

Transonic Aeroelastic Flutter Computations with Mesh Adaptation

Bret K. Stanford*, Steven J. Massey†, Kevin E. Jacobson‡ and Pawel Chwalowski§

NASA Langley Research Center, Hampton, VA, 23681

This paper considers multiscale mesh adaptation of aeroelastic wing models in transonic flow. Workflows are developed which allow for high-fidelity static aeroelastic equilibrium conditions and a matched-point dynamic aeroelastic flutter mechanism, to evolve alongside the mesh adaptation mechanics, via a staggered procedure. Rather than enforcing static equilibrium or matched-point constraints for a given mesh, these conditions are allowed to be more gradually met across the mesh adaptation process. These cost-effective workflows are successfully demonstrated for a wing-body semispan wind-tunnel model. For this configuration, the results show substantial interdependencies between the mesh-adapted static and dynamic aeroelastic physics.

I. Introduction

Transonic aeroelastic flutter predictions remain a challenging computational problem. The state-of-the-practice is typically based on the unsteady Reynolds-averaged Navier-Stokes (uRANS) equations, coupled to a modal representation of the structural model [1–3]. This system is integrated in the time-domain, at progressively higher dynamic pressures, until an oscillatory instability is observed. This flutter computation process is expensive and cumbersome, in addition to challenges with mesh convergence, time step convergence, and subiterative convergence [4]. The uRANS modeling is only valid for flutter mechanisms driven by moderate amounts of flow separation [5], but aeroelastic computations with scale-resolving turbulence models are still relatively rare [6, 7].

One promising tool for more efficient uRANS-based flutter computations is the linearized frequency domain (LFD) method [8–10]. LFD linearizes about a steady background flow via infinitesimal oscillatory perturbations of each mode shape, resulting in a complex valued set of generalized aerodynamic forces (GAFs). These GAFs are then used to inexpensively compute the flutter onset point via some appropriate eigen-solver. An LFD solver eliminates issues related to time step and subiterative convergence, at the expense of multiple (one per mode shape and reduced frequency combination of interest) complex-valued linear solves. A second issue is matched-point considerations: the LFD flutter dynamic pressure should match the dynamic pressure used to compute the steady background flow, potentially necessitating an iterative process [2, 11]. A matched-point solution is automatically obtained with the time-domain process outlined above.

A second tool that has shown promise for flutter computations is multiscale mesh adaptation techniques [12], which reduce spatial interpolation error in the flow field based on some scalar output of the flow. Past work by the authors [13–15] has successfully linked these mesh refinement tools to LFD, to obtain high-quality transonic flutter predictions with no temporal errors (as these solutions are in the frequency domain) and minimal/quantifiable spatial meshing errors. The meshes created by this refinement tool (namely, *refine*^a) produce elements with a high degree of anisotropy near the wing surface. This type of grid topology produces solutions of high quality for the aeroelastic cases cited above, in addition to a wide-range of rigid-model applications [16–19].

*Research Aerospace Engineer, Aeroelasticity Branch, bret.k.stanford@nasa.gov, AIAA Associate Fellow.

†Research Aerospace Engineer, Aeroelasticity Branch, AIAA Senior Member.

‡Research Aerospace Engineer, Aeroelasticity Branch, AIAA Member.

§Research Aerospace Engineer, Aeroelasticity Branch, AIAA Senior Member.

^aAvailable via <https://github.com/NASA/refine>.

These *refine*-generated volume meshes cannot easily be deformed in a finite manner due to surface motions of the flexible structure [20] without the formation of negative volumes. This failure is potentially due to the inability of highly-skewed surface elements to withstand bending curvature of cantilevered wing structures. LFD is still viable, as this tool only considers infinitesimal wing deformations, but any finite static aeroelastic deformations due to the steady flowfield (about which LFD is linearized) must be ignored. As such, the authors’ past refinement-based LFD work has considered pitch-and-plunge airfoil cases where static aeroelastic pitch deformations can be set as a change in the flow angle of attack [14], or wings with symmetric airfoils at 0° angle of attack (the AGARD 445.6 configuration) where there is no static aeroelastic deformation [15].

The goals of the current work are to explore mesh-adapted transonic flutter workflows that are able to include finite static aeroelastic deformations. This is done with a hybrid workflow of *refine* and the commercial software HeldenMesh^b, which can produce mesh-adapted grids with a sustained prismatic topology in the boundary layer (a similar workflow aimed to preserve the prisms in the boundary layer can be found in Ref. [21]), and as such, can undergo finite surface deformations.

In addition to demonstrating mesh adaptation driven by static aeroelastic considerations, this paper will also utilize LFD-based flutter computations, linearized about these mesh-adapted deformed wing shapes. A final goal is to demonstrate the role of matched-point dynamic pressure requirements. These results will be exercised on a generic wind-tunnel model with cantilevered bending/torsion wing deformations.

