbol{0}, \tag{2.5}$$

for which the optimal choice of test basis $\boldsymbol{\Psi} \in \mathbb{R}^{N \times p}$ is

$$\boldsymbol{\Psi} = \left. \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} \right|_{\tilde{\boldsymbol{x}}} \boldsymbol{\Phi}. \tag{2.6}$$

In practice, Eqn. 2.5 is solved using an optimizer, such as a Gauss-Newton optimizer. Each Gauss-Newton update is

$$\hat{\boldsymbol{x}}^{k+1} = \hat{\boldsymbol{x}}^k - (\boldsymbol{\Psi}^T \boldsymbol{\Psi})^{-1} \boldsymbol{\Psi}^T \boldsymbol{f}(\boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{x}}) \tag{2.7}$$

2.2. Steady-state Basis Adaptation

The linear basis assumption in Eqn. 2.3 has been shown to give accurate solutions in a variety of applications such as heat transfer, low-speed flow, and solid mechanics, but is known to result in poor accuracy in shock-dominated problems. In problems where the shock location can vary, linear combinations of the basis vectors can only predict solutions with shocks at the locations of the training data set shocks. Basis adaptation was developed by Ref. [15] to allow for solutions unconstrained by the linear basis assumption. However, the methods developed in Refs. [15] and [14] only work for transient problems. Here we develop two basis adaptation methods for steady flows.

2.2.1. Method 1: Gradient-based adaptation

This method uses gradients to adapt the basis so the residual ℓ^2 norm is decreased. Equation 2.4 can be rewritten as

$$\hat{\boldsymbol{x}} = \underset{\hat{z} \in \mathbb{R}^p}{\arg\min} \boldsymbol{f}(\tilde{z})^T \boldsymbol{f}(\tilde{z}). \tag{2.8}$$

where $\tilde{z} = \boldsymbol{\Phi}\hat{z} + \boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu})$. The gradient of $f(\tilde{\boldsymbol{x}})^T f(\tilde{\boldsymbol{x}})$ is

$$\frac{\partial \boldsymbol{f}^T \boldsymbol{f}}{\partial \tilde{\boldsymbol{x}}} = \left(\left. \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} \right|_{\tilde{\boldsymbol{x}}} \right)^T \boldsymbol{f}(\tilde{\boldsymbol{x}}) \tag{2.9}$$

We note that both the Jacobian $\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}$ and the residual $\boldsymbol{f}(\tilde{\boldsymbol{x}})$ are computed each step of the Gauss-Newton update (Eqn. 2.7), and therefore there is little additional computational expense to compute this update when no hyperreduction is used. We discuss more on hyperreduction at the end of this section.

A gradient-based optimizer takes a step Δ in the direction of decreasing gradient, leading to the following update:

$$\tilde{\boldsymbol{x}}^{k+1} = \tilde{\boldsymbol{x}}^k - \Delta \left(\left. \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} \right|_{\tilde{\boldsymbol{x}}} \right)^T \boldsymbol{f}(\tilde{\boldsymbol{x}}) \tag{2.10}$$

The convergence of the optimizer is highly dependent on the choice of step size Δ, and there is no obvious choice for the step size. Rather than attempting to guess the optimal step size, we choose to add the gradient computed in Eqn. 2.9 to the basis as an additional basis vector. Since the basis has an additional vector, the pROM state $\hat{\boldsymbol{x}}$ must also be assigned an additional value corresponding to the appended basis vector. The additional value can be chosen to be zero, but we suggest to choose a small negative value as this corresponds to a small positive Δ. A large negative value can lead to overshoot and high pROM errors.

Using this method in combination with hyperreduction adds to the complexity of the setup. During hyperreduction, the FOM governing equations are solved at a small number of points in the domain and those solutions are used to approximate the residual ℓ^2 norm. Adding this method would require computing the residual and Jacobian at the additional points where basis adaptation is to be applied.

In cases with shocks, the number of points near the shocks is generally small relative to the mesh size, so adapting only the points near the shock would not add excessively to the computational cost.

2.2.2. Method 2: Selected points adaptation

In problems with a wide range of scales gradient-based methods can be problematic because the gradient will vary by orders of magnitude, and taking a step in the gradient direction has negligible effect at most points. Typically, gradient-based methods work best at points far from a minima, whereas methods such as Gauss-Newton work best close to a minima. Here we use the Levenberg–Marquardt optimizer, which is a method that combines Gauss-Newton and gradient descent. The Levenberg-Marquardt algorithm solves least-squares problems of the form of Eqn. 2.8. The update step β is given by

$$(J^T J + \lambda I)\beta = J^T f \tag{2.11}$$

where $J = \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}$ and λ is the damping factor. A small damping factor brings the method closer to Gauss-Newton, and a large damping factor brings it closer to gradient-descent.

