Accelerating PDE-Constrained Optimization using Adaptive Reduced-Order Models

Investigators:
Matthew J. Zahr
Charbel Farhat

Army High Performance Computing Research Center

The desire to solve optimization problems governed by partial differential equations exists in all fields of science and engineering. These PDE-constrained optimization problems inevitably require a large number of solutions of the partial differential equation of interest and become prohibitively expensive if a fine discretization is required. We introduce a fast algorithm for solving such optimization problems that leverages adaptive reduced-order models and is provably globally convergent.

Motivation

- The problem of **designing a hull** to maximize energy absorption of the armor in order **protect the occupants** in the event of an underbody blast is of utmost importance
- Testing a given design incurs substantial cost ($5M): turn to **computational tools** to **analyze and design** such systems
- Computational analysis of a single design may take **many hours on a supercomputer** due to the complex geometry and physics involved
- Optimization of such a system may require **simulation of thousands of designs**, rendering the problem practically infeasible

- We introduce a framework for solving such PDE-constrained optimization problems using **reduced-order models** with the goal of substantial CPU saving.

Goal

Accelerate solution of a PDE-constrained optimization problem using a Reduced-Order Model (ROM) as a surrogate for the PDE

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

$\mathbf{R}(\mathbf{w},\boldsymbol{\mu})$ – discretized PDE
$\mathbf{w}$ – state vector
$\boldsymbol{\mu}$ – parameter

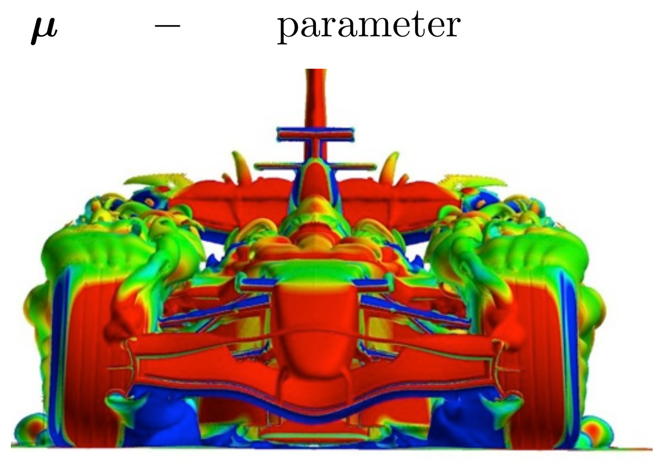

Optimization via Adaptive Reduced-Order Model

Assume state vector lies in *r*-dimensional *trial* subspace where $r \ll N$, defined by the Reduced Basis (RB) and project equations into *r*-dimensional *test* subspace.

$$\mathbf{w} = \boldsymbol{\Phi}\mathbf{y} \qquad \boldsymbol{\Psi}^T\mathbf{R}(\boldsymbol{\Phi}\mathbf{y},\boldsymbol{\mu}) = 0$$

Non-quadratic trust-region model problem: *error-aware ROM-constrained optimization*

