Accelerating PDE-Constrained Optimization using Adaptive Reduced-Order Models

Matthew J. Zahr

Institute for Computational and Mathematical Engineering
Farhat Research Group
Stanford University

Sandia National Laboratories
July 8, 2015

Outline

Application I: Shape Optimization of Vehicle in Turbulent Flow

- Volkswagen Passat
- Shape optimization
  - Minimum drag configuration
  - Unsteady effects
- Simulation
  - 4M vertices, 24M dof
  - Compressible Navier-Stokes
  - Spalart-Allmaras
- Single forward simulation
  - $\approx$ 1 day on 2048 CPUs

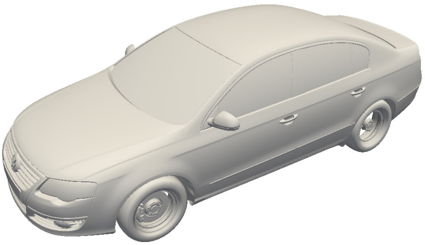

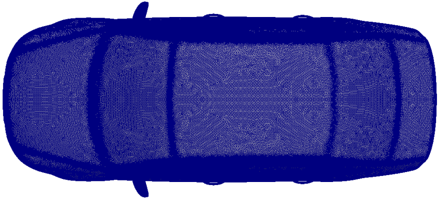

Application II: Optimal Control Flapping Wing

- Biologically-inspired flight
  - Micro Aerial Vehicles (MAVs)
- Mesh
  - 43,000 vertices
  - 231,000 tetra ($p = 3$)
  - 2,310,000 DOF
- CFD
  - Compressible Navier-Stokes
  - Discontinuous Galerkin
- Shape optimization, control
  - unsteady effects
  - min energy, const thrust

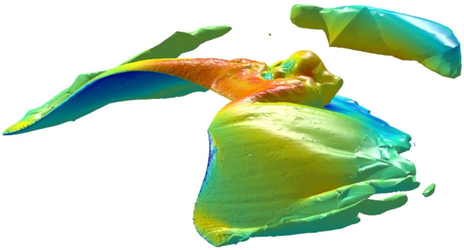

Figure: Flapping Wing (?)

Application III: Topology Optimization

- Design of new lacrosse head [1]
- Mesh
  - 96,247 vertices
  - 475,666 tetra
  - 276,159 DOF
- Single forward simulation
  - $\approx$ 5 minutes on 1 core

- Desired: topology optimization
  - Finer mesh (10-100x)
  - Realistic material model

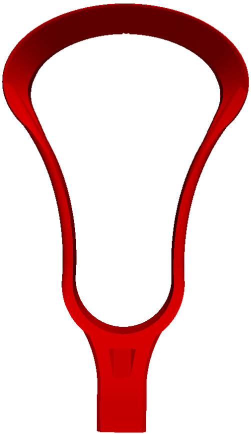

[1] Collaboration with K. Washabaugh

Application III: Topology Optimization

- Design of new lacrosse head [1]
- Mesh
  - 96,247 vertices
  - 475,666 tetra
  - 276,159 DOF
- Single forward simulation
  - $\approx$ 5 minutes on 1 core

- Desired: topology optimization
  - Finer mesh (10-100x)
  - Realistic material model

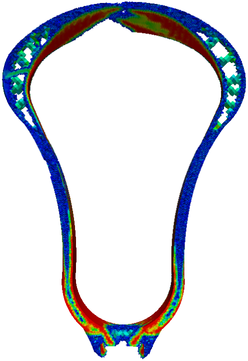

[1] Collaboration with K. Washabaugh

Application III: Topology Optimization

- Design of new lacrosse head [1]
- Mesh
  - 96,247 vertices
  - 475,666 tetra
  - 276,159 DOF
- Single forward simulation
  - $\approx$ 5 minutes on 1 core

- Desired: topology optimization
  - Finer mesh (10-100x)
  - Realistic material model

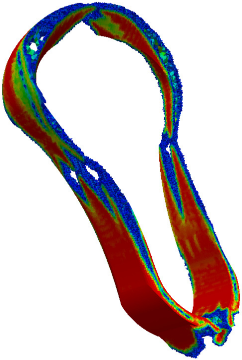

[1] Collaboration with K. Washabaugh

Reduced-Order Models (ROMs)

ROMs as Enabling Technology

- Optimization: design, control
  - Single objective, single-point
  - Multiobjective, multi-point
  - Unsteady effects
- Uncertainty Quantification
- Optimization under uncertainty

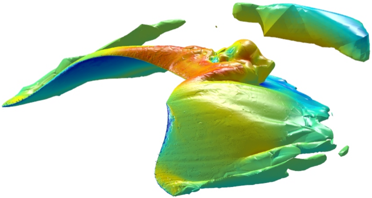

Figure: Flapping Wing (?)

