

Leveraging a Doublet Lattice Correction Method to Expedite Linearized Frequency-Domain Optimizations in MPhys

Andrew S. Thelen,^{*} Bret K. Stanford,[†] Kevin E. Jacobson,[‡]
NASA Langley Research Center, Hampton, VA, 23681

The Linearized Frequency-Domain method was recently implemented in FUN3D’s Stabilized Finite Element solver to allow for gradient-based optimization subject to aeroelastic flutter and gust constraints. This method, which linearizes about a nonlinear flow solution (thereby being more accurate than panel-based aerodynamics), makes computing these constraints much more tractable than equivalent responses with the full nonlinear time-domain solution. Computing these responses and their analytic derivatives, however, can still be quite costly depending on the number of structural modes and frequencies of interest. Correction techniques for Doublet Lattice Method aerodynamics have long been used in an effort to make panel-based aerodynamics more accurate, while not incurring the cost of a time- or frequency-domain Computational Fluid Dynamics solution. This work presents an implementation of such a correction procedure using the authors’ toolset developed at NASA Langley Research Center, including FUN3D’s Linearized Frequency-Domain solver and an in-house Doublet Lattice Method implementation. The toolset leverages the MPhys library for OpenMDAO, which helps facilitate solver modularity and the computation of analytic design derivatives, among other advantages. Optimizations with flutter and gust (both continuous and discrete) constraints, and with both shape and sizing design variables, are demonstrated for the AGARD 445.6 wing at a transonic flow condition.

I. Introduction

Recent work by the authors has demonstrated gradient-based flutter- [1–4] and gust-constrained [5–7] optimization using either panel aerodynamics or Computational Fluid Dynamics (CFD) solvers. The panel method is an in-house Doublet Lattice Method (DLM) implementation, while the Linearized Frequency Domain (LFD) method implemented in NASA’s FUN3D [8–11] Stabilized Finite Element (SFE) solver [12, 13] serves as the CFD solver. The MPhys library [14, 15] for OpenMDAO [16] is used to modularize the unsteady aeroelastic analyses, which eases the interchangeability between aerodynamic fidelities while using the same structural solver. While DLM-based design evaluations can generally be carried out in seconds or minutes, LFD evaluations may take hours. Both DLM and LFD require flow solutions for each structural mode at a predetermined list of reduced frequencies. Gradient evaluations for LFD-based flutter constraints typically require individual linear solutions for a subset of reduced frequencies (i.e., those involved in the flutter mechanism) and for each mode, in order to compute the cost function’s overall adjoint solution. This assumes the cost function is based on the flutter onset conditions (e.g., flutter dynamic pressure); an aggregated damping cost function would require linear solutions for the full set of reduced frequencies. This is also true for gradients of discrete or continuous gust responses.

DLM correction methods [17–21] have been heavily used for decades, as they allow the computationally cheap DLM solver to be improved in accuracy using data from CFD, experiments, or even flight tests. Corrected DLM remains a workhorse for industry (cf. Jadic et al. [19] or Moreno et al. [22]) because the cost of the higher fidelity methods and physical testing is intractable for a full flight vehicle flutter test matrix. The purpose of this paper is therefore to demonstrate gradient-based optimization with corrected DLM in order to reduce the cost of flutter- and gust-constrained optimizations with CFD. Such an implementation can improve the accuracy of DLM-based optimization while not incurring the full cost of LFD-based optimization. In this work, the accuracy is demonstrated by subsequently performing LFD-based optimization from the resulting optimum; if the correction method is sufficiently accurate, then the corrected DLM optimization will locate a design sufficiently close to the LFD optimum, and the final LFD-based optimization will generally require fewer design iterations. The performance of such a process is compared against starting the LFD-based

^{*}Research Aerospace Engineer, Aeroelasticity Branch, andrew.s.thelen@nasa.gov, AIAA Member.

[†]Research Aerospace Engineer, Aeroelasticity Branch, AIAA Associate Fellow.

[‡]Research Aerospace Engineer, Aeroelasticity Branch, AIAA Member.

optimization from the uncorrected DLM optimum. For comparison, LFD-based optimization from an arbitrary baseline design generally requires much more iterations than either approach, at least for the cases considered.

In the following sections, method and implementation details are first given in Sect. II, including a description of the correction procedure in Sect. II.A followed by a brief summary of flutter and gust analyses in Sect. II.B. Section III then shows sample DLM correction results for the AGARD 445.6 wing [23]. This includes, in Sect. III.A, flutter analysis results using both steady pitch corrections as well as the full list of modes at a subset of frequencies. This is followed by flutter-constrained optimizations in Sect. III.B, where steady pitch corrections are initially used followed by the non-zero frequency corrections with all modes. Section III.C adds discrete and continuous gust constraints to the flutter-constrained optimization problem, demonstrating optimization convergence with and without shape design variables. The paper ends with conclusions and comments on future work in Sect. IV.

II. Methodology

This section provides an overview of the analysis methods used. Section II.A first provides a description of the DLM correction formulation. Selected details of the flutter and gust analyses are provided in Sect. II.B.

