High-fidelity Multidisciplinary Design Optimization of Aircraft Configurations

Joaquim R. R. A. Martins (PI), Gaetan K. W. Kenway, David Burdette, Eirikur Jonsson
University of Michigan, Ann Arbor, Michigan

Graeme J. Kennedy
Georgia Institute of Technology, Atlanta, Georgia

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

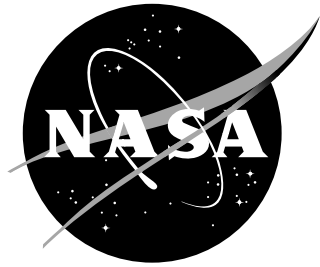- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at *http://www.sti.nasa.gov*

- E-mail your question to help@sti.nasa.gov

- Phone the NASA STI Help Desk at 757-864-9658

- Write to:
  NASA STI Information Desk
  Mail Stop 148
  NASA Langley Research Center
  Hampton, VA 23681–2199

High-fidelity Multidisciplinary Design Optimization of Aircraft Configurations

Joaquim R. R. A. Martins (PI), Gaetan K. W. Kenway, David Burdette, Eirikur Jonsson
University of Michigan, Ann Arbor, Michigan

Graeme J. Kennedy
Georgia Institute of Technology, Atlanta, Georgia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681

August 2017

Acknowledgments

Abstract

To evaluate new airframe technologies we need design tools based on high-fidelity models that consider multidisciplinary interactions early in the design process. The overarching goal of this NRA is to develop tools that enable high-fidelity multidisciplinary design optimization of aircraft configurations, and to apply these tools to the design of high aspect ratio flexible wings. We develop a geometry engine that is capable of quickly generating conventional and unconventional aircraft configurations including the internal structure. This geometry engine features adjoint derivative computation for efficient gradient-based optimization. We also added overset capability to a computational fluid dynamics solver, complete with an adjoint implementation and semi-automatic mesh generation. We also developed an approach to constraining buffet and started the development of an approach for constraining flutter. On the applications side, we developed a new common high-fidelity model for aeroelastic studies of high aspect ratio wings. We performed optimal design trade-offs between fuel burn and aircraft weight for metal, conventional composite, and carbon nanotube composite wings. We also assessed a continuous morphing trailing edge technology applied to high aspect ratio wings. This research resulted in the publication of 26 manuscripts so far, and the developed methodologies were used in two other NRAs.

Contents

List of Figures

3

List of Tables

1   Introduction

With increasing interest in unconventional aircraft, there is a need to develop design tools and processes that can handle next-generation concepts. Due to the lack of experience and historical data for these concepts, low-fidelity aircraft conceptual design tools are not sufficient to resolve the design trade-offs and accurately quantify the performance benefits. This NRA addressed these issues under two main parts:

1. Geometry and meshing

2. High aspect ratio flexible wings

The geometry and meshing are currently two major bottlenecks in the creation of high-fidelity models for aircraft configurations. Therefore, we focused our efforts on the development of tools that facilitate the rapid generation of geometry and computational meshes for both CFD and structural finite element models. Two main research questions arise under this topic:

- How can we rapidly construct aircraft geometries with a parametrization that is compatible with high-fidelity models and gradient-based optimization?

- How can we handle the meshing of arbitrary aircraft geometries and speed up the meshing process for both the external flow and the internal structure?

The high aspect ratio wing project focused on developing and applying high-fidelity aerostructural optimization tools to explore new high-performance wings. While there were many possibilities under this overall goal, we prioritized the following questions:

- What is the optimum wing planform when considering new materials?

- What is the optimum use of morphing technologies?

- How can we constrain buffet margin in high-fidelity wing design optimization?

- What should a common benchmark for aeroelastic studies of high aspect ratio transonic wings look like?

- How can we constrain flutter when optimizing both structural sizing and aerodynamic shape?

This report is structured around the answers to the above research questions. Due to the extent of the work, this report does not repeat the details that have already been published as part of this project, and instead we cite the relevant publications. Therefore, sections with material whose corresponding papers have not yet been published (Sections 5, 6, and 8) are much longer than the other sections.

This NRA resulted in the following 26 publications:

1. Kenway, G. K. W.; and Martins, J. R. R. A.:, Buffet onset constraint formulation for aerodynamic shape optimization. *AIAA Journal*, 55(6):1930–1947, June 2017. doi:10.2514/1.J055172

2. Jonsson, E.; Kenway, G. K. W.; Kennedy, G. J.; and Martins, J. R. R. A.:, Development of flutter constraints for high-fidelity aerostructural optimization. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2017

3. Kenway, G. K. W.; Secco, N.; Martins, J. R. R. A.; Mishra, A.; and Duraisamy, K.:, An efficient parallel overset method for aerodynamic shape optimization. In *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, January 2017. doi:10.2514/6.2017-0357

4. Martins, J. R. R. A.:, *Encyclopedia of Aerospace Engineering*, volume Green Aviation, chapter Fuel burn reduction through wing morphing, pages 75–79. October 2016, pages 75–79. doi:10.1002/9780470686652.eae1007

5. Kenway, G. K. W.; and Martins, J. R. R. A.:, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. *AIAA Journal*, 54(1):113–128, January 2016. doi:10.2514/1.J054154

6. Kenway, G. K. W.; and Martins, J. R. R. A.:, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. *AIAA Journal*, 54(1):113–128, January 2016. doi:10.2514/1.J054154

7. Hwang, J. T.; and Martins, J. R. R. A.:, An unstructured quadrilateral mesh generation algorithm for aircraft structures. *Aerospace Science and Technology*, 59:172–182, December 2016. doi:10.1016/j.ast.2016.10.010

8. Lyu, Z.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of an adaptive morphing trailing edge wing. *Journal of Aircraft*, 52(6):1951–1970, November 2015. doi:10.2514/1.C033116

9. Lyu, Z.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Aerodynamic shape optimization investigations of the Common Research Model wing benchmark. *AIAA Journal*, 53(4):968–985, April 2015. doi:10.2514/1.J053318

10. Hwang, J. T.; Lee, D. Y.; Cutler, J. W.; and Martins, J. R. R. A.:, Large-scale multidisciplinary optimization of a small satellite's design and operation. *Journal of Spacecraft and Rockets*, 51 (5):1648–1663, September 2014. doi:10.2514/1.A32751

11. Burdette, D. A.; Kenway, G. K.; and Martins, J. R. R. A.:, Performance evaluation of a morphing trailing edge using multipoint aerostructural design optimization. In *57th AIAA Structures, Structural Dynamics, and Materials Conference*, January 2016. doi:10.2514/6.2016-0159

12. Burdette, D. A.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Aerostructural design optimization of a continuous morphing trailing edge aircraft for improved mission performance. In *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2016. doi:10.2514/6.2016-3209

13. Kenway, G. K.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of the CRM configuration including buffet-onset conditions. In *54th AIAA Aerospace Sciences Meeting*, January 2016. doi:10.2514/6.2016-1294

14. Burdette, D.; Kenway, G. K. W.; Lyu, Z.; and Martins, J. R. R. A.:, Aerostructural design optimization of an adaptive morphing trailing edge wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2015. doi:10.2514/6.2016-1294

15. Kenway, G. K. W.; and Martins, J. R. R. A.:, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2015. doi:10.2514/6.2015-0264

16. Kenway, G. W. K.; and Martins, J. R. R. A.:, High-fidelity aerostructural optimization considering buffet onset. In *Proceedings of the 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2015

17. Gray, J.; Hearn, T.; Moore, K.; Hwang, J. T.; Martins, J. R. R. A.; and Ning, A.:, Automatic evaluation of multidisciplinary derivatives using a graph-based problem formulation in OpenMDAO. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-2042

18. Hwang, J. T.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Geometry and structural modeling for high-fidelity aircraft conceptual design optimization. In *Proceedings of the 15th AIAA Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-2041

19. Kennedy, G. J.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Towards gradient-based design optimization of flexible transport aircraft with flutter constraints. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-2726

20. Kenway, G. K. W.; Kennedy, G. J.; and Martins, J. R. R. A.:, Aerostructural optimization of the Common Research Model configuration. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-3274

21. Lyu, Z.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of an adaptive morphing trailing edge wing. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-3275

22. Lyu, Z.; Kenway, G. K. W.; and Martins, J. R. R. A.:, RANS-based aerodynamic shape optimization investigations of the common research model wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2014. doi:10.2514/6.2014-0567

23. Kennedy, G. J.; Kenway, G. K. W.; and Martins, J. R. R. A.:, High aspect ratio wing design: Optimal aerostructural tradeoffs for the next generation of materials. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2014. doi:10.2514/6.2014-0596

24. Hwang, J. T.; Lee, D. Y.; Cutler, J. W.; and Martins, J. R. R. A.:, Large-scale MDO of a small satellite using a novel framework for the solution of coupled systems and their derivatives. In *Proceedings of the 54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, April 2013. doi:10.2514/6.2013-1599

25. Kennedy, G. J.; and Martins, J. R. R. A.:, An adjoint-based derivative evaluation method for time-dependent aeroelastic optimization of flexible aircraft. In *Proceedings of the 54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, April 2013

26. Hwang, J. T.; and Martins, J. R. R. A.:, GeoMACH: Geometry-centric MDAO of aircraft configurations with high fidelity. In *Proceedings of the 14th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, September 2012. doi:10.2514/6.2012-5605

Part 1: Geometry and meshing

We now describe the development of GeoMACH, an open source geometry engine that is capable of generating geometries for conventional and unconventional aircraft configurations. We then detail the development of the capability to mesh these configurations for CFD analysis and optimizations using the overset technique.

2 Differentiable geometry parametrization engine

Under this project, we developed an open-source package for high-fidelity geometry parametrization and automatic structural mesh generation, called Geometry-centric MDO of Aircraft Configurations with High fidelity (GeoMACH) [22, 30]. GeoMACH enables the use of high-fidelity computational tools at the conceptual design level through a centralized geometry modeler. By employing a robust, general, and flexible geometry representation, GeoMACH supports a configuration-level multidisciplinary analysis and optimization environment in which the entire process, is streamlined, from geometry creation to visualization of results from high-fidelity tools. One key feature of GeoMACH is the ability to efficiently compute derivatives. This capability is especially valuable when used in conjunction with a gradient based optimization.

In GeoMACH, the user creates an aircraft configuration as a set of components such as wings and fuselages, and selects the desired shape design variables. These shape variables can be as high-level as aircraft design parameters like sweep, or as low-level as individual NURBS control points. Based on the structural layout defined by the user, GeoMACH automatically creates the structural mesh as well using a novel unstructured mesh generation algorithm. Fig. 1 shows the structural meshes created by GeoMACH for several unconventional configurations. To enable gradient-based optimization, GeoMACH efficiently computes numerically exact derivatives of the discretized geometry and the structural nodes with respect to the shape design variables.