Applying the Levenberg-Marquardt algorithm on the entire domain is infeasible because Eqn. 2.11 requires a linear solve with dimension N. However, it is possible to apply the Levenberg-Marquardt algorithm to a limited number of points in the domain, specifically the N_{adapt} points with high residual. This can be formulated by replacing J with $J\,S$ and replacing f with $S^T f$, where $S \in \mathbb{R}^{N \times N_{adapt}}$ is a sampling matrix where each column has a single 1 in the location corresponding to the point being sampled, and all other values are 0. The algorithm then becomes

$$(S^T J^T J\,S + \lambda I)\beta = S^T J^T S^T f \tag{2.12}$$

This requires solving an linear system of dimension N_{adapt}. In hypersonic pROMs, the residuals at the shock are typically much higher than throughout the rest of the domain, so adapting only points close to the shock is expected to improve results.

3. 1D BURGER'S EQUATION

3.1. Setup

The first test case is the 1D steady Burgers' equation with a point source. A shock develops near the point source, so the shock location and strength can be varied by changing the point source's location and strength. The Burgers' equation and boundary conditions are given in Eqn. 3.1.

$$\frac{1}{2}\frac{\partial u^2}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2} + s(x) \tag{3.1}$$

$$s(x) = \xi_{mag}\delta(x - \xi_{loc}) \tag{3.2}$$

$$u(0) = 1 \tag{3.3}$$

$$u(1) = 10 \tag{3.4}$$

with dynamic viscosity $\nu = 0.1$, δ is a delta function, and ξ_{mag} and ξ_{loc} determine the magnitude and location of the point source. The Burgers' equation is solved using the finite volume method with 1st order fluxes over a domain from x=0 to x=1 with 100 equally spaced cells. Fig. 3-1 shows full order solutions with parameters given in Table 3-1.

3.2. pROMs

The pROMs are trained on three FOM solutions given in Table 3-1. It is not necessary to perform POD for dimensionality reduction as a pROM dimension of p is sufficiently small, so the basis $\mathbf{\Phi}$ simply consists of the FOM solution vectors concatenated together. The pROMs use RBF interpolation and are run until a relative pROM residual of 1e-5 is reached or 20 iterations.

Figure 3-2 shows the state resulting from the initial guess using inverse distance interpolation and from a standard linear pROM. Instead of having a single shock, both the initial guess and pROM solution have a triple shock structure because the training cases have shocks at slightly different

Table 3-1. Parameters of the 1D Burgers' equation case

	$\xi_{\mathbf{mag}}$	$\xi_{\mathbf{loc}}$
Training FOM 0	100	0.56
Training FOM 1	50	0.60
Training FOM 2	80	0.63
Test FOM	60	0.58

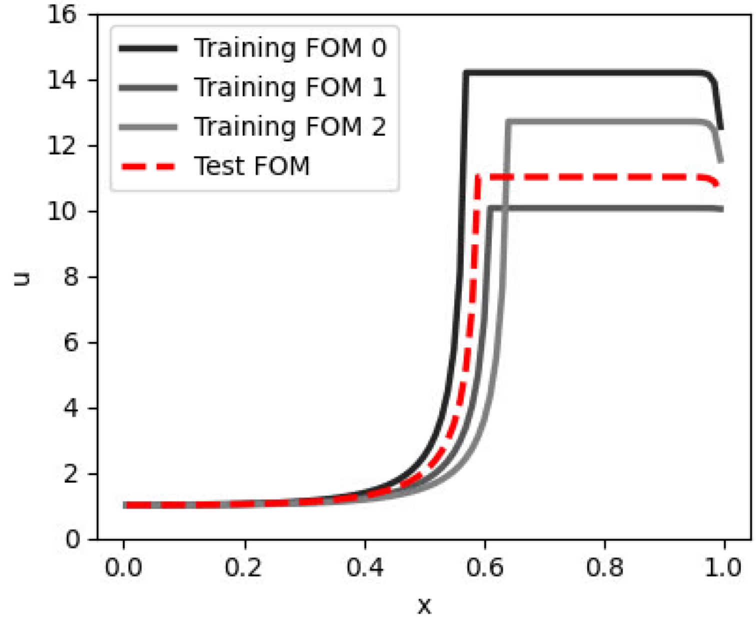

Figure 3-1. Full order solutions of 1D Burgers' equation

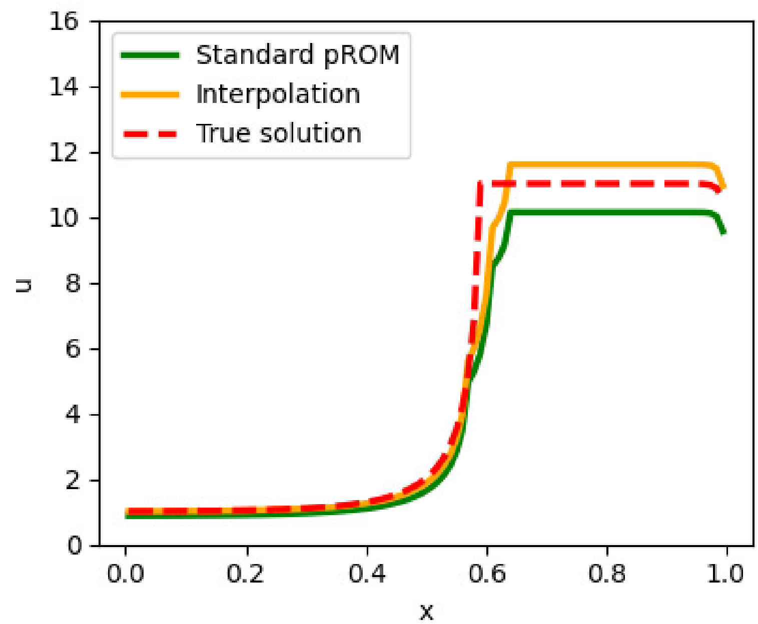

Figure 3-2. State resulting from inverse distance interpolation and standard linear pROM of 1D Burgers' equation