A. DLM Correction Procedure

Derived loosely from the method in Moreno et al. [22], a full, diagonally-dominant correction factor matrix, $[\mathbf{C}_F]$, is computed as

$$[\mathbf{C}_F] = [\mathbf{I}] + \left[\left(\{C_{P,dlm}\}^T \right)^\dagger (\{C_{P,ref}\} - \{C_{P,dlm}\}) \right]^T. \quad (1)$$

Here, $(\cdot)^\dagger$ is the Moore-Penrose pseudoinverse operator [24, 25] while $\{C_{P,dlm}\}$ and $\{C_{P,ref}\}$ (both of size $N_{panels} \times N_{ref}$) are the DLM- and CFD-computed oscillatory pressures on the DLM mesh panels. These pressures are computed at a particular reduced frequency and for an arbitrary set of N_{ref} oscillatory modes. The corrected DLM pressures for all N_{modes} oscillatory modes used in the analysis are then computed by pre-multiplying $[\mathbf{C}_F]$ by the DLM-computed Aerodynamic Influence Coefficient (AIC) matrix,

$$\{C_P\} = [\mathbf{C}_F] [\mathbf{AIC}]^{-1} \{w\}, \quad (2)$$

where $\{w\}$ contains the downwash vector for each oscillatory mode included in the analysis. Overall, this correction ensures that the corrected DLM pressure exactly matches the CFD-computed pressure (projected onto DLM panels) for all modes used in the correction data, $\{C_{P,ref}\}$. The resulting structure-on-structure Generalized Aerodynamic Force (GAF) matrix (of size $N_{modes} \times N_{modes}$) is then

$$[\mathbf{Q}_{ss}] = [\mathbf{S}] \{C_P\}, \quad (3)$$

where $[\mathbf{S}]$ is the integration matrix corresponding to the mode shapes. The same operation may be used to compute corrected gust GAFs as well, using the gust-induced oscillatory pressures [5]. Some testing by the authors found that fewer issues are encountered if the structure-on-structure and gust-on-structure correction matrices are computed separately, rather than combining them into the above oscillatory pressures.

Both the DLM and CFD analyses in the present work use the MPhys library for OpenMDAO. The DLM and structural analyses are performed using in-house solvers pyDLM and pyShell. The CFD reference data, $\{C_{P,ref}\}$, is generated using FUN3D's LFD solver, which computes the complex aerodynamic force vector at each CFD surface node. Structural mode shapes, which are the same among DLM and LFD analyses, are transferred to the different aerodynamic meshes using the Matching-based Extrapolation of Loads and Displacements (MELD) [26] load and displacement transfer scheme. While not demonstrated in the present work, MELD may optionally transfer static aeroelastic deflections from the FUN3D mesh to the panel mesh (though transferring from the structural mesh may yield a comparable result). Once an LFD analysis is complete, MELD is used to transfer the complex forces from the CFD surface nodes to the DLM panel centroids. Dividing the box forces by box areas subsequently yields the desired oscillatory pressures.

Note that when transferring LFD loads to DLM panels, MELD simply conserves energy through the principle of virtual work without knowledge of the mode shapes. As a result, the inner product of the loads and mode shapes will not necessarily be conserved, subsequently causing a mismatch between GAFs computed by LFD and the GAFs computed using the MELD-projected complex pressures, $[\mathbf{S}] \{C_{P,ref}\}$. Because this mismatch can shift the GAFs and thus affect

flutter and gust responses, two remedies have been developed by the authors. The simpler option is to apply a shift in each corrected DLM GAF such that it matches the LFD-computed GAF at the reference conditions. Alternatively, the projected pressure may itself be corrected before being used to compute $[C_F]$. This is done for an individual mode i in a similar manner as the DLM correction approach,

$$[C_{F,proj}]^i = [I] + \left[\left(\{C_{P',proj}\}^T \right)^\dagger (\{C_{P',ref}\} - \{C_{P',proj}\}) \right]^T, \quad (4)$$

where $[C_{F,proj}]^i$ is a diagonally-dominant correction matrix to be applied to the i^{th} mode's projected pressure,

$$\{C_{P,ref}\}^i = [C_{F,proj}]^i \{C_{P,proj}\}^i. \quad (5)$$

The so-called approximate panel pressures

$$\{C_{P',proj}\}^i = ([S]^T)^\dagger \{Q_{ss,proj}\}^i \quad (6)$$

and

$$\{C_{P',ref}\}^i = ([S]^T)^\dagger \{Q_{ss,lfd}\}^i \quad (7)$$

represent “pseudo-pressures” of sorts from the MELD-projected and original LFD analyses, as integration by mode shapes would result in their respective GAF columns, $\{Q_{ss,proj}\}^i$ and $\{Q_{ss,lfd}\}^i$. This projected pressure correction process, overall, ensures that the reference pressure used in the correction, $\{C_{P,ref}\}^i$, produces the same column of GAFs as LFD when integrated by the DLM mode shapes. This is achieved generally with no visible differences in the projected DLM panel pressures. A caveat to this process, however, is that modes that consist almost entirely of in-plane bending (thus being nearly zero in $\{w\}$ and similarly $[S]$) will yield ill-conditioned pressure correction matrices rather than ones close to identity, subsequently causing the same ill-conditioning issue with the DLM correction matrix. This can be avoided in an ad-hoc manner by checking the off-diagonal terms, and when the maximum is larger than a specified tolerance (e.g. 1), throwing out the mode and simply applying a GAF shift instead (so as to still ensure consistency with the reference conditions).

Note that although analytic derivatives of corrected DLM have been implemented (i.e., sensitivities of GAFs with respect to nodal coordinates, Mach, reference length, and mode shapes), sensitivities of the correction matrix have not been. As a result, the current implementation may only reliably run a corrected DLM optimization with a fixed correction matrix. In a case where CFD data are updated for each new design, the current lack of correct derivatives may slow down or completely prevent convergence to an optimum. This potential development may be addressed in the future, although such an approach would be significantly more costly due to the need for CFD at every new design.