GeoMACH is described in much more detail by Hwang et al. [22], and is currently being used in a another NASA Research Announcement (NRA) for the design optimization of a TBW configuration [31].

The development of automatic structural mesh generation was not originally proposed, but we determined together with NASA it would be valuable to add this to GeoMACH, so we developed it. The structural mesh generation algorithm is described by Hwang and Martins [11].

This project resulted in another serendipitous collaboration stream. Since GeoMACH was to be tightly integrated with NASA's OpenMDAO software, our team ended up contributing to OpenMDAO with an algorithmic framework we named Modular Analysis and Unified Derivatives (MAUD) [32], which is now the core algorithm and data structure in OpenMDAO.
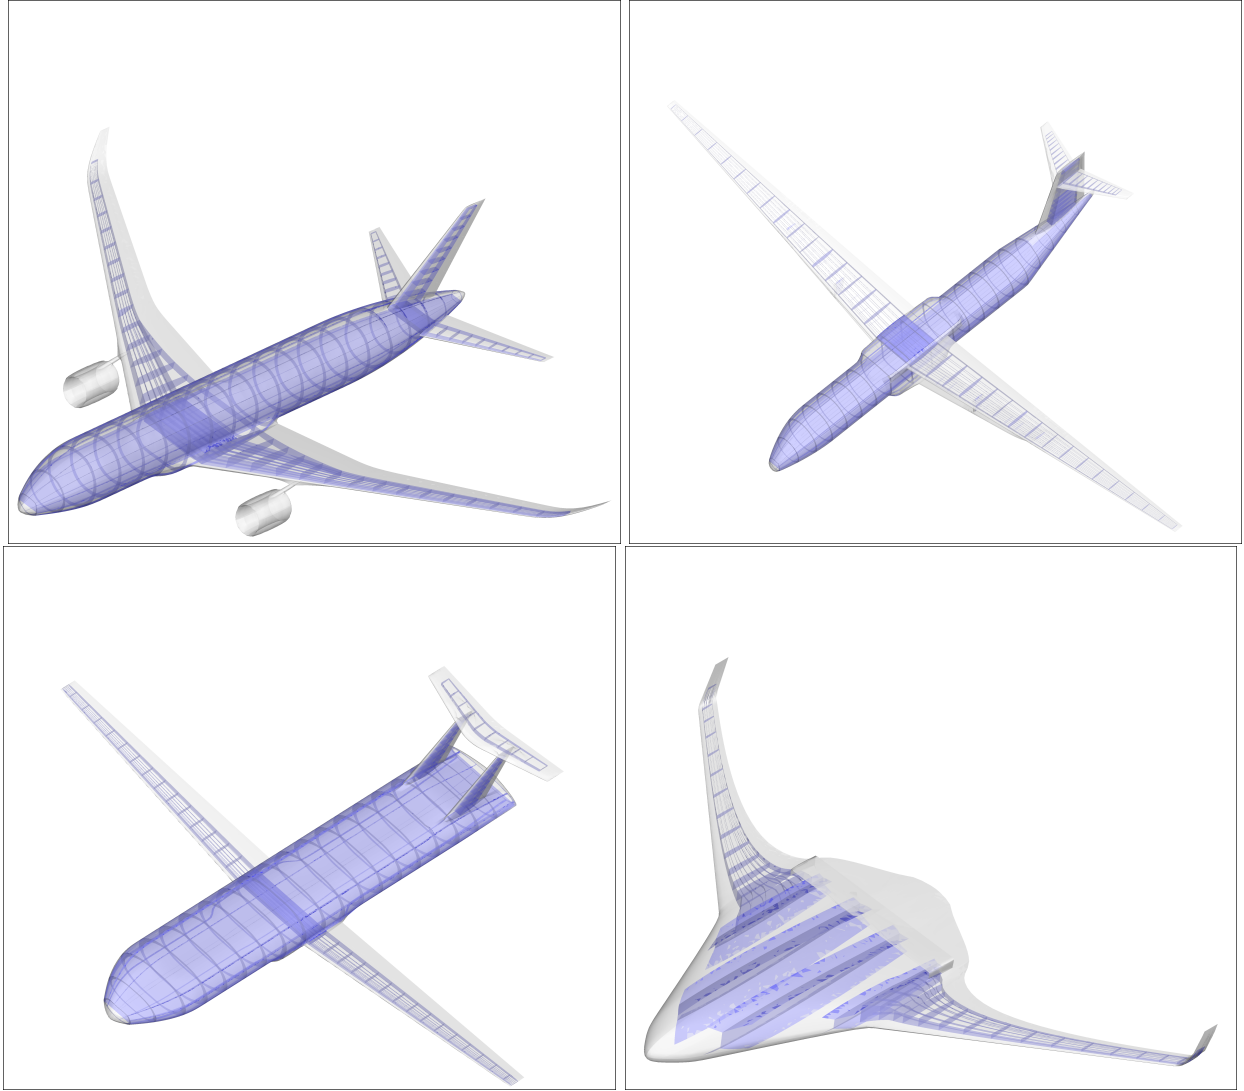
Figure 1: Geometry and structural models for various configurations generated using Geo-MACH.

3   Overset CFD for aerodynamic shape optimization

Having the capability to quickly generate aircraft geometries suitable for high-fidelity aerodynamic and structural analysis is not enough. Creating meshes for the aerodynamic and structural solvers is another major bottleneck in the overall workflow. The generation of the structural meshes was addressed within the GeoMACH project [11], but the automatic generation of CFD meshes is much more complex and was not proposed or addressed in that project.

Our CFD solver, ADflow, is a Reynolds-averaged Navier–Stokes (RANS) structured multiblock solver with an efficient adjoint solver that enables aerodynamic shape optimization with respect to large numbers of variables; it has been used extensively for both aerodynamic shape [10, 12, 13, 33, 34] and aerostructural design optimization of aircraft [15, 24, 35, 36].

A fully automatic generation of multiblock RANS grids for full aircraft configurations is currently not possible, and would be extremely challenging. Furthermore, in many practical applications, creating a face-matched multiblock mesh is not possible, or results in a low quality mesh. Even a relatively simple wing-body-tail configuration without engines, flap track fairings, or deployed high-lift system requires substantial user effort to generate a structured multiblock mesh of acceptable quality. This is the motivation to add overset capability to ADflow. In addition, to retain the efficient design optimization capability, the adjoint gradient computation should also be implemented in the overset version of ADflow. While a few examples of aerodynamic shape optimization exist [37, 38], those efforts did not result in further published work, there have not been published demonstrations on full aircraft configurations using RANS.

The overset approach consists in combining multiple structured meshes in a unstructured way. Interpolation is performed at the overlap regions to transfer information between clusters of structured meshes. One important advantage is that we can generate specialized meshes for individual components which can be then overlapped to constitute complex configurations. We use hyperbolic extrusion methods [39, 40] to automatically generate high-quality meshes for each geometry component, which speeds up the mesh generation.

One key component in the overset capability is the computation of the overset mesh connectivity, i.e., which cells are actually used for computation (compute cells), which cells should interpolate values from other mesh clusters (interpolate cells), what the interpolation weights are for those cells, and which cells should be excluded from the analysis (blanked cells), either because they are outside of the domain of the problem or because they have no impact on the overall solution.

The implicit hole cutting (IHC) technique [41] is an automated way of overset connectivity generation. This methodology preserves smaller cells, which usually are close to viscous walls, while larger cells become either interpolate or blanked cells. We implemented the IHC functionality in ADflow to minimize effort on the user side when setting up aerodynamic analysis cases. Even though IHC is well known in the CFD community, we spent substantial effort to efficiently parallelize the algorithm, making it feasible for optimization applications.

Figure 2 illustrates the overset mesh approach and IHC, where we can see that there are active cells inside the bodies. These cells should be blanked since they are outside of the domain of interest and could contaminate the solution in the exterior cells. We implemented a flooding algorithm [42] in ADflow to blank these cells. A cell of a given mesh cluster is flagged as a flood seed if it intersects walls defined in other clusters. Then, during the flooding process, the flood seed cells blank adjacent compute cells. The flooding propagates from compute cell to compute cell within the set of interpolated cells, and all flooded cells become blanked cells. The final mesh of the example after the flooding process is shown in Figure 3. Additional details on the implemented IHC are described by Kenway et al. [8].

