Oral presentation | Fluid-structure interaction
Fluid-structure interaction-II
Mon. Jul 15, 2024 2:00 PM - 4:00 PM Room A

[2-A-01] Multi-Fidelity Gradient-based Aerostructural Optimization

*Markus Peer Rumpfkeil[1], Phil Beran[2] (1. University of Dayton, 2. U.S. Air Force Research Laboratory)

Keywords: Multi-fidelity Constrained Optimization, Flutter, Aeroelastic Simulations

Twelfth International Conference on Computational Fluid Dynamics (ICCFD12), Kobe, Japan, July 14-19, 2024

ICCFD12-xxxx

Multi-Fidelity Gradient-based Aerostructural Optimization

Markus Peer Rumpfkeil[1] and Philip Beran[2]

[1] Department of Mechanical and Aerospace Engineering, University of Dayton, USA
[2] U.S. Air Force Research Laboratory, Wright-Patterson Air Force Base, USA
Corresponding author: Markus.Rumpfkeil@udayton.edu

Abstract: The traditional design process relies heavily on lower fidelity models for expedience and resource savings. However, the reduced accuracy and reliability of low-fidelity tools often lead to the late discovery of design defects or inadequacies. These deficiencies result either in costly changes or the acceptance of configurations that do not meet expectations. Multi-fidelity methods attempt to blend the increased accuracy and reliability of high-fidelity models with the reduced cost of low-fidelity models. In this paper, a gradient-based multi-fidelity constrained optimization framework is applied to an aeroelastic drag minimization of an efficient supersonic air vehicle (ESAV). Constraints are imposed on the lift and pitching moments. The high- and low-fidelity analysis levels considered are Euler and panel solutions, respectively, all combined with a modal structural solver. The coupling of the two solvers is accomplished with FUNtoFEM, a Python-based framework developed for both high-fidelity aeroelastic analysis and adjoint-based aeroelastic optimization. This work is a further step towards developing the capabilities of a multi-fidelity, multidisciplinary analysis and optimization of this type of vehicle.

Keywords: Multi-fidelity, Aeroelastic Simulations, Gradient-based optimization, MDO

Nomenclature

AR	=	aspect ratio
C_L	=	lift coefficient
C_D	=	drag coefficient
C_M	=	pitching moment coefficient
D	=	shape design variables
f	=	objective or cost function
k	=	stiffness
Φ	=	mode-shapes
M	=	Mach number
X_A	=	aerodynamic mesh coordinates
X_S	=	structural mesh coordinate
α	=	angle of attack ($\circ$)
Λ	=	sweep angle ($\circ$)
λ	=	taper ratio

1 Introduction and Background

To push the bounds of aircraft performance, aircraft makers must also be willing to push the conventional notion of what a well performing aircraft looks like. This requires the analysis of unconventional aircraft configurations for which little to no historical data exists. Thus, new aircraft configurations have to be analyzed with computational approaches in the early stages of development to predict their performance

and feasibility and it is very important to employ modeling methods that capture the complicated and often coupled physics of flight as accurately as possible with minimal computational cost.

The efficient supersonic air vehicle (ESAV) is an example for such an unconventional aircraft since it features swept lambda wings with no tail, that is envisioned to perform in the subsonic and supersonic regimes [1]. Flying a tailless vehicle in a complex flight envelope presents challenges in trim and stability, with the accurate prediction of aerodynamic forces and moments being critical to its success; this necessity for accurate physics modeling is further exacerbated by the fact that such a vehicle must both dwell in and fly through the transonic regime, for which aerodynamic modeling is known to be difficult.

As a great motivation for the development of multi-fidelity methods, one can take a look at the Department of Defense life cycle acquisition process. According to a study on pre-milestone A and early-phase systems engineering [2], 70-75% of a system's life cycle cost is locked in during the analysis of alternatives and down-selection to a single configuration for maturation. It is during this phase that designers have the greatest freedom to make critical design changes. However, it is also at this stage that they have the lowest confidence in performance predictions due to the predominantly use of low-fidelity tools such as historical data and simulations with lack of interdisciplinary couplings. However, as configurations deviate further from the envelope of historical data, designers run the risk of missing critical, design-driving phenomena stemming from individual disciplines or their couplings [3]. The discovery of these so-called "late defects" requires costly changes late in the design cycle and/or acceptance of a product with worse-than-advertised performance. Multi-fidelity techniques may help address these problems by enabling the injection of higher fidelity analysis earlier in the design process for performance checks and the identification of design-driving phenomena that would not otherwise be discovered until late in design.

A model's level of fidelity may be defined as its accuracy in determining a quantity or behavior of interest of a real system. Fidelity levels fall on a multi-dimensional spectrum. Variations may include, for example, changes in physics modeled, inclusion or coupling of different disciplines, variation in geometric detail, faithfulness of boundary and loading conditions, or reduction of numerical error in a solution process [4]. What constitutes a change in fidelity depends on the intended use and the user. Multi-fidelity methods strive to combine these various levels. Lower-cost (time, money, personnel, computing, etc.), less-accurate or reliable predictions may be combined with higher-cost, more accurate or reliable predictions to identify design-driving physics before final down-selection. Lower fidelity methods may be leveraged to speed the refinement of the final design, and even accelerate the final analysis at the highest fidelity. Multi-fidelity methods are not a panacea but they present an approach that, in concert with advances in other design tools and processes, may help to design better aircraft at lower cost. Ultimately, designers should be able to access and combine any necessary tools during any phase of the design process to aid and guide their decision making.

The overall goal of this research is to set up a geometrically parameterized, multi-fidelity model of an efficient supersonic air vehicle (ESAV) so that aeroelastic and control analyses can be performed to enable multi-fidelity and multidisciplinary design optimization. In a recent paper [5] we showed multi-fidelity parameter studies of this vehicle by varying flow conditions and geometric parameters including flap deflections to enable moment controls. We also conducted an inverse design optimization to demonstrate optimization capabilities, however, the gradients were finite-differenced as no adjoint was available at the time. Adjoint gradients were implemented via FUNtoFEM and demonstrated for a single-fidelity optimization of this vehicle in Rumpfkeil and Beran [6].

Multi-fidelity approaches are frequently used in engineering design when high-fidelity models are too expensive to use directly and lower fidelity models of reasonable accuracy exist. In an optimization context, corrected low-fidelity data is typically used in a series of sequential optimizations bounded by trust regions around the approximate model. Trust Region Model Management (TRMM) frameworks [7, 8] allow for approximate models to be used in place of more expensive high-fidelity models over a limited region of the design space. Each sub-optimization is independent of the previous one, except for the starting design point and trust region size since for every outer loop iteration, the trust region is re-centered about the current design, and the sizes of the move limits or subproblem bounds are determined via a heuristic [7, 9]. Thus, a design optimization may be carried out as a sequence of optimizations of the approximate subproblem which is provably convergent to a local optimum of the high-fidelity problem provided that the approximate problem maintains first-order consistency (i.e., matching of function and gradient values) with the high-fidelity problem [8].

Examples of applying multi-fidelity concepts to gradient-based optimization in the literature are based on the Trust Region Model Management (TRMM) approach presented by Lewis [7] and Alexandrov et al. [8] for unconstrained optimization. Design constraints are explicitly considered by Rodríguez, Renaud,

and Watson [10] and Alexandrov et al. [11]. Alexandrov et al. [12] also extended the TRMM concepts to a general Approximation and Model Management Optimization framework, demonstrating it for augmented Lagrangian optimization, the multilevel algorithms for large-scale constrained optimization (MAESTRO) framework, and a trust-region Sequential Quadratic Programming (SQP) method.

While its simplicity makes TRMM easy to implement its convergence rate toward the high-fidelity optimum depends on the accuracy of the successive approximate models for the high-dimensional subspaces [13, 14, 15, 16]. This is particularly problematic early on in the optimization when usually most of the reduction in the objective function is achieved but data for improving the approximate models is limited. This inadequacy of early approximate models is compounded by the fact that information about the high-fidelity problem is not propagated from one sub-optimization to the next, aside from the scaling of the trust region. Ideally, once the expensive high-fidelity truth function is evaluated, its value and gradient would be used to guide further progress by the optimizer. Also, while TRMM is provably convergent to an optimum of the high-fidelity problem, the rate of convergence in the neighborhood of the optimum can be slow [17]. Finally, each iteration of TRMM may require many low-fidelity function evaluations. While the lower fidelity prediction is cheaper than that of the high-fidelity, its cost may not be negligible. To the contrary, in order to obtain better approximations of the truth function, a better low-fidelity approximation may be required, making each evaluation more costly.