B. Description of Flutter and Gust Analyses

Once structure-on-structure GAFs are computed, the p - k method implementation described in Jacobson and Stanford [2] and Stanford et al. [3] (among others) is used to evaluate flutter onset conditions. For these analyses, mass- and stiffness-proportional Rayleigh structural damping values are set to 2 and 5×10^{-4} , respectively. The optimizer enforces, in optimizations shown later, that this onset dynamic pressure must not be less than the arbitrary value of 1 psi. This value was chosen simply based on the baseline flutter response and preliminary optimizations, such that the constraint is active yet not overly limiting, which may remove any feasible design space. When applicable, gust-on-structure GAFs, in addition, may be used to compute continuous (also known as stochastic [27]) and discrete gust constraints. Both analyses take the root bending moment to be the quantity of interest to be limited, although future work by the authors will instead limit gust-induced stresses.

The continuous gust analysis implementation is described in Jacobson et al. [5] and Thelen et al. [6]. To illustrate this analysis, Fig. 1 shows the modal solution for the AGARD 445.6 wing across a range of frequencies, where LFD aerodynamics are used at a Mach number and angle of attack of 0.95 and 6 degrees, respectively. The flow velocity and density are 12,730 in/s and 1.0027×10^{-8} lb/in³. When using the scaling vector corresponding to root bending induced by unit modal displacements, and after scaling by the Dryden atmospheric spectrum [28] (with turbulence

scale and velocity Root Mean Square (RMS) of 30,000 in and 300 in/s), the Power Spectral Density (PSD) shown in Fig. 2 is obtained. The corresponding RMS of the root bending moment is then constrained by the optimizer; later optimizations set this upper limit to be the rather arbitrary value of 300 lb-in (where, for reference, the PSD in Fig. 2 has an RMS of 284 lb-in). This value was chosen in a similar manner as the flutter constraint limit (i.e., based on preliminary optimizations). Note that the RMS in this case is largely driven by the zero-frequency response because (a) the exclusion of rigid body modes, which would yield modal solutions approaching zero at zero frequency (with only rigid body mode responses being nonzero), and (b) the Dryden gust spectrum drops off rather quickly with frequency. Other work [6, 7] has demonstrated the effects of rigid body modes with the current continuous gust analysis tools using the Common Research Model (CRM) wing.




Fig. 1 Sample modal displacements due to a continuous gust input.

The discrete gust responses, which use the one-minus-cosine gust profile, are computed for four different gust lengths (all with the same vertical gust amplitude) in the time-domain using the inputted structure-on-structure and gust-on-structure GAFs, among other inputs (such as atmospheric velocity and density, which are taken to be fixed). A Kreisselmeier-Steinhauser (KS) [29, 30] aggregate of the four maximum root bending moments then forms a scalar response to be limited by the optimizer (with 300 lb-in later used, to be consistent with the continuous gust constraint), so that analytic derivatives must only be computed once for all gust lengths. An illustration of this analysis is given in Fig. 3.

III. Results for the AGARD 445.6 Wing Case

This section provides results using the described DLM correction procedure. Section III.A first shows sample flutter analysis results using a steady pitch correction. An extension of the method to all structural modes and a particular nonzero frequency is then demonstrated. Flutter-constrained optimization results are then given in Sect. III.B, where DLM is corrected first using steady pitching followed by all modes at multiple frequencies. Section III.C then demonstrates flutter- and gust-constrained optimizations with full corrections, where the effect of including or excluding shape design variables is also shown. In all sections, a 6 degree angle of attack and Mach number of 0.95 was chosen to ensure the presence of transonic nonlinearity, which can significantly alter the DLM and LFD optimal designs. In contrast, earlier testing with lower angles of attack (e.g., 0 or 2 degrees) caused the optima to be relatively close in the




Fig. 2 Resulting root bending moment PSD using the Dryden gust spectrum. The RMS, related to the area under the PSD (284 lb-in in this case), is limited by the optimizer to be no greater than an arbitrary limit of 300 lb-in.




Fig. 3 One-minus-cosine gust profiles with four different gust lengths (all amplitudes of 300 in/s) and the resulting root bending moment. A KS aggregate of maximum gust-induced root bending moment is constrained by the optimizer to be no greater than 300 lb-in.

design space, thus there was not much to gain by correcting the DLM aerodynamics. This is expected as the AGARD wing shows weak transonic effects at low angles of attack due to the thinness of the wing; therefore, the uncorrected DLM is sufficient to predict the compressibility effects [31].

A. Flutter Analysis Results at the Baseline Design

Figure 4 illustrates the correction matrix for this Mach 0.95, 6 degree angle-of-attack freestream condition, where steady pitching, simulated using FUN3D's Euler LFD solver, is used as a reference condition. For reference, the pressure perturbations for this pitching condition are shown in Fig. 5, where the LFD pressure has been transferred from the CFD mesh to DLM panels. In Fig. 4, box numbering goes chord-wise from leading to trailing edge, and from root to tip, with the panel mesh having 20 and 40 chord- and span-wise boxes, respectively. Shown in the left plot, off-diagonal terms at higher column numbers, which correspond to the outboard region of the wing, are larger but still relatively small with a range of about -0.034 to 0.039 (corresponding to a log-scale magnitude of roughly -1.5 in the figure). In

the right plot, diagonal terms are shown to be close to 1, with the outboard boxes (i.e., those with higher box numbers) deviating more from 1.




Fig. 4 Visualization of correction matrix, where steady pitching motion is used for the LFD reference data.




Fig. 5 DLM- and LFD-computed pitching panel pressures at zero frequency (i.e., perturbation per unit change in angle of attack) used for computing the correction matrix.