Numerical integration of flow variables over surfaces is a key component required in CFD to obtain important quantities such as lift and drag. However, we do not get a watertight surface mesh
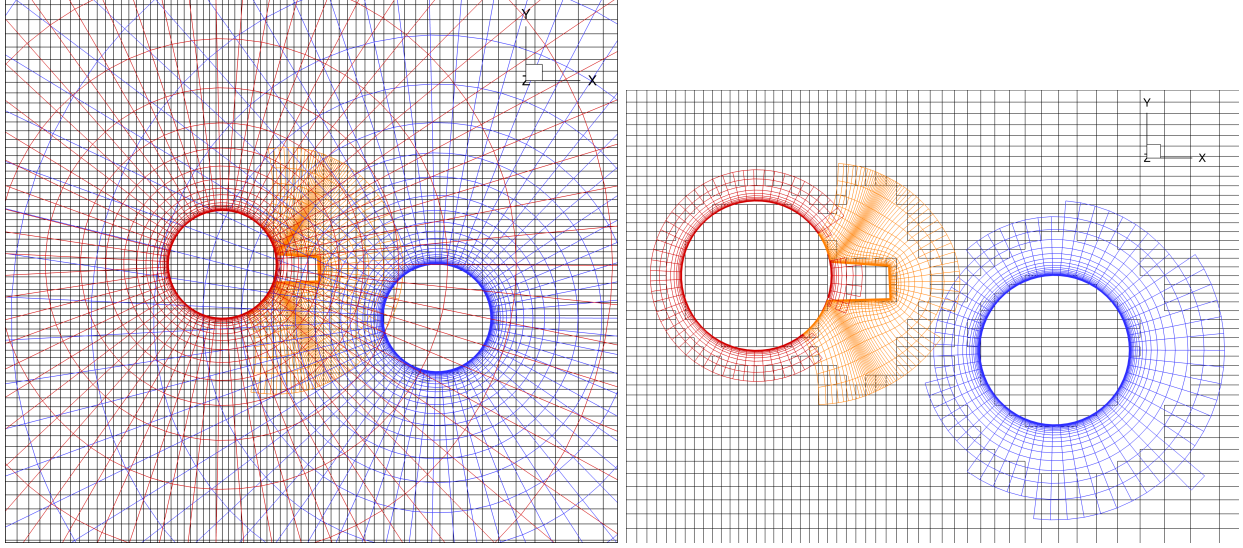
Figure 2: Example of overset connectivity computation starting with four different meshes (left), and after IHC is applied (right). The four different meshes are: background (black), left circle (red), rectangular appendage (yellow), and right circle (blue). The IHC method preserves the small cells near walls. Note that some cells of the background mesh that are inside the cylinders are still present, but will be blanked by the flooding procedure.
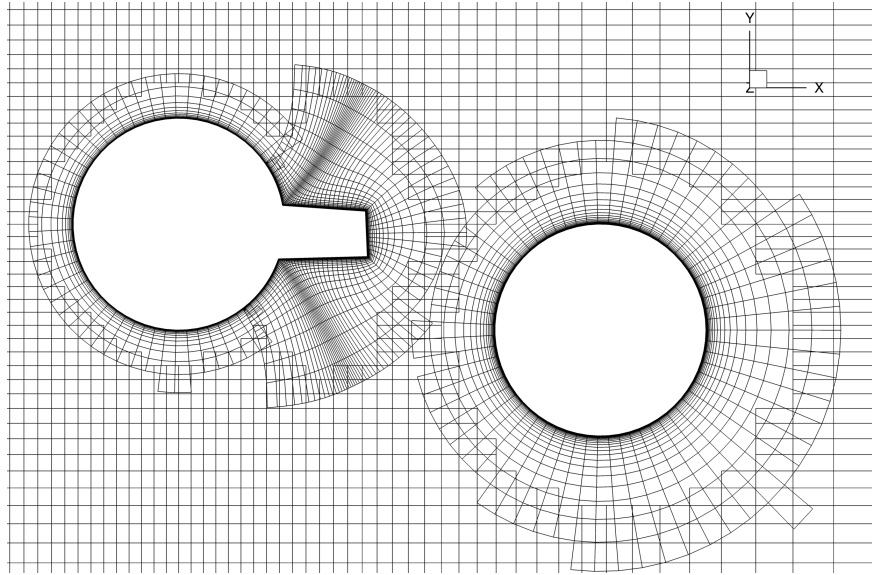
Figure 3: Final overset connectivity; only compute cells are shown.

(which is required for the surface integration) after the IHC process due to cell overlap requirements for the interpolations. To be able to perform integrations over the surfaces, we implemented a fully automated *zipper mesh* approach [43], which creates a watertight surface by triangulating gaps between overlapping surfaces. The main steps of the fully automatic zipper mesh algorithm implemented in ADflow are shown in Figure 4 for a simple sphere example. To illustrate the use of the overset capability in ADflow, we show the wing-body Common Research Model configuration used for DPW6 in Figure 5.
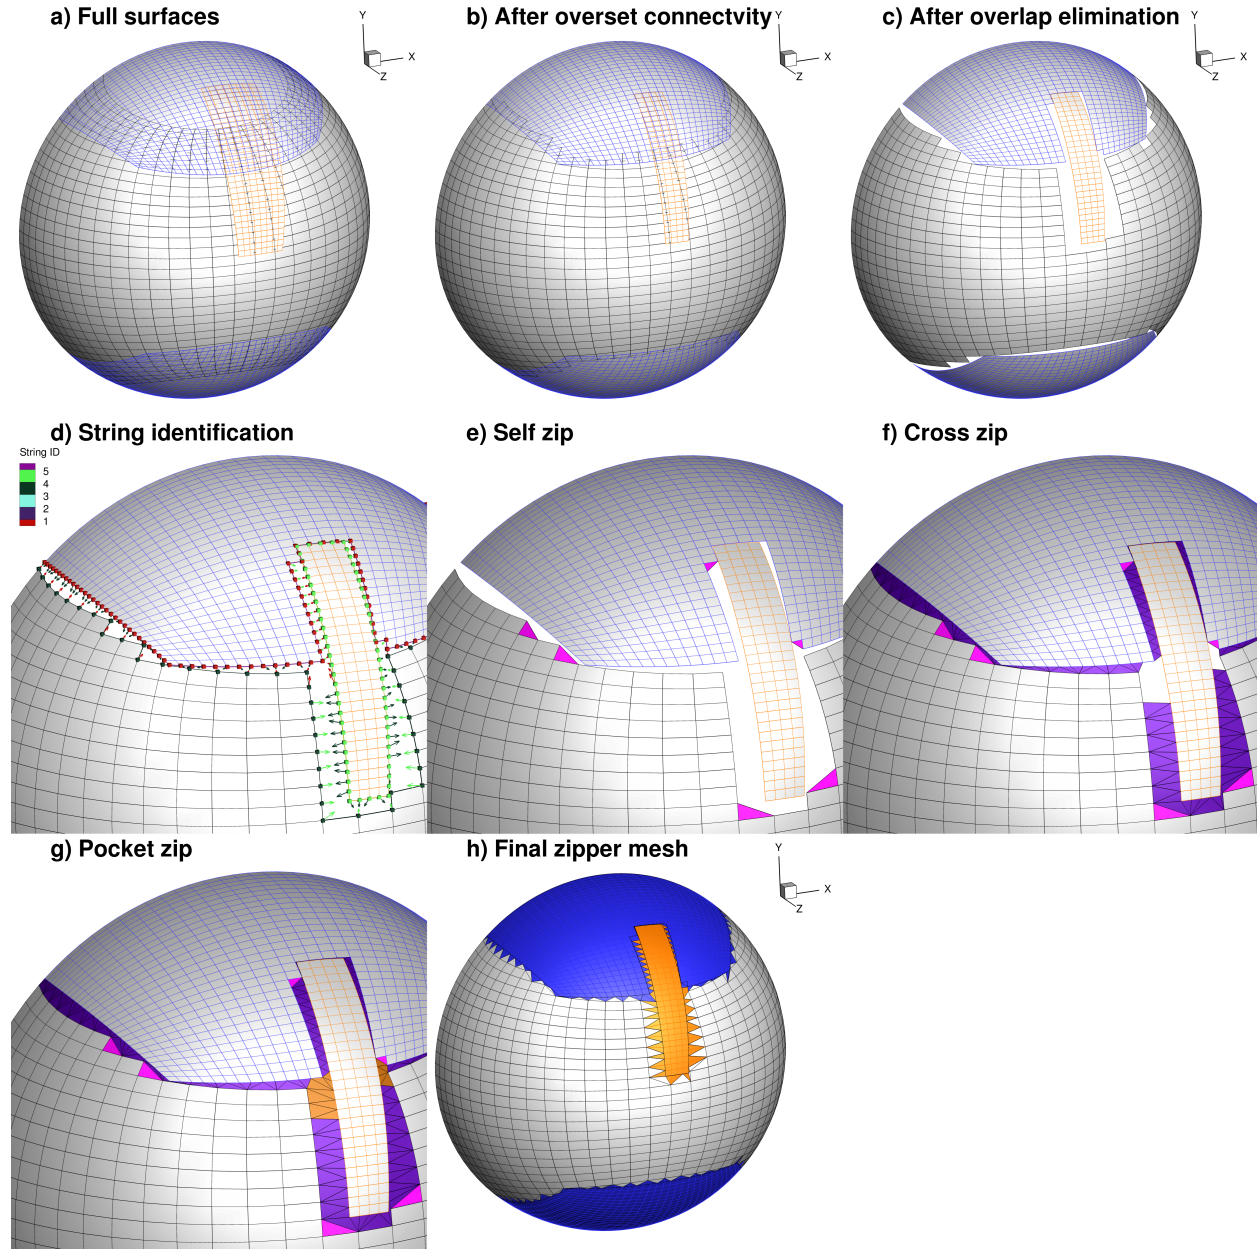
Figure 4: Overview of the zipper mesh procedure.

After implementing and testing the new ADflow capability for solving using overset grids, we added the adjoint derivative computation capability in order to be able to perform design optimization. Previous work already demonstrated the feasibility of overset-based aerodynamic shape
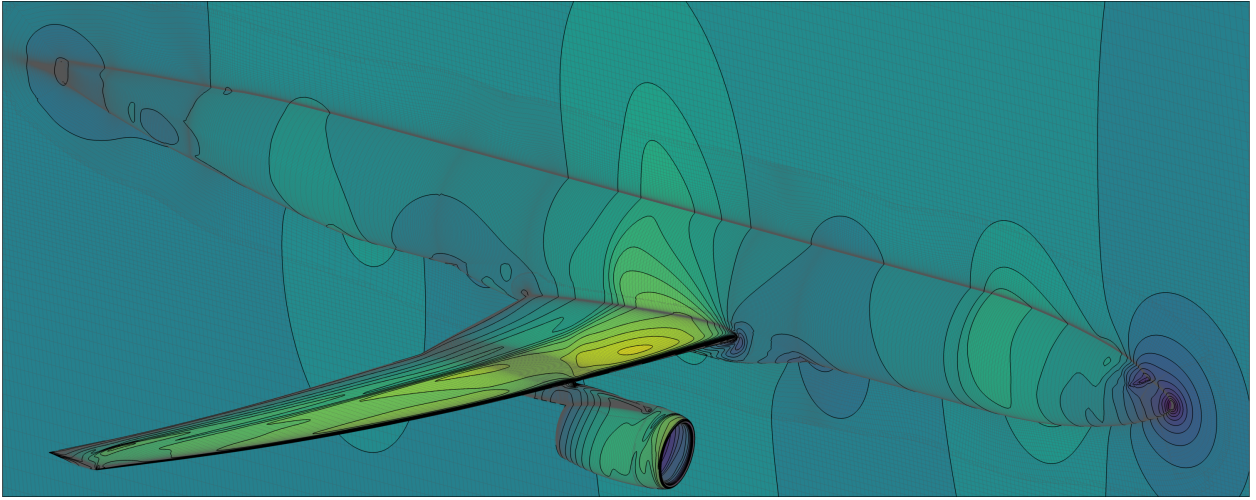
Figure 5: Wing-body overset mesh application including engine nacelles.

optimizations. Lee and Kim [37] used the overset formulation while solving the Euler equations to optimize the DLR-F6 wing, while computing derivatives with the discrete adjoint method. Liao and Tsai [38] performed aerodynamic shape optimization of dual airfoils and turbine stators with overset meshes using continuous adjoint formulation for Euler equations. The overset connectivity was recomputed on every optimization iteration.

For the first implementation of the adjoint method for overset meshes in ADflow, we employed what we call *frozen overset connectivity*, that is, the overset connectivities and interpolation weights are computed before starting the optimization and assumed to remain fixed. This version was used for the optimizations shown by [8].