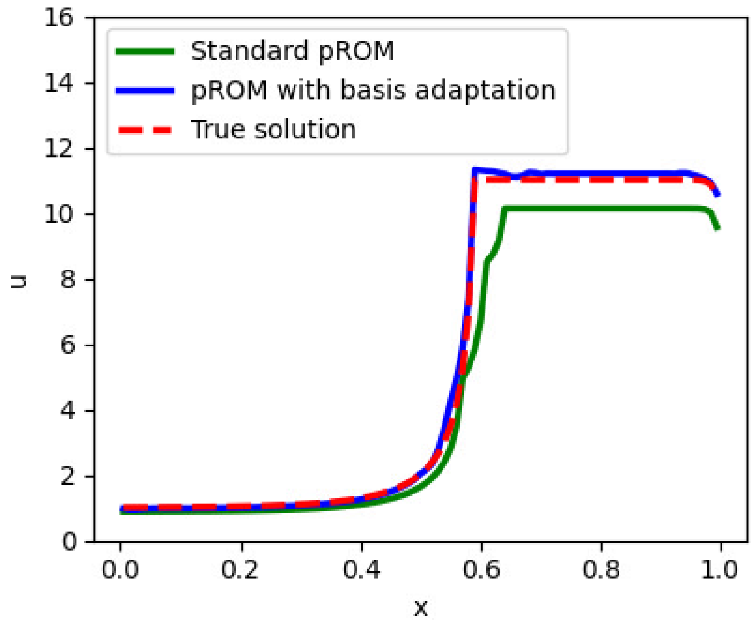

Figure 3-3. Solution resulting from pROM with basis adaptation method 1

locations. In addition, both the initial guess and standard linear pROM significantly mispredict the state to the right of the shock.

Both basis adaptation methods were applied to the 1D steady Burgers' equation. For both methods, basis adaptation was applied every two iterations between the fourth and twelfth iterations, for a total of five basis adaptations. While applying the basis adaptation repeatedly may lead to convergence to the true solution, the goal is to apply basis adaptation a small number of times to keep the additional computational cost small, and we therefore apply the adaptation only five times. Figure 3-3 shows the pROM state resulting from basis adaptation method 1. The adapted pROM has a single shock rather than the triple shock from the standard linear pROM, and the solution to the right of the shock is more accurate. However, a slight oscillation is observed immediately to the right of the shock and there are slight errors in the shock, but overall the adapted pROM is much more accurate than the standard linear pROM.

Fig. 3-4 shows the result from adaptation method 2. The adaptaion was applied to the 5 points of highest residual and their neighbors, for a total of $N_{adapt} = 6 - 9$ points each adaptation. The neighboring points are included because the residual at a single point is affected by the neighboring points, so reducing the residual at a single point may require modifying the neighbors. The basis adapted pROM is much more accurate than the standard linear pROM, as both the shock structure and the values to the right of the shock are much closer to the true solution.

The full-order residual ℓ^2-norm $\left\|\boldsymbol{f}(\tilde{\boldsymbol{x}})\right\|_2$ shown in Fig. 3-5 demonstrates the residual decrease for the 3 pROMs. The standard linear pROM relative residual decreases to 0.98, while the basis adaptation methods decrease the residual to 0.08 and 0.02, demonstrating that the adapted pROMs are solving the governing equations much more accurately than the standard linear pROM.

Table 3-2 shows the state ℓ^2-norm errors for each case. While the standard linear pROM increases