Using this correction matrix computed based on steady pitching motion, Fig. 6 shows the uncorrected and corrected DLM pressure perturbations at zero frequency for the first flexible mode compared to the LFD pressure transferred to the DLM panels. This first mode shape is also shown for reference. This is provided because it shows how the correction performs for a mode not included in the reference LFD data; ultimately, six modes will later be used for flutter analyses. In this figure, the corrected DLM pressure appears to capture the pressure perturbation at the left piece of the lambda-shaped shock fairly well, but underpredicts the magnitude of pressure on the right leg of the shock. The pressure perturbation near the tip appears to predict the LFD pressure fairly well, although the magnitude at the shock is underpredicted. Overall, the corrected DLM pressure appears significantly closer to LFD pressure when compared to DLM, which has no capability to compute the lambda shock.

For reference, Fig. 7 shows the corresponding steady pressure and zero-frequency mode 1 pressure perturbations on the CFD mesh. The mesh is also shown to give a sense for mesh resolution. As can be seen in the previous figures,




Fig. 6 Zero-frequency pressure perturbations due to unit perturbations of the first mode, where steady pitching motion is used for the LFD reference data.

there is a strong lambda shock on the upper surface that moves as the first mode is perturbed. The rounded tip also has large pressure perturbations near the trailing edge, which is reflected in the panel pressures in the previous figure.

Figure 8 shows the resulting GAFs corresponding to the first three flexible modes, where the corrected DLM analyses use the same steady pitch data described previously. In this figure, the zero-frequency correction shifts the DLM GAF real parts closer to those of LFD (e.g., dropping zero-frequency error by a median factor of three, relative to the LFD value). The imaginary parts are left uncorrected as the LFD reference data are only evaluated at zero frequency. This decision can readily be changed in the implementation; however, limited testing has suggested that this only has a small effect on predicted flutter onset conditions for the AGARD 445.6. This is perhaps because the flutter reduced frequency is relatively low at this flow condition: approximately 0.08.

The GAFs are used to compute the p - k damping and frequency curves in Fig. 9. While only four modes are shown for clarity, six are used in the flutter analysis. Table 1 summarizes the predicted flutter conditions, comparing corrected and uncorrected DLM to LFD. Note that in spite of the transonic Mach number, the DLM result is conservative as the LFD flutter dynamic pressure is larger; this is likely because the increased angle of attack shifted the transonic dip to lower Mach numbers and Mach 0.95 is past this dip. Nonetheless, the correction methodology succeeds in shifting damping and frequency curves closer to those of LFD, which is taken to be the truth data here. The errors in the resulting flutter dynamic pressure and frequency drop from around 12% with uncorrected DLM to less than 2% with corrected DLM. For perspective, the added computational cost to achieve this is the time it takes for one nonlinear flow solution in FUN3D, followed by one LFD linear problem. The LFD analysis, on the other hand, requires 66 LFD linear problems, as each of the 6 modes are solved for 11 reduced frequencies.

The same correction methodology is next demonstrated using all six modes, in addition to pitching, with Fig. 10 showing the terms of the correction matrix at a reduced frequency of 0.1. The LFD-projected pressure correction procedure described in Sect. II.A is used here, although it likely has no visible effect on the matrix terms. In this figure, note that whereas the real part of the correction matrix is close to an identity matrix, the imaginary part has




Fig. 7 FUN3D-computed steady pressure and zero-frequency pressure perturbation due to unit perturbation of the first mode.

Table 1 Flutter onset results.

Description	q_f (% error) [psi]	f_f (% error) [Hz]
LFD	0.447755	14.849003
DLM	0.393733 (-12.06%)	13.151406 (-11.43%)
Corrected DLM	0.456073 (+1.86%)	14.607825 (-1.62%)

no diagonal dominance. Overall, while using the full matrix is more accurate (with flutter q_f error even driven down to machine precision, if all modes and frequencies are used), the zero-frequency pitch correction shown previously would of course be cheaper to evaluate. Computing this full correction, however, is still cheaper than a typical LFD optimization iteration, which would require derivative computations (perhaps several, depending on the number of constraints); using a subset of modes and frequencies would also lessen the computational expense of the correction.

B. Flutter-constrained Optimizations

In the previous section, the zero-frequency pitch correction is shown to improve agreement of the DLM with LFD, using a relatively small set of CFD data from one steady nonlinear flow solution and one LFD linear solution. Also demonstrated is that the method can be extended to include nonzero frequencies and an arbitrary set of modes, which would be more costly but improve the agreement with LFD. This section is now concerned with leveraging the correction method for optimization with both shape and structural sizing design variables. The general approach to carry this out is to iteratively perform sub-optimizations, where the correction matrices are recomputed at the start of each sub-optimization:

- Compute the correction matrix at the baseline design.
- Run the corrected DLM optimization.
- At the new design, recompute the correction matrix and run the corrected DLM optimization; repeat until the change in optimized objective function is less than a certain tolerance.
- Finally, run the LFD optimization starting from the final design.

The goal here is to utilize the corrected DLM sub-optimizations to obtain a design closer to the LFD optimum than uncorrected DLM. If the designs are indeed close, then this can lead to a reduced number of required iterations in the final LFD optimization. For reference, in this work, each sub-optimization requires wall time similar to one LFD




Fig. 8 GAFs of the first three modes, where steady pitching is used for the LFD reference data.

optimization iteration (including function and gradient call). This relative cost is of course dependent on the DLM and CFD mesh resolutions, with the DLM mesh here being relatively fine and the CFD mesh being a moderately fine Euler mesh (as visualized in the previous section). Meanwhile, each new correction matrix computation takes some fraction of an LFD optimization iteration, as requiring derivatives can increase the cost of the iteration by several times, while the correction data can also be limited to a subset of modes and frequencies.

As an optimization test case, the AGARD 445.6 wing optimization from Jacobson and Stanford [2] is selected. The optimization problem seeks to minimize structural mass subject to a flutter inequality constraint (flutter dynamic




Fig. 9 p - k computed damping and frequency using the different aerodynamic fidelities.




