Multifidelity Optimization with Transonic Flutter Constraints

Bret K. Stanford[*], Andrew S. Thelen[†], Kevin E. Jacobson[‡]

NASA Langley Research Center, Hampton, VA, 23681

This work considers static and dynamic aeroelastic optimization of a cantilevered plate wing in transonic flow. Low- and high-fidelity aeroelastic predictions are fed into a standard trust region model management scheme in order to efficiently solve this expensive optimization problem with multifidelity methods. Differences in fidelity are entirely driven by different aerodynamic solvers: linear panel methods for the low-fidelity method, and inviscid CFD-based solvers for the high-fidelity response. The optimization problem utilizes shape, sizing, and trim variables to satisfy stress, trim, and flutter constraints, and is solved across a range of subsonic and transonic Mach numbers. The multifidelity method is able to provide a speed-up for all cases considered here, despite sizable inaccuracies in the low-fidelity response for transonic flows.

I. Introduction

Recent literature has shown progress in the ability to handle high-fidelity unsteady aeroelastic physics during design optimization: see Refs. 1–8. In addition to the computation of CFD (computational fluid dynamics)-based flutter and gust mechanisms, the tools in these papers also compute analytical derivatives for gradient-based optimization, the most economical choice for high-fidelity large-scale design.[9] There is a particular interest in the literature on transonic flows: gust and flutter constraints are typically most aggressive in this regime,[10] as the aerodynamics become very sensitive to structural deformations.[11]

Flutter predictions that rely on linearized frequency domain methods[12,13] are more computationally efficient than those in the time domain, but still the computational cost of the linearized flutter analysis (and the flutter derivatives) is relatively high. Multifidelity optimization schemes may be used to decrease the overall computational cost, where the brunt of the aeroelastic optimization is conducted with a low-fidelity tool, with occasional course-correction via a high-fidelity tool. Specifically, the aerodynamic fidelities of the aeroelastic computations are of concern here, where the CFD computations (both for nonlinear steady analysis and linearized unsteady analysis) may be replaced with a low-fidelity panel tool, such as the doublet lattice method (DLM).[14]

The majority of recent multifidelity optimization research relies on global surrogate approximations,[15,16] but these methods are not amenable to optimization problems with large numbers of design variables. For multifidelity optimization schemes that rely entirely on local gradient information (and thus scale to large design spaces), the most popular is the trust region model management (TRMM) idea first developed in Ref. 17. TRMM develops a corrected low-fidelity model that satisfies first-order consistency with a high-fidelity model at an expansion point. The corrected low-fidelity model is optimized over a trust region, and then the resulting optimum is evaluated with the high-fidelity model. Depending on the improvement of the high-fidelity response relative to the previous high-fidelity data point, the optimum is accepted or rejected, and the trust region is expanded or contracted. Reference 18 expands the TRMM scheme to include second-order corrections via a quasi-Newton approximation, and further extensions can be found in Refs. 19–22, where Ref. 22 is particularly concerned with handling different fidelity combinations in multidiscplinary problems.

This work will use the second-order TRMM scheme of Ref. 18 to solve multifidelity optimization problems of an aeroelastic wing with transonic trim, stress, and flutter constraints, under both shape and structural

[*]Research Aerospace Engineer, Aeroelasticity Branch, bret.k.stanford@nasa.gov, AIAA Associate Fellow.
[†]Research Aerospace Engineer, Aeroelasticity Branch, AIAA Member.
[‡]Research Aerospace Engineer, Aeroelasticity Branch, AIAA Member.

sizing design variables. Results in this paper are demonstrated on the AGARD 445.6 wing,[23] a well-studied aeroelastic validation test case, retrofitted here as a design optimization demonstration. Aeroelastic computations and derivatives are computed with the assistance of OpenMDAO[24] and MPhys.[25,26] MPhys provides libraries to help set up and solve coupled aeroelastic problems in OpenMDAO, among other types of multidisciplinary couplings. MPhys also provides the ability to easily swap fidelities within a coupled aeroelastic workflow (i.e., replace a CFD aerodynamic solver for a linear DLM solver), a capability that can be utilized in the development of TRMM algorithms.

II. Aeroelastic Solvers

For a given fidelity, the aeroelastic workflow as assembled by OpenMDAO and MPhys is shown in the extended design structure matrix (XDSM) diagram[27] of Fig. 1. In this diagram, each rectangle represents an OpenMDAO subsystem, which can be a single OpenMDAO component or a group of components associated with a discipline.

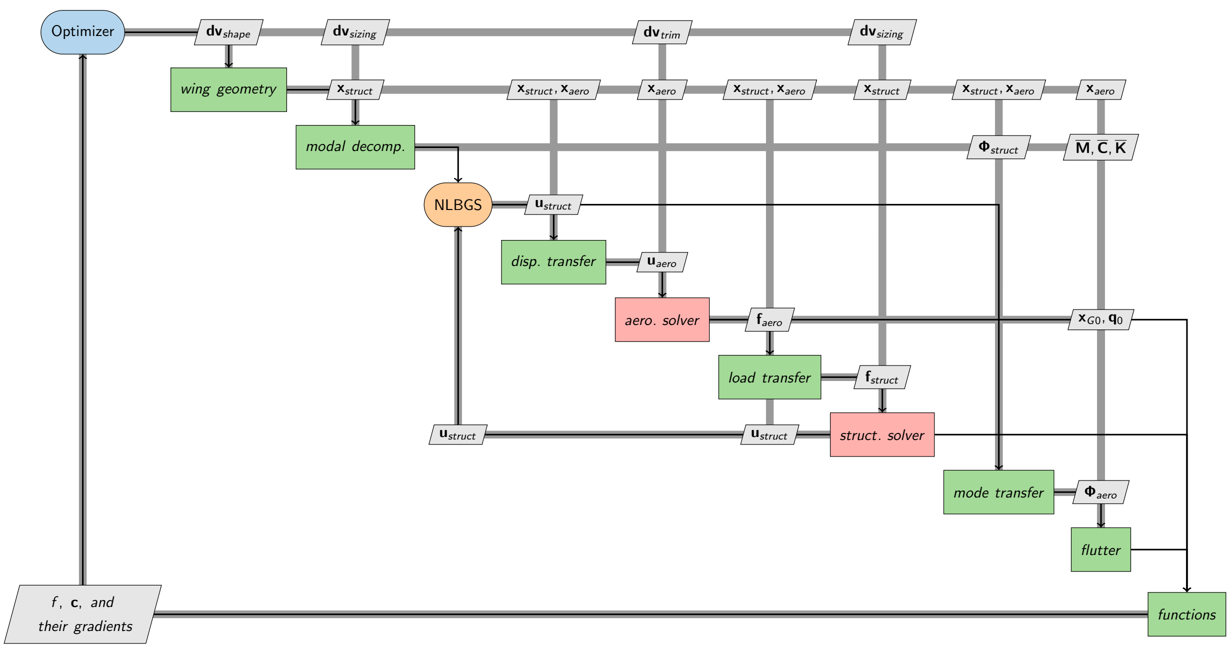

Figure 1. XDSM diagram of a single-discipline aeroelastic optimization process.