II. Aeroelastic Modeling

All RANS-based simulations are conducted in the SFE [22] flow path of the FUN3D solver [23], a Streamline Upwind Petrov-Galerkin formulation with an SA-neg turbulence model. Steady solutions can be obtained over the rigid wing, or alternatively, the wing surface may be deformed into some prescribed shape, the volume mesh is subsequently deformed (using an elasticity-based method [20]), and steady solutions can be obtained over that deformed wing shape. The prescribed shape is obtained from a modal-based structural analysis, where FUN3D reads the amplitude η of each mode shape to construct the wing shape. An iterative static aeroelastic process could be used to couple FUN3D to the modal structural representation, but an alternative process is used here, where the coupling is closely embedded within the mesh adaptation process, described below.

After computing the steady flow over the wing (using either the undeformed or deformed shape), the LFD solver [10] is used to compute the GAFs, which finally enables computation of the flutter onset dynamic pressure, flutter- Q , via the $p - k$ eigen-equation [24].

III. Multiscale Mesh Adaptation

Mesh adaptation mechanics are handled with a hybrid workflow of *refine* and the commercial software HeldenMesh [25]. The adaptations reduce spatial interpolation error in the flow field based on a particular multiscale metric [12] scalar field, which in this work is the Mach number. For a non-aeroelastic simulation of a rigid wing, the overall workflow is summarized as follows:

1. An initial coarse mesh is manually constructed in HeldenMesh with a prismatic topology within the boundary layer of the model.
2. A steady flow simulation is conducted with FUN3D/SFE.
3. Using the scalar Mach field from the steady flow solution, *refine* generates a multiscale metric field. A set of anisotropic sources are derived from this field, and the sources within the boundary layer (a fixed wall distance defined by the user) are removed.
4. This altered set of sources is provided to HeldenMesh to create the next mesh. This subsequent mesh will preserve prismatic layers near the viscous surfaces, owing to the deletion of the near-body sources. The effect of the sources outside the boundary layer still penetrate the prism layer however, resulting in adaptation near the surface.
5. The process repeats to step 2, and continues until the computational budget is met.

^b<https://heldenaero.com/heldenmesh>

For static aeroelasticity, step 2 could be expanded to include an iteratively-coupled workflow: the steady pressures on the wing are integrated into steady GAFs, and the modal amplitudes computed as:

$$\omega_i^2 \cdot \eta_i = q \cdot GAF_i \quad (1)$$

where i is the mode number and ω is the natural vibration frequency of the mode (and so ω^2 is the generalized stiffness). These modal amplitudes could be fed back into FUN3D, which would deform the mesh as described above. Steady FUN3D/SFE could be solved on the deformed mesh, the GAFs recomputed, and the process repeated to convergence. Finally, the *refine*/HeldenMesh workflow could be exercised on this final converged static aeroelastic Mach field, a new mesh would be created, and the process repeats. Clearly, this workflow is extremely expensive, necessitating numerous steady flow simulations for each refinement step.

Alternatively, the static aeroelastic iterations may be implemented alongside the mesh adaptation steps:

1. An initial coarse mesh is manually-constructed.
2. A single steady flow simulation is conducted with FUN3D/SFE.
3. The GAFs are computed corresponding to this steady flow, and the modal amplitudes are computed via Eq. 1.
4. Using the scalar Mach field from the steady flow solution, *refine* and HeldenMesh create the next mesh.
5. This new mesh is deformed using the modal amplitudes found from the old mesh.
6. The process repeats to step 2, and continues until the computational budget is met.

This is a staggered procedure: modal amplitudes are computed for mesh j , but implemented on mesh $j + 1$. With this process, only a single steady FUN3D/SFE solve is required for each mesh, and the static aeroelastic solution should converge alongside the mesh adaptation process. There is no guarantee of convergence however. For transonic cases with large deformations, there may be a strong coupling between the wing shape and shock location (a primary driver of the grid clustering location), and this process may struggle or fail to converge. All of the results shown in this work do largely converge however (as do results in Ref. [13], which use a similar staggered idea for static pitch deformations of an airfoil), even for cases with strong transonic nonlinearities. Aitken relaxation [26] (with a relaxation factor of 0.5 between adjacent mesh convergence steps) is also used to improve convergence.

It is also noted that, for this work, the HeldenMesh anisotropic sources do not account for the deformation of the grid. The scalar Mach field from the deformed mesh is transferred to the undeformed mesh, then *refine* computes the anisotropic metric on the undeformed mesh, and HeldenMesh finally creates a new grid in the undeformed state. This new grid is then deformed by FUN3D based on the past vector of modal amplitudes, η . For the mesh adaptation software used here, it is possible to consider a case where the underlying geometry changes from step to step [13], but this has not been exercised. For moderate wing deformations, this disconnect is unlikely to have a strong impact.