To address these issues a unified multi-fidelity quasi-Newton approach has been previously developed by Bryson and Rumpfkeil [9, 18, 19] that preserves an estimate of the inverse Hessian between iterations and which determines search directions from high-fidelity data using multi-fidelity surrogate models for the line searches where unified, hybrid additive-multiplicative PCE corrections are employed. It was demonstrated that the resulting algorithm produced better search directions, maintained larger step sizes, and required significantly fewer low-fidelity function evaluations than TRMM approaches for a large variety of analytical test functions [18]. In that algorithm the approximate high-fidelity inverse Hessian and the high-fidelity gradient are used to calculate the expected (bound constrained) optimal point as in typical quasi-Newton methods. While the search direction is forced to be along the quasi-Newton step the line search itself is performed on the general low-fidelity model, corrected to match high-fidelity data at two points using a unified, hybrid additive-multiplicative polynomial chaos expansion (PCE) approach also developed by Bryson and Rumpfkeil [20]. The first-order consistency condition is satisfied by fitting the models to the function values and gradients of only two points, one of which is the current center point of the trust region. The second point is the most recently evaluated point if the step was rejected; otherwise, the previous accepted design is used. However, the convexity of the approximate response is not guaranteed as it is in the classical trust region method. As a safeguard, a trust region imposes move limits on the design step to prevent the line search from exploiting extrapolatory errors in the correction functions beyond the expected optimal point. The trust region is adjusted according to the same rules as in a TRMM, depending on the performance of the approximate model relative to the high-fidelity function [9, 18, 19].

This framework is only able to accommodate constraints indirectly through penalty functions. Thus, it was extended by Rumpfkeil and Beran [21] to include constraints directly and its effectiveness was demonstrated on a variety of benchmark problems. For this, a mono-fidelity sequential least-squares programming (SLSQP) [22] algorithm has been adapted and incorporated into the unified multi-fidelity quasi-Newton framework. For constraint evaluations during the line-search phase multi-fidelity PCE surrogate models as described above with all available low- and high-fidelity data points are constructed and evaluated for each constraint separately.

2 Aeroelastic Analysis Solver and Adjoint Derivatives

In an aeroelastic analysis it is necessary to capture the complex and coupled physical phenomena present in the operating environment and flight regime of modern aircraft. The process begins with the selection of a set of design parameters via a parametric study, surrogate model training point selection, optimizer, etc. A model configuration and geometry generator interprets and maps the given set of design parameters into the required aerodynamic and structural analysis models which also usually require an adequate mesh. For a complex aeroelastic application, a tailless, efficient supersonic air vehicle (ESAV) illustrated in Figure 1(a), will be employed here. This type of configuration has been the subject of many studies at AFRL [5, 6, 23, 24, 25, 26, 27] and is of interest due to potential savings in weight and drag since the empennage is eliminated.

2.1 Geometry

A shared geometric representation of the vehicle is paramount for any multi-fidelity and multi-disciplinary analysis and optimization. Using a single source ensures that the inputs given to each analysis are consistent and aids in the transfer of data between disciplines or fidelity levels. This objective is achieved by using the Computational Aircraft Prototype Synthesis (CAPS) [28] geometry program. Within CAPS exists a parametric, attributed model of the vehicle. The attributes provide logical information required for the generation of analysis inputs. For example, attributes identify the vehicle skins where aeroelastic data transfers take place, symmetry planes for the application of boundary conditions, and bodies to which material properties should be applied. When a shape design parameter is changed, the geometry is regenerated, and analysis models (meshes, properties, etc.) may be requested for various disciplinary analyses at varying levels of fidelity. An additional benefit of having a fully parametric model is that it is differentiable all the way to geometric primitives which means that very accurate gradients can be obtained for gradient-based optimization and other purposes.

The analysis model generation proceeds as follows. Using the current design parameters, the airfoil cross-sections and the planform shape are determined. Lofting these airfoils provides a solid body representing the outer mold line (OML). These same airfoils also provide the boundaries for defining mid-surface aerodynamic panel models. The CFD domain is generated by subtracting the OML solid from a bounding box. The internal structure results from intersecting the OML body with a grid representing the structural layout. The wing skins are extracted from the outer surface of the OML body. Sample geometric entities used for building various analysis models for the ESAV are shown in Figure 1.

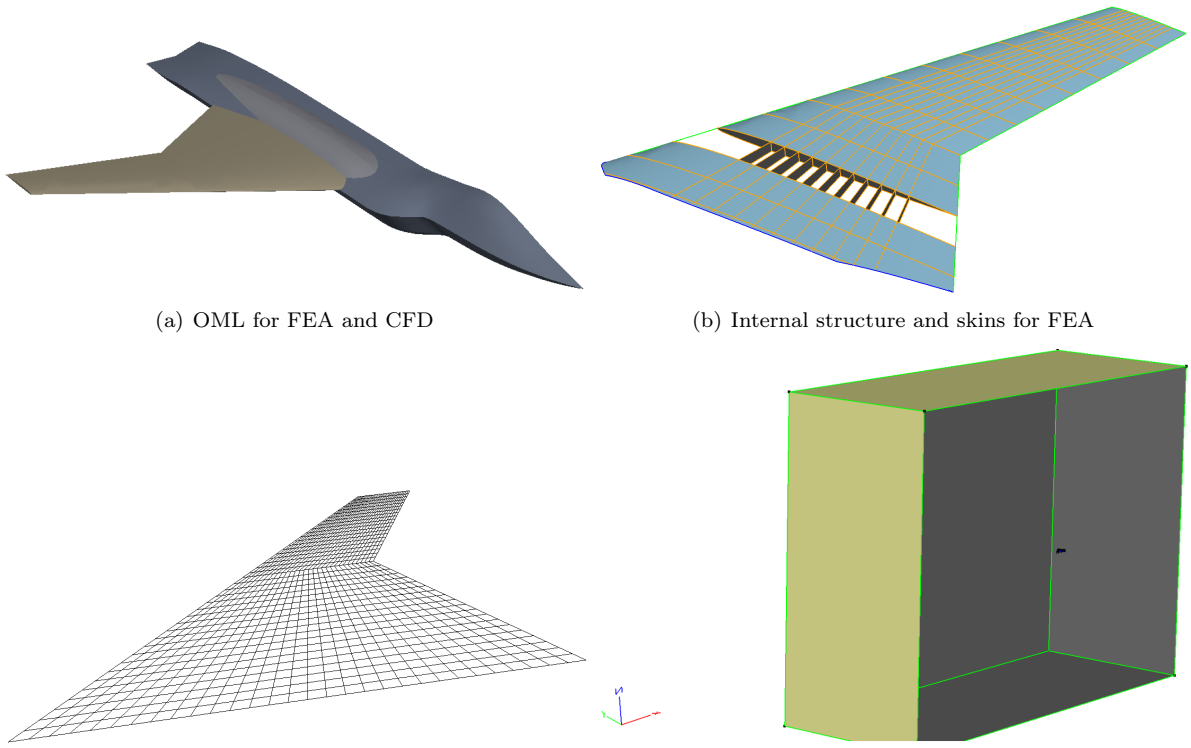

(a) OML for FEA and CFD

(b) Internal structure and skins for FEA

(c) Panels for low-fidelity aerodynamics

(d) Fluid domain for CFD

Figure 1: Representation of ESAV for multi-fidelity and multi-disciplinary analyses.

2.2 Structural Solver

The employed structural solver is the Automated STRuctural Optimization System (ASTROS) [29]. ASTROS can perform static, modal, and transient linear finite-element analysis (FEA), and has an internal aerodynamics capability for static and dynamic aeroelastic analyses (see Subsection 2.3.1). The same structural analysis model feeds all subsequent multi-fidelity aeroelastic analyses described in the next subsection. The generic structural model shown in Figure 1(b) is employed here consisting of 11 spars and 7 ribs each in the inboard and outboard section. The ribs/spars as well as skins are represented

by shell elements made out of aluminum and are 2 *in* and 0.1 *in* thick, respectively. To account for non-structural mass, the material density is scaled by a factor of 1.5. Although important to the structure's behavior, panel stiffeners and spar caps are not currently accounted for, as their geometry would need to be modeled explicitly in CAPS.