Fig. 10 Correction matrix using all six modes at a reduced frequency of 0.1.

pressure $q_f \geq 1$ in this case) and a planform area equality constraint (must equal the baseline design's area). Gust constraints explained previously are thus far omitted. Design variables include 100 structural thicknesses and four planform variables: sweep, span, root chord, and tip chord (all defined as linear differential distances in inches, as illustrated in Stanford et al. [3]). Following the previous section, a Mach number of 0.95 and angle of attack of 6 degrees are used. Likewise, the DLM correction is again initially based on steady pitching. Figures 11 and 12 show the DLM- and LFD-based convergence metrics for this case.

In the left DLM optimization figure (Fig. 11a), three corrected DLM sub-optimizations are shown, with uncorrected DLM optimization also included for reference. The mass objective changed very little after the first sub-optimization, and in fact does not change at all during the last one. This is because shape design variables did not change during the final outer-loop optimization; although structural design variables did change slightly, this has no effect on the rigid pitching motion used for computing the DLM correction (which therefore remained identical between the final sub-optimizations). This suggests that just the first sub-optimization would have brought the optimizer sufficiently close to the LFD optimum. Nonetheless, the right DLM optimization figure (Fig. 11b) shows the corresponding flutter constraint values for these optimizations.

Meanwhile, Fig. 12a shows the subsequent LFD optimizations, starting from the corrected or uncorrected DLM optima, while Fig. 12b shows the corresponding flutter constraint convergence histories. An incomplete optimization from the baseline design is shown to give a sense for how long it may eventually take if no DLM optimization is available.




Fig. 11 DLM optimization convergence, where the corrected DLM optimizations are stopped once the change in objective function is less than 10^{-5} .




Fig. 12 LFD optimization convergence starting from the uncorrected and corrected DLM optima, denoted as $\mathbf{x}^*$. An incomplete optimization from the baseline design is also shown for reference.

Ultimately, the optimization from the corrected DLM optimum required fewer iterations than that of uncorrected DLM (11 versus 18), as the LFD optimization was able to start closer to the LFD optimum. Also noteworthy are the initial values of the flutter constraint: at the uncorrected DLM optimum, the LFD-computed constraint is -0.8, indicating an LFD-computed flutter dynamic pressure of 1.8 compared to the DLM-predicted value of 1. Optimization from the corrected DLM optimum, on the other hand, starts with a flutter constraint of close to zero, indicating close agreement between corrected DLM and LFD at the corrected DLM optimum.

Finally, Fig. 13 shows the optimized planforms and thickness distributions of the three optima: corrected and uncorrected DLM and LFD. The baseline design is also included for reference. With this particular case, most of the shape design variables are at their bounds but thickness distributions differ significantly. Visually, the DLM optimum places more mass towards the root trailing edge; though not shown here, this optimum is at least visually similar to the LFD optimum at zero angle of attack. At this higher angle of attack, however, mass is no longer needed at the root trailing edge but the optimizer strengthens the mid-span leading edge area. Visually, the corrected DLM appears to match this trend well.




Fig. 13 Baseline and optimal designs, where the black dotted lines show the initial planform for reference.

As a step beyond this zero-frequency pitch result, the same correction procedure is now applied using all modes (in addition to pitching) at a subset of (six out of eleven) reduced frequencies. When a subset of frequencies is used, the implementation uses linear interpolation of the correction matrix at reduced frequencies not included in the reference data. In addition, while the previous correction leaves the imaginary GAFs uncorrected (as the reference CFD data has no knowledge of nonzero frequencies), the imaginary parts are now corrected when using non-zero frequencies. Also note that the LFD panel correction procedure explained in Sect. II.A is used to avoid a shift between LFD GAFs on the CFD mesh and the GAFs corresponding to the MELD-projected LFD pressures on the DLM panels. This is not necessary when using pitching (which involves a rigid rotation with no static deformation) as a uniform mode shape ensures the conservation of the product of the load and mode shape.

Figure 14 illustrates the corrected DLM sub-optimizations with the full corrections in blue. For reference, the black curve shows the zero-frequency pitch-corrected DLM presented previously in Fig. 11. Uncorrected DLM is also shown, as is the LFD optimum's mass for reference. Scatter points on the sub-optimizations indicate where LFD corrections were recomputed with the corrected DLM optimization subsequently restarted. Overall, the full correction process required two more LFD evaluations before stopping (based on the same objective tolerance). Thus, this new correction required slightly more outer-loop iterations, with each new correction costing somewhat more than the previous steady pitch correction, which used only one LFD evaluation rather than one for each mode and frequency combination. The full correction, however, was able to reach the LFD optimum, and a subsequent LFD optimization (synonymous to Fig. 12 but not shown here) completed in only one optimization iteration. In addition, derivative calls with LFD can be quite costly, especially when including gust constraints, so each full correction calculation was still significantly faster than a typical LFD optimization iteration (as the correction procedure does not require derivatives).

For reference, Fig. 15 shows convergence of the four shape design variables during these sub-optimizations. For this case, three out of four design variables are at their bounds of -5 or 5 inches, while the span differential design variable ends up being zero. Span being at exactly zero (within roughly 10^{-8}) may suggest that it is more weakly influential to the design problem, and a value of zero satisfies the area constraint when the other design variables are at their bounds (noting that area is unaffected by sweep). Nonetheless, in theory, having this LFD optimum mostly at design variable bounds may help stabilize the process (presuming corrected DLM shape derivatives still point in the right direction) as the strong shocks would not be changing in position and strength between correction iterations.




Fig. 14 DLM optimization convergence with non-zero frequency corrections, where projected LFD pressures are initially corrected as described in Sect. II.A.