A final addition to this workflow is the LFD-based flutter computations. This can be simply added onto step 2 above: the steady flow simulation is used as the background flow for LFD, which then computes unsteady GAFs and finally the flutter- Q . Because the refinement metrics are still completely driven by the steady flow features, LFD is only computed as a side-process to the mesh adaptation process and has no direct impact on it (Ref. [13] uses a mesh adaptation driven by unsteady snapshots reconstructed from the fluttering mode shape, but this idea is not used here).

However, if a matched-point analysis is desired, the flutter- Q found at mesh j would be used as the static aeroelastic dynamic pressure in Eq. 1, for mesh $j + 1$. In this workflow, both the static aeroelastic analysis and the matched-point flutter computations would converge alongside the mesh adaptation steps, and LFD would have an indirect impact on the mesh adaptation.

IV. Test Case

The test case used here is an aeroelastic wind-tunnel model tested in the NASA Langley Transonic Dynamics Tunnel (TDT) in January 2024. This model was originally tested in the TDT in 1993, 1994, and 2000 as the MAVRIC model (Models for Aeroelastic Validation Research Involving Computation) [27]. This is a wing-body semispan sidewall-mounted model. The wing has no twist or dihedral, a sweep angle of 23°, a taper ratio of 0.29, and a span of 53.17 in. The original model construction was a stepped thickness

aluminum plate with balsa to form the airfoil shape, but was rehabilitated in 2023 with a high-density foam to form a NACA0012 airfoil shape (the previous airfoil spanwise distribution was supercritical in nature).

This rehabilitated model, re-termed as the FFTM (Facility Flutter Training Model, Fig. 1) and tested in the TDT in 2024, is used as a demonstration case for this work.




Figure 1. FFTM model installed in the TDT (credit: NASA).

V. Results

For all of the work considered here, the initial manually-constructed semispan mesh for this case consists of 2.5 million nodes and can be seen on the left of Fig. 2. The mesh adaptation process begins with this 2.5 million mesh and proceeds based on a source term scaling factor, where the number of desired nodes in the discrete mesh is proportional to the value of this scaling factor. For this work, the scaling factor is doubled if the number of nodes in two consecutive meshes is less than a ratio of 1.25. All of the mesh adaptation processes shown here are terminated after 7 steps, which corresponds to between 10 and 11 million nodes: most of the aeroelastic metrics are reasonably converged at this point. The static aeroelastic results could reasonably be run at finer meshes, but the LFD analyses become very expensive and the associated linear complex-value solves are very difficult to converge for highly-adapted meshes.

Six structural mode shapes (tuned to match the experimental data) are used for aeroelastic computations, and have natural frequencies ranging from 4.3 Hz (first bending mode) to 61.4 Hz (second torsion mode).




Figure 2. Initial 2.5 million node coarse mesh of the rigid wing (left), and 11.8 million node adapted mesh of the deformed wing (right): Mach=0.8, AoA=3°, Q=110 psf.

A. Static Aeroelastic Results

In this section, static aeroelastic mesh-adapted FUN3D/SFE simulations are run at a Mach number of 0.8, a Reynolds number of $8.4 \cdot 10^5$, and a fixed dynamic pressure of 110 psf. No LFD-based flutter computations are considered in this section. As discussed above, the staggered mesh adaptation process proceeds across 7 steps: modal amplitudes are computed for mesh j but implemented on mesh $j + 1$. The final mesh on the

right of Fig. 2 has 11.8 million nodes. This mesh is in the deflected static aeroelastic shape (the initial coarse mesh on the left of the figure is undeflected) due to a 3° model AoA. This mesh displays grid clustering due to the presence of a strong shock but retains relatively low aspect ratio cells along the wing surface. A detail of the wingtip meshes is shown in Fig. 3, and an overlay of the surface pressures is shown in Fig. 4.




Figure 3. Detail of the two meshes in Fig. 2 at the wingtip.

Additional mesh and flow information can be seen Figures 5 and 6, which show Mach contours sliced over the wing at 50% and 95% semispan, respectively. This is done for AoAs of 0° , 1° , 2° , and 3° , where the 3° result corresponds to the results shown previously in Figures 2 - 4. Each of the results at these four AoA values are the product of four separate mesh adaptation processes. For the initial mesh results in the top row of Figures 5 and 6, a strong but under-resolved shock develops over the wing as the AoA is increased, and the flow is separated at the foot of the shock for the outboard station in Fig. 6. The final mesh-adapted results in the bottom row of the two figures reflects the deflected wing state, which increases for higher AoA and also for the outboard 95% station (compared to the 50% station in Fig. 5). The shocks for the deflected adapted cases are well-resolved due to the grid clustering, but weaker relative to the undeflected results, largely due to the passive bend-twist washout coupling of the swept wing deformation which reduces the effective angle of attack along the cantilevered wing.




Figure 4. Overlayed surface pressure contours of the initial undeflected mesh and the final adapted deformed mesh in Fig. 2.