2.3 Flow Solver

Two different fidelity levels are employed for the flow solver which will be explained in the following subsections.

2.3.1 Low-fidelity: ASTROS

The low-fidelity analysis is performed with the ASTROS package. The static aeroelastic analysis features in ASTROS provide the capability to analyze and design linear structures in the presence of steady aerodynamic loading. The USSAERO (Unified Subsonic and Supersonic Aerodynamic Analysis) algorithm [30] is employed in ASTROS which can determine the pressure distributions on lifting wing-body-tail combinations for steady subsonic or supersonic flow. The solid boundaries are represented by a number of discrete panels and the flow around the solid boundaries can be estimated by the superposition of source type singularities for non-lifting bodies and vortex singularities for wing-like structures. The aerodynamic forces and influence coefficients are applied to the structure through built-in splines [31]. Boundary conditions can account for the angle of attack, control surface settings (through the bulk data entry AESURF) as well as airfoil camber and thickness. The number of panels was varied until the solution was "panel-converged" and the resulting 1200 aerodynamic panels are shown in Figure 1(c).

2.3.2 High-fidelity: Fun3D

The high-fidelity level considered in this work utilizes NASA's Fully-Unstructured Navier-Stokes 3D (FUN3D) [32] code in Euler mode. FUN3D is a node-centered, implicit, upwind-differencing finite-volume solver. Inviscid wall boundary conditions are applied to the wing outer mold line, and the symmetry plane is modeled with a symmetric boundary condition. The initial surface and volume grids are generated by AFLR4 and AFLR3 [33, 34], respectively, and a grid convergence study yielded satisfactory volume meshes with approximately 950,000 nodes and about five and a half million elements. The underlying surface mesh is shown in Figure 2.

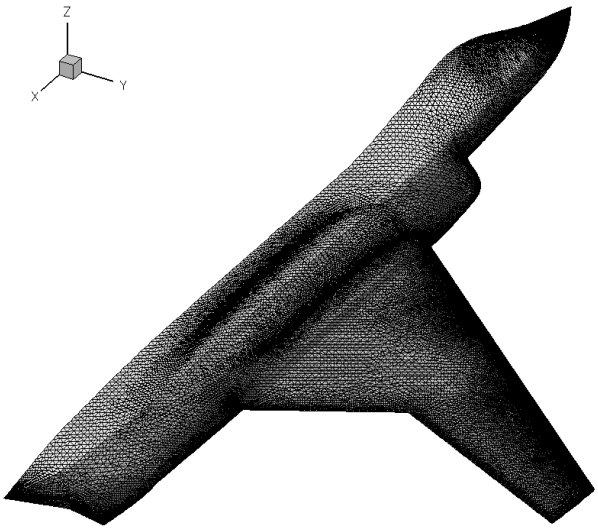

Figure 2: CFD surface mesh.

Since no aeroelastic and transonic experimental data exists for this vehicle for validation purposes, the rigid aerodynamic case for a low subsonic Mach number has been analyzed in Lickenbrock et al. [27] through a sweep of the angle of attack and compared to wind tunnel data obtained by Lockheed Martin [1]. The inviscid FUN3D results closely matched the lift trends and values of the wind tunnel results. When adjusted for viscous effects, by shifting the inviscid FUN3D results so that the drag

matches the Wind Tunnel drag at zero angle of attack, the drag was very close to the wind tunnel results as well [27].

FUN3D's internal aeroelastic capability [35] utilizes a modal structural decomposition approach and FUNtoFEM [36, 37] has an identical solver module implemented as well. The utilized linear structural dynamic equations in both cases are appropriate for small deflections [35]. The deflections are represented as linear combination of eigenmodes and typically only a limited set of the "important" eigenmodes (ten here) are transferred. The transfer of mode shapes from the structural mesh to the fluid surface mesh is handled by CAPS internally and given to FUNtoFEM. Starting from free-stream conditions 250 coupling iterations are employed to yield the final static aeroelastic deflection.

Figure 3 exhibits sample pressure contours for the baseline vehicle at a Mach number of 0.9 and angle of attack of 5°.

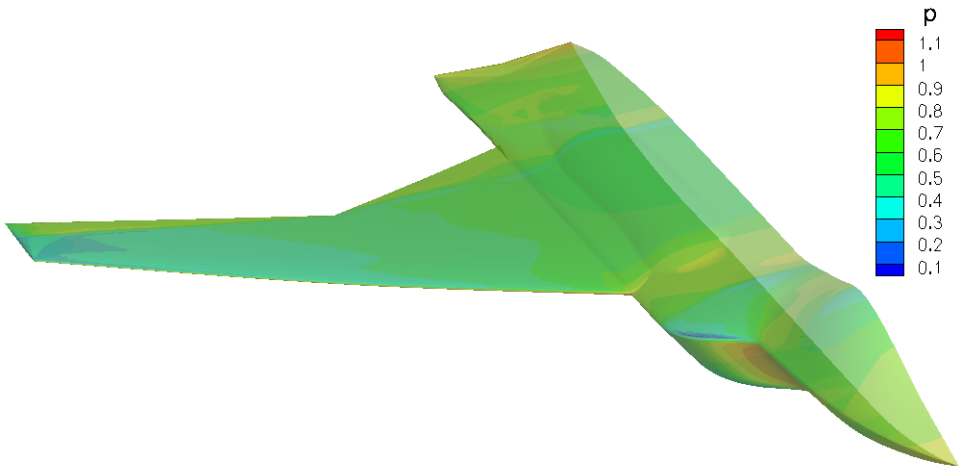

Figure 3: Pressure distribution at a transonic Mach number and moderate angle of attack.

Figure 4 shows the original and the deformed shape at the end of the aeroelastic simulation. The

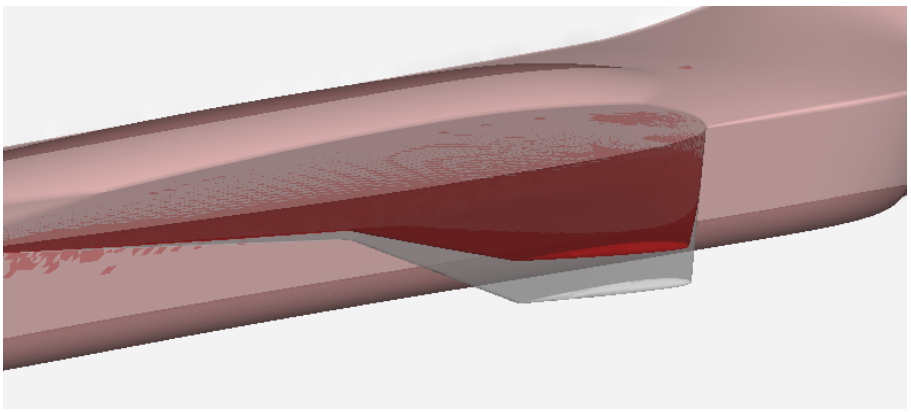

Figure 4: Aeroelastic deformed shape (red) and original shape (white).

corresponding plot of C_L versus the number of iterations as well as the convergence history of the conservation of mass residual is shown in the left of Figure 5 and the displacement response for the first four modes versus number of iterations is shown to the right in the same figure. One can observe a deep convergence of the residual which is important for an accurate adjoint solution for the gradient. The first mode is clearly the most important contribution to the aeroelastic deformation shown in Figure 4.