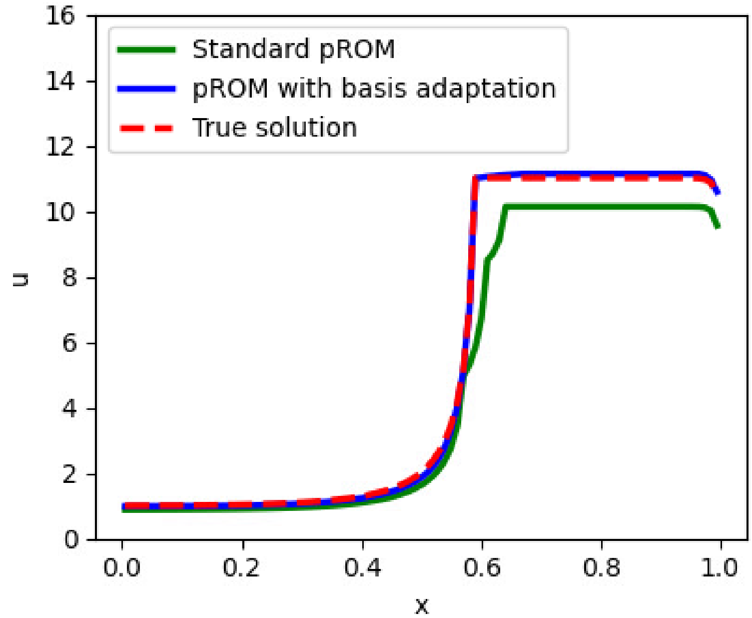

Figure 3-4. Solution resulting from pROM with basis adaptation method 2

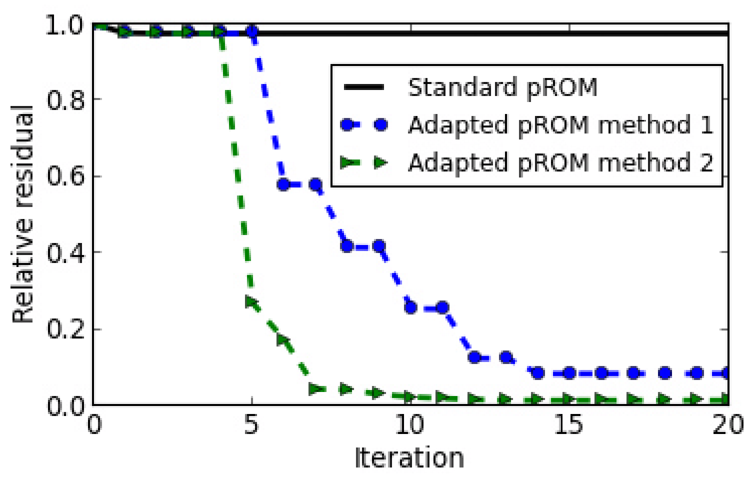

Figure 3-5. Relative full order residual ℓ^2-norm vs iteration for standard linear pROM and adapted pROMs.

Table 3-2. Errors from the 1D Burgers' equation case

	Error ℓ^2-norm
Inverse distance interp	6.9
Standard linear pROM	9.7
pROM with adaptation method 1	1.9
pROM with adaptation method 2	1.0

the overall state error, both adaptation methods dramatically decrease the error by factors of over 3.

4. 2D HIFIRE

4.1. Setup

The second flow configuration is a two-dimensional simulation of the HIFiRE-1 (Hypersonic International Flight Research Experiment) nose cone [5]. The nose cone has a rounded nose and a 7° taper angle and is 11.8 cm long from the nose to the end of the cone. The geometry is axisymmetric with zero angle of attack, so a 2D mesh on one side of the centerline shown in Figure 4-1 is sufficient to simulate the flow. The mesh has 32,512 cells, with a maximum y+ of 0.81.

Simulations are run using the S-A turbulence model [18], which has six conserved variables. The simulations are solved in the Sandia Parallel Aerodynamics and Reentry Code (SPARC) with an interface to Pressio to run the pROMs. [8, 17]. The freestream conditions are given in Table 4-1, and simulations are run across various freestream densities and velocities. All cases are run with grid tailoring applied twice.

This case was chosen because Ref. [3] showed that pROMs have very poor accuracy on this case without grid-tailoring, but even with grid-tailoring the pROMs had poor accuracy.

4.2. pROMs

The standard linear pROMs use 9 FOM solutions for the basis and RBF interpolation for the initial guess, after which they are solved until the pROM residual ℓ^2-norm decreases by 5 orders of magnitude. The basis adapted pROMs are then applied over 5 consecutive iterations. The FOM grid tailoring algorithm computes displacements for each mesh node, and during the pROM setup

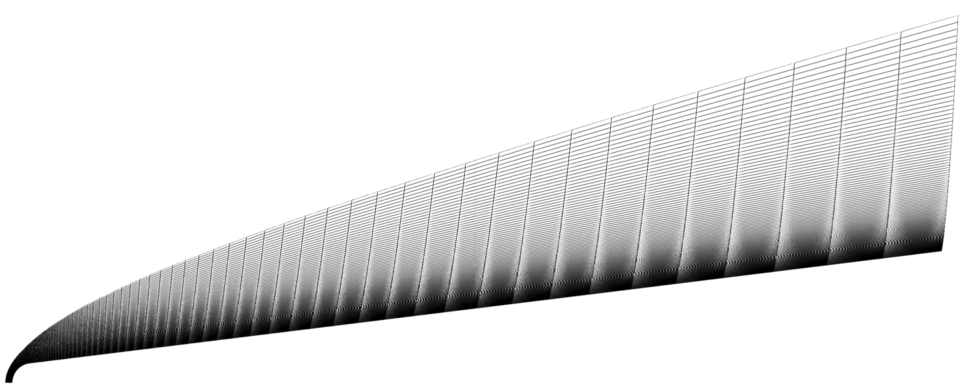

Figure 4-1. Baseline mesh for 2D HIFiRE-1 without grid tailoring.