Convergence of the static aeroelastic residual (Eq. 1) is shown in Fig. 7, across the 7 steps of the mesh adaptation process. The figure shows over two orders-of-magnitude drop in this residual: residuals are lower for the smaller AoA's (because the deflections are smaller), but the relative drop in residual across the process is roughly the same for each AoA result. A further drop could be expected with additional grid refinement steps, but it can be seen in Fig. 8 that the modal amplitudes (modes 1, 3, and 4 are shown) have largely




Figure 5. Mach contours at 50% semispan, for the initial undeformed mesh (top row), and the final adapted deformed mesh (bottom row), at AoA's of 0° - 3°.




Figure 6. Mach contours at 95% semispan, for the initial undeformed mesh (top row), and the final adapted deformed mesh (bottom row), at AoA's of 0° - 3°.

stabilized. Another option would be to decrease the aggressiveness of the refinement schedule, such that the mesh would grow at a smaller rate. The static aeroelastic process would then have more time to reach equilibrium before the mesh becomes prohibitively large, but this was not considered here.

It is noted that the mode shapes shown in Fig. 8 are scaled for visualization and are not plotted at the amplitudes indicated in the bottom figure. Beyond mode 1, the deformational contributions of the higher modes are relatively small. Mode 1 amplitude convergence in Fig. 8 displays a mostly monotonic behavior for each AoA, and this is due to the use of Aitken relaxation. Mesh adaptation without Aitken relaxation,

which is not shown here, is highly non-monotonic due to the fact that positive deflection unloads the swept wing, and the resulting modal amplitudes are not well-converged within the 7 mesh adaptation steps.




Figure 7. Convergence of the static aeroelastic residual (Eq. 1) across the mesh adaptation process.




Figure 8. Convergence of select modal amplitudes 1, 3, and 4, across the mesh adaptation process.

B. Non-Matched-Point Flutter Results

An initial flutter mapping is shown in Fig. 9, comparing experimental TDT results, results computed with linear aeroelastic tools [28], and LFD results computed with the initial 2.5 million node mesh. The LFD results in Fig. 9 are linearized about the rigid wing at 1° and thus ignore static aeroelastic considerations. The experimental and LFD flutter boundaries show the expected ‘dip’ and recovery in the transonic flow regime, though the computed flutter dip underpredicts the measurement by 10-20 psf. The linear aeroelastic result misses these flow features due to an inability to predict shocked flow. LFD-based eigenvalue migration plots corresponding to the flutter results in Fig. 9 are shown in Fig. 10 for Mach numbers 0.5 (subsonic flow), 0.825 (bottom of the transonic dip), and 0.875 (back side of the dip). These results show a binary coupling of modes 1 and 2, where mode 1 is a first bending mode (shown in Fig. 8), and mode 2 is a second bending mode with some torsional vibration. The frequency coalescence is weaker for the transonic flows in Fig. 10, but the rate at which the model destabilizes with increased dynamic pressure is not appreciably different for any of the Mach numbers shown.

There are numerous possible reasons for why the LFD solver underpredicts the experimental flutter boundary in Fig. 9, some of which may be the spatial gridding errors and static aeroelastic load redistribution themes considered in this paper. As will be seen below, however, exercising these tools does not improve the correlation with the TDT experiment, so some other source of error is likely the cause. This comparison with experiment is not a core theme of this paper, but some additional discussion can be found in Appendix A.




Figure 9. Computed and measured flutter boundaries across the transonic regime at 1° angle of attack. Conditions above a given flutter boundary are unstable.




Figure 10. Eigenvalue migration plots at three Mach numbers corresponding to the data in Fig. 9.

A general sense of the role of mesh adaptation in the transonic flutter dip computation can be seen in Fig. 11, where the initial mesh result (as above, from the manually-constructed 2.5 million node mesh) is directly copied from Fig. 9, and the adapted result is from the final step in the mesh refinement process. These results are again linearized about the rigid wing at 1° AoA. Because the refinement metrics are still completely driven by the steady flow features, LFD is only computed as a side-process to the mesh adaptation process and has no direct impact on it. For the result in Fig. 11, the mesh adaptation only shows a relevant shift in the flutter- Q above Mach 0.75, with an increase in 50 psf at the highest considered Mach value of 0.875.

The role of mesh-adapted static aeroelasticity in the LFD-based flutter computations is considered next, focusing solely on Mach 0.8 and a fixed static Q value of 110 psf (the same condition used in the previous section). These results are not matched-point, and this concept will be explored in the next section. As such, again LFD is only computed as a side-process to the mesh adaptation process and has no direct impact on it. The adapted meshes, the steady modal aeroelastic deflections, and the steady transonic flows are all identical to what was found previously in Figures 2 - 8.