1. First, an optimizer dictates shape $\boldsymbol{dv}_{shape}$, sizing $\boldsymbol{dv}_{sizing}$, and trim $\boldsymbol{dv}_{trim}$ (i.e., angle of attack, AoA) design variables.
2. A geometry component uses the shape variables to determine coordinates of the aerodynamic ($\boldsymbol{x_{aero}}$) and structural ($\boldsymbol{x_{struct}}$) nodes on the wing.
3. A structural modal analysis computes the matrix of structural mode shapes $\boldsymbol{\Phi}_{struct}$, and the modally-reduced mass, damping, and stiffness matrices ($\overline{\boldsymbol{M}}$, $\overline{\boldsymbol{C}}$, and $\overline{\boldsymbol{K}}$).
4. A nonlinear block Gauss-Seidel algorithm (NLBGS) is used to compute the static aeroelastic deformations. Structural displacements $\boldsymbol{u}_{struct}$ are mapped onto the aerodynamic surface mesh to form $\boldsymbol{u}_{aero}$. The aerodynamic solver computes the aerodynamic forces $\boldsymbol{f}_{aero}$ (and other metrics of interest, such as aerodynamic lift, drag, etc.), which is then mapped onto the structural mesh to form $\boldsymbol{f}_{struct}$. The structural solver subsequently computes the resulting $\boldsymbol{u}_{struct}$. This NLBGS loop iterates until convergence, which can take 10-30 iterations, depending on the flexibility of the configuration.
5. The structural mode shapes are mapped on to the aerodynamic surface mesh to form $\boldsymbol{\Phi}_{aero}$.
6. The flutter analysis first linearizes about the steady flow (defined by the steady flow solution $\boldsymbol{q}_0$ and the steady volume mesh coordinates $\boldsymbol{x}_{G0}$) to compute oscillatory flows in response to the surface motion

of $\boldsymbol{\Phi}_{aero}$. Generalized aerodynamic forces (GAFs) are then computed, and used (along with $\overline{\boldsymbol{M}}$, $\overline{\boldsymbol{C}}$, and $\overline{\boldsymbol{K}}$) to compute the flutter boundary via a nonlinear eigenvalue p-k solver.

Additional XDSM charts for the steady aerodynamic solver and the flutter solver can be found in Ref. 6. The portion of the diagram in Fig. 1 between the modal decomposition and the flutter analysis is an aeroelastic scenario in MPhys. The coupling group solved by the NLBGS scheme is a coupling subsystem, the modal computation is a precoupling subsystem, and the mode transfer and flutter computations are a postcoupling subsystem.[26]

For the high-fidelity aeroelastic workflow, steady aerodynamics are computed with a stabilized finite element (SFE) solver[28] in the FUN3D software,[29] and unsteady aerodynamics are computed with a linearized frequency domain (LFD) solver written within FUN3D/SFE.[13] Load, displacement, and mode shape mapping is done with Matching-based Extrapolation for Loads and Displacements (MELD).[30] Finally, a linear shell finite element analysis (FEA) is used for both the modal decomposition and the structural portion of the NLBGS. For the low-fidelity aeroelastic workflow, the steady and unsteady FUN3D-based CFD solvers are swapped out for linear DLM[14] solvers, though the other solvers (MELD, FEA, $p-k$) remain in-place. In this case any dependence between the steady NLBGS and the unsteady flutter problem is removed (represented by $\boldsymbol{q}_0$ and $\boldsymbol{x}_{G0}$ for the CFD problem in Fig. 1), because steady and unsteady linear aeroelasticity are entirely separable.

For either high- or low-fidelity aeroelasticity, five functions are pulled from the workflow for use in optimization. First, the structural mass is set as the objective function f. Next, the planform area, the steady aerodynamic lift, the maximum structural von Mises stress seen in the structure during the static aeroelastic deformations (approximated with a Kreisselmeier-Steinhauser (KS) aggregation function[31]), and the flutter dynamic pressure, are all grouped into a vector of constraints, $\boldsymbol{c}$. Furthermore, analytically-computed derivatives of f and $\boldsymbol{c}$ are available with respect to $\boldsymbol{dv}_{shape}$, sizing $\boldsymbol{dv}_{sizing}$, and trim $\boldsymbol{dv}_{trim}$.

III. Multifidelity Optimization

The TRMM[17] scheme used here is relatively well known, and so will only be briefly described here. An XDSM diagram for this TRMM process is shown in Fig. 2.

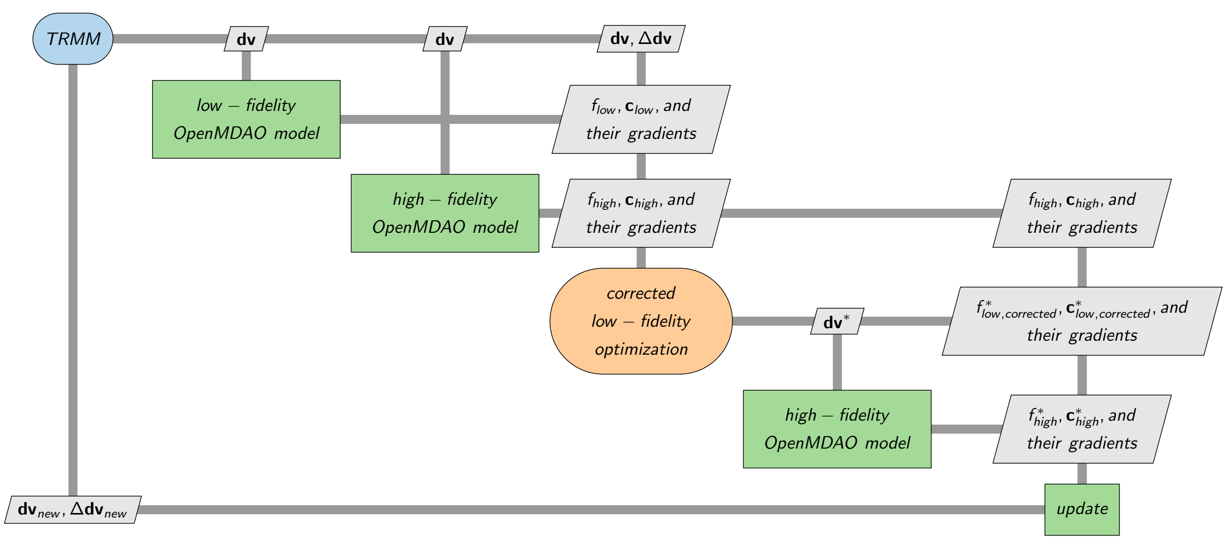

Figure 2. XDSM diagram of the TRMM process.

The overall process is described as:

1. First, both the low- and high-fidelity OpenMDAO models are queried at a location in the design space, $\boldsymbol{dv}$. This produces f_{low}, $\boldsymbol{c}_{low}$, f_{high}, $\boldsymbol{c}_{high}$, their derivatives with respect to $\boldsymbol{dv}$, and approximations to the second-derivative Hessian.

2. This information is provided to a subproblem optimization process that takes place within a prescribed trust region of size $\Delta \boldsymbol{dv}$. An XDSM diagram of this subproblem optimization is shown in Fig. 3. An optimizer obtains a design within the trust region, $\boldsymbol{dv}^*$, that will optimize a corrected low-fidelity model. The corrected low-fidelity model satisfies first- and second-order consistency with the high-fidelity model at the expansion point $\boldsymbol{dv}$. The low-fidelity model is queried at the potential new design point, $\boldsymbol{dv}^*$, to provide f^*_{low} and $\boldsymbol{c}^*_{low}$, and then additive second-order corrections are applied as documented in Ref. 18. The resulting $f^*_{low,corrected}$ and $\boldsymbol{c}^*_{low,corrected}$ are provided back to the optimizer.
3. Once the optimization process of Fig. 3 converges (which happens very quickly, because only the inexpensive low-fidelity model is iteratively queried here), the new candidate design point, $\boldsymbol{dv}^*$, is evaluated with the high-fidelity model.
4. Scalar merit functions are computed by adding the constraint violations to the objective function via a penalty term. Three merit functions are utilized here, one each for the high-fidelity data at $\boldsymbol{dv}^*$, the corrected low-fidelity data at $\boldsymbol{dv}^*$, and the old high-fidelity data at $\boldsymbol{dv}$.
5. Finally, an improvement ratio is computed, which quantifies how much better the new design point $\boldsymbol{dv}^*$ is compared to the old point $\boldsymbol{dv}$, relative to how much better the corrected low-fidelity model thought it would be. Based on the value of this improvement ratio, the new point $\boldsymbol{dv}^*$ can be accepted or rejected, and the trust region $\Delta \boldsymbol{dv}$ can be expanded or contracted. Rules governing these decisions are followed from Ref. 17.