Table 4-1. Parameters of the 2D HIFiRE-1 case

Freestream density	0.04 - 0.08 kg/m^3
Mach number	5.5 - 9.5
Freestream temperature	226.46 K
Wall temperature	296.7 K

the displacements are computed using RBF interpolation and fixed, identical to the method used in Ref. [4].

The work done in this study was performed over a short time period and under time constraints, and there are therefore a number of unfinished aspects to this work. The previous work found that pROMs with primitive variables had better accuracy than conserved variables, but all results here use conserved variable pROMs. Furthermore, no equation scaling was used although it is known that a dimensionally consistent inner product gives better accuracy [13]. Only basis adaptation method 1 was applied, as there was no time to apply method 2. Finally, the methods are applied inefficiently and take significant computational expense, but the objective of this study was to demonstrate the accuracy of the basis adaptation methods, and an efficient implementation is left for future work.

We examine three QoIs error metrics for this case. The first is the axial force error, defined as

$$\varepsilon_{F_x} = \frac{|F_x(\boldsymbol{\mu}) - \tilde{F}_x(\boldsymbol{\mu})|}{|F_x(\boldsymbol{\mu})|}, \tag{4.1}$$

where $F_x(\boldsymbol{\mu})$ and $\tilde{F}_x(\boldsymbol{\mu})$ are the integrals of axial force computed with the FOM and some corresponding approximation with (2.3), respectively. Similarly, the integrated heat flux error is defined as

$$\varepsilon_{Q_{wall}} = \frac{|Q_{wall}(\boldsymbol{\mu}) - \tilde{Q}_{wall}(\boldsymbol{\mu})|}{|Q_{wall}(\boldsymbol{\mu})|}, \tag{4.2}$$

where $Q_{wall}(\boldsymbol{\mu})$ and $\tilde{Q}_{wall}(\boldsymbol{\mu})$ are integrated heat fluxes calculated with the FOM and an approximation. Finally, the state ℓ^2 error is defined as

$$\varepsilon_{\boldsymbol{x}} = \frac{\|\boldsymbol{x}(\boldsymbol{\mu}) - \tilde{\boldsymbol{x}}(\boldsymbol{\mu})\|_2}{\|\boldsymbol{x}(\boldsymbol{\mu})\|_2}, \tag{4.3}$$

where $\boldsymbol{x}(\boldsymbol{\mu})$ and $\tilde{\boldsymbol{x}}(\boldsymbol{\mu})$ are the full state computed with the FOM and some approximation with (2.3), respectively.

Figs. 4-2 shows the force error and integrated heat flux errors resulting from the initial guess, standard linear pROM, and basis adapted pROM. The RBF initial guess is clearly the most accurate for both QoIs, but the basis adapted pROMs have generally better accuracy than the standard pROMs, particularly for heat flux. The basis adapted pROMs are still significantly less accurate than the initial guess, but it is hoped that incorporating primitive variables and residual norm scalings would improve the results.