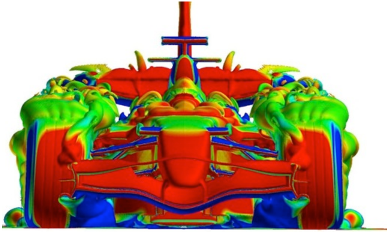

Problem Formulation

Goal: Rapidly solve PDE-constrained optimization problems of the form

$$\begin{aligned} \underset{\mathbf{w}\in\mathbb{R}^N,\ \boldsymbol{\mu}\in\mathbb{R}^p}{\text{minimize}} \quad & f(\mathbf{w},\boldsymbol{\mu}) \\ \text{subject to} \quad & \mathbf{R}(\mathbf{w},\boldsymbol{\mu}) = 0 \end{aligned}$$

Discretize-then-optimize

where $\mathbf{R} : \mathbb{R}^N \times \mathbb{R}^p \rightarrow \mathbb{R}^N$ is the discretized (steady, nonlinear) PDE, $\mathbf{w}$ is the PDE state vector, $\boldsymbol{\mu}$ is the vector of parameters, and N is **assumed to be very large**.

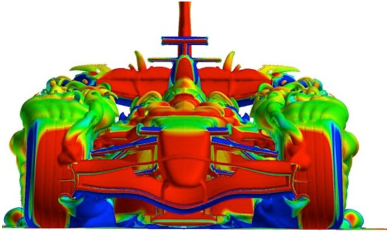

Outline

Reduced-Order Model

- Model Order Reduction (MOR) assumption: *state vector lies in low-dimensional affine subspace*

$$\mathbf{w} \approx \mathbf{w}_r = \bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y} \qquad \Longrightarrow \qquad \frac{\partial \mathbf{w}}{\partial \boldsymbol{\mu}} \approx \frac{\partial \mathbf{w}_r}{\partial \boldsymbol{\mu}} = \mathbf{\Phi}\frac{\partial \mathbf{y}}{\partial \boldsymbol{\mu}}$$

  where $\mathbf{y} \in \mathbb{R}^n$ are the reduced coordinates of $\mathbf{w}_r$ in the basis $\mathbf{\Phi} \in \mathbb{R}^{N \times n}$, and $n \ll N$

- Substitute assumption into High-Dimensional Model (HDM), $\mathbf{R}(\mathbf{w}, \boldsymbol{\mu}) = 0$

$$\mathbf{R}(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu}) \approx 0$$

- Require projection of residual in low-dimensional *left subspace*, with basis $\mathbf{\Psi} \in \mathbb{R}^{N \times n}$ to be zero

$$\boxed{\mathbf{R}_r(\mathbf{y}, \boldsymbol{\mu}) = \mathbf{\Psi}^T\mathbf{R}(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu}) = 0}$$

Reduced Optimization Problem

ROM-Constrained Optimization

$$\begin{aligned} \underset{\boldsymbol{\mu}\in\mathbb{R}^p}{\text{minimize}} \quad & f(\bar{\mathbf{w}} + \boldsymbol{\Phi}\mathbf{y}(\boldsymbol{\mu}), \boldsymbol{\mu}) \\ \text{subject to} \quad & \boldsymbol{\Psi}^T\mathbf{R}(\bar{\mathbf{w}} + \boldsymbol{\Phi}\mathbf{y}, \boldsymbol{\mu}) = 0 \end{aligned}$$

- Issues that must be considered
  - Construction of bases
  - Speedup potential
  - Sensitivity analysis (adjoint method)
  - Training

Offline-Online Approach

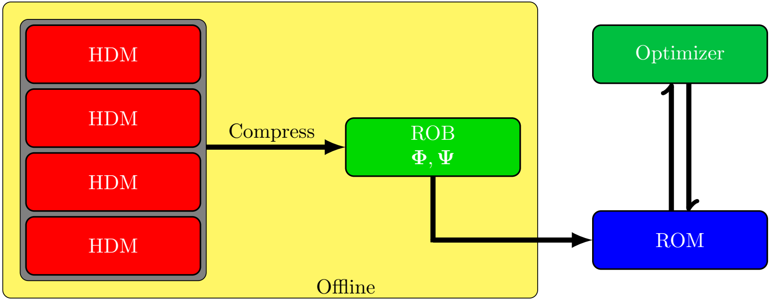

Figure: Schematic of Algorithm

Offline-Online Approach

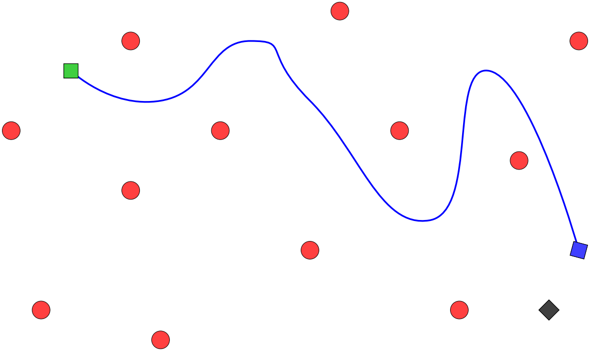

(a) Idealized Optimization Trajectory: Parameter Space

Offline-Online (Database) Approach

Offline-Online Approach to ROM-Constrained Optimization

- Identify samples in *offline* phase to be used for training
  - Space-fill sampling (i.e. latin hypercube)
  - Greedy sampling
- Collect snapshots from HDM
- Build ROB $\mathbf{\Phi}$
- Solve optimization problem

$$\begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^n,\ \boldsymbol{\mu}\in\mathbb{R}^p}{\text{minimize}} \quad & f(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu}) \\ \text{subject to} \quad & \mathbf{\Psi}^T\mathbf{R}(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu}) = 0 \end{aligned}$$

(?), (?), (?), (?)