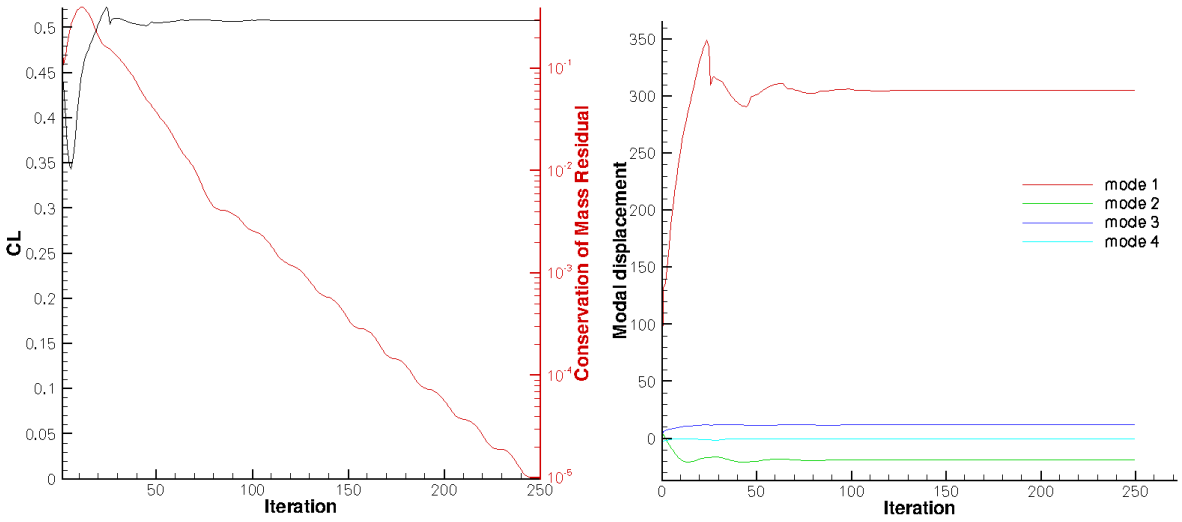

Figure 5: Representative plot of C_L and conservation of mass vs number of iterations (left) and displacement response for the first four modes vs iterations (right).

2.4 FUNtoFEM Coupling and Adjoint Derivative

In order to enable high-fidelity adjoint capabilities the coupling of FUN3D and ASTROS is performed with FUNtoFEM [36, 37], a Python-based framework developed for both high-fidelity aeroelastic analysis and adjoint-based aeroelastic optimization. This work has been already presented by Rumpfkeil and Beran [6] but is repeated here for completeness. The modularity of FUNtoFEM was designed with the goals of being able to couple any adjoint-enabled flow and structural models and to permit its extension to other disciplines such as thermal analysis [36, 37]. FUNtoFEM also includes several load and displacement transfer techniques for aeroelastic coupling. Figure 6 displays a flowchart of the basic data flow and capabilities of FUNtoFEM.

For a modal-based aeroelastic simulation, the objective or cost function, f, is a function of the shape variables, D, through both the structural and aerodynamic mesh coordinates, X_S and X_A [38]:

$$\frac{df}{dD} = \frac{\partial f}{\partial X_A}\frac{dX_A}{dD} + \frac{\partial f}{\partial X_S}\frac{dX_S}{dD} + \frac{\partial f}{\partial k}\frac{\partial k}{\partial X_S}\frac{dX_S}{dD} + \frac{\partial f}{\partial \Phi}\frac{\partial \Phi}{\partial X_S}\frac{dX_S}{dD} \tag{1}$$

where k is the stiffness and Φ are the mode-shapes. The green terms are computed by FUNtoFEM and the orange terms have to be supplied by the employed structural solver. The blue and red terms have to be computed by the geometry generator and aerodynamic or structural mesher, respectively. The mode shape term (last summand) can likely be ignored in the fixed mode assumption used here. In addition, ASTROS does not supply $\frac{\partial \Phi}{\partial X_S}$ and it would be computationally very expensive to finite-difference. Furthermore, it is also not easy to simply finite-difference $\frac{\partial \Phi}{\partial X_S}\frac{dX_S}{dD}$ since the structural mesh and hence mode-shape locations might change with shape variable perturbations. For the second summand one would need $\frac{dX_S}{dD}$ which is currently not available in CAPS (though it is under development) so this summand is ignored here. Thus, the adjoint derivative is approximated with

$$\frac{df}{dD} \approx \frac{\partial f}{\partial X_A}\frac{dX_A}{dD} + \frac{\partial f}{\partial k}\frac{dk}{dD} \tag{2}$$

where $\frac{dX_A}{dD}$ is computed by an ESP body class that was written within the FUNtoFEM framework and $\frac{dk}{dD} = \frac{\partial k}{\partial X_S}\frac{dX_S}{dD}$ is obtained by using a fourth-order accurate central finite-difference of ASTROS. The latter is computationally not too expensive if not too many shape variables are used since ASTROS is running pretty fast for the employed structural mesh.

In order to verify the correctness of this aeroelastic adjoint the derivative of the baseline static lift coefficient (cf. Figure 5) with respect to five design variables (three shape and two flow variables) is compared to finite-difference approximations of various orders (1st, 2nd, and 4th order accurate) and with various stepsizes (relative to each design variable value) in Figure 7. Note that lines between the data points are for visual purposes only. The step size after the adjoint refers to the relative FD step

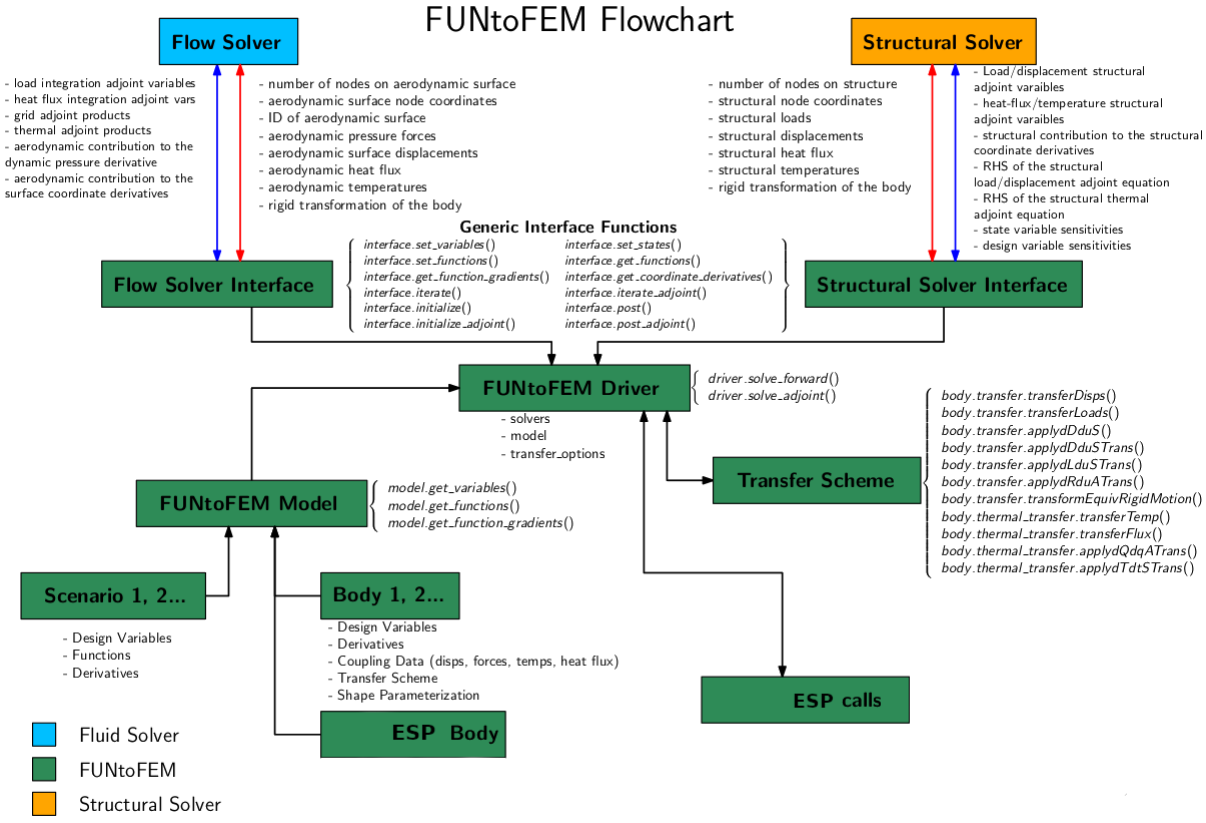

Figure 6: FUNtoFEM flowchart adapted with permission from Jacobson et al. [36, 37].