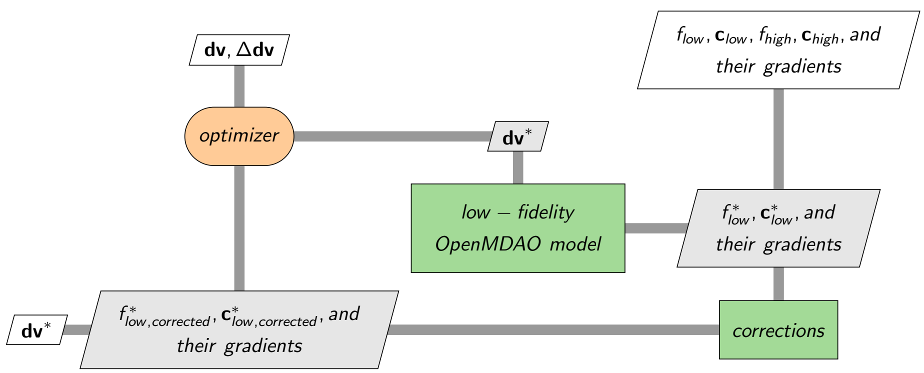

Figure 3. XDSM diagram of the corrected low-fidelity optimization process.

IV. Results

A. Problem Definition

The AGARD demonstration case is shown in Fig. 4. The AGARD configuration has a root chord of 22 inches, a 30 inch span, a 45° quarter-chord sweep angle, and a taper ratio of 0.66. The wing consists of a symmetric NACA65A004 airfoil extruded from root to tip. The plate-like wing structure, clamped along a portion of the root, consists of a weakened orthotropic mahogany material with a nonuniform thickness distribution shown in Fig. 4, ranging between 0.3 and 0.8 inches. An inviscid Euler mesh with 200,000 nodes is used for the CFD aerodynamic modeling, a planar mesh with 128 quadrilaterals is used for the linear DLM aerodynamic modeling, and a planar mesh with 12,800 triangular shell elements is used for the FEA structural modeling. A range of subsonic and transonic Mach numbers are considered here, but the static aeroelastic scenario is always run at a dynamic pressure of 0.42 psi.

For optimization, there are three types of design variables:

1. Shape design: there are 4 planform shape design variables (also shown in Fig. 4), attached to the root chord, the tip chord, the tip sweep (defined here as the streamwise distance between the root leading edge and the tip leading edge), and the span. Each additive design variable ranges from -5 to 5, with a baseline value of 0: the wingspan can therefore range between 25 to 35 in, for example.

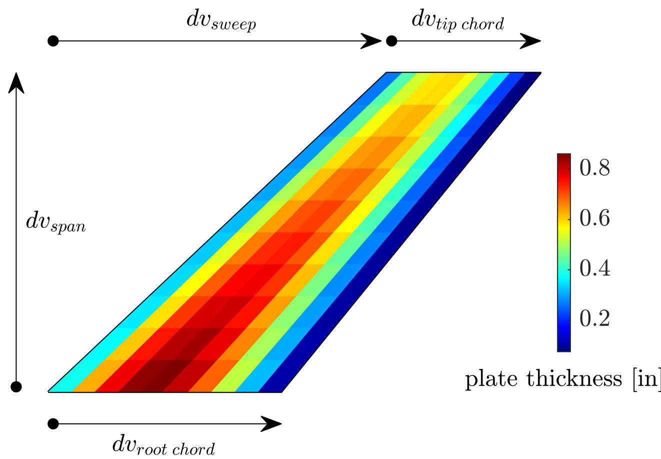

Figure 4. Shape and sizing design variables for the AGARD wing.

2. Sizing design: for the initial results in this paper, there is a single structural sizing variable that uniformly scales the plate thickness values in Fig. 4. The baseline value for this sizing variable is 1, and can range from 0.2 to 1.8 during optimization. Later results in this paper apply a separate thickness scaling design variable to each of the patches shown in Fig. 4, for a total of 100 thickness design variables.
3. Trim design: the angle of attack used during the static aeroelastic NLBGS routine (and the subsequent flutter linearization) is a design variable, with a baseline value of 3°, and side bounds of 0° and 5°.

If a single scalar thickness design variable is used, there are 6 total design variables. If the patchwise thickness variables are used, there are a 105 total design variables.

The objective function f is to minimize structural mass, and the design constraints $\boldsymbol{c}$ are:

1. The planform area must remain equal to the baseline value of 546 in^2.
2. The steady aerodynamic lift coefficient (C_L) must remain equal to 0.1.
3. The stress based KS function must be less than 1 (i.e., all of the plate stresses must be less than the von Mises stress).
4. The flutter dynamic pressure (q_f) must be greater than 1 psf.

B. Baseline Design Behavior

First, comparisons between low- and high-fidelity predictions (and sensitivities) are presented for the baseline wing configuration. These comparisons can help place multifidelity optimization results (presented in later sections) in context, as the TRMM scheme may be expected to struggle for situations where there are large differences in the low- and high-fidelity design spaces. Conversely, situations where the TRMM algorithm excels, but where low- and high-fidelity design spaces are very similar, should not necessarily be taken as a referendum on the efficacy of TRMM.

Steady CFD-based pressures are shown in Fig. 5 for Mach numbers of 0.5, 0.9, 0.95 and 0.98, using the baseline design variables values listed above. The wings in this figure are deflected under static aeroelastic loads, but the deflections are generally small (1-2 inches of upward deflection at the wingtip), and not clearly obvious in the figure. A weak shock appears at Mach 0.9, then strengthens and shifts rearward for higher Mach numbers. The integrated C_L is shown in Fig. 6, for comparison with values computed from the low-fidelity DLM-based model. C_L derivatives with respect to AoA and the wing sweep variable are also shown in the figure.

Differences between low- and high-fidelity C_L metrics in Fig. 6 are moderate, but grow with increased transonic Mach numbers, as expected. The DLM-based C_L metrics monotonically increase with Mach, whereas the CFD-based metrics peak near Mach 0.95, and then attenuate above that value. This is due to the growth and movement of the transonic shock[11] in Fig. 5, which the DLM is unable to capture. Similar results are shown in Fig. 7, for the stress-based KS function computed from the static aeroelastic deformation, and its derivatives with respect to AoA and sweep. Here the agreement between low- and high-fidelity is much better at subsonic Mach numbers, but again shows the expected differences in the transonic regime.

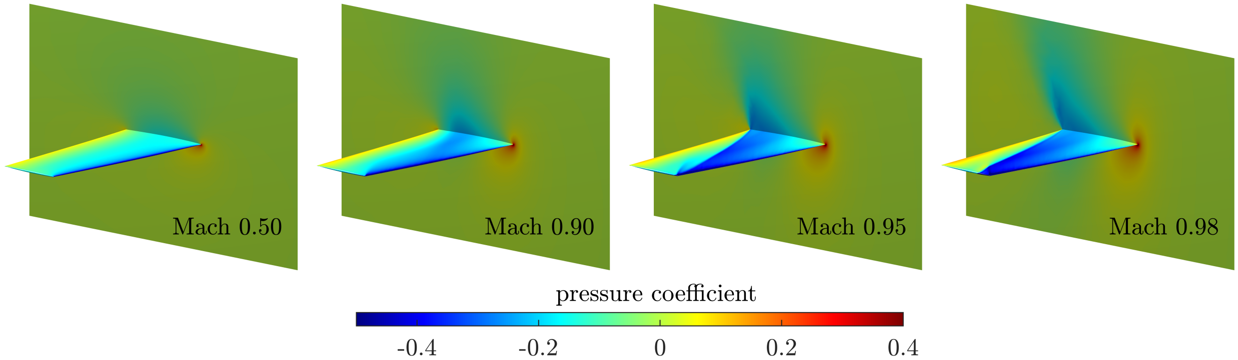

Figure 5. Steady CFD-based pressure distributions.