The frozen overset connectivity approach has the advantages that the computational cost is low due to the elimination of the connectivity updates, and that it yields a smoother design space since it eliminates discontinuities that arise when the connectivity is updated. This makes the frozen connectivity approach particularly attractive for optimization. However, for frozen connectivity approach to give accurate solutions, it is necessary to deform the entire mesh in a consistent manner to ensure that overlapped cells remain in the same relative location to each other. This can be accomplished by using the unstructured-inverse distance mesh deformation approach that we recently developed based on the method described by [44]. This approach modifies the nodal positions of all overlapping meshes simultaneously, thus preserving their relative positions.

The implementation of the adjoint method required three main additions: 1) Differentiation of the zipper mesh surface integration routines, 2) an overset halo exchange routine to the differentiated code, and 3) the explicit linearization of the interpolated cells with respect to their real donor cells in the Jacobian matrix of the residuals with respect to the flow state variables. The resulting adjoint derivative computation was verified against the complex-step method and shown to be as accurate as the original multiblock adjoint.

This overset version of ADflow already proved invaluable in another NASA project, where we performed aerostructural and engine integration optimization for the D8 aircraft configuration, as shown in Figure 6.

The assumption of frozen overset connectivity might be unfeasible for cases in which we need to move geometry components and warp their associated meshes in an independent manner, such as engine-airframe integration, or the placement of the struts in a strut-braced wing. This situation requires a complete overset connectivity update, and led us to implement a second version of the
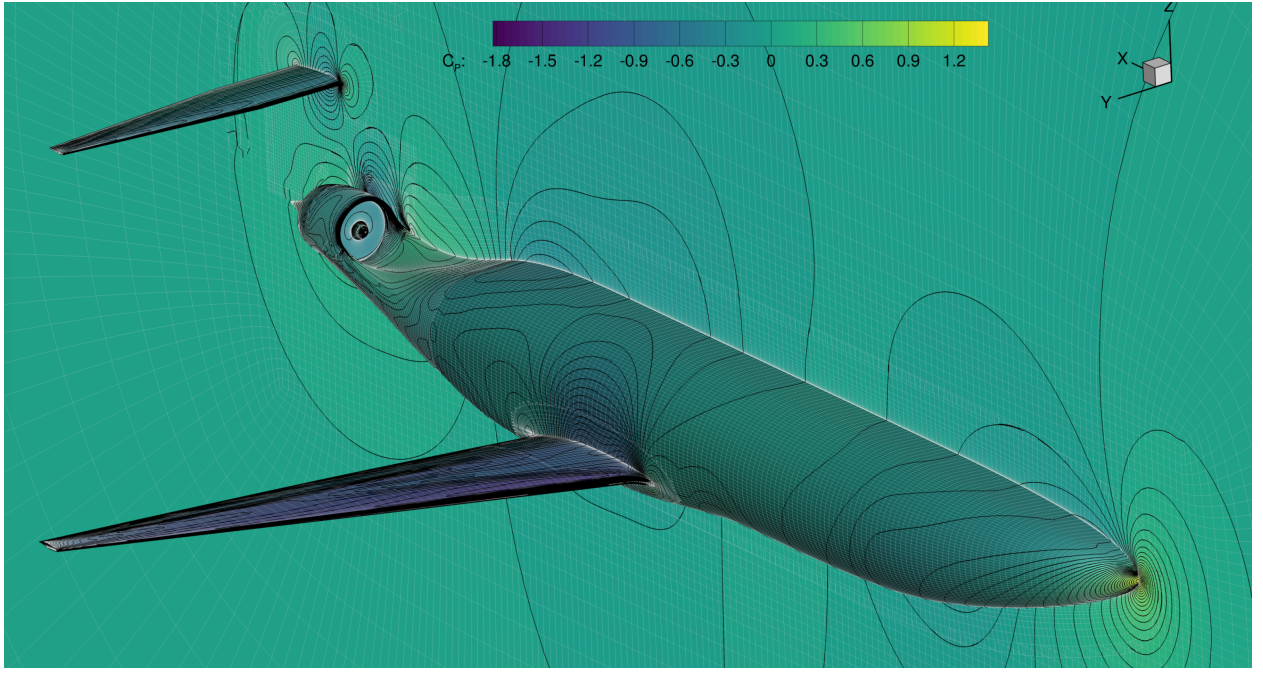
Figure 6: Overset mesh and ADflow solution for the D8 aircraft configuration.

overset method adjoint in which we allow connectivity updates during each optimization iteration. As mentioned above, this poses some challenges to gradient-based optimization, since it introduces discontinuities in the functions of interest. For instance, if we allow overset connectivity updates during the optimization process, the interpolation stencils of the interpolated cells might change. In addition, changes in overset connectivity in viscous surfaces modify the pattern of the zipper mesh, which is used for surface forces integration. These effects manifest themselves as noise in the functions of interest, such as drag coefficient. Figures 7, 8, and 9 illustrate how drag coefficient changes for a vertical wing translation with respect to the fuselage on a representative transonic airplane configuration.

The noise amplitude observed in some optimization cases is in the order of hundreds of drag counts, which is within the expected range of model and discretization errors. Furthermore, gradient-based optimizers can still make progress because the gradients consistently point towards the improvement direction. Discontinuities might hinder the optimizer performance near the optimum, because the noise generates trends inconsistent with the small magnitude gradients.

A third version of the adjoint overset code is still under development. This version will take into account the sensitivity of the interpolation weights with respect to the nodal coordinates during the derivative computation. This will make the linearized code closer to the original nonlinear code, which should improve the accuracy of the derivatives.
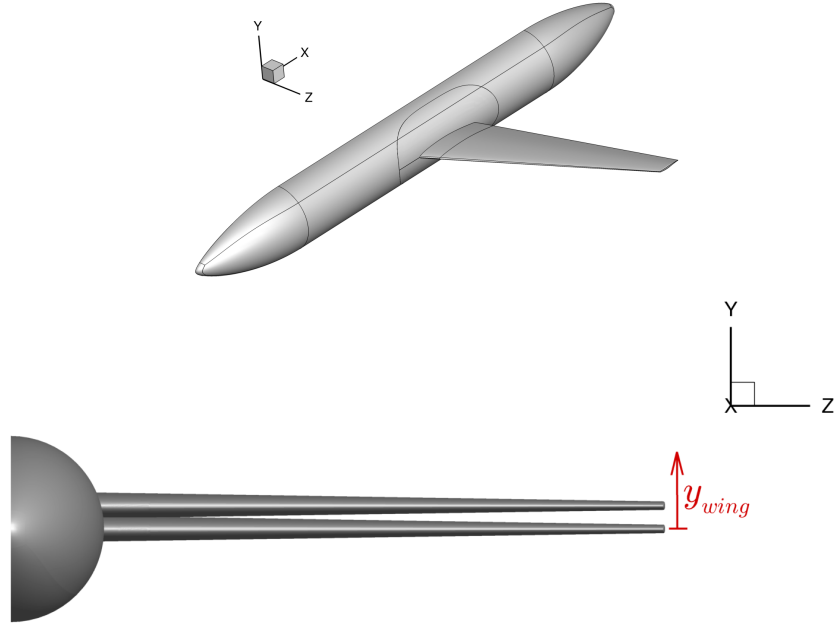
Figure 7: Transonic configuration used for wing translation study.
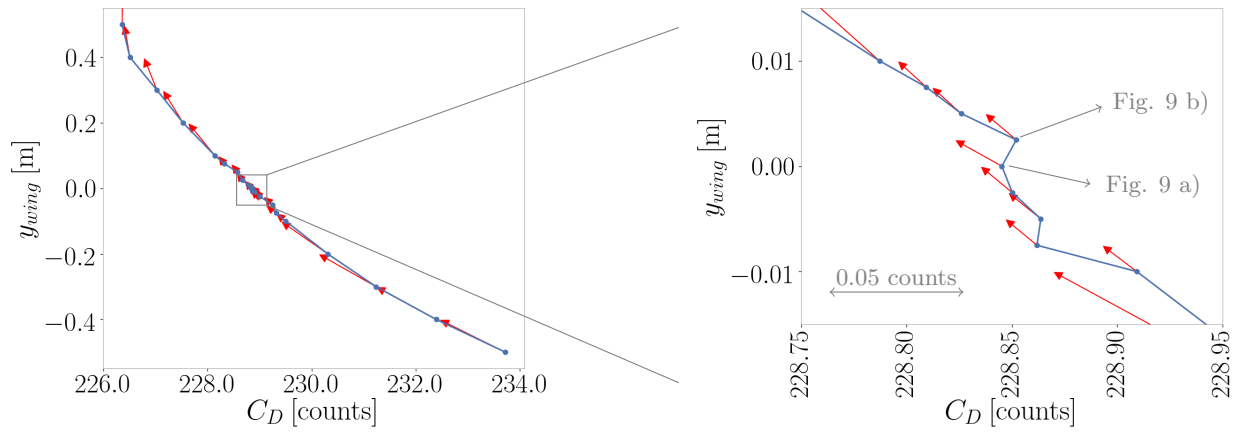
Figure 8: Drag variations due to the vertical translation of the wing on a transonic airplane configuration. The red arrows represent gradients. Noise is observed at small steps, but the gradients are still consistent with the overall trend. The noise is mainly associated with changes in the zipper mesh. The zipper mesh change illustrated in Fig. 9 caused the gradient inconsistency between the two indicated points.
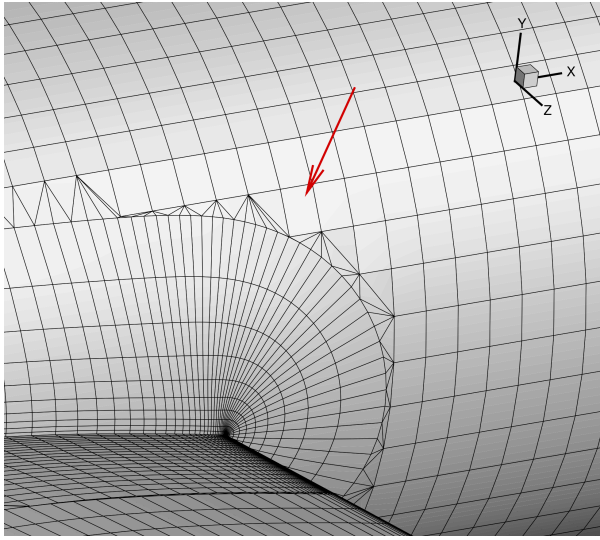
a) $y_{wing} = 0.0000$ m
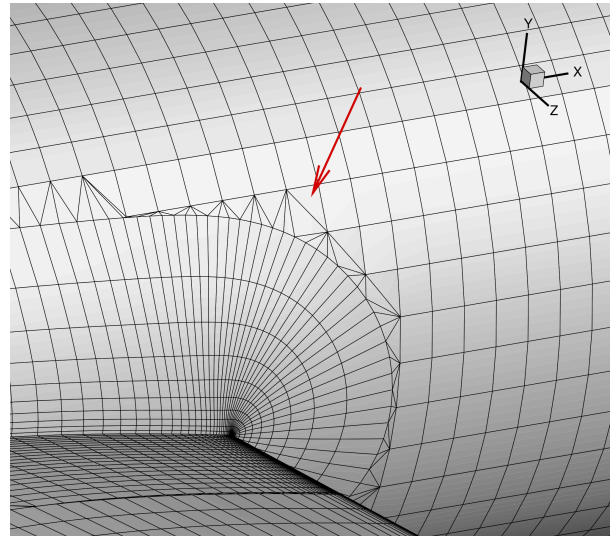
b) $y_{wing} = 0.0025$ m

Figure 9: Changes in the zipper mesh configuration causes noise in the drag curve shown in Fig. 8.