Results are shown in Fig. 12 for flutter- Q as a function of AoA, for the initial 2.5 million node mesh, the adapted mesh without any static aeroelastic deformations, and the adapted mesh with static aeroelastic deformations. Results at 1° AoA are replicated from Fig. 11. Both the initial mesh and the case which




Figure 11. Mesh-adapted flutter results when linearizing about the rigid wing shape, at 1° angle of attack.

ignores static aeroelastic deflections (so LFD is linearized about the rigid wing) show a substantial increase in flutter- Q with increased AoA, in Fig. 12. This is likely due to the leftward shift of the transonic dip with increased AoA (i.e., a shift towards smaller Mach values) [29], and so conditions on the front side of the flutter dip at lower AoA, shift to the backside of the dip at higher AoA. Furthermore, the shift in flutter- Q with mesh adaptation is more pronounced at higher AoA, as would be expected from the more aggressive transonic flows for those cases.

The more important result in Fig. 12 is with the static aeroelasticity included in the mesh adaptation, since refining about an undeflected rigid wing at higher AoAs is a fictional scenario. Here the change in flutter- Q with increased AoA is very small: the passive bend-twist washout coupling of the wing dilutes the transonic flows, as seen previously in Figures 5 and 6. The stronger transonic nonlinearities nominally seen at higher AoA are nearly offset by the associated increased static aeroelastic deflections. This result may have important ramifications for the design of aeroelastic wind-tunnel models for transonic flutter benchmarking, as discussed in more detail in Appendix B.

The mesh refinement convergence details for the cases in Fig. 12 are shown in Fig. 13. The flutter- Q metric converges well for the two cases shown with static aeroelastic deformations; for the LFD cases that are linearized about the undeflected rigid wing, LFD results could not be obtained for the final 7th mesh. The transonic flows are very aggressive here (as seen in Figures 5 and 6), and the linear complex-value solves could not be converged for these highly-adapted meshes. Furthermore, the 3° AoA undeflected result in Fig. 13 is not obviously converged by mesh level 6 (7.7 million nodes), though the 2° result is.




Figure 12. Mesh-adapted flutter results as a function of angle of attack.




Figure 13. Convergence of flutter- Q across the mesh adaptation process.

C. Flutter Results (Matched-Point)

All of the previous results fixed the dynamic pressure used for static aeroelastic deformations at 110 psf. The results in this final section enforce a matched-point flutter condition, where the static aeroelastic dynamic pressure will be driven (during the mesh adaptation process) to be equal to the flutter dynamic pressure. In contrast with the previous workflows, the LFD does now have an indirect impact on the mesh adaptation, due to the matched-point constraint. Both the static aeroelastic analysis and the matched-point flutter computations converge alongside the mesh adaptation steps. A form of Aitken relaxation is also used for the matched-point restriction, where the static aeroelastic dynamic pressure is updated as: $Q_{j+1} = \delta \cdot Q_j^{flutter} + (1 - \delta) \cdot Q_j$. The results shown here use a relaxation value of 0.75, though a value of 1.0 was tested also with only a minor degradation in convergence.

The mesh-adapted matched-point flutter results are shown in Fig. 14, for comparison with the results that do not include the matched-point restriction. These latter results are copied directly from Fig. 12, where the static aeroelastic deformations are driven by a fixed Q of 110 psf. For this case, the matched-point flutter boundary is only slightly different from the non-matched-point flutter boundary, with a peak difference of 2 psf at the highest 3° AoA. The static aeroelastic deformations for the matched-point case are smaller (evidenced by the mode 1 amplitude shown on the right of Fig. 14), because they are driven by a dynamic pressure in the 90 psf range, as opposed to the non-matched-point deformations under a fixed 110 psf load.

The minor difference in the flutter boundaries in Fig. 14 is, as above, likely due to the washout of the transonic flow features when the swept wing statically deforms, which dilutes the importance of matched-point considerations (just as the overall importance of AoA is also diluted). Despite this, the matched-point controller must overcome large initial differences in the two Q values, as seen in Fig. 15: at 3° AoA, an initial 55 psf difference is driven to less than 0.1 psf after the 7 mesh adaptation steps. It is understood that the static aeroelastic residual is simultaneously being driven down during this mesh adaptation process, but these residuals are not meaningfully different from those in Fig. 7 and 8, and are not shown here.

VI. Conclusions

This paper has demonstrated various mesh adaptation strategies for transonic aeroelasticity. Previous mesh adaptation work by the authors had been limited to cases without finite static aeroelastic deformation, as the mesh refinement software produced elements with a high degree of anisotropy near the wing surface, and these meshes could not be deformed in a finite manner without the formation of negative volumes. The current work uses the commercial software HeldenMesh, whose mesh refinement process can sustain a prismatic topology in the boundary layer: these meshes are therefore able to withstand finite surface deformations.