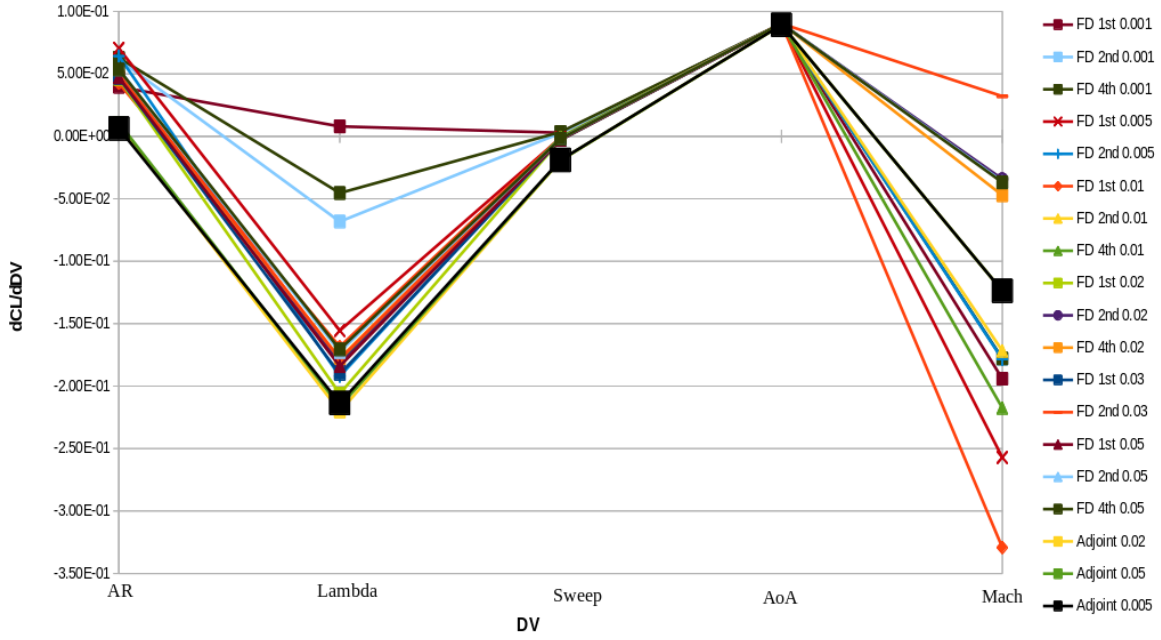

Figure 7: Verification of aeroelastic adjoint gradient with various finite-difference approximations.

size for the fourth-order accurate central finite-difference of ASTROS but as one can infer the adjoint derivatives are pretty close to one another. For the remainder of this paper a step size of 0.005 was chosen for the fourth-order finite-differencing of ASTROS for $\frac{dk}{dD}$.

Overall, one can observe that the FD trends match the adjoint ones though the magnitude can be

off likely due to truncation or round-off error in the finite-differences and the neglection of adjoint terms in Equation (2) as discussed earlier. For flow design variables such as Mach number or angle of attack, FUNtoFEM can compute the full adjoint derivative internally and thus should be correct [38]. Figure 7 shows very good agreement with finite-differences for the angle of attack (AoA) but not for the Mach number. The reason for the relatively large scatter of the various finite-differences for the Mach number derivative is the fact that the lift coefficient is near a (local) maximum at this Mach number and thus depending on which neighbors are involved in the FD stencil very different results are achieved.

When total derivatives of the low-fidelity model's (ASTROS) output quantities of interest (lift, drag and moment coefficients) with respect to design variables are required a fourth-order accurate central finite-difference with a step size of 0.01 is employed which again is not too expensive for the few shape design variables employed here since ASTROS is running pretty fast.

3 Constrained Gradient-based Aerostructural Design Optimization

The ultimate goal of creating a geometrically-parameterized multi-disciplinary analysis model is to use it in a gradient-based optimization framework. Four bounded design variables are considered here as an application example; three geometric parameters namely aspect ratio, $AR \in [2, 6]$, taper, $\lambda \in [0.3, 0.7]$, and sweep, $\Lambda \in [30°, 60°]$ as well as the angle of attack, $\alpha \in [0°, 6°]$. The extreme cases of each geometric parameter can be seen in comparison to the baseline ESAV model in Figure 8.

To bring consistency to the parameter space, the model maintains constant planform areas in the outboard and inboard portions of the wing, and the y-directional span is set to a constant value. This results in longer root chord lengths at smaller aspect ratios, and narrow outboard sections of the wing at higher aspect ratios, as seen in Figure 8. While some of these configurations do not seem feasible at first glance, the extreme values for each parameter will allow for greater design freedom for the optimizer, and appropriately applied constraints should drive the aircraft to a feasible optimal configuration. Since the correct trends of the aeroelastic adjoint gradient have been established in Subsection 2.4 a gradient-based optimization can be attempted.

A lift- and moment-constrained drag minimization with $M = 0.9$ at sea-level is performed. The baseline vehicle has an aspect ratio of, $AR = 4.0003$, taper, $\lambda = 0.492$, sweep, $\Lambda = 41.324°$, and $\alpha = 2.0°$ yielding $C_L^* = 0.17278$ and that lift coefficient will be tried to be maintained (or exceeded) while reducing the drag coefficient, C_D, subject to a pitching moment, C_M, constraint to make sure that the vehicle can be trimmed. The optimization problem can be cast as

$$\begin{aligned} \underset{AR,\lambda,\Lambda,\alpha}{\text{minimize}} \quad & C_D \\ \text{subject to} \quad & c_1 = C_L - C_L^* \geq 0 \\ & c_2 = 0.03 - |C_M| \geq 0 \\ & 2 \leq AR \leq 6 \\ & 0.3 \leq \lambda \leq 0.7 \\ & 30° \leq \Lambda \leq 60° \\ & 0° \leq \alpha \leq 6° \end{aligned} \tag{3}$$

In order to decrease the overall runtimes a coarser mesh with only about 260,000 nodes and 1.4 Million elements is employed for the high-fidelity function and gradient evaluations. The total computational cost estimate assumes that one high-fidelity function evaluation has a cost of one (yielding all 3 required coefficients) and each adjoint gradient costs about the same (one needs to compute three; one each for the lift, drag and moment coefficient). The low-fidelity aeroelastic analysis with ASTROS runs about 100 times faster than FUN3D. Figure 9 shows the objective function convergence history of this optimization and Figures 10 and 11 display the feasibility. Each data point symbolizes an optimization iteration. HF indicates an optimization performed with only high-fidelity function and gradient evaluations while MF indicates a multi-fidelity optimization both using the framework described above. Various starting points in the design space are considered where "Center" implies that all initial design variables are set to the mid-points of their respective intervals. Likewise, "1 Quarter", "1 Third", and "2 Fifth" imply all initial design variables are set to one-quarter, one-third, or two-fifth of their respective intervals (scaled to unity), respectively.