Part 2: High aspect ratio flexible wings

This part of the project focuses on enabling high aspect ratio flexible wings. Large spans decrease induced drag, but also incur a structural weight penalty. The overall goal of this project is to find ways to mitigate this weight penalty. Two airframe technologies are considered to enable this: advanced materials and continuous morphing trailing edge. In addition, we developed the formulation of design optimization constraints that are critical in the MDO of flexible wings: buffet onset, and flutter boundary constraints.

4 Optimal high aspect ratio wings considering material trade-offs

Current and future composite material technologies have the potential to greatly improve the performance of large transport aircraft. However, the coupling between aerodynamics and structures makes it challenging to design optimal flexible wings, and the transonic flight regime requires high fidelity computational models.

We addressed these challenges by solving a series of high-fidelity aerostructural optimization problems that explore the design space for the wing of a large transport aircraft. We considered three different materials: aluminum, carbon-fiber reinforced composites and an hypothetical composite based on carbon nanotubes.

Table 1 lists the properties for the materials. Note that the properties for the CNT-based composite material are hypothetical; they are based on a tensile modulus of 1.2 TPa and a tensile strength of 6 GPa, while the remaining properties are scaled to match the corresponding composite material data. A composite material with such properties is currently not available, but we use it here to quantify the impact on wing design of improving material properties by an order of magnitude.

The methodology and results are only briefly described here; much more detail is provided in the paper presented in a special session organized by NASA [27], and a contract report written exclusively about this study [45].

For this study, we used both a medium-fidelity and a high-fidelity aerodynamic analysis coupled with a detailed finite element model of the wing. The medium-fidelity aerodynamics consisted of TriPan, an unstructured, three-dimensional parallel panel code for calculating the aerodynamic forces, moments and pressures for inviscid, incompressible, external lifting flows using the Prandtl–Glauert equation [46]. TriPan also includes an estimate of the viscous drag that uses sectional viscous drag data based on Mach number, Reynolds number, and airfoil thickness. The sectional viscous drag is integrated across the span to obtain the total viscous drag. This estimate models the effect of wetted area, t/c, and taper ratio on viscous drag. For the high-fidelity aerodynamics, we used the CFD solver ADflow, which is a second-order structured block-based finite-volume solver for the Euler, Navier–Stokes and RANS equations [47].

The structural analysis is performed using the Toolkit for the Analysis of Composite Structures (TACS), a parallel, finite-element code designed specifically for the analysis of stiffened, thin-walled, composite structures using either linear or geometrically nonlinear strain relationships [46].

The load and displacement transfer scheme follows the work of Brown [48]. The displacements from the structures are extrapolated to the aerodynamic nodes using rigid links. These rigid links are formed by locating the closest point on the structural surface to each of the aerodynamic nodes. The structural surface is determined by interpolating between structural nodes using the finite-element shape functions.

Efficient gradient-based optimization requires the accurate and efficient evaluation of gradients of the objective function and constraints. To achieve this, we use the coupled adjoint approach [49,

Table 1: Mechanical properties of the metallic, conventional composite and hypothetical CNT-based composite used in this study.

Parameter	Value	Units	Parameter	Value	Units
Aluminum material data					
E	70.0	GPa	ν	0.3	
σ_{YS}	420	MPa	ρ	2780	kg/m^3
Composite material data					
E_1	128	GPa	E_2	11	GPa
G_{12}	4.5	GPa	G_{13}	4.5	GPa
G_{23}	3.2	GPa	ν_{12}	0.25	
X_t	1170	MPa	X_c	1120	MPa
Y_t	40	MPa	Y_c	170	MPa
S	48	MPa	ρ	1522	kg/m^3
Carbon nanotube-based composite material data					
E_1	1200	GPa	E_2	120	GPa
G_{12}	45	GPa	G_{13}	45	GPa
G_{23}	32	GPa	ν_{12}	0.25	
X_t	6000	MPa	X_c	5000	MPa
Y_t	400	MPa	Y_c	1600	MPa
S	500	MPa	ρ	1522	kg/m^3

50] We have developed a coupled aerostructural adjoint that is based on analytic derivatives without the use of finite-difference computations. Further details of the approach for the medium-fidelity case are given in Kennedy and Martins [46], and the high-fidelity case is similar to Kenway et al. [24].

For this study we have developed what we called the Quasi-CRM (QCRM) geometry to serve as a baseline. The QCRM is a wing and wing-tail geometry that has a planform roughly the same as the Common Research Model (CRM) used in the Drag Prediction Workshops (DPW) 4 and 5 [51]. However, it should be emphasized that the QCRM geometry is not derived from the CRM geometry directly. The QCRM wing geometry is shown in Figure 10. The aircraft parameters are similar to those of a next-generation 777-sized aircraft.

We minimized a weighted combination of fuel burn and take-off gross weight (TOGW) over four prescribed missions. The contribution to the objective from each mission was weighted based on the proportion of flights flown for each distance. We computed the fuel consumption for each mission based on the Breguet range equation

$$\text{FB} = \text{LGW} \left(\exp \left(\frac{R\, \text{TSFC}}{V\, (L/D)} \right) - 1 \right), \tag{4.1}$$

where FB is the fuel burn for the entire mission, LGW is the landing gross weight, R is the mission range, TSFC is the thrust-specific fuel consumption, V is the cruise speed and L/D is the lift to drag ratio.

The structural design must satisfy a series of failure and buckling constraints at the two critical loading conditions such that the yield stress and buckling constraints at each loading condition are within the admissible failure envelope. In addition, we imposed steady-state lift and trim constraints at each of the 6 operating conditions. This is an idealization for the maneuver conditions since, in reality, a steady pull-up condition would require a finite turning radius.
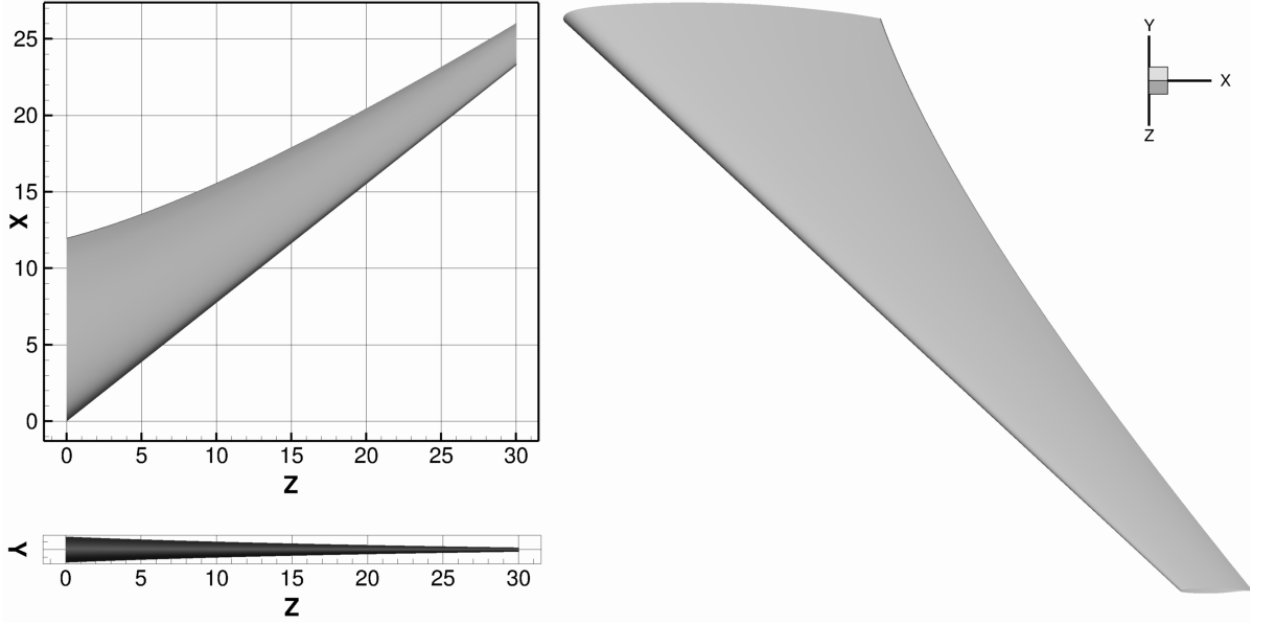
Figure 10: The Quasi-CRM (QCRM) wing geometry.

The design variables consist of geometric variables, structural variables, and aerodynamic design variables. The aerodynamic design variables consist of the angles of attack at each of the 6 operating conditions. The geometric variables consist of 4 twist variables distributed along the wing span, with a fixed root twist, one span-scaling variable, 5 vertical scaling variables, which modify the thickness to chord ratio, and a single chord scaling variable for the entire wing. Together the span and chord scaling variables admit a series of planforms that are stretched in the chord-wise and span-wise directions, but share similar geometric features. In addition, we employ 6 tail rotation angle variables in order to trim the aircraft at each flight condition.

The structural design variables consist of a top stiffener pitch and bottom stiffener pitch variable. Each panel has 2 thickness variables for the skin and stiffener, respectively as well as a stiffener height. The stiffener base width is fixed based on the stiffener height. For the metal wing parametrization, each material is isotropic and no further information is necessary. For the composite wing parametrization, there are an additional 3 variables that define the ply-fraction for the skin. We fixed the ply-fractions within the stiffeners to a 0°-ply dominant laminate.