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

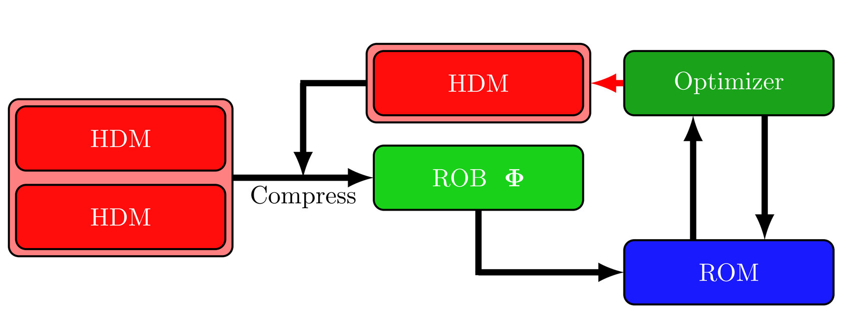

Workflow Schematic

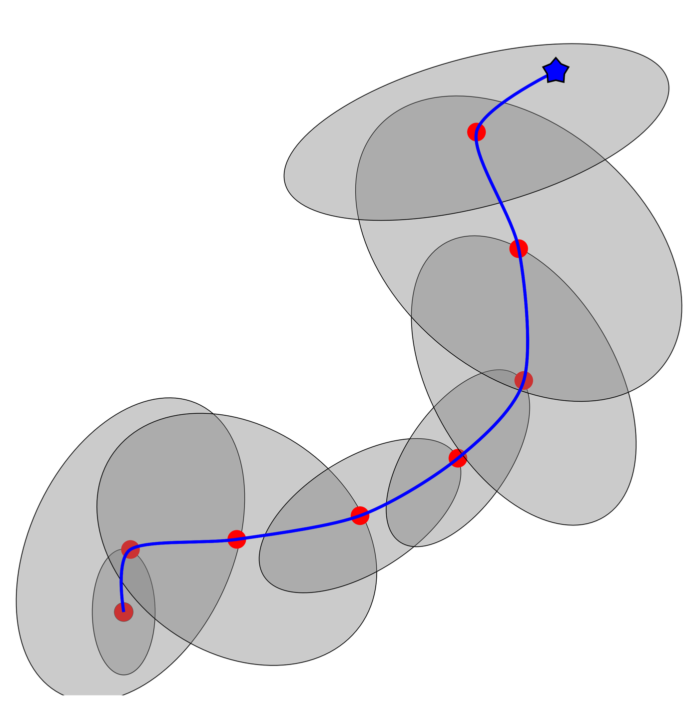

Parameter Space

Global Convergence Theorem

Define the implicit functions

$$g(\boldsymbol{\mu}) := f(\mathbf{w}(\boldsymbol{\mu}),\boldsymbol{\mu}) \qquad m_k(\boldsymbol{\mu}) := f(\boldsymbol{\Phi}_k\mathbf{y}(\boldsymbol{\mu}),\boldsymbol{\mu})$$

and assume they are continuously differentiable with bounded Hessian. If the following relaxed first-order conditions are met (guaranteed by the proposed minimum-residual primal and sensitivity reduced-order model framework): $\exists\, \xi > 0$

$$m_k(\boldsymbol{\mu}_k) = g(\boldsymbol{\mu}_k) \qquad \|\nabla g(\boldsymbol{\mu}_k) - \nabla m_k(\boldsymbol{\mu}_k)\| \leq \xi \min\{\|\nabla m_k(\boldsymbol{\mu}_k)\|, \Delta_k\}$$

then the proposed trust-region algorithm produces a sequence of iterates that satisfies

$$\liminf_{k\to\infty} \|\nabla m_k(\boldsymbol{\mu}_k)\| = \liminf_{k\to\infty} \|\nabla g(\boldsymbol{\mu}_k)\| = 0$$

Thus the algorithm converges to a local minimum from *any starting point*.

Aerodynamic Shape Optimization

In this section, the proposed optimization algorithm that leverages adaptive reduced-order models is compared to a standard technique for PDE-constrained optimization on the problem of recovering a RAE2822 geometry from a NACA0012 geometry by considering only the discrepancy in the pressure distribution.

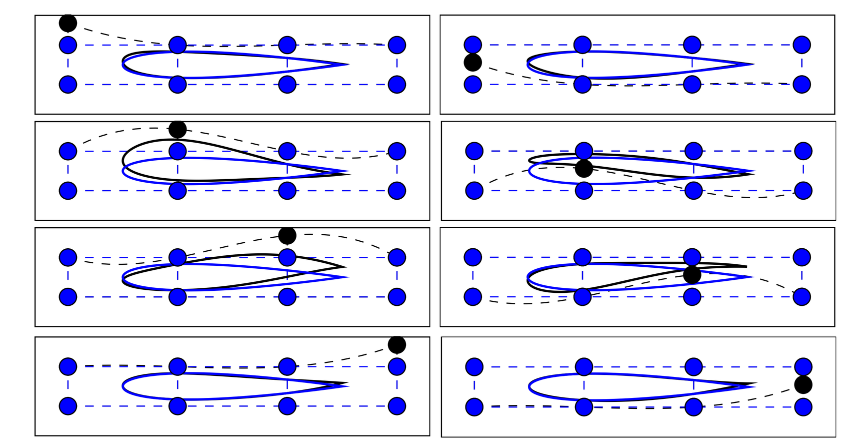

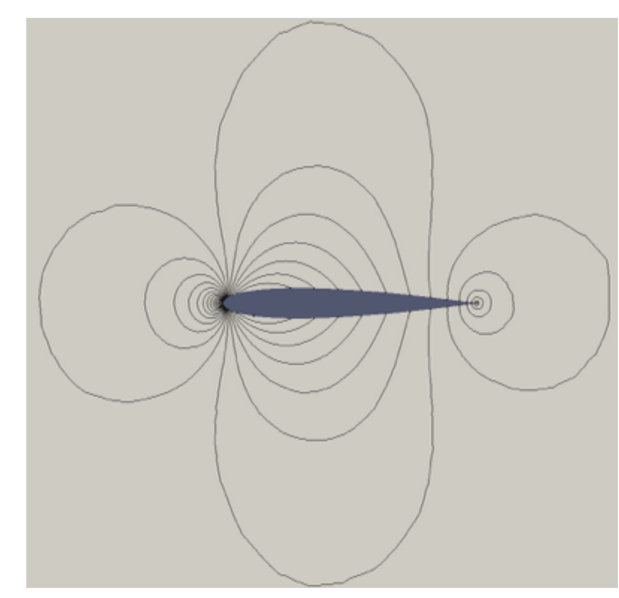

NACA0012
(M = 0.5, AOA = 0)