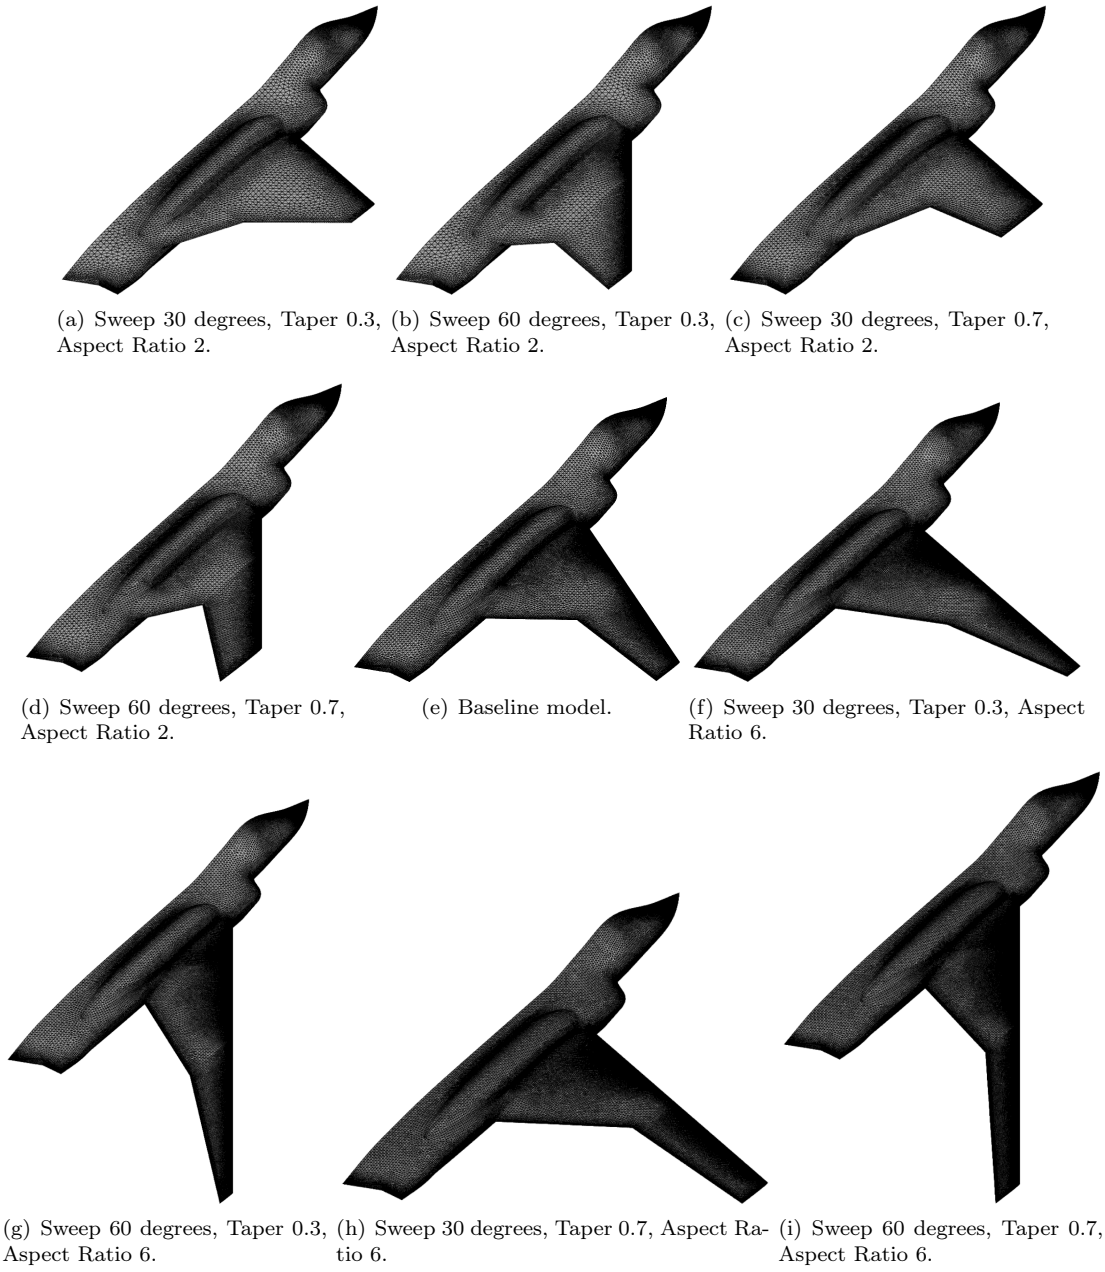

(a) Sweep 30 degrees, Taper 0.3, Aspect Ratio 2.

(b) Sweep 60 degrees, Taper 0.3, Aspect Ratio 2.

(c) Sweep 30 degrees, Taper 0.7, Aspect Ratio 2.

(d) Sweep 60 degrees, Taper 0.7, Aspect Ratio 2.

(e) Baseline model.

(f) Sweep 30 degrees, Taper 0.3, Aspect Ratio 6.

(g) Sweep 60 degrees, Taper 0.3, Aspect Ratio 6.

(h) Sweep 30 degrees, Taper 0.7, Aspect Ratio 6.

(i) Sweep 60 degrees, Taper 0.7, Aspect Ratio 6.

Figure 8: Sample FUN3D models of ESAV with varying geometric parameters.

The more successful optimizations are very close to or actually feasible while reducing the (inviscid) drag coefficient to about 75 counts. The baseline shape and the MF optimized shape starting from the baseline are shown in Figure 12 (both including their aeroelastic deformation) and resulting pressure contours are displayed in Figure 13 showing small differences between the initial and optimized design.

The MF optimized design variable values are: aspect ratio, $AR = 3.98$, taper, $\lambda = 0.488$, sweep, $\Lambda = 40.671°$, and $\alpha = 1.97°$, thus the optimized vehicle has slightly less sweep back, a smaller span, and a skinnier tip than the baseline one.

4 Conclusions

Robust computational analysis models have been created for the multi-fidelity aeroelastic analysis of a relatively complex vehicle including high-fidelity adjoint derivatives. The integration of these analysis models into a gradient-based constrained multi-fidelity optimization framework has been described and a lift- and moment-constrained drag minimization has been performed. This work is the foundation for

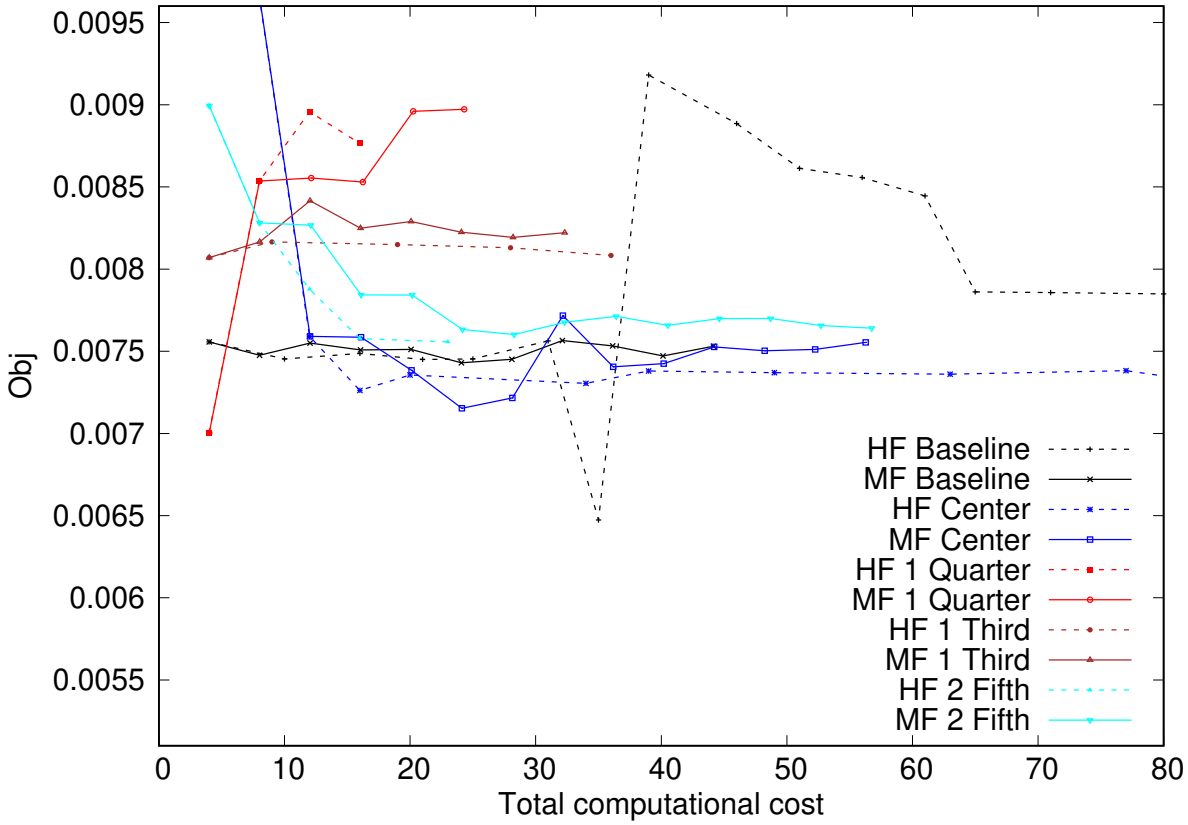

Figure 9: Objective function history for constrained drag minimization with four design variables.