The aerostructural optimization problem can be summarized as follows:

$$
\begin{aligned}
& \text{minimize} && \sum_{i=1}^{4} w_i \left(\beta \text{FB} + (1-\beta)\text{TOGW}\right) \\
& \text{w.r.t.} && \mathbf{x} \\
& \text{such that} && \text{KS}_{\text{fail}}^{(i)} \leq 1 && i = 1, \dots, 6 \\
& && \text{KS}_{\text{buckling}}^{(i)} \leq 1 && i = 1, \dots, 4 \\
& && L^{(i)} = n_{(i)} W^{(i)} && i = 1, \dots, 6 \\
& && M^{(i)} = 0 && i = 1, \dots, 6 \\
& && \mathbf{c}(\mathbf{x}) = 0 \\
& && \mathbf{l} \leq \mathbf{A}\mathbf{x} \leq \mathbf{u}
\end{aligned}
\tag{4.2}
$$

where $\text{KS}_{\text{fail}}^{(i)}$ and $\text{KS}_{\text{buckling}}^{(i)}$ are the Kreisselmeier–Steinhauser aggregation functions [52?] for failure and buckling constraints, respectively; $L^{(i)} = n_{(i)}W^{(i)}$ is the lift constraint for each analysis condition; $n_{(i)}$ is the load factor; $M^{(i)} = 0$ represents the trim condition for each load condition; and $\mathbf{c}(\mathbf{x}) = 0$ represents all the consistency constraints.

We solved this problem for several values of β for all three materials using the medium-fidelity solver TriPan as the aerodynamic model. Figures 11–13 shows the skin thickness distributions for the top and bottom skins of the metallic, conventional composite, and CNT-based composite designs. The panel thickness for the metallic design exhibits slightly larger thicknesses for the lower skin compared to the upper skin. For the composite designs, the lower skin thickness is significantly higher than the upper skin thickness due to a combination of the buckling and failure criteria at the -1 g condition and the failure criterion at the 2.5 g condition. Note that the designs for each of the metallic, conventional composite and CNT-based composite exhibit consistent design trends between the TOGW ($\beta = 0$) and fuel burn ($\beta = 1$) objectives.
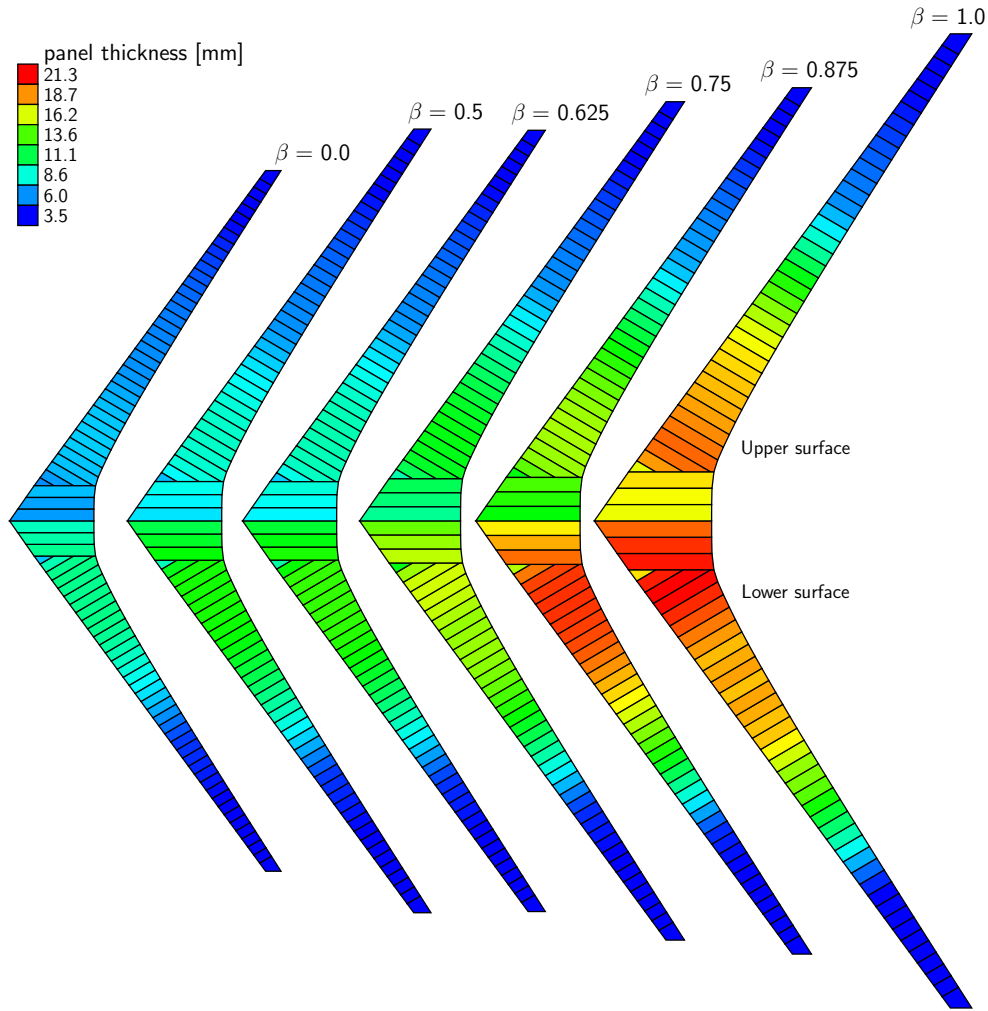
Figure 11: Skin thickness distributions for the metallic wing designs.

Figures 14 and 15 show the ply fractions in the wing skins for the composite and CNT-based designs. The reference axis associated with the ply angles is aligned along leading edge span of the wing. The laminate is fixed throughout the panel and the panel thickness is constant. A
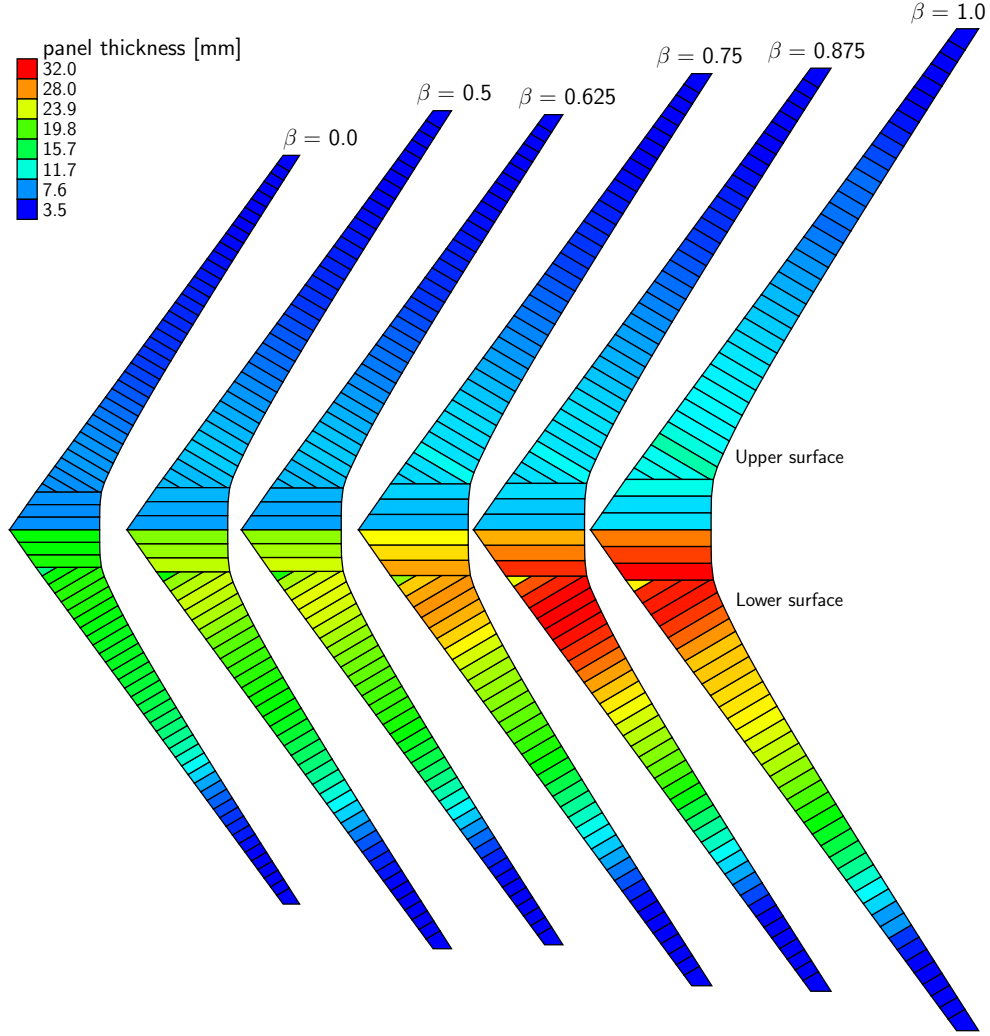
Figure 12: Skin thickness distributions for the conventional composite wing designs.

minimum ply-fraction constraint of 12.5% is enforced for any ply. Note that the plot shows the relative composition of the panel as a proportion of the panel area. Since the panel area reduces in the spanwise direction, this is not meant to represent absolute values. For both the composite and CNT-based composite wings, the designs with larger wing spans incorporate greater fractions of 0° plies. The spanwise variation of the ply fractions is also significant. All designs tend to utilize greater fractions of 0° plies near the root and larger fractions of ±45° plies towards the wing tip. This design feature reduces the tip twist of the wing under load and offsets the loss of torsional stiffness from the smaller, thinner tip sections.