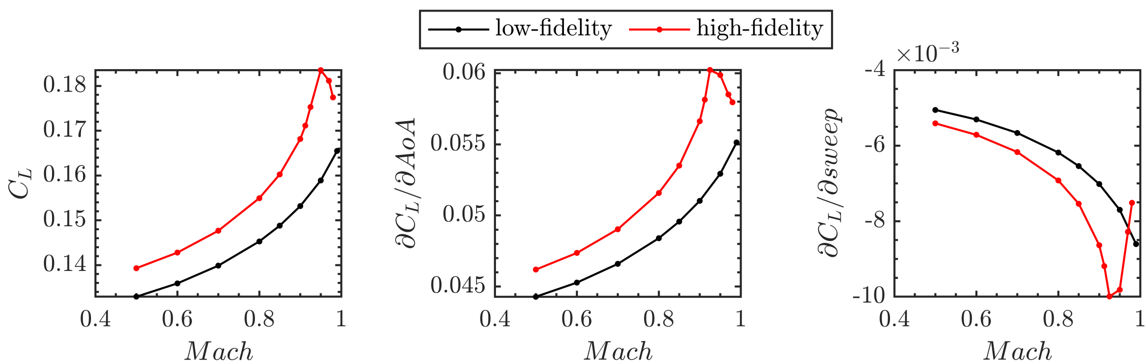

Figure 6. Dependence of lift and lift derivatives on Mach number and analysis fidelity.

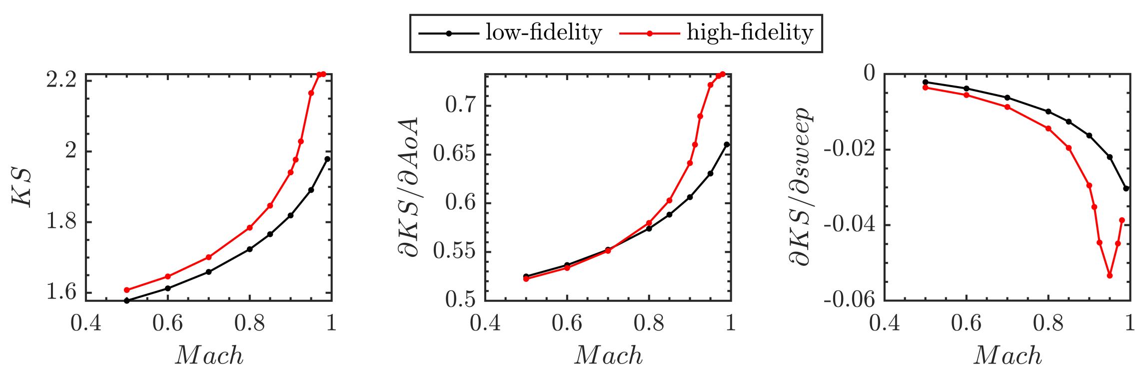

Figure 7. Dependence of stress and stress derivatives on Mach number and analysis fidelity.

Next, dynamic aeroelastic predictions are considered. Figures 8 and 9 show the real and imaginary parts of the CFD-based flutter mode, respectively, at Mach values of 0.5, 0.9, 0.95 and 0.98, again using the baseline design variables values listed above. The flutter q_f value itself, and its derivatives with respect to AoA and sweep, are shown in Fig 10. Both DLM- and CFD-based flutter q_f drop with increased Mach number due to compressibility effects, but a more severe drop is seen in the CFD-based flutter q_f in the transonic regime (the flutter dip), again increasing the gap between low- and high-fidelity predictions. This dip is likely driven by the transonic amplification of the C_L-AoA derivative[11] in Fig. 6 (though change in the pitching moment with shock movement is also a likely contributor[32]). Both lift and pitching moment are computed from the oscillatory pressures of Figs. 8 and 9, and these quantities are largest across the transonic shocks. Near the bottom of the transonic flutter dip, the flutter mechanism approaches a single-degree-of-freedom mechanism,[33] and this is reflected in the flutter mode shapes of Figs. 8 and 9. These mode shapes have large real and imaginary deformation components for subsonic flows, but the oscillatory deformations are largely real-valued for transonic flows.

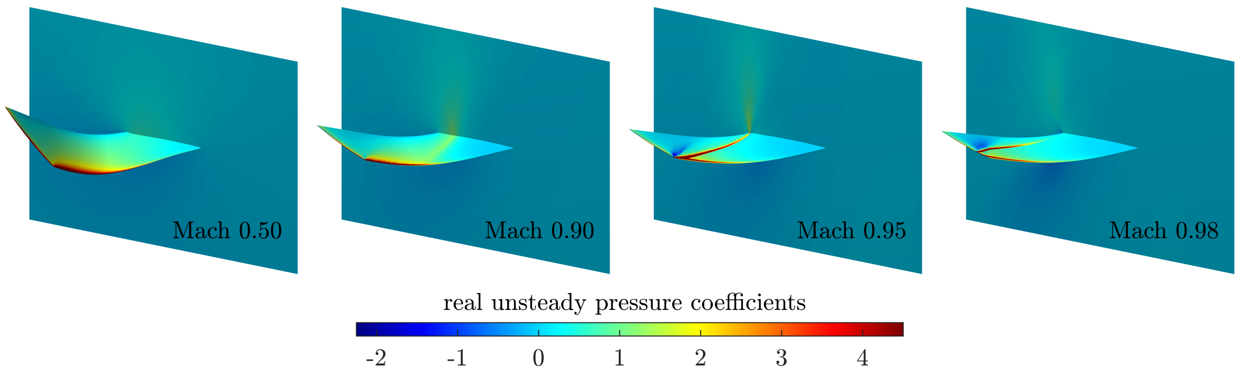

Figure 8. Oscillatory CFD-based pressure distributions (real part).

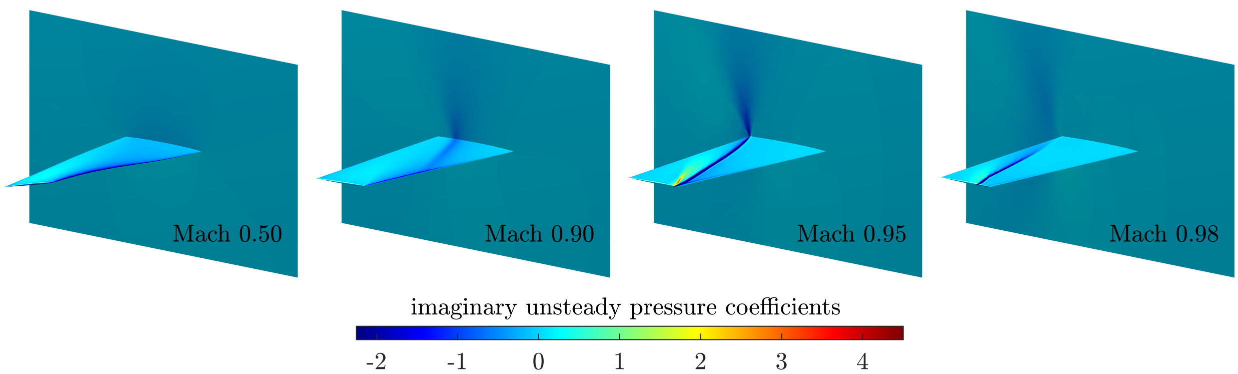

Figure 9. Oscillatory CFD-based pressure distributions (imaginary part).