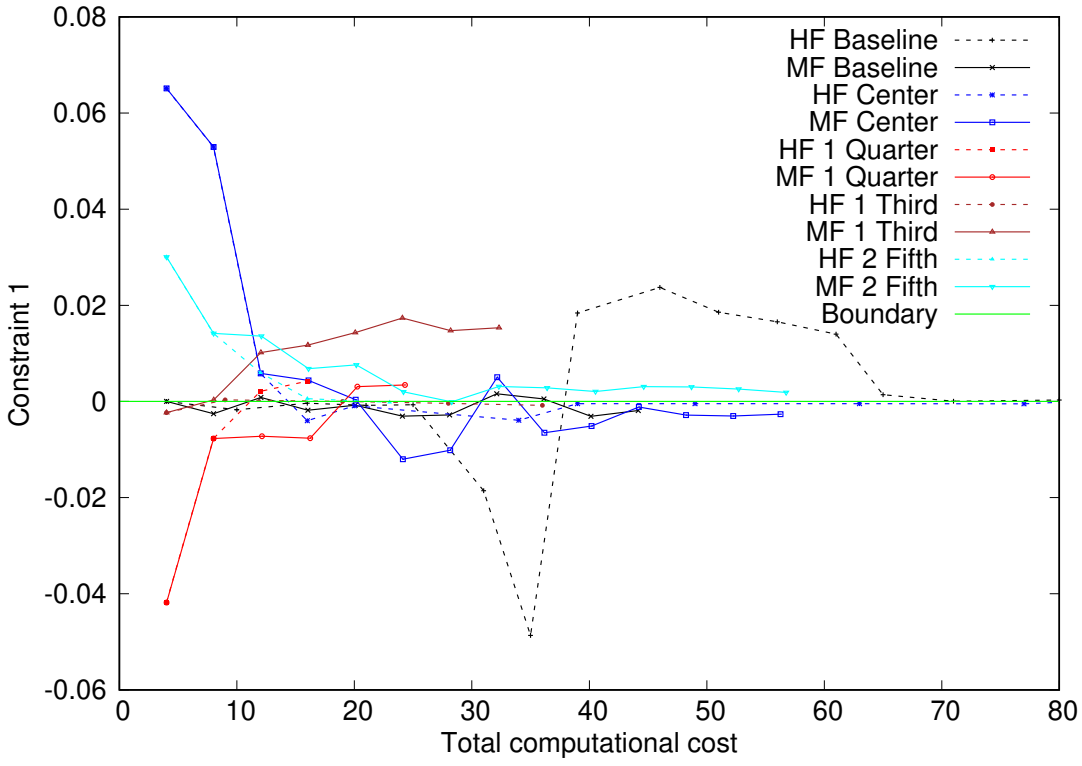

Figure 10: Lift constraint history for constrained drag minimization with four design variables.

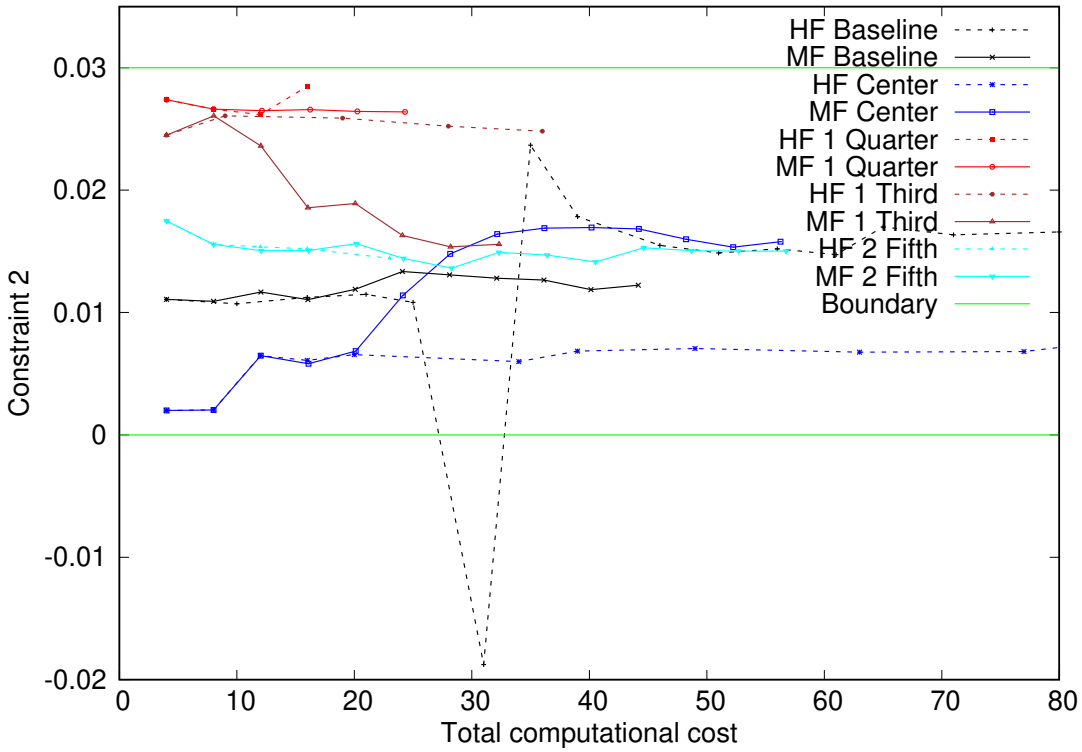

Figure 11: Moment constraint history for constrained drag minimization with four design variables.

Figure 12: Initial (white) and optimal (red) aeroelastically deformed shapes for the lift- and moment-constrained drag minimization.

a planned multi-fidelity and multidisciplinary maximization of the ESAV range where trim, deflection, flutter [39, 40], and other necessary structural constraints will be implemented to maintain the integrity of the model.

Acknowledgments

The second author acknowledges the support of the Air Force Office of Scientific Research (Dr. Fariba Fahroo, Computational Mathematics Program Officer). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

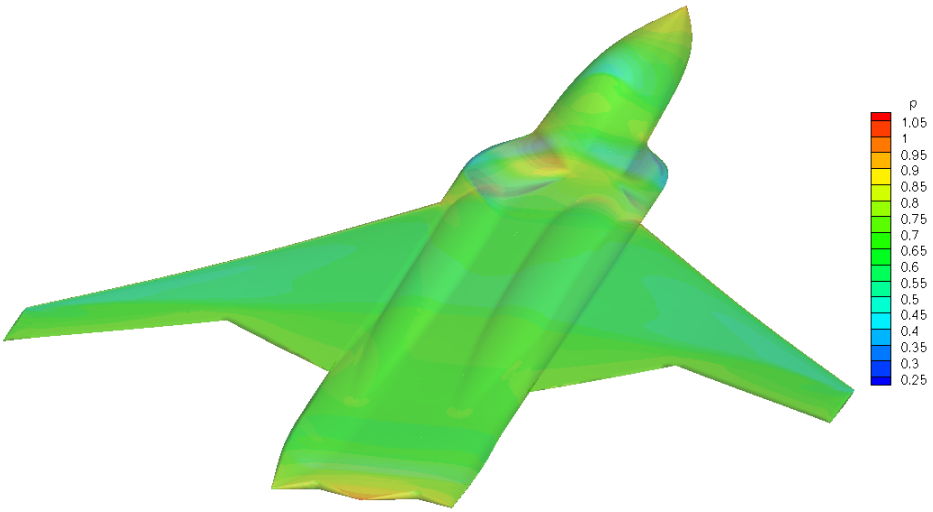

Figure 13: Non-dimensionalized pressure on initial (left) and optimal (right) aeroelastically deformed shapes for the lift- and moment-constrained drag minimization.

References

[1] Clifton Davies, Marc Stelmack, Scott Zink, Antonio De La Garza, and Peter Flick. High Fidelity MDO Process Development and Application to Fighter Strike Conceptual Design. AIAA Paper, 2012-5490, 2012.

[2] N. R. Council. Pre-Milestone A and Early-Phase Systems ENgineering: A Retrospective Review and Benefits for Future Air Force Systems Acquisition. Technical Report National Academies Press, Washington, DC, 2008.