Fig. 15 Shape design variable convergence using full corrections for the flutter-constrained optimization.

C. Flutter- and Gust-constrained Optimizations With and Without Shape Design Variables

The discrete and continuous gust constraints described in Sect. II.B are now added to the optimization problem, and the same process is performed with and without shape design variables. The structure-on-structure and gust-on-structure correction matrices are computed separately using all frequencies and modes. The case with shape design variables uses an option in the implementation that accumulates DLM and LFD pressure data at all previous evaluations, so that the correction matrix ensures a pressure match at not just the current LFD design but all previous designs as well. This was added in an effort to help stabilize the outer-loop convergence, as normally a new correction would have no knowledge of correction data at previous designs.

Figure 16 shows the objective function convergence for the corrected DLM sub-optimizations as well as the subsequent LFD optimizations. Uncorrected DLM with shape design variables, which appears to have converged on a completely different design as LFD, is also shown. The sizing-only uncorrected DLM case is not shown as it did not converge. This appeared to be due to the fidelity level's inability to satisfy the constraints (indicating a lack of feasible design space), with thickness design variables all reaching their upper bounds at which point the optimizer was unable to proceed. While not shown here, the flutter and discrete gust constraints are both active for all cases, while the continuous gust constraint is inactive.

While the sizing-only uncorrected DLM optimization case had no feasible design space, the same case with DLM corrections was able to find feasible space during each sub-optimization. Unfortunately, the mass objective tolerance





Fig. 16 Objective convergence with gust and flutter constraints, with and without shape design variables.


between sub-optimizations did not meet the same mass tolerance used in flutter-only optimizations and simply continued until the maximum correction iterations of 10. The subsequent sizing-only LFD optimization, while not reducing the mass any more, also became stuck in a line search and was halted manually. This issue, which was not observed in other design cases, needs further investigation and likely affected both the sub-optimization and the LFD optimization convergence. One possible factor is that the fifth mode shape when computing the second sub-optimization's correction is almost entirely in-plane, and its MELD-transferred surface displacement on the wing (in combination with the modal reference amplitude) essentially passes through the symmetry plane due to small unclamped regions at the root leading and trailing edges. As a result, this seems to cause negative volumes when computing the corresponding volume mode shape, although mesh deformation and LFD residuals still converge to machine precision. Presumably, this mode has very little effect on the flutter or gust responses, although the negative volumes may perhaps be a source of noise when changing the design slightly. Nonetheless, partly because of this difficulty converging, the case with shape design variables was halted manually as the objective function between outer-loop iterations appeared to have adequately converged. The subsequent LFD optimization reduced the objective function only slightly more, indicating that corrected DLM located a design close to the true optimum.

Figure 17 shows the corresponding shape design variable convergence histories. As shown in this figure, the uncorrected DLM optimization converged on a completely different planform as corrected DLM and LFD. Corrected DLM, after a few outer-loop iterations, hovers around the LFD-optimal planform but struggles to converge on it exactly. This appears to be mainly due to the sweep design variable, which is no longer at one of its bounds like in the flutter-only optimizations. In addition, while sweep does not affect the area constraint, the span and tip chord design variables do; corrected DLM indicates that these values are balanced to try to satisfy this constraint, rather than keeping tip chord all the way at its bound like with LFD.


The derivatives at the final design are compared between DLM, corrected DLM, and LFD to help gauge how well the aeroelastic constraints trend among the three analyses. Figure 18 illustrates the shape derivatives at this design. In this figure, corrected DLM shape derivatives appear to sometimes match those of LFD better than uncorrected DLM, but there are cases where uncorrected DLM has more accurate derivatives. This is perhaps not surprising, as while the current DLM correction may ensure a zero-order consistency with LFD, first-order consistency is in no way guaranteed as a single LFD evaluation does not inform DLM on how shocks move with geometry changes. It is unclear if the pressure accumulation option can eventually yield an “over-training” effect of sorts; while likely not a direct cause of inaccurate shape derivatives, this will be further investigated in the future. In addition, a less challenging flow condition where DLM and LFD start out with more similar trends may yield better shape derivatives when using the correction procedure. Unfortunately, however, there will always be shortcomings with shape design with the current approach as DLM inherently neglects thickness effects; airfoil shape, for instance, could not be designed at all as DLM does not model the geometry at this level of detail. For such a design problem, one may need to update the correction matrix at every design (and likewise implement analytic derivatives of the correction matrix), likely using a relatively cheap



(a) Uncorrected DLM




(b) Corrected DLM




(c) LFD from corrected DLM optimum

Fig. 17 Shape design variable convergence with both gust and flutter constraints.


training set such as steady pitching.



(a) Flutter




(b) Discrete gust




(c) Continuous gust

Fig. 18 Shape derivatives at the final corrected DLM optimum with flutter and gust constraints.


Figure 19 illustrates the structural thickness design variable derivatives at the final design with shape design variables. To give a sense for how derivatives are distributed, Fig. 20 shows the derivative contours on the final design using LFD aerodynamics. Unlike the shape derivatives, these sizing derivatives show a much closer match between corrected DLM and LFD. This seems intuitive, as sizing variables, when excluding static aeroelastic deformation like in this case, would only affect mode shapes and not the steady flow field.



(a) Flutter



(b) Discrete gust



(c) Continuous gust

Fig. 19 Thickness derivatives at the final corrected DLM optimum with shape design variables.




Fig. 20 LFD thickness derivative contours at the final corrected DLM optimum with shape design variables.