The derivatives of flutter q_f shown in Fig. 10 display larger low- and high-fidelity differences than seen previously. DLM-based flutter q_f is independent of static AoA (due to the separation of steady and unsteady physics seen in linear aeroelasticity predictions), and so this derivative is zero. The CFD-based AoA derivative is non-zero, owing to the fact that the LFD solver is linearized about the static aeroelastic deformation in Fig. 1. Changing the steady AoA will change the static aeroelastic deformation, shifting the shock locations, which finally impacts the flutter mechanism. This AoA derivative shows multiple rapid sign changes for transonic Mach values. The DLM-based tool *is* able to predict the derivative of q_f with respect to sweep, and this value compares well with the CFD-based value for subsonic flows. Transonically, again, the fidelity match is poor, with DLM mispredicting the sign of the sweep derivative for Mach values between 0.85 and 0.95. For both AoA and sweep derivatives, the inability of the low-fidelity solver to predict the correct derivative sign may signify a particularly challenging situation for TRMM.

C. Optimization Results: Small Design Space

Design optimization is first shown for cases where only a single sizing design variable is utilized. There are then a total of 6 design variables: 4 shape variables, 1 sizing variable, and 1 trim variable. Optimization convergence is shown in Figs. 11 and 12 for the objective function and the maximum constraint violation, respectively. Results are shown for both the single-fidelity solver (i.e., high-fidelity CFD-based solver) and the multifidelity TRMM scheme, for Mach values of 0.5, 0.9, 0.95 and 0.98. Convergence is specifically

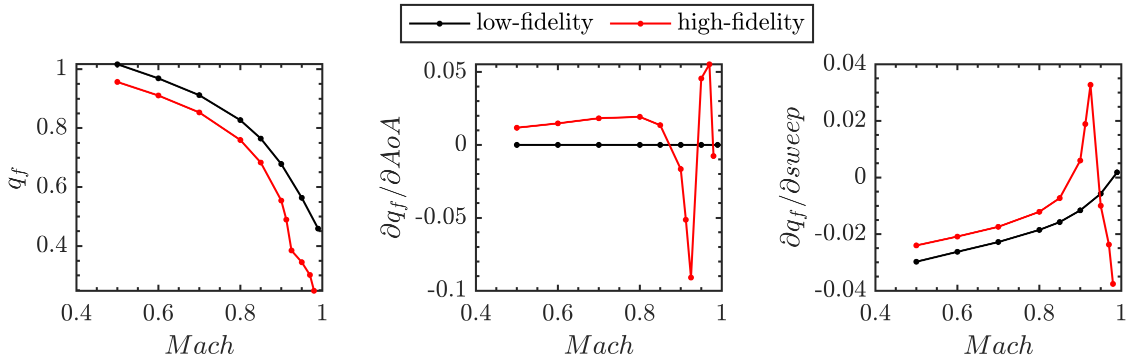

Figure 10. Dependence of flutter and flutter derivatives on Mach number and analysis fidelity.

plotted as a function of the number of high-fidelity samples, where each sample requires both a function call and a derivative call. On 400 Intel Skylake CPU cores, a single function call takes roughly 95 minutes and the derivative call takes 155 minutes. A complete low-fidelity (or corrected low-fidelity) optimization can typically be completed in less than 5 minutes.

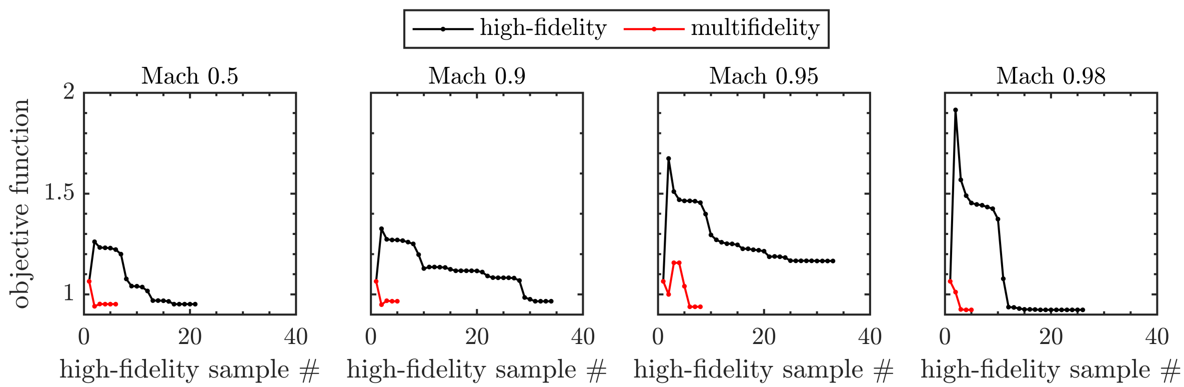

Figure 11. Objective function convergence across four Mach numbers.

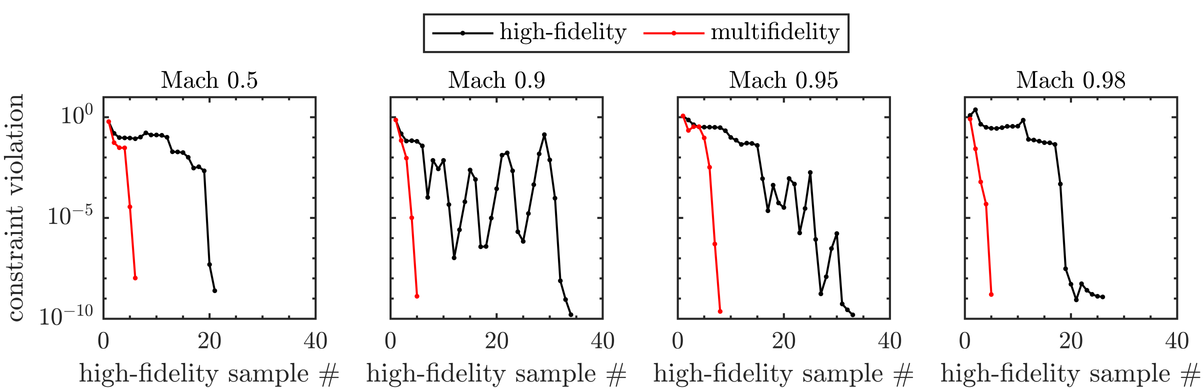

Figure 12. Constraint violation convergence across four Mach numbers.

All three aeroelastic constraints of the initial baseline design become more violated with increased Mach number. The C_L of Fig. 6 increases above the constraint boundary of 0.1, the KS of Fig. 7 increases above the constraint boundary of 1.0, and the flutter q_f of Fig. 10 decreases below the constraint boundary of 1.0 psf (via the flutter dip). The baseline constraint violation therefore increases with Mach, and the high-fidelity optimizer reacts by increasing the mass objective function at design cycle two. This however is not a predictor for how long the high-fidelity optimization takes to converge, with Mach 0.9 converging more slowly than 0.95 and 0.98. The high-fidelity convergence at Mach 0.9 is erratic, indicative of a complex design space with strong curvature, though alternative scaling of the optimization problem may alleviate this. At convergence, for each Mach number in Fig. 11, every constraint is active.

The TRMM performs well for each case in Figs. 11 and 12; a hypothesis was stated above that TRMM would struggle for increased transonic flows (given the fact that low-fidelity solvers mis-predict the magnitude and sign of certain design derivatives, specifically in Fig. 10), and this proves not to be true in the results. Taking an arbitrary convergence tolerance as when the maximum constraint violation drops below (and stays below) 10^{-8}, the TRMM provides a speed-up (again in terms of the number of high-fidelity samples required) of 3.5 (Mach 0.5), 6.6 (Mach 0.9), 3.8 (Mach 0.95), and 4.4 (Mach 0.98) relative to the high-fidelity optimization. These TRMM results all use an initial trust region size ($\Delta \boldsymbol{dv}$) of 1: i.e., the TRMM uses the entire design space for the first subproblem optimization. The penalty parameter to compute the scalar merit function is set to a default value of 100 for all results also, a value that was found to work well for simple analytical optimization example problems.