The aerostructural optimizations for various values of β enable us to plot Pareto fronts with respect to structural weight and fuel burn as shown in Figure 16. This is the weighed objectives approach. Another alternative to finding points on the Pareto front would be to constrain one of the objectives while minimizing the other. We chose to use the weighed objectives approach method for two reasons: 1) Both fuel burn and takeoff gross weight (mass) have the same units so it is natural to try to combine them, and 2) we did not know a priori which TOGW and fuel burn values would be of interest along the Pareto front, so constraining one and minimizing the other would have been challenging. Note that this is really a smaller portion of a Pareto front between
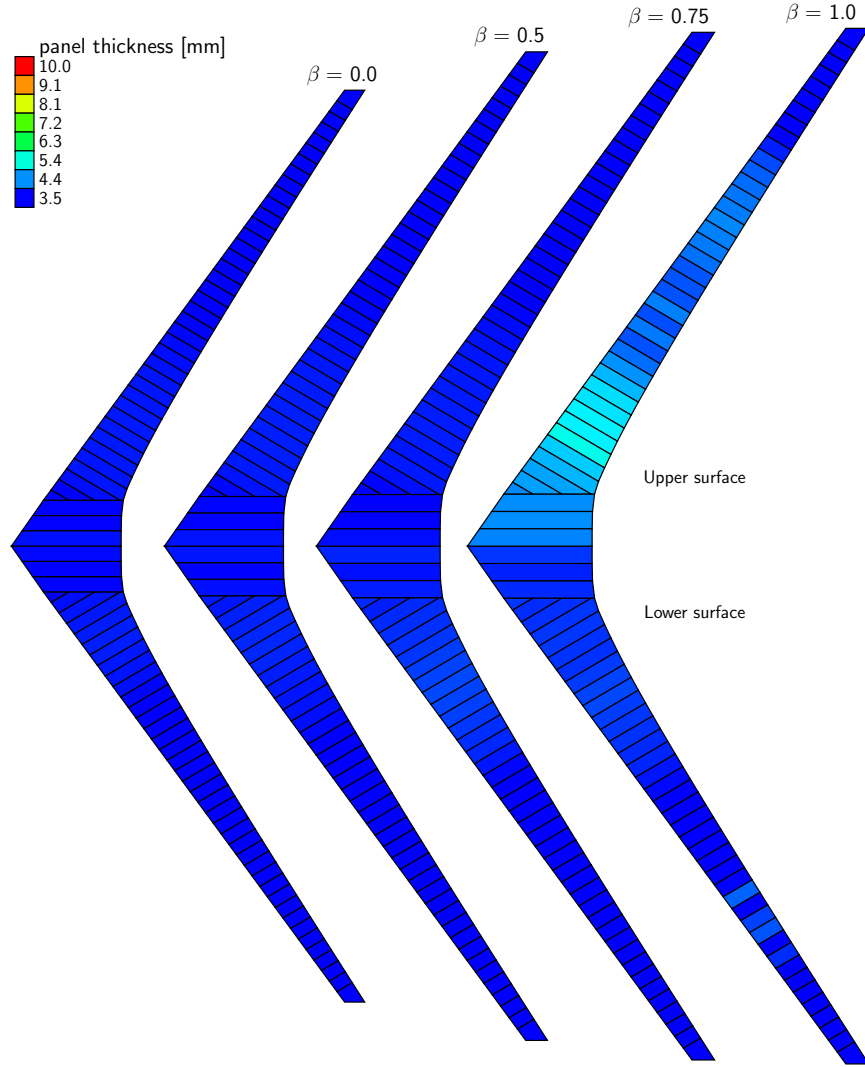
Figure 13: Skin thickness distributions for the CNT wing designs.

drag and structural mass.

From these Pareto front trends, it is clear that the use of advanced materials leads to simultaneous fuel burn and structural weight improvements. The uneven spacing in the Pareto front points is due in part to optimizer convergence. As more advanced materials are used, the difference in performance between fuel burn minimization and TOGW minimized designs becomes smaller and the Pareto fronts become more compact. This is due to the larger weight savings, which reduces both fuel consumption and TOGW directly. The composite TOGW design ($\beta = 0$) is 8.4% lighter than the metallic baseline, and the CNT-based composite design is 5.6% lighter relative to the $\beta = 0$ composite design. For the fuel burn optimized designs ($\beta = 1$) the composite design exhibits a 7.7% fuel burn advantage compared to the metallic design, and the CNT-based composite exhibits a further 5.2% fuel burn reduction compared to the composite $\beta = 1$ design. Thus, the improvements in performance of the CNT-based composite wings are relatively modest given that the assumed stiffness and strength are an order of magnitude higher than the conventional composite. This is in part because the weight savings were limited by minimum thickness constraints. Trends of the fuel burn and weight versus span, and the details of the resulting wing designs are
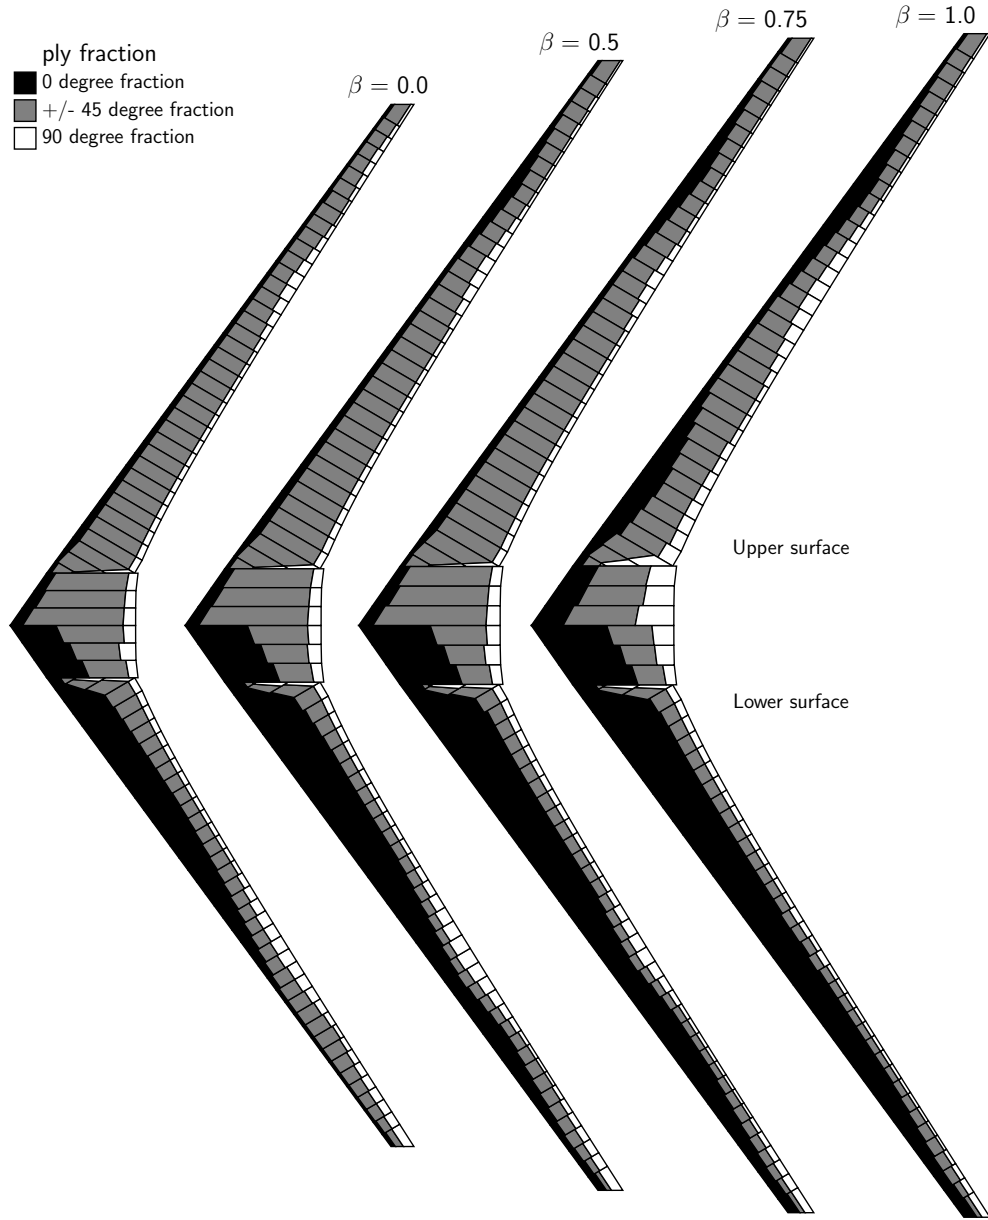
Figure 14: Ply fractions for the conventional composite wing designs.

presented in Kennedy et al. [27].

In Figure 16 we also show the results from a sequential design method for a series of fixed spans with metallic wings. In the sequential approach, we repeatedly perform design iterations where the structure is sized at fixed aerodynamic loads to obtain a structural weight estimate, followed by an aerodynamic optimization at fixed structural weight to minimize the fuel burn. This design process ignores the impact of aeroelastic deformation, but provides a useful comparison of the fully integrated optimization approach with a sequential design method. The span for a given β was fixed to the corresponding aerostructural optimum value because an aerodynamic optimization would tend to increase the span as much as possible, while a structural optimization would tend to decrease it as much as possible. Therefore, the span variable would not converge using this sequential process. There is a significant difference between the performance of the wing obtained
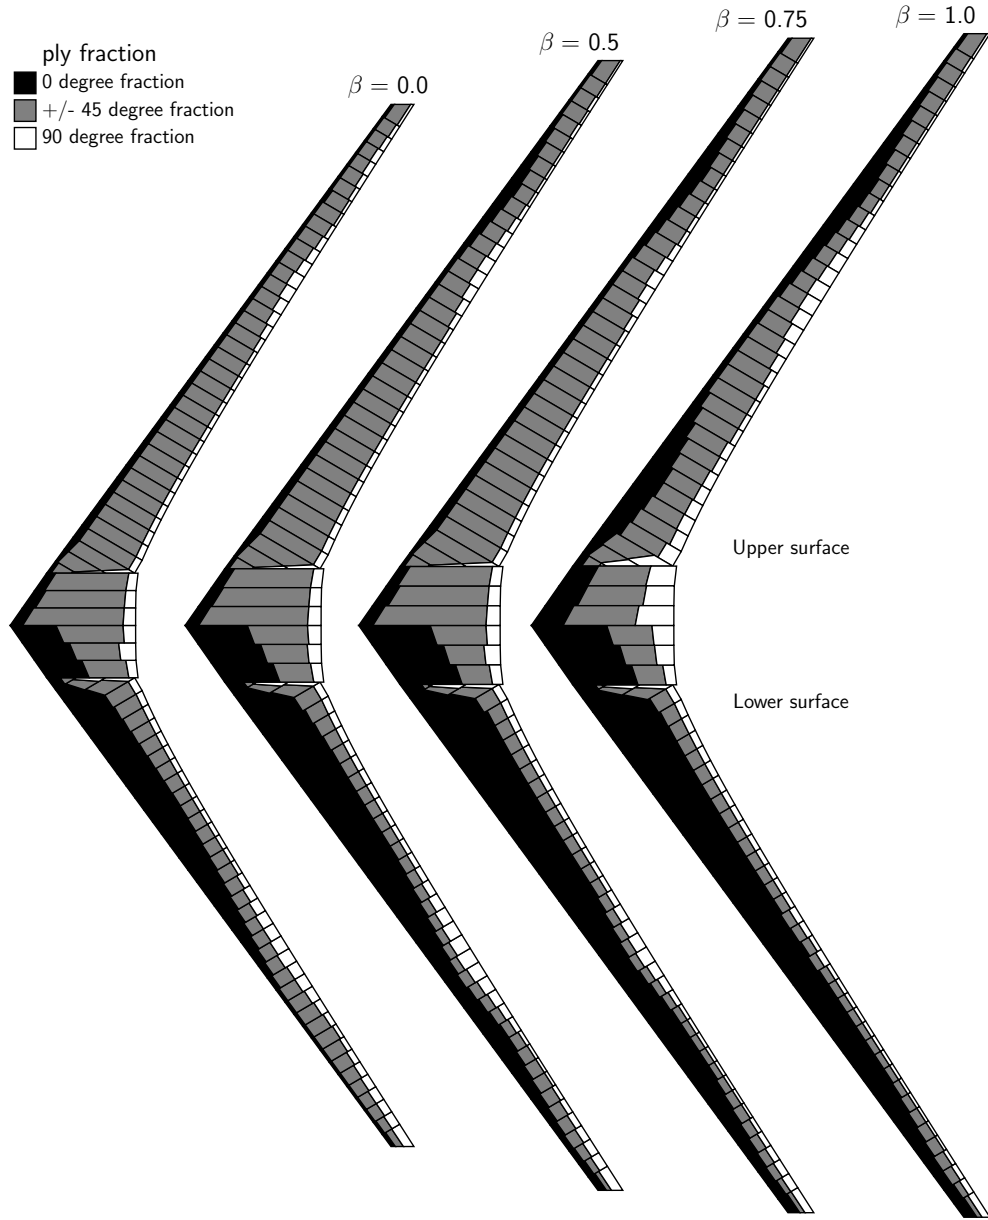
Figure 15: Ply fractions for the CNT wing designs.