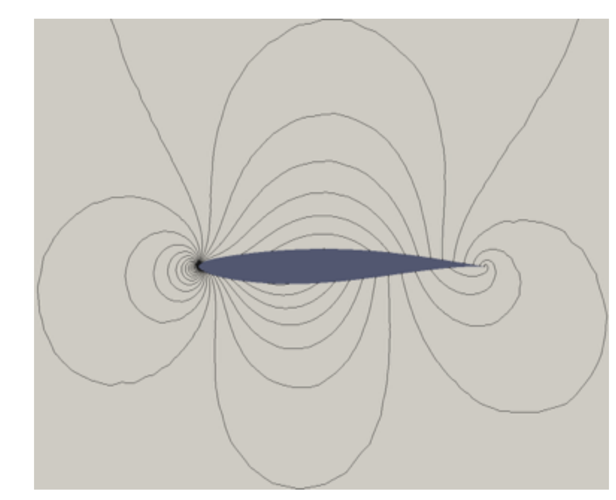

RAE2822
(M = 0.5, AOA = 0)

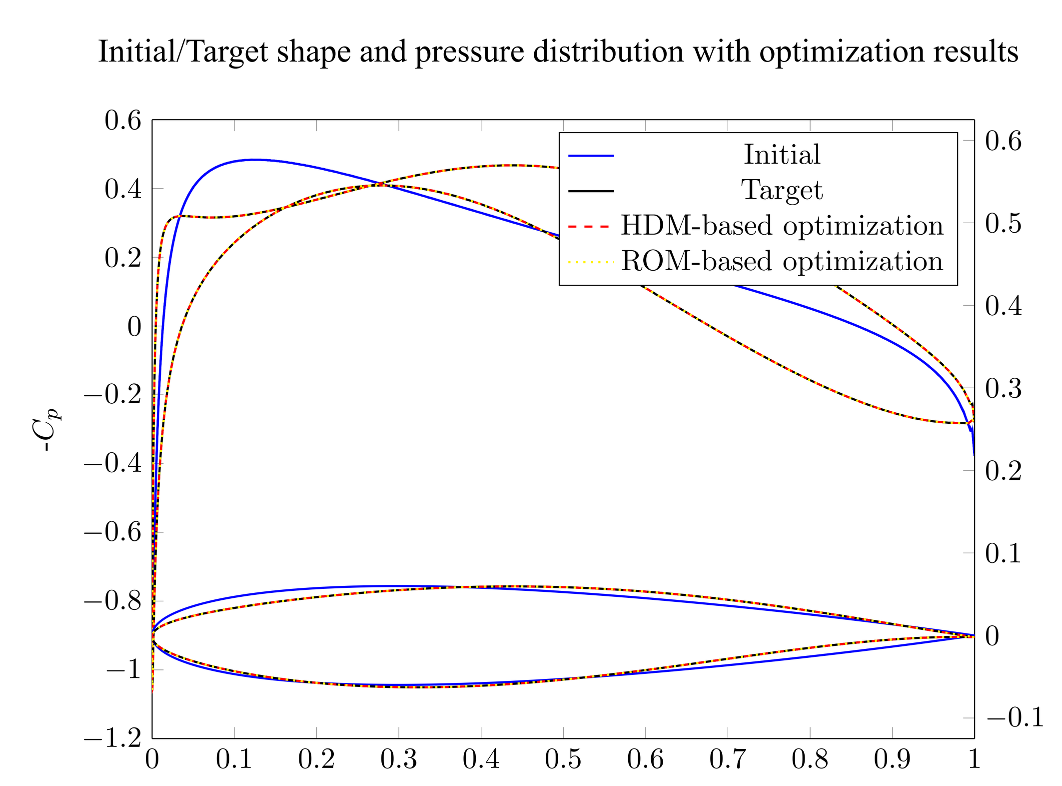

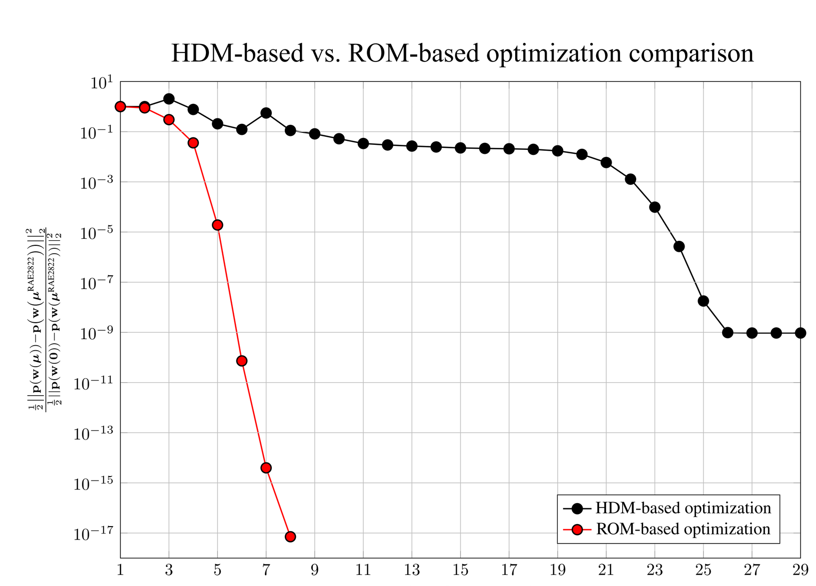

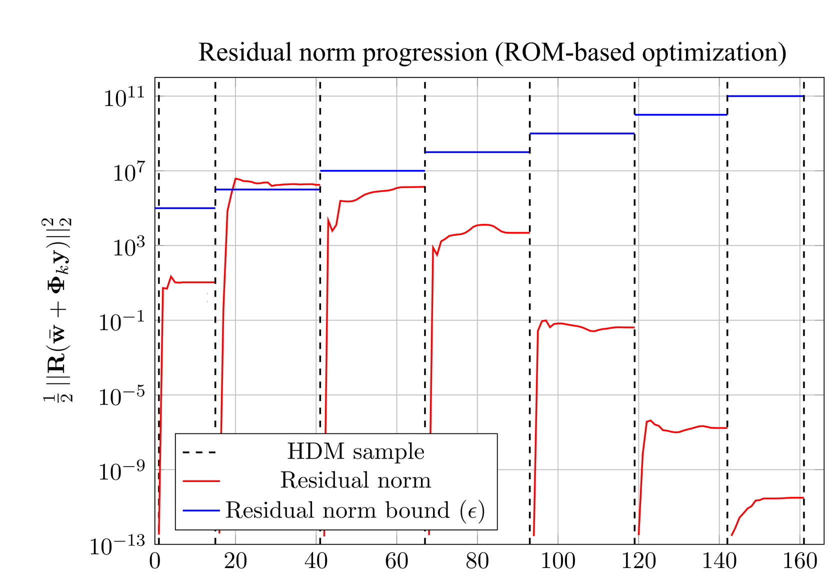

Conclusion

- Introduced new globally convergent method for accelerating PDE-constrained optimization using Reduced-Order Models.
- **Factor of 4 fewer HDM queries** observed on aerodynamic shape optimization problem where the optimal solution was recovered to machine precision.