At Mach 0.95 in Figs. 11 and 12, it can be seen that the multifidelity optimization converges to a different (better) design than the high-fidelity optimization. This result is indicative of multiple local minima in the design space. TRMM convergence, as a function of the initial trust region size, is shown in Fig. 13, where the high-fidelity result and the TRMM result with an initial $\Delta \boldsymbol{dv}$ equal to 1, are repeated from Figs. 11 and 12. Decreasing the initial trust region size to 0.5 converges the TRMM to the same inferior local optimum as the high-fidelity result. Further decreasing the trust region to 0.1 reverts the TRMM convergence to the superior optimum, and also slows down the TRMM convergence considerably, as would be expected. TRMM convergence to various local optimal may also be triggered by changing the penalty parameter, but this is not explored here. Changing the starting design of the high-fidelity optimization will also obviously force a convergence to different local optima, but this is not explored here either.

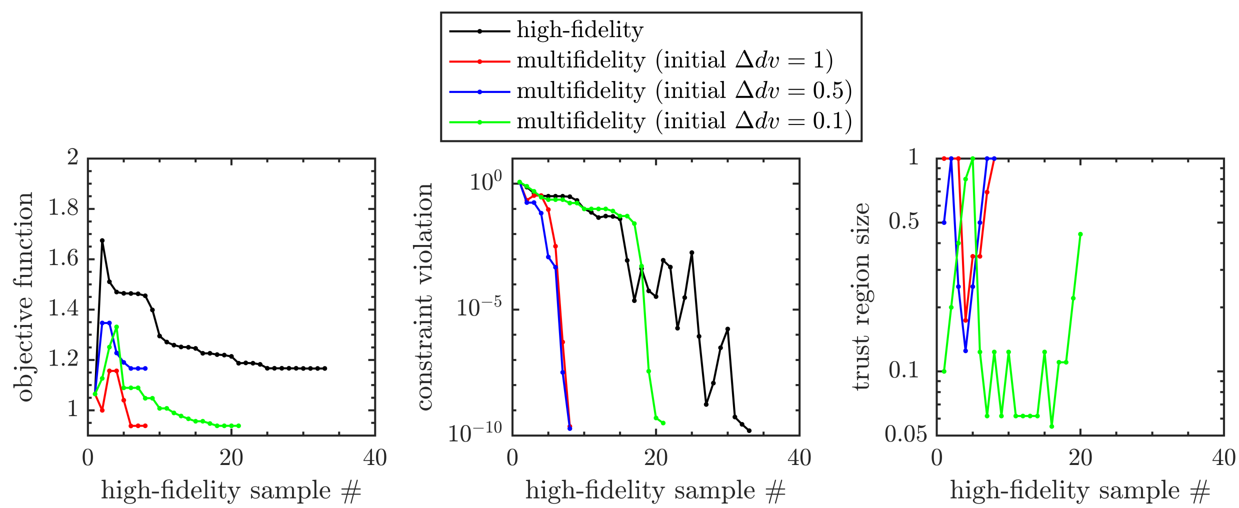

Figure 13. Dependence of the TRMM convergence upon the baseline trust region size, at Mach 0.95.

Figures 6, 7, and 10 had shown substantial differences in the transonic aeroelastic response between high- and low-fidelity predictions, but despite this, the TRMM solver performs relatively well in Figs. 11 and 12 at these Mach numbers. It may be that the aeroelastic response for the baseline design is quite different between the fidelities, but the optimizer pushes the design into a direction where the fidelities *do* match well, and this is the reason for the favorable TRMM performance (as opposed to the ability of the TRMM scheme to ably handle a large mismatch in low- and high-fidelity prediction).

This idea is further explored, and largely disproven, in Fig. 14, which shows the optimal solution as a function of Mach number for both the low- and high-fidelity solvers. Optimal objective function and design variables are shown in this figure, but not constraint values, as every constraint is active for every optimal design. It is understood that the multifidelity optimization will return the same result as the high-fidelity optima, except for the local minima observed in Fig. 13. These local minima are better quantified in Fig. 14, where the DLM-based low-fidelity solver also sees a local minimum above Mach 0.91. This local high-fidelity optimum is only plotted at Mach 0.95; it is likely that the local optima also exists at Mach 0.98, and simply was not found in any of the optimizations in Figs. 11 and 12. The superior designs in Fig. 14 are labeled as "global": attempts to find any superior optima via different optimization starting points (or, below Mach 0.91, any other optima at all) have failed.

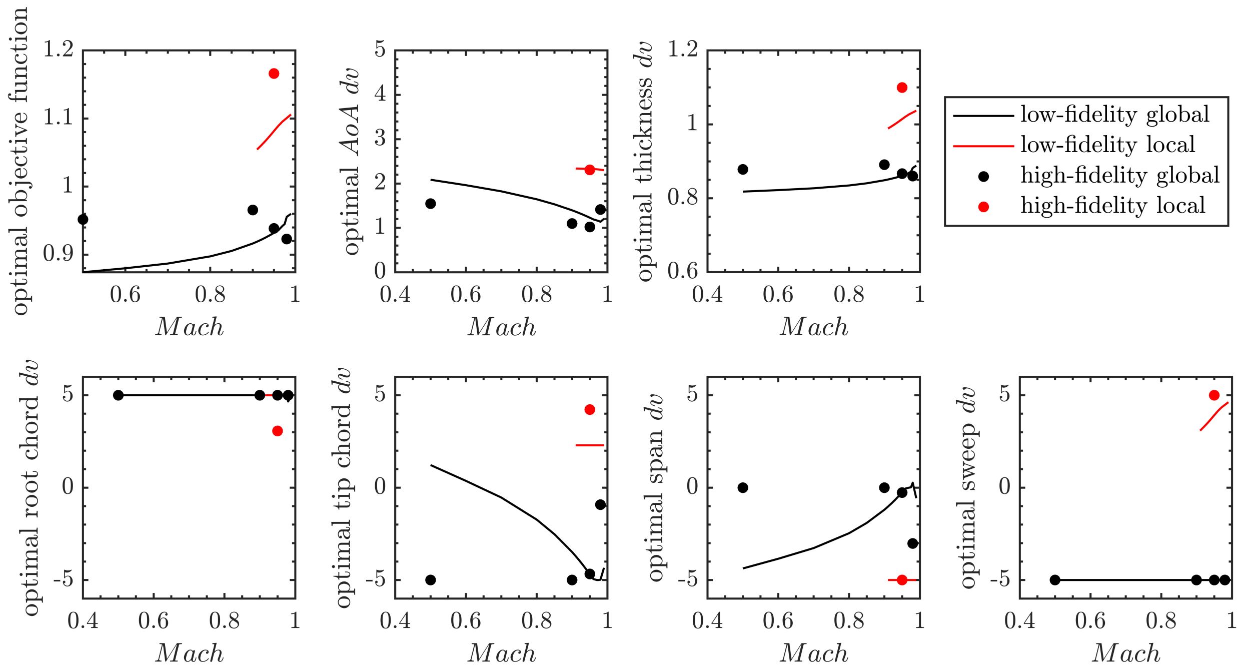

Figure 14. Low- and high-fidelity optima across different Mach numbers.

The low- and high-fidelity optima do show similarities for transonic flows, in terms of the trim AoA, the plate thickness, the root chord, and the wing sweep. The wing sweep in particular is pushed to the lower bound for both low- and high-fidelity optimization, despite the differing flutter sweep derivative signs shown in Fig. 10. The optimal tip chord and wing span are fairly different at Mach 0.98, however. Interestingly, these two design variables are very dissimilar at the subsonic Mach 0.5 also, where low- and high-fidelity design similarity would be expected. As already noted, no local optima have been found below Mach 0.91, and so this result would indicate that the optimal design is very sensitive to constraint boundary locations, even for benign aeroelastic physics.

D. Optimization Results: Large Design Space

Next, design optimization is conducted for cases where 100 sizing design variables are utilized, one for each of the thickness patches shown in Fig. 4. There are then a total of 105 design variables: 4 shape variables, 100 sizing variables, and 1 trim variable. Optimization results are shown in Fig. 15 for a single Mach number of 0.95. As expected, both TRMM and the high-fidelity optimization take much longer to converge than results with a single thickness design variable in the previous section, but TRMM still provides a sizable speed-up relative to the single-fidelity result.