The test case for this work is a wing-body semispan sidewall-mounted model, tested in the NASA LaRC Transonic Dynamics Tunnel. To better understand the transonic aeroelastic physics of this configuration, three levels of mesh adaptation are considered here:




Figure 14. Mesh-adapted matched-point results as a function of angle of attack, compared to results without matched-point restrictions (static aeroelastic Q fixed at 110 psf).




Figure 15. Convergence of the matched-point constraint across the mesh adaptation process.

1. Static aeroelasticity is ignored, and the mesh is adaptively refined on the steady flow features of the rigid undeformed wing. A flutter analysis (via the linearized frequency domain method) is conducted by linearizing about the rigid wing, but this flutter analysis has no impact on the mesh adaptation.
2. Some realistic static aeroelastic dynamic pressure is used to deform the wing, and the mesh is adaptively refined on the steady flow features of the deformed wing. Static aeroelasticity is computed with a staggered procedure, where structural deflections computed at mesh level j , are applied at mesh level $j + 1$. A flutter analysis is again computed by linearizing about the deformed wing shape, and again the flutter analysis has no impact on the mesh adaptation.
3. The static aeroelastic dynamic pressure is driven, during the mesh adaptation, to be equal to the flutter dynamic pressure for a matched-point analysis. Here the flutter analysis does implicitly impact the mesh adaptation process, again via a staggered procedure.

For the aeroelastic wind-tunnel model studied here, the differences between the first two mesh adaptation procedures are substantial. The bend-twist washout coupling of the swept wing leads to a large reduction in the strength of the transonic flow nonlinearities, significantly altering both the mesh refinement mechanics and the flutter mechanism. This shift is particular strong at higher model angles of attack. The third mesh adaptation strategy (matched-point analysis) is less significant, in that it produces results fairly similar to the second mesh adaptation workflow. This is partially due to the fact that the fixed dynamic pressure chosen for the second workflow is fairly close to the computed flutter dynamic pressures, and partially due to the aforementioned dilution of the transonic physics when the swept wing bends upwards.

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. Don Keller and Ian Giles of the NASA Langley Aeroelasticity Branch led the rehabilitation and TDT testing of the FFTM.