Adaptive Approach

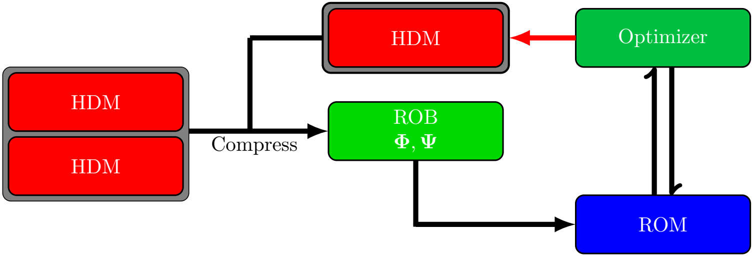

Figure: Schematic of Algorithm

Adaptive Approach

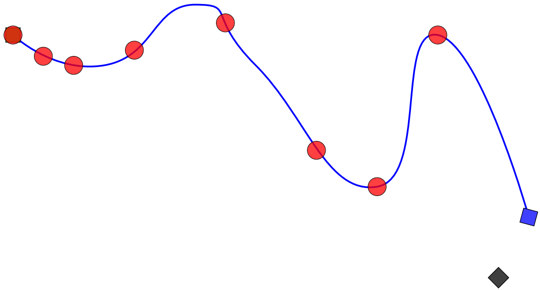

(a) Idealized Optimization Trajectory: Parameter Space

Adaptive Approach

Adaptive Approach to ROM-Constrained Optimization

- Collect snapshots from HDM at *sparse sampling* of the parameter space
  - Initial condition for optimization problem
- Build ROB $\mathbf{\Phi}$ from sparse training
- Solve optimization problem

$$
\begin{aligned}
\underset{\mathbf{y}\in\mathbb{R}^n,\ \boldsymbol{\mu}\in\mathbb{R}^p}{\text{minimize}} \quad & f(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu}) \\
\text{subject to} \quad & \mathbf{\Psi}^T\mathbf{R}(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu}) = 0 \\
& \frac{1}{2}||\mathbf{R}(\bar{\mathbf{w}} + \mathbf{\Phi}\mathbf{y}, \boldsymbol{\mu})||_2^2 \leq \epsilon
\end{aligned}
$$

- Use solution of above problem to enrich training and repeat until convergence

(?), (?), (?), (?), (?), (?), (?)

Difficulty of Breaking Offline-Online Barrier

Offline-Online Approach

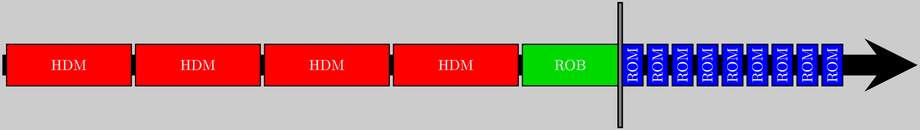

Figure: Offline-Online Approach

- Offline/Online Barrier
  - + Enables *large online* speedups
  - - Difficult to construct accurate, robust ROM
- Minimize ROM!

Difficulty of Breaking Offline-Online Barrier

Progressive Approach

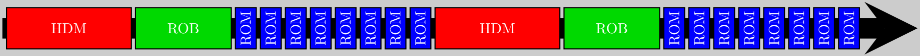

Figure: Progressive Approach

- Requires minimizing HDM, ROB, and ROM!
    - Cost and Quantity

Progressive Approach

Ingredients of Proposed Approach (?)

- Minimum-residual ROM (LSPG) and minimum-residual sensitivities
  - $f_r(\boldsymbol{\mu}) = f(\boldsymbol{\mu})$ and $\dfrac{\mathrm{d}f_r}{\mathrm{d}\boldsymbol{\mu}}(\boldsymbol{\mu}) = \dfrac{\mathrm{d}f}{\mathrm{d}\boldsymbol{\mu}}(\boldsymbol{\mu})$ for training parameters $\boldsymbol{\mu}$
- Reduced optimization (sub)problem

$$\begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^n,\ \boldsymbol{\mu}\in\mathbb{R}^p}{\text{minimize}} \quad & f(\bar{\mathbf{w}} + \boldsymbol{\Phi}\mathbf{y}, \boldsymbol{\mu}) \\ \text{subject to} \quad & \boldsymbol{\Psi}^T\mathbf{R}(\bar{\mathbf{w}} + \boldsymbol{\Phi}\mathbf{y}, \boldsymbol{\mu}) = 0 \\ & \frac{1}{2}||\mathbf{R}(\bar{\mathbf{w}} + \boldsymbol{\Phi}\mathbf{y}, \boldsymbol{\mu})||_2^2 \leq \epsilon \end{aligned}$$

- Efficiently update ROB with additional snapshots or new translation vector
  - Without re-computing SVD of entire snapshot matrix
- Adaptive selection of $\epsilon \rightarrow$ trust-region approach