The high-fidelity optimal design is shown in Fig. 16, for comparison with the baseline design, and the high-fidelity design found in the previous section with only a single thickness design variable. The larger design space allows the optimizer to decrease the objective function (structural mass) from 0.938 to 0.773, by decreasing the span and the tip chord. Structural thicknesses are redistributed to target high stiffness near the root, where the stresses are highest. As above, all of the aeroelastic constraints are active for this optimal design. The static aeroelastic flow conditions for these same three designs are shown in Fig. 17, where the baseline result is repeated from Fig. 5. For the optimal designs, the shock at the root is shifted slightly upstream, and the outboard shock is shifted much closer to the trailing edge, as compared to the baseline case.

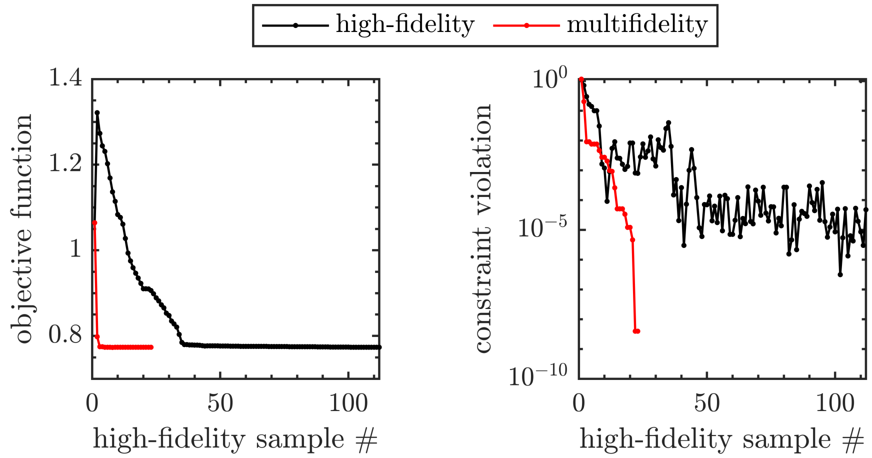

Figure 15. TRMM performance at Mach 0.95, for the case with 105 design variables.

The TRMM results in the previous section at Mach 0.95 returned different designs compared to the high-fidelity result, owing to the existence of local minima. This is not the case in Fig. 15, as both solvers converge to nearly the same design. Using DLM-based low-fidelity simulations (not shown here), that inferior local minima *is* still present in this larger design space, and presumably exists for the high-fidelity model also. The high-fidelity optimization in Fig. 15 does not converge to that inferior local minima, for reasons that are unclear.

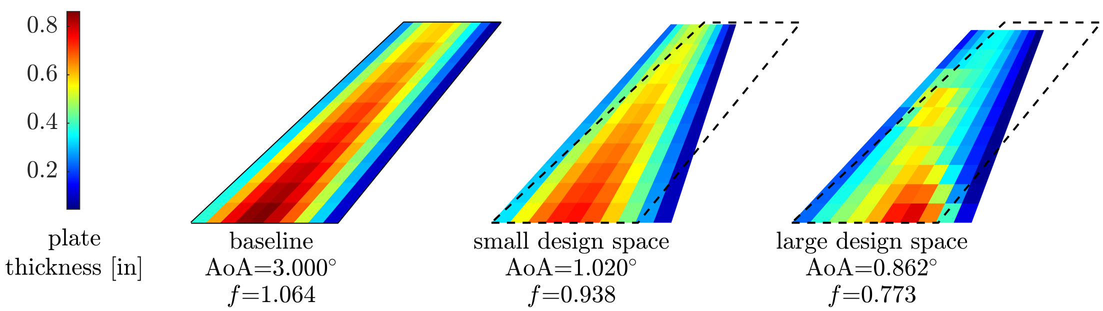

Figure 16. Baseline and optimal wing designs computed with the high-fidelity model at Mach 0.95.

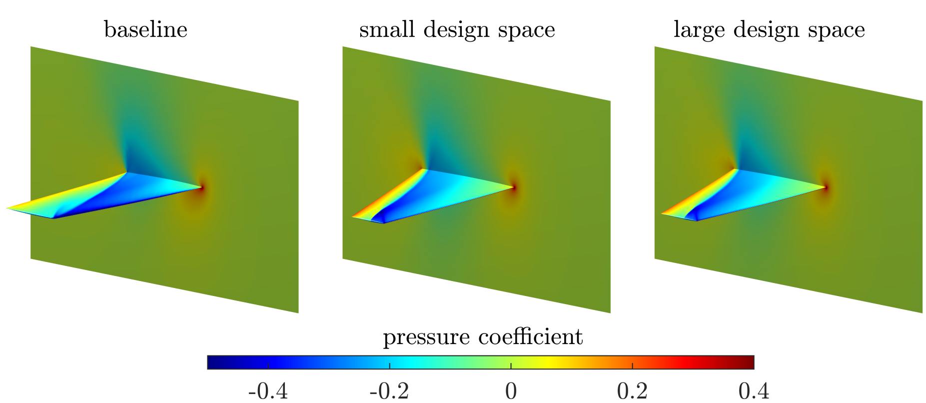

Figure 17. Steady CFD-based pressure distributions of the baseline and optimal wing designs at Mach 0.95.

V. Conclusions

This paper has demonstrated the well-known trust region model management scheme for multifidelity optimization of a transonic aeroelastic wing. OpenMDAO-based static and dynamic aeroelastic solvers are used to create high- and low-fidelity computations (and analytical design derivatives), where high-fidelity tools use inviscid computational fluid dynamics for the aerodynamics, and low-fidelity tools use linear panel methods. The aeroelastic scenario for these cases is a static trim condition at a nonzero angle of attack, and then a flutter condition that linearizes about the static result. The optimization problem is to minimize structural mass of a cantilever plate wing under trim, stress, and flutter constraints, with planform shape, structural sizing, and trim design variables.

Optimization results are tracked as a function of increasing Mach number: for transonic flows, shocks develop over the wing, increasing the aerodynamic sensitivity of the wing. This in turn leads to large differences between the low- and high-fidelity aeroelastic responses, as well as aeroelastic derivatives with respect to design parameters. Despite these differences, the multifidelity tool has been shown to be effective at increasing optimization convergence rates relative to a high-fidelity optimization, for all cases considered here: subsonic and transonic conditions, relatively small design spaces (6 design variables), and relatively large design spaces (105 design variables).

A key theme of this work is to understand the degree to which the multifidelity algorithm is being stressed, relative to its success at optimization speed-up. If the multifidelity tool shows benefits compared to a single-fidelity optimization, but the low-fidelity and the high-fidelity information being fed to the multifidelity algorithm are very similar to each other, then this is not a meaningful result. Conversely, if multifidelity convergence is strong despite inaccurate low-fidelity information, then this is a meaningful result. For the design cases considered here, this conclusion is difficult to decisively demonstrate, even for the small design space, given the high sensitivity of the optimal result to constraint boundary locations (i.e., subsonic low- and high-fidelity optimizations converge to different designs, despite very-similar baseline responses), and also the existence of multiple local optima for transonic flows. However, at the highest transonic Mach number considered here (0.98) the baseline low- and high-fidelity design sensitivities are very different, and the low- and high-fidelity optimizations converge to fairly different designs also. These factors do suggest that the multifidelity tool can efficiently solve the CFD-based aeroelastic trim and flutter optimization problems considered here, despite relying on inaccurate and inexpensive low-fidelity information.

Acknowledgments

This work is supported by the Transformational Tools and Technologies (TTT) project of the NASA Transformative Aeronautics Concepts Program (TACP). Computational resources for this work are provided by the NASA Langley K cluster.