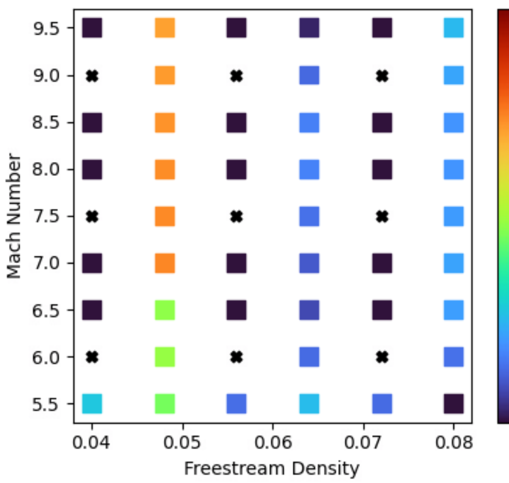

(a) Initial guess force error ε_{F_x}

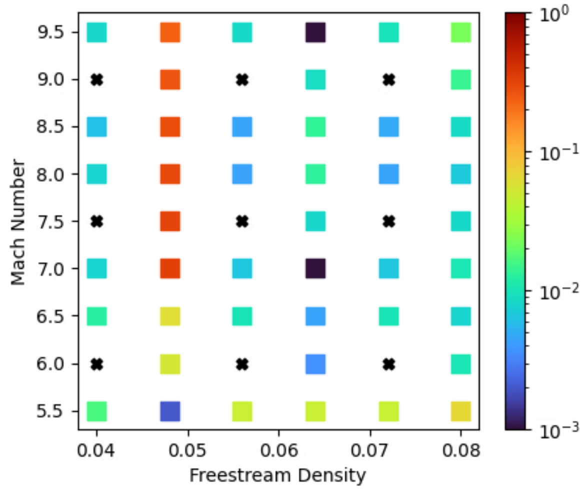

(b) Initial guess heat flux error $\varepsilon_{Q_{wall}}$

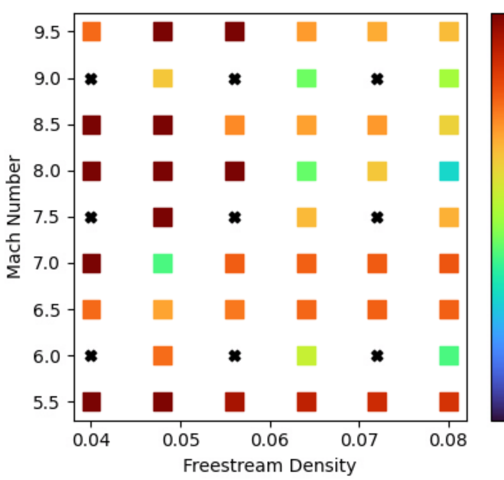

(c) Standard linear pROM force error ε_{F_x}

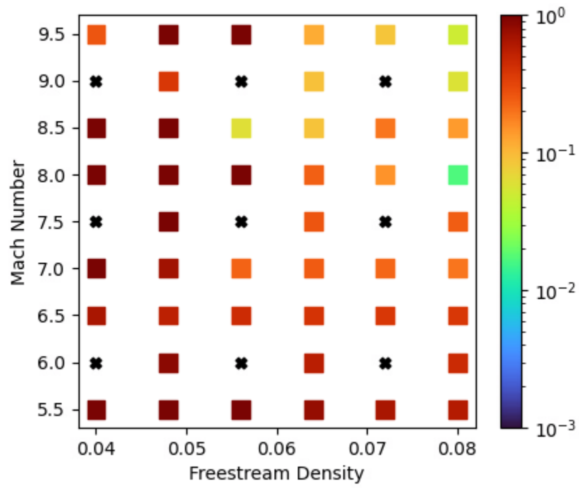

(d) Standard linear pROM heat flux error $\varepsilon_{Q_{wall}}$

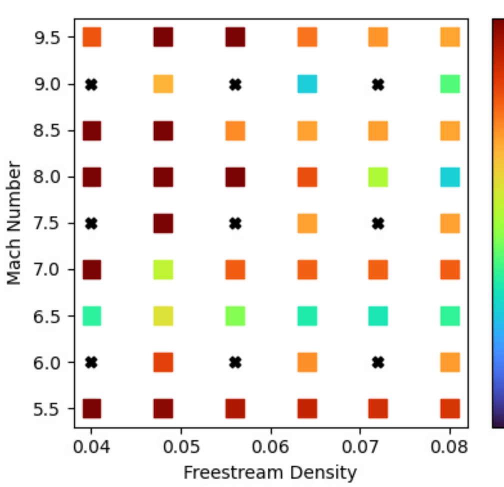

(e) Adapted pROM force error ε_{F_x}

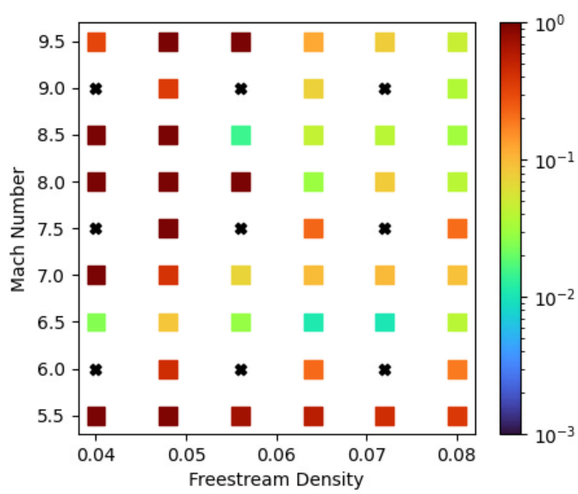

(f) Adapted pROM heat flux error $\varepsilon_{Q_{wall}}$

Figure 4-2. Errors for 2D HIFiRE. Training cases are indicated by black X's

The difference between standard linear pROMs and basis adapted pROMs is demonstrated by considering an error ratio, defined as

$$Error\ ratio = \frac{Basis\ adapted\ pROM\ error}{Standard\ linear\ pROM\ error} \tag{4.4}$$

The error ratios for streamwise force, integrated heat flux, and state error are shown in Fig. 4-3. The force error ratio varies widely across parameter space, with an average force error decrease of 8%. At the highest Mach numbers where pROMs are extrapolating, the basis adaptation makes the force error worse, but at lower Mach numbers the force error generally decreases or stays the same, with the exception of several cases where the standard linear pROM has very low error and the basis adaptation increases the force error significantly. However, the integrated heat flux error ratio shows that basis adaptation lowers the heat flux error dramatically across nearly the entire parameter space, and the average heat flux error decreases by 69%. The state error ratio is also improved for most of the pROMs, although at the highest Mach numbers the state error generally increases. However, the state errors only decrease slightly, with an average decrease of 14%.