Outline

Compressible, Inviscid Airfoil Inverse Design

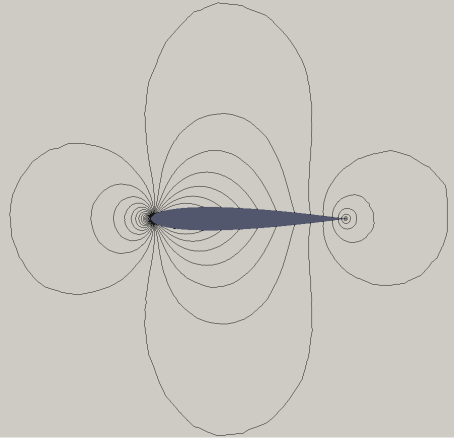

(a) NACA0012: Pressure field ($M_\infty = 0.5$, $\alpha = 0.0°$)

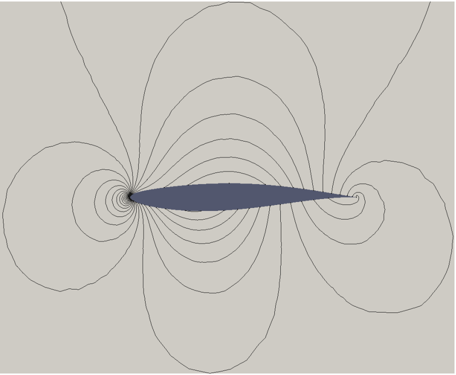

(b) RAE2822: Pressure field ($M_\infty = 0.5$, $\alpha = 0.0°$)

- Pressure discrepancy minimization (Euler equations)
  - Initial Configuration: NACA0012
  - Target Configuration: RAE2822

Initial/Target Airfoils: Scaled

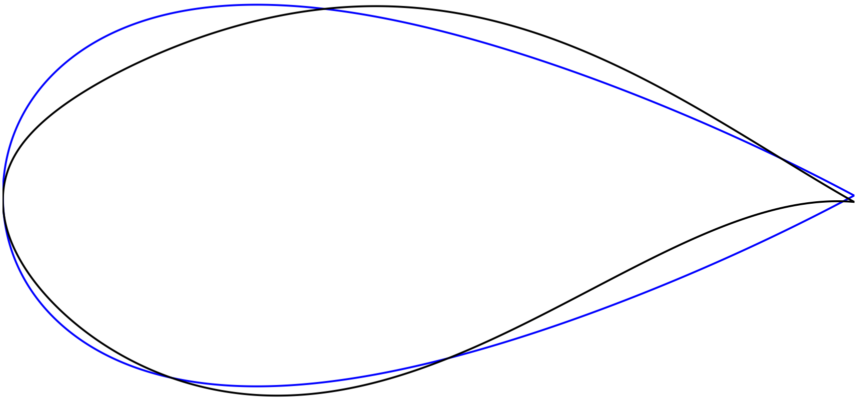

Shape Parametrization

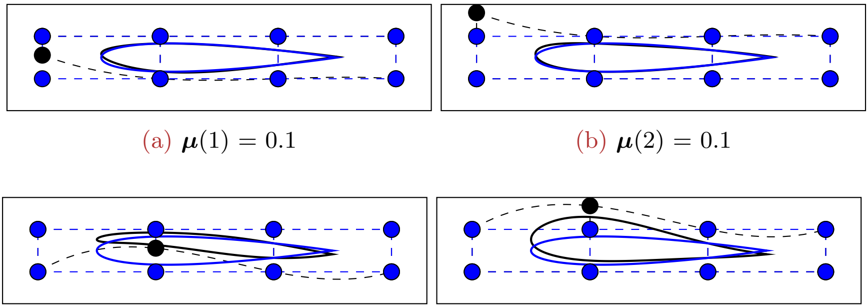

(a) $\boldsymbol{\mu}(1) = 0.1$

(b) $\boldsymbol{\mu}(2) = 0.1$

(c) $\boldsymbol{\mu}(3) = 0.1$

(d) $\boldsymbol{\mu}(4) = 0.1$

Figure: Shape parametrization of a NACA0012 airfoil using a *cubic* design element

Shape Parametrization

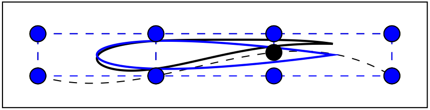

(a) $\boldsymbol{\mu}(5) = 0.1$

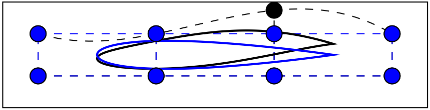

(b) $\boldsymbol{\mu}(6) = 0.1$

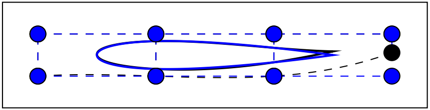

(c) $\boldsymbol{\mu}(7) = 0.1$

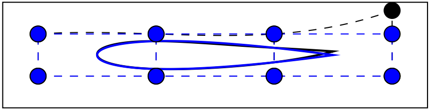

(d) $\boldsymbol{\mu}(8) = 0.1$

Figure: Shape parametrization of a NACA0012 airfoil using a *cubic* design element