[3] M. Love, R. Yoakum, and R. Britt. Identification of Critical Flight Loads. 44th AIAA Structures, Structural Dynamics, and Materials Conference, 2003.

[4] Philip S Beran, Dean Bryson, Andrew S Thelen, Matteo Diez, and Andrea Serani. Comparison of multi-fidelity approaches for military vehicle design. In *AIAA AVIATION 2020 FORUM*, page 3158, 2020.

[5] M. P. Rumpfkeil, M. Lickenbrock P. Beran, and R. Kolonay. Multi-fidelity, Aeroelastic Analysis and Optimization with Control Surface Deflections of an Efficient Supersonic Air Vehicle. AIAA Paper, 2021-0732, 2021.

[6] M. P. Rumpfkeil and P. Beran. Aeroelastic Analysis and Optimization using FUNtoFEM of an Efficient Supersonic Air Vehicle. AIAA Paper, 2022-2094, 2022.

[7] Robert M Lewis. A trust region framework for managing approximation models in engineering optimization. In *AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, pages 1053–1055, 1996.

[8] N M Alexandrov, J E Jr Dennis, R M Lewis, and V Torczont. A trust-region framework for managing the use of approxima- tion models in optimization. *Structural Optimization*, 15:16–23, 1998.

[9] D.E. Bryson and M.P. Rumpfkeil. Aero-Structural Design Optimization using a Multi-fidelity Quasi-Newton Method. *AIAA Journal of Aircraft*, 56(5):2019–2031, 2019.

[10] J. F. Rodriguez, J. E. Renaud, and L. T. Watsen. Convergence of Trust Region Augmented Lagrangian Methods Using Variable Fidelity Approximation Data. *Structural Optimization*, 14(3-4):141–156, 1998.

[11] N. M. Alexandrov, R. M. Lewis, C. R. Gumbert, L. L. Green, and P. A. Newman. Optimization with Variable-Fidelity Models Applied to Wing Design. AIAA-2000-0841, 2000.

[12] N. M. Alexandrov, R. M. Lewis, C. R. Gumbert, L. L. Green, and P. A. Newman. Approximation and Model Management in Aerodynamic Optimization with Variable-Fidelity Models. *Journal of Aircraft*, 38(6):1093–1101, 2001.

[13] A. Forrester, A. Sobester, and A. Keane. *Engineering Design via Surrogate Modelling: A Practical Guide*. John Wiley & Sons, 2008.

[14] G. Xu, K. Yamazaki, and G. D. Cheng. A New Two-Point Approximation Approach for Structural Optimization. *Structural and Multidisciplinary Optimization*, 20(1):22–28, 2000.

[15] J. W. Bandler, Q. S. Cheng, S. A. Dakroury, A. S. Mohamed, Bakr, K. M. H., Madsen, and J. Sondergaard. Space mapping: the state of the art. *IEEE Transactions on Microwave Theory and Techniques*, 52(1):337–361, 2004.

[16] B. Peherstorfer, K. Willcox, and M. Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. Technical report, Massachusetts Institute of Technology Aerospace Computational Design Laboratory, 2016.

[17] M. S. Eldred, A. A. Giunta, S. S. Collis, N. A. Alexandrov, and R.M. Lewis. Second-order corrections for surrogate-based optimization with model hierarchies. AIAA Paper, 2004-4457, 2004.

[18] D. E. Bryson and M. P. Rumpfkeil. Comparison of Unified and Sequential-Approximate Approaches to Multifidelity Optimization. 2017.

[19] D. E. Bryson and M. P. Rumpfkeil. A Multifidelity Quasi-Newton Method for Design Optimization. *AIAA Journal*, 56, No. 10:4074–4086, 2018.

[20] D. E. Bryson and M. P. Rumpfkeil. All-at-Once Approach to Multifidelity Polynomial Chaos Expansion Surrogate Modeling. *Aerospace Science and Technology*, 70C:121–136, 2017.

[21] M. P. Rumpfkeil, A. Serani, and P. Beran. Multi-fidelity Constrained Optimization Methods Applied to Benchmark Problems. AIAA Paper, 2024-0385, 2024.

[22] Dieter Kraft. A software package for sequential quadratic programming. Technical report, Oberpfaffenhofen: Institut für Dynamik der Flugsysteme, 1988.

[23] S. A. Burton, E. J. Alyanak, and R. M. Kolonay. Efficient Supersonic Air Vehicle Analysis and Optimization Implementation using SORCER. AIAA Paper, 2012-5520, 2012.

[24] C. M. Meckstroth, E. J. Alyanak, N. J. Lindsley, M. Gabbard, and A. Hunton. Aerodynamic Modeling Techniques for Efficient Supersonic Air Vehicle Multidisciplinary Design Optimization. AIAA Paper, 2014-3254, 2014.

[25] Chris M. Meckstroth and William B. Blake. Control Focused Multidisciplinary Design Optimization of Tailless Fighter Aircraft. AIAA Paper, 2015-2324, 2015.

[26] Chris M. Meckstroth. Parameterized, Multi-fidelity Aircraft Geometry and Analysis for MDAO Studies using CAPS. AIAA Paper, 2019-2230, 2019.

[27] M. Lickenbrock, M. P. Rumpfkeil, P. Beran, and R. Kolonay. Multi-fidelity, Multidisciplinary Design Analysis of an Efficient Supersonic Air Vehicle. AIAA Paper, 2020-2223, 2020.

[28] E. Alyanak, R. Durscher, R. Haimes, J. F. I. Dannenhoffer, N. Bhagat, and D. Allison. Multi-fidelity, Gemoetry-centric Multi-disciplinary Analysis for Design. AIAA Paper, 2016-4007, 2016.

[29] E.H. Johnson and V.B. Venkayya. *Automated Structural Optimization System (ASTROS)*, volume Volume I - Theoretical Manual. 1988.

[30] F. A. Woodward. An Improved Method for the Aerodynamic Analysis of Wing-Body-Tail Configurations in Subsonic and Supersonic Flow: Part I - Theory and Application. Technical Report CR-2228, NASA, 1973.

[31] R. L. Harder and R. N. Desmarais. Interpolation using Surface Splines. *Journal of Aircraft*, 9 No. 2:189–191, 1972.

[32] R.T. Biedron, J.-R. Carlson, J.M. Derlaga, P.A. Gnoffo, D.P. Hammond, W.T. Jones, B. Kleb, E.M. Lee-Rausch, E.J. Nielsen, M.A. Park, C.L. Rumsey, J.L. Thomas, and W.A. Wood. *FUN3D Manual: 13.2*. 2017.

[33] D. L. Marcum and N. P. Weatherill. Unstructured Grid Generation Using Iterative Point Insertion and Local Reconnection. *AIAA Journal*, 33 No. 9:1619–1625, 1995.

[34] D. L. Marcum. *Unstructured Grid Generation Using Automatic Point Insertion and Local Reconnection*. The Handbook of Grid Generation, edited by J.F. Thompson, B. Soni, and N.P. Weatherill, CRC Press, 1998.

[35] R. T. Biedron and J. L. Thomas. Recent Enhancements to the FUN3D Flow Solver for Moving-Mesh Applications. AIAA Paper, 2009-1360, 2009.

[36] J. Kiviaho, K. Jacobson, M. Smith, and G. Kennedy. A Robust and Flexible Coupling Framework for Aeroelastic Analysis and Optimization. AIAA Paper, 2017-4144, 2017.

[37] K. Jacobson, J. Kiviaho, M. Smith, and G. Kennedy. An Aeroelastic Coupling Framework for Time-accurate Analysis and Optimization. AIAA Paper, 2018-0100, 2018.

[38] Kevin E. Jacobson. *ADJOINT-BASED AEROELASTIC OPTIMIZATION WITH HIGH-FIDELITY TIME-ACCURATE ANALYSIS*. PhD thesis, Georgia Institute of Technology, 2019.

[39] M. P. Rumpfkeil and P. Beran. Multi-Fidelity Surrogate Models for Flutter Database Generation. *Computer and Fluids*, 197, 2020.

[40] M. P. Rumpfkeil and P. Beran. Multi-Fidelity Sparse Polynomial Chaos Surrogate Models Applied to Flutter Databases. *AIAA Journal*, 58, No.3:1292–1303, 2020.