For reference, Fig. 21 shows the same derivative comparison for the case without shape design variables, where the agreement in fact seems somewhat better. It is unclear if this better agreement is a possible sign of an “over-training” effect of the pressure accumulation option used in the case with shape design variables (as it was not used for this sizing-only case), or if other factors are more responsible for this better fit (e.g., flow-field, mode shapes, LFD-to-panel pressure correction procedure, or solver settings). The option appeared to stabilize convergence somewhat in earlier testing, but the implications of such an option has not extensively been tested thus far due to the numerous possible combinations of settings and the computational cost of each full run. In the future, performing optimizations with only the four shape design variables may shed light on what may work best. Nonetheless, this result indicates that the current DLM correction procedure is perhaps best suited for structural sizing optimizations; while cases with shape design variables may ultimately be expedited as well, further developments may be needed to achieve reliable convergence to the LFD optimum.




Fig. 21 Thickness derivatives at the final corrected DLM optimum without shape design variables.

IV. Conclusion

This work has demonstrated a gradient-based optimization process with flutter and gust constraints based on corrected DLM. Being implemented using the MPhys library for OpenMDAO allows for easier interchangeability between aerodynamic solvers while using the same structural solver, and also facilitates the computation of analytic derivatives (although derivatives of the correction matrix itself are still lacking). The corrected DLM method was demonstrated using steady pitch data, as well as data from several modes and frequencies. Flutter-constrained optimization of the AGARD 445.6 wing (using steady or unsteady corrections) was shown as well as results including gust constraints with and without shape design variables.

The approach, which iteratively performs corrected DLM sub-optimizations, each with a new pre-multiplied AIC correction matrix, tends to speed up the search for the true LFD optimum by locating a design closer to it (when compared to uncorrected DLM). This is especially true for the flutter-constrained case with full non-zero frequency corrections, which located the optimum exactly (thanks to the optimum being at shape design variable bounds). The presence of

shape design variables, however, adds difficulty to the design problem, as indicated in the gust- and flutter-constrained problem. This is perhaps not surprising, as the correction procedure does not ensure first-order consistency with LFD, and DLM inherently assumes a lifting body is made up of two-dimensional panels. The Mach 0.95, 6 degree angle-of-attack flow condition selected may be especially challenging due to the presence of a strong lambda-shaped shock that moves with planform geometry, presumably causing a poorer correlation between LFD and the underlying DLM physics. Nonetheless, the approach seems particularly well suited for structural sizing optimizations, as flutter and gust constraint derivatives with respect to these design variables are shown to closely match LFD derivatives. This is at least true when neglecting static aeroelastic deflection; a case where sizing variables affect static deflection and thus shocks may prove to be somewhat more difficult.

While a different correction matrix formulation may perform better than the current implementation, reliably obtaining an LFD optimum with shape design variables may require a constantly varying correction matrix (as opposed to sub-optimizations each with fixed correction matrices). This would also increase the cost of each analysis immensely when compared to a corrected DLM analysis. If one uses only a limited data set for computing the correction matrix (e.g. steady pitching, limited modes and frequencies), however, this may still prove fruitful as it would still drive down the cost of a full LFD optimization.

Nonetheless, with the current sub-optimization approach, several aspects could be further investigated to gauge performance and reliability of the implementation. A case with only the limited set of shape design variables used here, perhaps with a lower Mach and/or angle of attack, could help identify when the method starts to have difficulties. Such a test case may also help identify alternative correction matrix formulations that work better when the geometry is varying. The pressure accumulation option, which continuously builds up reference DLM and LFD pressure data, could also be better understood to see if there are limitations relating to factors such as “over-training”. A case with static aeroelastic deflection, likely excluding shape design variables, may also help confirm whether this flexibility is a significant issue for the implementation. While not presumed to be a major factor, a case that uses Reynolds-averaged Navier-Stokes aerodynamics rather than Euler equations may also be a useful data point. Finally, while somewhat out of scope, similar corrections could be applied to the authors’ in-house Vortex Lattice Method solver to recover a similar improvement in accuracy of steady aerodynamics.

Acknowledgments

This work is supported by the Transformational Tools and Technologies (TTT) project of the NASA Transformative Aeronautics Concepts Program (TACP).

References

- [1] Jacobson, K. E., Stanford, B. K., Wood, S. L., and Anderson, W. K., “Adjoint-based Sensitivities of Flutter Predictions based on the Linearized Frequency-domain Approach,” *59th AIAA Aerospace Sciences Meeting*, Virtual Event, 2021.
- [2] Jacobson, K. E., and Stanford, B. K., “Flutter-Constrained Optimization with the Linearized Frequency-Domain Approach,” *AIAA SciTech Forum*, San Diego, California, USA, 2022.
- [3] Stanford, B., Thelen, A., and Jacobson, K., “Fixed- and Free-Mode Flutter Derivatives During Aeroelastic Optimization,” *AIAA Journal*, Vol. 62, No. 12, 2024, pp. 4645–4653. doi:10.2514/1.J064166.
- [4] Stanford, B., Thelen, A., and Jacobson, K., “Multifidelity Optimization with Transonic Flutter Constraints,” *AIAA AVIATION Forum and ASCEND*, Las Vegas, NV, 2024. doi:10.2514/6.2024-4025.
- [5] Jacobson, K., Thelen, A., and Stanford, B., “Linearized Frequency-Domain Gust Analysis and Adjoint-Based Sensitivities,” *AIAA SciTech*, Orlando, Florida, USA, 2024. doi:10.2514/6.2024-2591.
- [6] Thelen, A., Stanford, B., and Jacobson, K., “Flutter and Stochastic Gust Constraint Studies of the Undeformed CRM using MPhys,” *AIAA SciTech Forum*, Orlando, Florida, USA, 2024. doi:10.2514/6.2024-2774.
- [7] Thelen, A. S., Stanford, B. K., and Jacobson, K. E., “Gradient-based Optimization of the Common Research Model Wing Subject to CFD-based Gust and Flutter Constraints,” *International Forum on Aeroelasticity and Structural Dynamics (IFASD)*, The Hague, The Netherlands, 2024.
- [8] Anderson, W. K., and Bonhaus, D. L., “An Implicit Upwind Algorithm for Computing Turbulent Flows on Unstructured Grids,” *Computers & Fluids*, Vol. 23, No. 1, 1994, pp. 1–21.