References

- ¹Bartels, R., Rumsey, C., Biedron, R., “CFL3D Version 6.4 - General Usage and Aeroelastic Analysis,” NASA TM 2006-214301.
- ²Bendiksen, O., “Review of Unsteady Transonic Aerodynamics: Theory and Applications,” *Progress in Aerospace Sciences*, Vol. 47, No. 2, pp. 135-167, 2011, <https://doi.org/10.1016/j.paerosci.2010.07.001>.
- ³Antimirova, E., Jung, J., Zhang, Z., Machuca, A., Gu, G., “Overview of Computational Methods to Predict Flutter in Aircraft,” *Journal of Applied Mechanics*, Vol. 91, No. 5, 2024.
- ⁴Kamakoti, R., Shyy, W., “Fluid-Structure Interaction for Aeroelastic Applications,” *Progress in Aerospace Sciences*, Vol. 40, No. 8, pp. 535-558, 2004.
- ⁵Edwards, J., “Calculated Viscous and Scale Effects on Transonic Aeroelasticity,” *Journal of Aircraft*, Vol. 45, No. 6, pp. 1863-1871, 2008.
- ⁶Chwalowski, P., et al. “Flutter Prediction Report in support of the High Angle Working Group at the Third Aeroelastic Predictions Workshop” *AIAA SciTech Forum*, AIAA Paper 2024-0418.
- ⁷Soneda, K., Tsushima, N., Hamidreza, K., Saitoh, K., “Flutter Evaluation in Transonic Regime Based on Detached Eddy Simulation” *AIAA SciTech Forum*, AIAA Paper 2025-0819.
- ⁸Thormann, 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.
- ⁹Bekemeyer, P., Thormann, R., Timme, S., “Frequency-Domain Gust Response using Computational Fluid Dynamics,” *AIAA Journal*, Vol. 55, No. 7, pp. 2154-2187, 2017.
- ¹⁰Jacobson, 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.
- ¹¹Gallman, J., Batina, J., Yang, T., “Computational Transonic Flutter Boundary Tracking Procedure” *Journal of Aircraft*, Vol. 25, No. 3, pp. 263-270, 1988.
- ¹²Loseille, A., Dervieux, A., Frey, P., Alauzet, F., “Achievement of Global Second Order Mesh Convergence for Discontinuous Flows with Adapted Unstructured Meshes,” *AIAA Computational Fluid Dynamics Conference*, Miami, Florida, 2007.
- ¹³Jacobson, K., Stanford, B., Kiviaho, J., Ozoroski, T., Park, M., Chwalowski, P., “Multiscale Mesh Adaptation for Transonic Aeroelastic Flutter Problems,” *AIAA Aviation Forum*, virtual event, 2021.
- ¹⁴Stanford, B., Jacobson, K., “Transonic Aeroelastic Modeling of the NACA 0012 Benchmark Wing,” *AIAA Journal*, Vol. 59, No. 10, pp. 4134-4143, 2021.
- ¹⁵Stanford, B., Jacobson, K., “Transonic Flutter Dips of the AGARD 445.6 Wing,” *Journal of Aircraft*, Vol. 61, No. 2, pp. 515-522, 2024.
- ¹⁶Kleb, B., Park, M., Wood, W., Bibb, K., Thompson, K., Gomez, R., Tesch, S., “Sketch-to-Solution: An Exploration of Viscous CFD with Automatic Grids,” *AIAA Aviation Forum*, Dallas, TX, June 17-21, 2019.
- ¹⁷Balan, A., Park, M., Anderson, W., Kamenetskiy, D., Krakos, J., Michal, T., Alauzet, F., “Verification of Anisotropic Mesh Adaptation for RANS Simulations over a ONERA M6 Wing,” *AIAA Journal*, Vol. 58, No. 4, pp. 1550-1565, 2020.
- ¹⁸Shehata, H., Korzun, A., Schoenenberger, M., “Uncertainty Modeling for Mars Ascent Vehicle’s Aerodatabase Development,” *AIAA SciTech Forum*, Orlando, FL, January 8-12, 2024.
- ¹⁹Park, M., “HLPW-5: Mesh Adaptation for RANS Technology Focus Group Workshop Summary,” *AIAA SciTech Forum*, Orlando, FL, January 6-10, 2025.
- ²⁰Biedron, R., Thomas, J., “Recent Enhancements To The FUN3D Flow Solver For Moving-Mesh Applications,” AIAA Paper 2009-1360.
- ²¹Nastac, G., Miller, F., Schneider, D., Jacobson, K., Jones, B., Opgenorth, M., “Toward Adaptive Mixed-Element Unstructured Grids for Simulations of Viscous Flows,” *AIAA SciTech Forum*, Orlando, FL, January 6-10, 2025.
- ²²Anderson, W., Newman, J., Karman, S., “Stabilized Finite Elements in FUN3D,” *Journal of Aircraft*, Vol. 55, No. 2, pp. 696-714, 2017.
- ²³Anderson, W., Biedron, R., Carlson, J., Derlaga, J., Diskin, B., Druyor, C., Gnoffo, P., Hammond, D., Jacobson, K., Jones, W., Kleb, B., Lee-Rausch, E., Liu, Y., Lohry, M., Nastac, G., Nielsen, E., Padway, E., Park, M., Rumsey, C., Thomas, J., Thompson, K., Walden, A., Wang, L., Wood, S., Wood, W., Zhang, X., “FUN3D Manual: 14.2,” NASA TM-2025-0002308.
- ²⁴van Zyl, L., Maserumule, M., “Divergence and the p-k Flutter Equation,” *Journal of Aircraft*, Vol. 38, No. 3, pp. 584-586, 2001, <https://doi.org/10.2514/2.2805>.
- ²⁵Wick, A., Hooker, R., “HeldenMesh Contributions and Grid Generation Lessons Learned from the 5th AIAA High-Lift Prediction Workshop,” *AIAA SciTech Forum*, Orlando, FL, January 6-10, 2025.
- ²⁶Kenway, G., Kennedy, G., Martins, J., “Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations,” *AIAA Journal*, Vol. 52, No. 5, pp. 935-951, 2014.
- ²⁷Edwards, J., Spain, C., Keller, D., Moses, R., Schuster, D., “Transport Wing Flutter Model Transonic Limit Cycle Oscillation Test,” *Journal of Aircraft*, Vol. 46, No. 4, pp. 1104-1113, 2009.

²⁸Rodden, W., Johnson, E., “MSC/NASTRAN Aeroelastic Analysis, User’s Guide, v68,” The MacNeal-Schwendler Corporation, 1994.

²⁹Edwards, J., Bennett, R., Whitlow, W., Seidel, D., “Time-Marching Transonic Flutter Solutions Including Angle-of-Attack Effects” *Journal of Aircraft*, Vol. 20, No. 11, pp. 527-534, 1983.

³⁰Wieseman, C., Bennet, R., “Wall Boundary Layer Measurements for the NASA Langley Transonic Dynamics Tunnel,” NASA TM 2007-214867.

³¹Ivanco, T., “Unique Testing Capabilities of the NASA Langley Transonic Dynamics Tunnel, an Exercise in Aeroelastic Scaling,” *AIAA Ground Testing Conference*, San Diego, CA, June 24-27, 2013.

³²Stanford, B., “Scaled Aeroelastic Wind Tunnel Model Design with Topology Optimization,” *AIAA Aviation Forum*, virtual event, August 2-6, 2021.