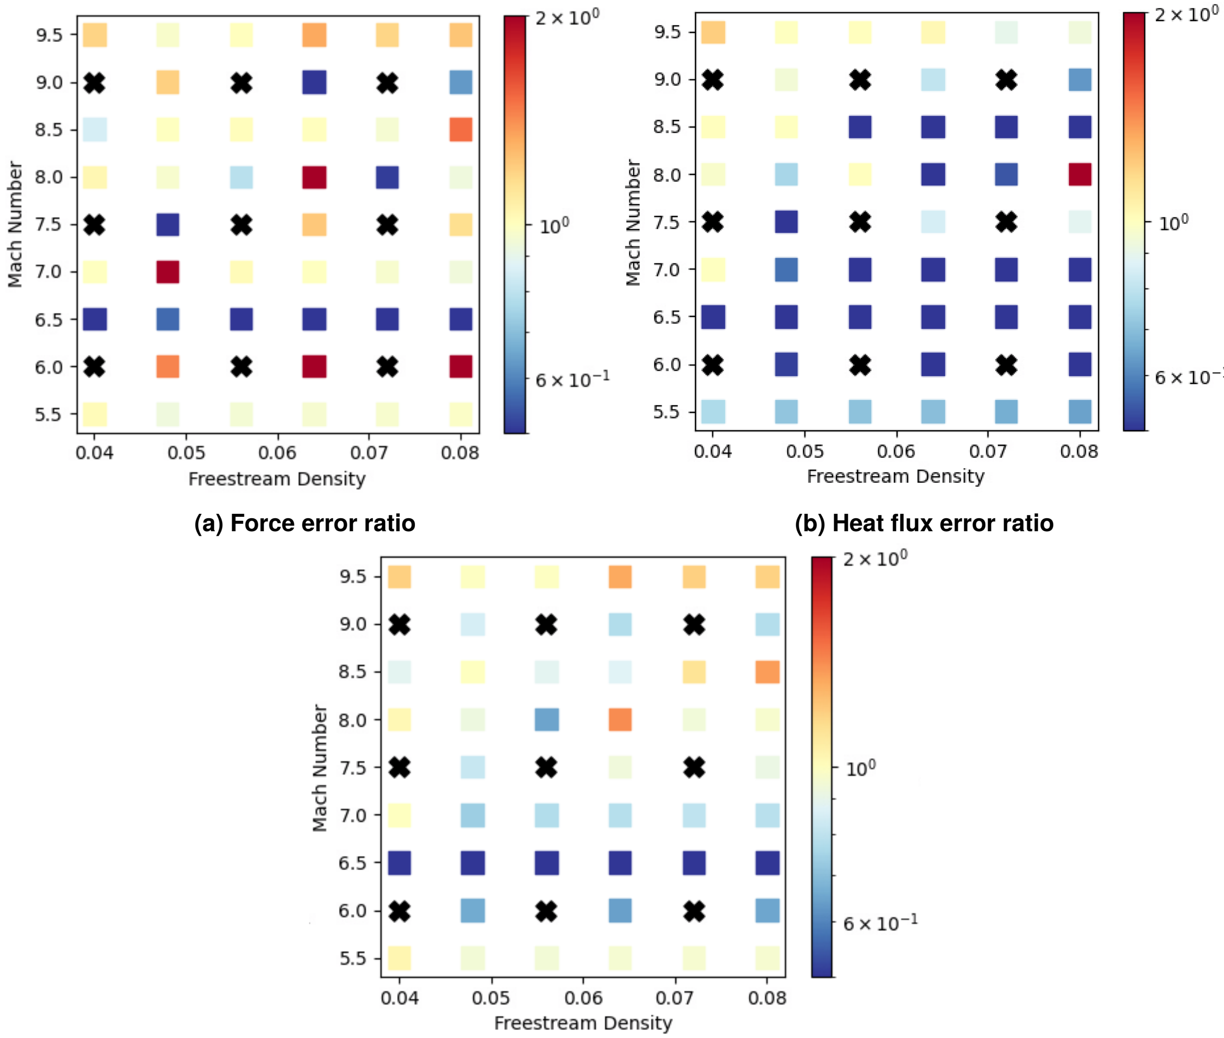

(a) Force error ratio

(b) Heat flux error ratio

(c) State ℓ^2 error ratio

Figure 4-3. Ratio of basis adapted to standard linear pROM errors for 2D HIFiRE. Training cases are indicated by black X's

5. CONCLUSIONS

This study developed two basis adaptation methods for steady problems, one using a gradient based approach, and the other using an optimizer on a limited number of points. Both methods were tested on a 1D Burgers' equation with spatially varying source, a case in which standard linear pROMs have poor accuracy, and both basis adaptation methods dramatically improved the results. The gradient based method was also tested on hypersonic flow around a 2D HIFiRE-1 vehicle and was shown to improve the results in comparison to a standard linear pROM.

The 2D HIFiRE-1 case used here was chosen because previous work showed that pROMs had poor accuracy for this case, largely due to the highly refined mesh [3]. While basis adaptation generally improved QoI accuracy, the force error was only slightly decreased on average, and there were cases where the errors increased. This indicates that the basis adaptation method used has potential to increase accuracy, but it does not unaccompanied make pROMs robust for hypersonics. Rather, a robust hypersonic pROM may only be obtained using a combination of pROM developments, including using grid-tailoring, primitive variables, entropy scaling, and domain decomposition.