- [9] Biedron, R. T., and Thomas, J. L., "Recent Enhancements to the FUN3D Flow Solver for Moving-Mesh Applications," *AIAA 47th Aerospace Sciences Meeting*, Orlando, Florida, 2009.
- [10] Nielsen, E. J., and Anderson, W. K., "Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes," *AIAA Journal*, Vol. 40, 2002, pp. 1155–1163.
- [11] Wang, L., Diskin, B., Biedron, R., Nielsen, E. J., and Bauchau, O., "Sensitivity Analysis of Multidisciplinary Rotorcraft Simulations," *55th AIAA Aerospace Sciences Meeting*, 2017, p. 1670.
- [12] Anderson, W. K., Newman, J. C., and Karman, S. L., "Stabilized Finite Elements in FUN3D," *Journal of Aircraft*, Vol. 55, No. 2, 2018. doi:10.2514/1.C034482.
- [13] Anderson, W. K., and Newman, J., "High-Order Stabilized Finite Elements on Dynamic Meshes," *2018 AIAA Aerospace Sciences Meeting*, 2018.
- [14] Yildirim, A., Jacobson, K. E., Stanford, B. K., Gray, J. S., Mader, C. A., Martins, J. R. R. A., and Kennedy, G. J., "MPhys: A Modular Multiphysics Simulation Package using the OpenMDAO Framework," *AIAA Aviation Forum*, Presentation Only, Virtual, 2020.
- [15] Yildirim, A., Jacobson, K., Anibal, J., Stanford, B., Gray, J., Mader, C., Martins, J., and Kennedy, G., "MPhys: A Modular Multiphysics Simulation Package using the OpenMDAO Framework," Submitted, 2024.
- [16] Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A., "OpenMDAO: An Open-Source Framework for Multidisciplinary Design, Analysis, and Optimization," *Structural and Multidisciplinary Optimization*, Vol. 59, 2019, pp. 1075–1104. doi:10.1007/s00158-019-02211-z.
- [17] Giesing, J., Kalman, T., and Rodden, W., "Correction Factor Techniques for Improving Aerodynamic Prediction Methods," Tech. Rep. CR 144967, NASA, May 1976.
- [18] Palacios, R., Climent, H., Karlsson, A., and Winzell, B., "Assessment of strategies for correcting linear unsteady aerodynamics using CFD or experimental results," *IFASD 2001-074*, 2001.
- [19] Jadic, I., Hartley, D., and Giri, J., "AN ENHANCED CORRECTION FACTOR TECHNIQUE FOR AERODYNAMIC INFLUENCE COEFFICIENT METHODS," 1999. URL <https://api.semanticscholar.org/CorpusID:55182192>.
- [20] Valente, C., Jones, D., Gaitonde, A., Cooper, J., and Lemmens, Y., "Doublet-Lattice Method Correction by Means of Linearised Frequency Domain Solver Analysis," *AIAA SciTech Forum*, 2016. doi:10.2514/6.2016-1575.
- [21] Chen, P., Silva, R., and Liu, D., "Transonic AIC Weighting Method using Successive Kernel Expansion," *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics & Materials Conference*, Austin, TX, 2005. doi:10.2514/6.2005-1991.
- [22] Moreno, R., Narisetti, R., von Knoblauch, F., and Taylor, P., "A Modification to the Enhanced Correction Factor Technique to Correlate With Experimental Data," *AIAA SciTech Forum*, Kissimmee, FL, 2015. doi:10.2514/6.2015-1421.
- [23] Yates, C. E., "AGARD Standard Aeroelastic Configurations for Dynamic Response. Candidate Configuration I - Wing 445.6," Tech. Rep. TM-100492, NASA, August 1 1987.
- [24] Moore, E., "On the reciprocal of the general algebraic matrix," *Bulletin of the American Mathematical Society*, Vol. 26, No. 9, 1920, pp. 394–395.
- [25] Penrose, R., "A generalized inverse for matrices," *Proceedings of the Cambridge Philosophical Society*, Vol. 51, No. 3, 1955, pp. 406–413.
- [26] Kiviaho, J. F., and Kennedy, G. J., "Efficient and Robust Load and Displacement Transfer Scheme Using Weighted Least Squares," *AIAA Journal*, Vol. 57, No. 5, 2019, pp. 2237–2243.
- [27] Stanford, B., *Aeroservoelastic Optimization under Stochastic Gust Constraints*, American Institute of Aeronautics and Astronautics, 2018. doi:doi:10.2514/6.2018-2837, URL <https://doi.org/10.2514/6.2018-2837>.
- [28] Federal Aviation Administration, "14 CFR §25.341," <https://www.ecfr.gov/current/title-14/chapter-I/subchapter-C/part-25/subpart-C/subject-group-ECFR3e855ea22ea15d0/section-25.341>, retrieved 14 November 2023.

- [29] Kreisselmeier, G., and 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.
- [30] Kennedy, G. J., and Hicken, J. E., "Improved Constraint-Aggregation Methods," *Computer Methods in Applied Mechanics and Engineering*, Vol. 289, 2015, pp. 332–354.
- [31] Cunningham, H., Batina, J., and Bennett, R., "Modern Wing Flutter Analysis by Computational Fluid Dynamics Methods," *AIAA Journal of Aircraft*, Vol. 25, No. 10, 1988.