A. Appendix: Numerical Comparison with TDT Flutter Data

The comparison between LFD and the TDT experiment in Fig. 9 is mixed: the overall location (in terms of Mach number) and relative severity of the flutter dip (in terms of Q) is captured, but the prediction is consistently 10-20 psf below the measurement. There are many possible culprits for this gap:

1. Flutter testing usually exhibits a certain amount of conservatism, so as not to risk the destruction of the wind-tunnel model. This, however, does not explain the discrepancy in Fig. 9, since the numerical model underpredicts the measurements, rather than vice-versa.
2. There may be uncertainties in the structural model, but this is not likely a large factor, since the structural model used here matches the experimentally-measured vibration properties well.
3. There are uncertainties in the ability of the URANS modeling to correctly predict transonic unsteady aerodynamics, but there are no pressure measurements on the model during the test, so this uncertainty is unknowable. Furthermore, the flutter prediction in the subsonic range is only slightly closer to the experiment, despite the fact that the aerodynamic modeling is well-trusted for subsonic flows. It is also noted that LFD and the linear aeroelastic tool do not agree as well as may be expected for the subsonic flows in Fig. 9, although the aerodynamics in this regime are nominally linear.
4. The model is attached directly to the tunnel wall, and therefore the inboard portion of the model is immersed within the boundary layer of the model. The thickness of the TDT boundary layer at the model location is unclear [30] but likely does extend past the fuselage and into some portion of the inboard wing. The modeling here assumes a symmetry plane as the tunnel wall; Ref. [15] had found that modeling the wall as viscous caused a large increase (15-25 psf) in the computed transonic flutter boundary of the AGARD 445.6 configuration, compared to a symmetry plane. The AGARD configuration is an extreme case however, since this wing is attached directly to the wall (as opposed to the fuselage used for the current case) and will therefore be more susceptible to the boundary layer of the tunnel.

B. Appendix: Design of Benchmarking Models for Transonic Flutter

A complete set of flutter boundaries is shown in Fig. 16 for AoA’s between 0° and 3° . This is done for the 2.5 million manually-constructed coarse mesh utilized above, as it would be too expensive to run mesh refinement cases for each of these conditions. These results then contain some mesh refinement errors (particularly for the strong transonic flows), but the basic trends shown in this figure should suffice for the discussion considered here. Flutter results are shown for LFD cases linearized about the rigid wing and also cases which are linearized about the matched-point static aeroelastic wing shape.

The flutter results linearized about the rigid wing shape (left side of Fig. 16) show the expected trend with increased AoA: a leftward shift in the transonic flutter dip, as transonic effects are encountered at earlier Mach numbers [29]. The bottom of the dip decreases from Mach 0.825 at 0° AoA, to Mach 0.725 at 3° AoA. For the results linearized about the static aeroelastic wing shape (right side of Fig. 16), the AoA effects are diluted as discussed above. There is still a change in flutter dip behavior with AoA, but it is more subtle, with the dip location only decreasing to Mach 0.8.

TDT testing of this FFTM model did not prioritize changes in AoA, but if this had been a testing priority, one could obviously expect that experimental results would mirror the static aeroelastic results in Fig. 16, rather than the rigid results (which are not physically realizable). Given the noise, uncertainty, and conservatism inherent in flutter testing, it could be difficult to experimentally realize such subtle trends

seen on the right of Fig. 16. On the other hand, an ideal transonic flutter case for prediction benchmarking would mirror the general trends on the left side of Fig. 16, with large quantifiable shifts in the flutter dip with increased AoA, where these shifts are driven by hard-to-predict nonlinearities in the flow. The FFTM, then, is likely a non-optimal transonic benchmarking case, given that the static aeroelasticity dilutes much of the key physics desired for code benchmarking.




Figure 16. Complete flutter results for the 2.5 million manually-constructed coarse mesh (no refinement), with and without static aeroelastic effects.

This would suggest design guidelines for future aeroelastic flutter models aimed at prediction benchmarking. For scaled-down models of flight aircraft, both static aeroelastic and dynamic aeroelastic (i.e., flutter) scaling laws must be observed [31]. But manufactured experiments (such as may be used for benchmarking/validation tests) can utilize any scaling, and typically these models are designed so that they flutter within the envelope of the tunnel but can still withstand static loads to the appropriate safety factor. This exercise requires a strong design conflict between low stiffness (for flutter) and high strength which can be difficult to satisfy [32]. The discussion in this Appendix would indicate a third requirement for low static aeroelastic deflections (high stiffness) to preserve the transonic nonlinearities in the flow field.

It can be expected that this third requirement would be in even greater conflict with the low-stiffness flutter requirement, further restricting the model design space. Adding some torsional inertial to the wingtip (in the form of a cylindrical rod, e.g.) can decrease the flutter boundary and still maintain high wing stiffness, but this is not a satisfactory solution for wind-tunnel models aimed to replicate subsonic transport aircraft. Similarly, considering an unswept wing model geometry would remove the passive bend-twist washout coupling responsible for the degraded aeroelastic nonlinearities, but this again is not satisfactory for a research thrust aimed at subsonic transport aircraft.