Optimization Results

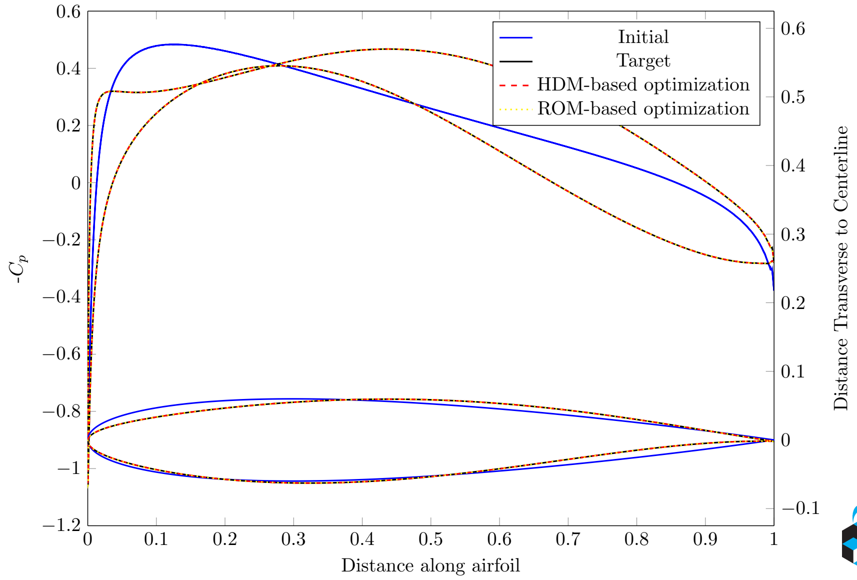

Optimization Results

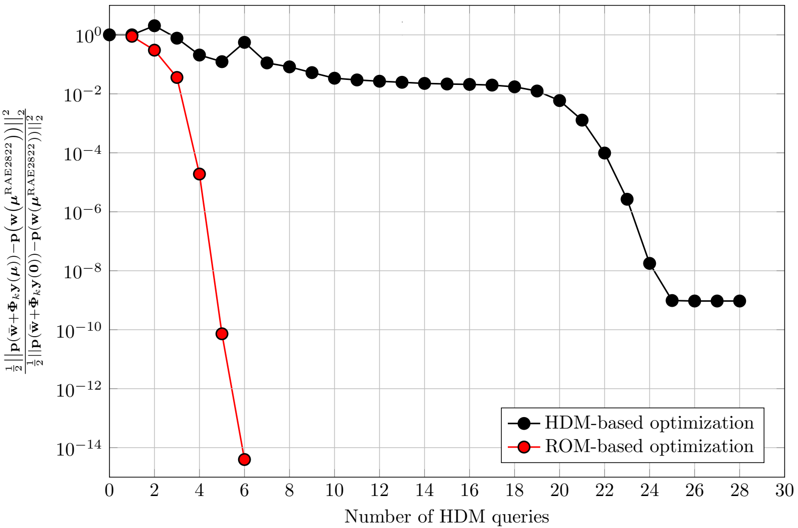

Optimization Results

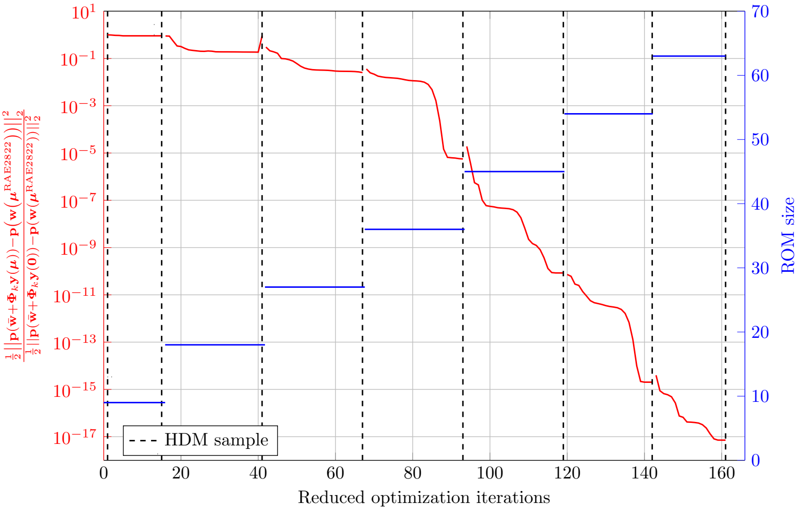

Optimization Results

	HDM-based optimization	ROM-based optimization
# of HDM Evaluations	29	7
# of ROM Evaluations	-	346
$\frac{\|\|\boldsymbol{\mu}^* - \boldsymbol{\mu}^{RAE2822}\|\|}{\|\|\boldsymbol{\mu}^{RAE2822}\|\|}$	$2.28 \times 10^{-3}\%$	$4.17 \times 10^{-6}\%$

Table: Performance of the HDM- and ROM-based optimization methods

Problem Setup

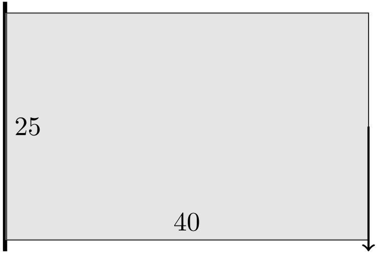

- 16000 8-node brick elements, 77760 dofs
- Total Lagrangian form, finite strain, StVK [2]
- St. Venant-Kirchhoff material
- Sparse Cholesky linear solver (CHOLMOD[3])
- Newton-Raphson nonlinear solver
- Minimum compliance optimization problem

$$\begin{aligned} \underset{\mathbf{u}\in\mathbb{R}^{n_\mathbf{u}},\ \boldsymbol{\mu}\in\mathbb{R}^{n_{\boldsymbol{\mu}}}}{\text{minimize}} \quad & \mathbf{f}_{\text{ext}}{}^T\mathbf{u} \\ \text{subject to} \quad & V(\boldsymbol{\mu}) \leq \frac{1}{2}V_0 \\ & \mathbf{r}(\mathbf{u},\ \boldsymbol{\mu}) = 0 \end{aligned}$$

- Gradient computations: Adjoint method
- Optimizer: SNOPT (?)
- Maximum ROM size: $k_\mathbf{u} \leq 5$

[2] (?), (?)
[3] (?)