This is a preliminary study to examine if basis adaptation for steady problems is feasible, and there is much remaining work to be done. Some potential future works are:

- Implementing the basis adaptation techniques efficiently and incorporating hyperreduction
- Combining with primitive variable pROMs
- Incorporating the entropy ℓ^2-norm
- Improving boundary condition treatment by excluding points adjacent to boundaries
- Implementing shock capturing to only adapt points near shocks

The code used to generate and run the cases is available at cee-gitlab.sandia.gov/dching/rom-basis-adaptation.

REFERENCES

[1] Nicholas Arnold-Medabalimi, Cheng Huang, and Karthik Duraisamy. Large-eddy simulation and challenges for projection-based reduced-order modeling of a gas turbine model combustor. *International Journal of Spray and Combustion Dynamics*, 14(1-2):153–175, 2022.

[2] Joshua Barnett and Charbel Farhat. Quadratic approximation manifold for mitigating the kolmogorov barrier in nonlinear projection-based model order reduction. *Journal of Computational Physics*, page 111348, 2022.

[3] Patrick J. Blonigan, Marco Arienti, David S. Ching, Jeffrey A. Fike, and Elizabeth Krath. Rapid high-fidelity aerothermal responses with quantified uncertainties via reduced-order modeling. *Sandia Report*, pages 2021–12316, 2021.

[4] David S Ching, Patrick J Blonigan, Francesco Rizzi, and Jeffrey A Fike. Model reduction of hypersonic aerodynamics with residual minimization techniques. In *AIAA SCITECH 2022 Forum*, page 1247, 2022.

[5] Douglas Dolvin. Hypersonic international flight research and experimentation (hifire) fundamental science and technology development strategy. In *15th AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, page 2581, 2008.

[6] Philip A Etter and Kevin T Carlberg. Online adaptive basis refinement and compression for reduced-order models via vector-space sieving. *Computer Methods in Applied Mechanics and Engineering*, 364:112931, 2020.

[7] Andrea Ferrero, Tommaso Taddei, and Lei Zhang. Registration-based model reduction of parameterized two-dimensional conservation laws. *Journal of Computational Physics*, 457:111068, 2022.

[8] Micah Howard, Andrew Bradley, Steven W. Bova, James Overfelt, Ross Wagnild, Derek Dinzl, Mark Hoemmen, and Alicia Klinvex. Towards performance portability in a compressible CFD code. In *23rd AIAA Computational Fluid Dynamics Conference*, volume 1, Denver, CO, 2017. AIAA.

[9] Cheng Huang, Jiayang Xu, Karthik Duraisamy, and Charles Merkle. Exploration of reduced-order models for rocket combustion applications. In *2018 AIAA Aerospace Sciences Meeting*, page 1183, 2018.

[10] Kookjin Lee and Kevin T Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.

[11] Rambod Mojgani and Maciej Balajewicz. Physics-aware registration based auto-encoder for convection dominated pdes. *arXiv preprint arXiv:2006.15655*, 2020.

[12] Nirmal J Nair and Maciej Balajewicz. Transported snapshot model order reduction approach for parametric, steady-state fluid flows containing parameter-dependent shocks. *International Journal for Numerical Methods in Engineering*, 117(12):1234–1262, 2019.

[13] Eric J Parish and Francesco Rizzi. On the impact of dimensionally-consistent and physics-based inner products for pod-galerkin and least-squares model reduction of compressible flows. *arXiv preprint arXiv:2203.16492*, 2022.

[14] Benjamin Peherstorfer. Model reduction for transport-dominated problems via online adaptive bases and adaptive sampling. *SIAM Journal on Scientific Computing*, 42(5):A2803–A2836, 2020.

[15] Benjamin Peherstorfer and Karen Willcox. Online adaptive model reduction for nonlinear systems via low-rank updates. *SIAM Journal on Scientific Computing*, 37(4):A2123–A2150, 2015.

[16] Donsub Rim, Benjamin Peherstorfer, and Kyle T Mandli. Manifold approximations via transported subspaces: Model reduction for transport-dominated problems. *arXiv preprint arXiv:1912.13024*, 2019.

[17] Francesco Rizzi, Patrick J. Blonigan, and Kevin T. Carlberg. Pressio: Enabling projection-based model reduction for large-scale nonlinear dynamical systems. Submitted to the SIAM Journal on Scientific Computing, Feb. 2020, 2020.

[18] Philippe Spalart and Steven Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th aerospace sciences meeting and exhibit*, page 439, 1992.

[19] Tommaso Taddei. A registration method for model order reduction: data compression and geometry reduction. *SIAM Journal on Scientific Computing*, 42(2):A997–A1027, 2020.

DISTRIBUTION

Hardcopy—Internal

Number of Copies	Name	Org.	Mailstop
1	David Ching	8751	9042
1	Francesco Rizzi	8739	9152
1	Jaideep Ray	8739	9152

Email—Internal (encrypt for OUO)

Name	Org.	Sandia Email Address
CA Technical Library	8551	cateclib@sandia.gov

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