References

[1]Mani, K., Mavriplis, D., "Adjoint-Based Sensitivity Formulation for Fully Coupled Unsteady Aeroelasticity Problems," *AIAA Journal*, Vol. 47, No. 8, pp. 1902–1915, 2009.

[2]Zhang, Z., Chen, P., Yang, S., Wang, Z., Wang, Q., "CFD-based Aeroelastic Adjoint Sensitivities for Flight Vehicle Weight Minimization using Sizing and/or Shape Design Variables," *AIAA Structures, Structural Dynamics, and Materials Conference*, Grapevine, TX, January 2017.

[3]Jacobson, K., Kiviaho, J., Kennedy, G., Smith, M., "Evaluation of Time-Domain Damping Identification Methods for Flutter-Constrained Optimization," *Journal of Fluids and Structures*, Vol. 87, pp. 174-188, 2019.

[4]He, S., Jonsson, E., Mader, C., Martins, J., "Aerodynamic Shape Optimization with Time Spectral Flutter Adjoint," *AIAA SciTech Forum*, San Diego, CA, January 2019.

[5]Thomas, J., Dowell, E., "Discrete Adjoint Approach for Nonlinear Unsteady Aeroelastic Design Optimization," *AIAA Journal*, Vol. 57, No. 10, pp. 4368-4376, 2019.

[6]Jacobson, K., Stanford, B., "Flutter-Constrained Optimization with the Linearized Frequency-Domain Approach" *AIAA Scitech Forum*, San Diego, CA, January 3-7, 2022.

[7]Prasad, R., Choi, S., Patil, M., "Aerodynamic Shape Optimization using a Time Spectral Coupled Adjoint for Nonlinear Aeroelastic Problems" *Aerospace Science and Technology*, Vol. 126, No. 107495, 2022.

[8]Thelen, A., Stanford, B., Jacobson, K., "Flutter and Stochastic Gust Constraint Studies of the Undeflected CRM using MPhys" *AIAA Scitech Forum*, San Diego, CA, January 8-12, 2024.

[9]Haftka, R., Gurdal, Z., *Elements of Structural Optimization*, Kluwer Academic Publishers, Dordecht, The Netherlands, 1992.

[10]Jonsson, E., Riso, C., Lupp, C., Cesnik, C., Martins, J., Epureanu, B., "Flutter and Post-Flutter Constraints in Aircraft Design Optimization," *Progress in Aerospace Sciences*, Vol. 109, 100537, 2019.

[11]Bendiksen, O., "Review of Unsteady Transonic Aerodynamics: Theory and Applications," *Progress in Aerospace Sciences*, Vol. 47, No. 2, pp. 135-167, 2011.

[12]Thormann, R., Widhalm, M., "Linear-Frequency-Domain Predictions of Dynamic-Response Data for Viscous Transonic Flows," *AIAA Journal*, Vol. 51, No. 11, pp. 2540-2557, 2013.

[13]Jacobson, K., Stanford, B., Wood, S., Anderson, W., "Frequency-Domain Flutter Analysis and Sensitivities with Stabilized Finite Elements," *AIAA SciTech Forum*, Orlando, FL, January 6-10, 2020.

[14]Blair, M., "A Compilation of the Mathematics Leading to the Doublet Lattice Method," WL-TR-92-3028, August,1992.

[15]Peherstorfer, B., Willcox, K., Gunzburger, M., "Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization," *SIAM Review*, Vol. 60, No. 3, pp. 550-591, 2018.

[16]Fernandez-Godino, G., Park, C., Kim, N., Haftka, R., "Issues in Deciding Whether to Use Multifidelity Surrogates," *AIAA Journal*, Vol. 57, No. 5, pp. 2039-2054, 2019.

[17]Alexandrov, N., Lewis, R., Gumbert, C., Green, L., Newman, P., "Approximation and Model Management in Aerodynamic Optimization with Variable-Fidelity Models," *Journal if Aircraft*, Vol. 38, No. 6, pp. 1093-1101, 2019.

[18]Eldred, M., Giunta, A., Collis, S., "Second-Order Corrections for Surrogate-Based Optimization with Model Hierarchies," *AIAA Multidisciplinary Analysis and Optimization Conference*, Albany, NY, August 30 - September 1, 2004.

[19]Gratton, S., Sartenaer, A., Toint, P., "Recursive Trust-Region Methods for Multiscale Nonlinear Optimization," *SIAM Journal of Optimization*, Vol. 19, No. 1, pp. 414-444, 2008.

[20]Olivanti, R., Gallard, F., Brezillon, J., Gourdain, N., "Comparison of Generic Multi-Fidelity Approaches for Bound-Constrained Nonlinear Optimization Applied to Adjoint-Based CFD Applications," *AIAA Aviation Forum*, Dallas, TX, June 17-21, 2019.

[21]Bryson, D., Rumpfkeil, M., "Aerostructural Design Optimization using a Multifidelity Quasi-Newton Method," *Journal of Aircraft*, Vol. 56, No. 5, pp. 2840-2855, 2019.

[22]Wu, N., Mader, C., Martins, J., "A Gradient-Based Sequential Multifidelity Approach to Multidisciplinary Design Optimization," *Structural and Multidisciplinary Optimization*, Vol. 65, No. 131, 2022.

[23]Yates, C., "AGARD Standard Aeroelastic Configurations for Dynamic Response. Candidate Configuration I - Wing 445.6," NASA TM-100492, August, 1987.

[24]Gray, J., Hwang, J., Martins, J., Moore, K., Naylor, B., "OpenMDAO: An Open-Source Framework for Multidisciplinary Design, Analysis, and Optimization," *Structural and Multidisciplinary Optimization*, Vol. 59, No. 4, pp. 1075-1104, 2019.

[25]Yildirim, A., Jacobson, K., Stanford, B., Gray, J., Mader, C., Martins, J., Kennedy, G., "A Modular Multiphysics Simulation Framework Using OpenMDAO," *AIAA Aviation Forum*, virtual event, June 15-19, 2020.

[26]Yildirim, A., Jacobson, K., Anibal, J., Stanford, B., Gray, J., Mader, C., Martins, J., Kennedy, G., "MPhys: A Modular Multiphysics Simulation Library Using the OpenMDAO Framework," *Structural and Multidisciplinary Optimization*, in review.

[27]Lambe, A., Martins, J., "Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes" *Structural and Multidisciplinary Optimization*, Vol. 46, pp. 273-284, 2012.

[28]Anderson, W., Newman, J., Karman, S., "Stabilized Finite Elements in FUN3D," *Journal of Aircraft*, Vol. 55, No. 2, pp. 696-714, 2017.

[29]Biedron, R., Carlson, J., Derlaga, J., Gnoffo, P., Hammond, D., Jacobson, K., Jones, W., Kleb, B., Lee-Rausch, E., Nielsen, E., Park, M., Rumsey, C., Thomas, J., Thompson, K., Walden, A., Wang, L., Wood, W., "FUN3D Manual: 13.7," NASA TM-2020-5010139.

[30]Kiviaho, J., Kennedy, G., "Efficient and Robust Load and Displacement Transfer Scheme using Weighted Least Squares" *AIAA Journal*, Vol. 57, No. 5., pp. 2238-2243, 2019.

[31]Kreisselmeier, G., Steinhauser, R., "Systematic Control Design by Optimizing a Vector Performance Index", *International Federation of Active Controls Symposium on Computer-Aided Design of Control Systems*, Zurich, Switzerland, 1979.

[32]Isogai, K., "On the Transonic-Dip Mechanism of Flutter of a Sweptback Wing," *AIAA Journal*, Vol. 17, No. 7, pp. 793-795, 1979, https://doi.org/10.2514/3.61226.

[33]Stanford, B., Jacobson, K., "Transonic Flutter Dips of the AGARD 445.6 Wing," *Journal of Aircraft*, Vol. 61, No. 2, pp. 515-522, 1979.