Optimal Solution Comparison

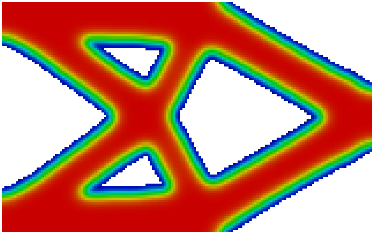

HDM

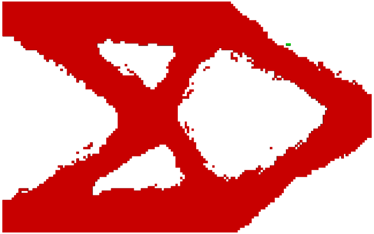

CTRPOD + Φ_μ adaptivity

HDM Solution	**HDM** Gradient	**HDM** Optimization
7458s (450)	4018s (411)	8284s

HDM
Elapsed time = 19761s

HDM Solution	**HDM** Gradient	**ROB** Construction	**ROM** Optimization
1049s (64)	88s (9)	727s (56)	39s (3676)

CTRPOD + Φ_μ adaptivity
Elapsed time = 2197s, Speedup ≈ 9x

Solution after 64 HDM Evaluations

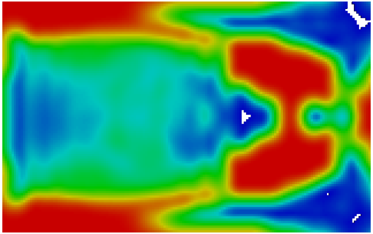

HDM

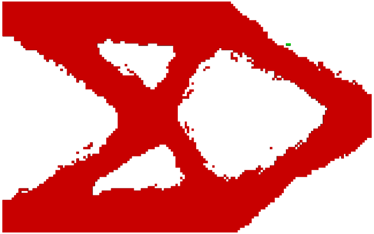

CTRPOD + Φ_μ adaptivity

- CTRPOD + Φ_μ adaptivity: superior approximation to optimal solution than HDM approach after fixed number of HDM solves (64)
- Reasonable option to *warm-start* HDM topology optimization

CTRPOD + Φ_μ adaptivity

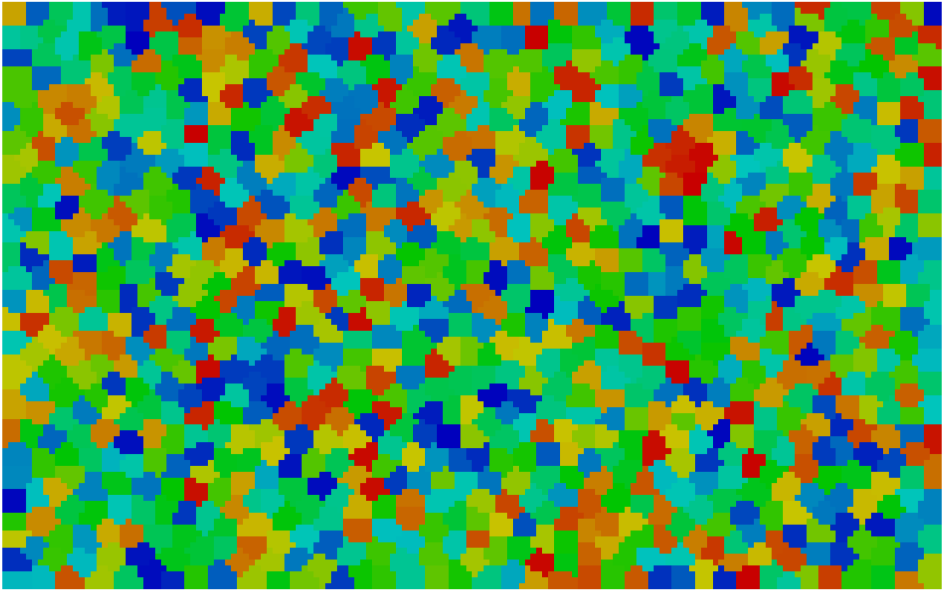

Problem Setup

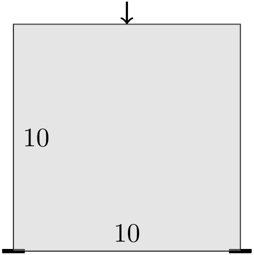

(a) XY view

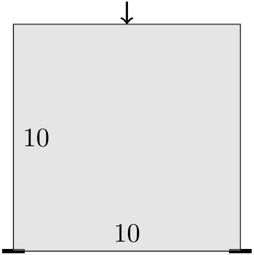

(b) XZ view

- 64000 8-node brick elements, 206715 dofs
- Total Lagrangian formulation, finite strain
- St. Venant-Kirchhoff material
- Jacobi-Preconditioned Conjugate Gradient
- Newton-Raphson nonlinear solver
- Minimum compliance optimization problem

$$\begin{aligned} \underset{\mathbf{u}\in\mathbb{R}^{n_{\mathbf{u}}},\ \boldsymbol{\mu}\in\mathbb{R}^{n_{\boldsymbol{\mu}}}}{\text{minimize}} \quad & \mathbf{f}_{\text{ext}}{}^T\mathbf{u} \\ \text{subject to} \quad & V(\boldsymbol{\mu}) \leq 0.15 \cdot V_0 \\ & \mathbf{r}(\mathbf{u},\ \boldsymbol{\mu}) = 0 \end{aligned}$$