- Ongoing work is focused on demonstrating the proposed approach on a large-scale problem - design of the Common Research Model.

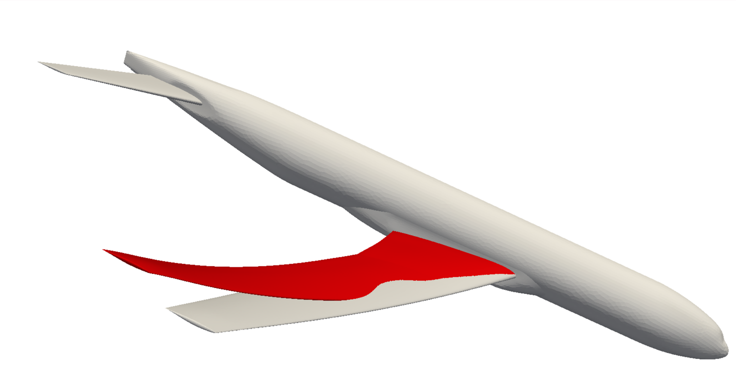

Contact: mzahr@stanford.edu

Collaborations

The following collaboration efforts are planned:

ARL/CSD: Pat Collins, on the CFD ROM component and its introduction at ARL/VTD where AERO-F is now known. Anticipated applications are design optimization of MAVs and flapping wings, among others

TARDEC: Matt Castanier, on the structural dynamics ROM component with applications to armor design optimization