- Gradient computations: Adjoint method
- Optimizer: SNOPT
- Maximum ROM size: $k_{\mathbf{u}} \leq 5$

Optimal Solution Comparison

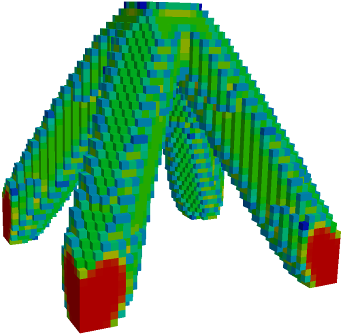

HDM

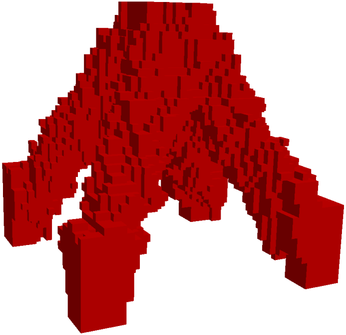

CTRPOD + Φ_μ adaptivity

- **HDM**, elapsed time = 179176s
- **CTRPOD**+Φ_μ adaptivity, elapsed time = 15208s
- Speedup $\approx 12\times$

Solution after 68 HDM Evaluations

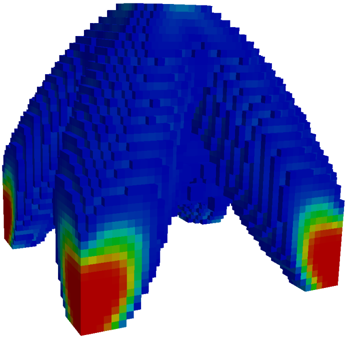

HDM

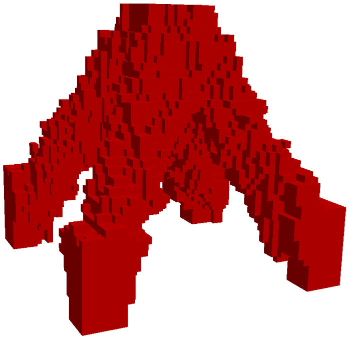

CTRPOD + Φ_μ adaptivity

- CTRPOD + Φ_μ adaptivity: superior approximation to optimal solution than HDM approach after fixed number of HDM solves (68)
- Reasonable option to *warm-start* HDM topology optimization

Outline

Problem Formulation

Goal: Rapidly solve PDE-constrained optimization problems of the form

$$\begin{aligned} \underset{\boldsymbol{U},\ \boldsymbol{\mu}}{\text{minimize}} \quad & \int_{T_0}^{T_f} f(\boldsymbol{U}(t), \boldsymbol{\mu}, t)\, dt \\ \text{subject to} \quad & \frac{\partial \boldsymbol{U}}{\partial t} + \nabla \cdot \boldsymbol{F}(\boldsymbol{U}, \nabla \boldsymbol{U}, \boldsymbol{\mu}) = 0 \end{aligned}$$

- Two-Phase approach
  - Develop *globally* high-order numerical scheme (HDM)
  - Adapt proposed trust-region approach with adaptive model reduction (ROM)
- Collaboration with P.-O. Persson (UCB)

Highlights

- Spatial discretization
  - High-order Discontinuous Galerkin Arbitrary-Lagrangian-Eulerian (DG-ALE)
  - GCL augmentation
- Temporal discretization
  - Diagonally-Implicit Runge Kutta
- Output integration
  - Solver-consistent
  - DG-ALE for spatial integrals
  - DIRK for temporal integrals
- Fully-discrete unsteady adjoint method

Energetically-Optimal Trajectory

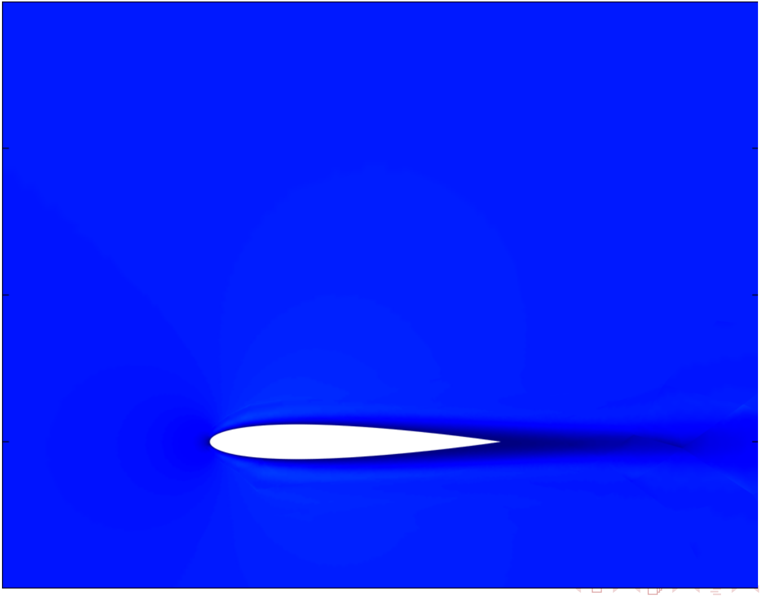

Coming soon(ish) ...

Collaboration with Kevin Carlberg and Drew Kouri

Outline

Summary

Summary

- Introduced nonlinear trust region framework for optimization using adaptive reduced-order models
- Demonstrated approach on canonical problem from aerodynamic shape optimization
    - Factor of 4 fewer queries to HDM than standard PDE-constrained optimization approaches
- Extension to problems with large-dimensional parameter space and constraints (topology optimization)
    - Order of magnitude speedup on canonical 2D/3D problems

Future Work

- **Convergence proof** for proposed progressive optimization framework
- Incorporate **hyperreduction** to realize speedups
- Application to **large-scale**, 3D problems

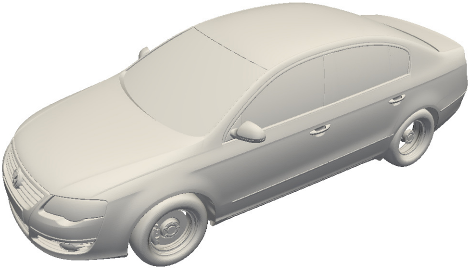

- Extension to **unsteady** PDE-constrained optimization
- Extension to **stochastic** PDE-constrained optimization

References I

Afanasiev, K. and Hinze, M. (2001).
Adaptive control of a wake flow using proper orthogonal decomposition.
Lecture Notes in Pure and Applied Mathematics, pages 317–332.

Arian, E., Fahl, M., and Sachs, E. W. (2000).
Trust-region proper orthogonal decomposition for flow control.
Technical report, DTIC Document.

Belytschko, T., Liu, W., Moran, B., et al. (2000).
Nonlinear finite elements for continua and structures, volume 26.
Wiley New York.

Bonet, J. and Wood, R. (1997).
Nonlinear continuum mechanics for finite element analysis.
Cambridge university press.

Chen, Y., Davis, T. A., Hager, W. W., and Rajamanickam, S. (2008).
Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate.
ACM Transactions on Mathematical Software (TOMS), 35(3):22.

Constantine, P. G., Dow, E., and Wang, Q. (2014).
Active subspace methods in theory and practice: Applications to kriging surfaces.
SIAM Journal on Scientific Computing, 36(4):A1500–A1524.

Fahl, M. (2001).
Trust-region methods for flow control based on reduced order modelling.
PhD thesis, Universitätsbibliothek.

Gill, P. E., Murray, W., and Saunders, M. A. (2002).
Snopt: An sqp algorithm for large-scale constrained optimization.
SIAM journal on optimization, 12(4):979–1006.

References II

Hinze, M. and Matthes, U. (2013).
Model order reduction for networks of ode and pde systems.
In *System Modeling and Optimization*, pages 92–101. Springer.

Kunisch, K. and Volkwein, S. (2008).
Proper orthogonal decomposition for optimality systems.
ESAIM: Mathematical Modelling and Numerical Analysis, 42(1):1.

Lassila, T. and Rozza, G. (2010).
Parametric free-form shape design with pde models and reduced basis method.
Computer Methods in Applied Mechanics and Engineering, 199(23):1583–1592.

LeGresley, P. A. and Alonso, J. J. (2000).
Airfoil design optimization using reduced order models based on proper orthogonal decomposition.
In *Fluids 2000 conference and exhibit, Denver, CO*.

Lieberman, C., Willcox, K., and Ghattas, O. (2010).
Parameter and state model reduction for large-scale statistical inverse problems.
SIAM Journal on Scientific Computing, 32(5):2523–2542.

Manzoni, A., Quarteroni, A., and Rozza, G. (2012).
Shape optimization for viscous flows by reduced basis methods and free-form deformation.
International Journal for Numerical Methods in Fluids, 70(5):646–670.

Maute, K. and Ramm, E. (1995).
Adaptive topology optimization.
Structural optimization, 10(2):100–112.

Persson, P.-O., Willis, D., and Peraire, J. (2012).
Numerical simulation of flapping wings using a panel method and a high-order navier–stokes solver.
International Journal for Numerical Methods in Engineering, 89(10):1296–1316.

Rozza, G. and Manzoni, A. (2010).
Model order reduction by geometrical parametrization for shape optimization in computational fluid dynamics.
In *Proceedings of ECCOMAS CFD*.

References III

Yue, Y. and Meerbergen, K. (2013).
Accelerating optimization of parametric linear systems by model order reduction.
SIAM Journal on Optimization, 23(2):1344–1370.

Zahr, M. J. and Farhat, C. (2014).
Progressive construction of a parametric reduced-order model for pde-constrained optimization.
International Journal for Numerical Methods in Engineering.