from a sequential design method and the aerostructural optimization results. These sequential optimization results are in agreement with previous low-fidelity aerostructural optimization results, where the sequential approach was shown to be inferior relative to the MDO approach [53, 54]. This indicates that when considering the next generation of aircraft it is important not only to invest in airframe technologies, but also invest in the development of new design methodologies.

As a second stage of this study, we performed RANS-based aerodynamic shape optimization of designs obtained from the medium-fidelity aerostructural design optimizations presented above. The goal was to optimize the airfoil cross sections to improve the transonic performance without compromising the previously optimized structural design. Two optimizations at the extremes of the Pareto front were performed: one for TOGW ($\beta = 0$) and one for fuel burn ($\beta = 1$). Table 2 shows the data for the TOGW and fuel burn optimizations, and Figure 17 compares the pressure
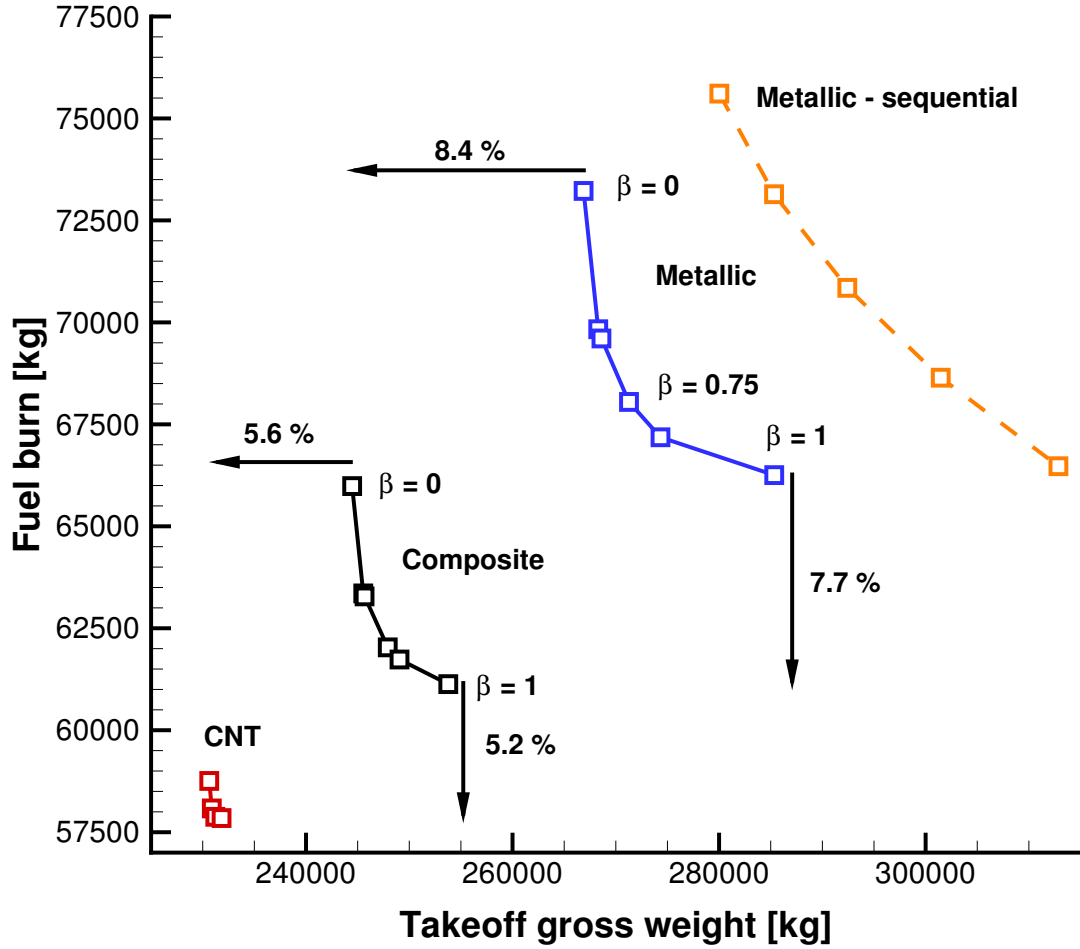
Figure 16: The Pareto fronts of fuel burn and takeoff gross weight shows the advantage of the conventional composite and CNT composite materials, as well as the advantage of aerostructural design optimization over sequential optimization.

contours, airfoil shapes, structural thicknesses, spanwise lift distributions, twist distributions, and t/c distributions for these two optimizations.

As shown in Table 2, the TOGW for the $\beta = 0$ optimization is lower than the fuel burn optimization, but converse is not true: The fuel-burn optimized design has 0.7% higher fuel burn. The reason for this discrepancy is that these optimizations were performed at a fixed altitude, and therefore the wing is not free to fly at the best point in its drag polar. This was an important lesson derived from this study: if large changes in wing area or mass are expected during the optimization, altitude variation, or a surrogate for altitude variation should be included as a design variable in the optimization problem. In fact, we have since been including altitude as a design variables (see Sections 6, 7, and [6, 15, 16, 36, 55, 56]).

The choice of objective function has a dramatic effect on nearly every aspect of the optimized design. The wing planform is the first major difference that we notice in Figure 17. Both designs have roughly the same aspect ratio, but the fuel burn design wing span is 7 m longer and has a 21.4% larger wing area. Since the fuel burn objective places less emphasis on the empty weight, the wing span extends to lower the span-loading and thus lower the induced drag. This increase

Figure 17: Comparison of TOGW and fuel burn optimized designs for the metallic wing.

Table 2: TOGW and fuel burn results for RANS-based aerostructural optimization

β	TOGW (kg)	Fuel Burn (kg)	Wing Mass (kg)	L/D	Span	Wing Area (m^2)	AR
0.0	268 630	74 242	26 567	20.13	73.8	431.348	12.62
1.0	290 206	74 754	47 631	21.87	80.8	523.468	12.45

in span comes at a substantial penalty in terms of structural weight, 21 064 kg or 79% of the TOGW optimized mass. This large wing mass is further explained by examining the distribution of thickness to chord ratio distributions. The average t/c for the TOGW minimization is 33% higher than for the fuel burn design. However, the difference in the wingbox depth is slightly lower than this due to the larger chord of the minimum fuel burn wing. Further confirmation of the increased wingbox mass can be seen in the upper skin thickness distribution contours in Figure 17. With the exception of the more lightly loaded wing tips, the fuel burn design skin is thicker than the minimum TOGW design.

The cruise lift distributions for both minimum fuel burn and minimum TOGW designs shown in Figure 17 are close to elliptical, but the TOGW design yields a closer match. However, there is a very large difference in the shape of the maneuver lift distributions. Both designs exhibit passive load alleviation—where the 2.5 g maneuver lift distribution is shifted inboard relative to the elliptical distribution—but this load alleviation is more pronounced for the minimum TOGW design, since it emphasizes weight reduction. For the minimum fuel burn design, the corresponding twist distribution shows the additional passive aeroelastic wash-out that occurs at the maneuver condition. This additional downward twisting reduces the tip load, shifting the distribution inboard and lowering the bending moment on the structure. This behavior is consistent with the medium fidelity results.

The 2.5 g maneuver lift distribution for the TOGW design shows a completely opposite twist behavior: the wing twist actually decreases under the higher loading condition. Even so, the maneuver lift TOGW maneuver lift distribution shifted further inboard and is even more favorable from a structural perspective. To explain this phenomenon, we examined the three-dimensional flow field of the TOGW 2.5 g maneuver condition. At this condition, we observed that a large portion of the wing has stalled, resulting in large region of separated flow. Although stall is not desirable within the flight envelope (especially tip stall), the optimizer exploited the fact that no stall constraint is imposed to implement an extremely effective way to alleviate the loads. This is an excellent example of the optimization exploiting a weakness in the problem formulation and provides a valuable lesson. Motivated in part by this issue, we developed the separation-based buffet constraint formulation described in Section 7 and in Kenway and Martins [6].

Overall, we conclude that composite wings consistently resulted in lower fuel burn and lower structural weight, and that the carbon nanotube composite did not yield the increase in performance one would expect from a material with such outstanding properties. This was in part due to the minimum structural thickness constraint. A future study that accounts for nanocomposite manufacturing architectures and more realistic structural properties is recommended once this technology is further developed and this information becomes available. In addition, we have since developed the capability to perform RANS-based aerostructural optimizations more efficiently, to include buffet constraints (see Section 7), and to include the altitude as a design variable, so we would be able to consider utilizing these developments in future studies. For all materials, the minimum fuel burn wings were found to be longer, heavier, thinner, more flexible, and more lightly loaded than their minimum TOGW counterparts. Finally, based on a comparison with sequential designs, we concluded that it is equally important to invest in new airframe technologies and in

new design optimization methodologies.

We have been using the knowledge gained in this project in a more recent NRA on high-fidelity aerostructural design optimization of tow-steered wings, which uses exclusively RANS for the aerodynamics, and includes the buffet constraint formulation described in Section 7 and by Brooks et al. [36, 55].

5 uCRM—High aspect ratio undeflected Common Research Model

5.1 Introduction

The NASA Common Research Model (CRM) has served as a useful aerodynamic benchmark for drag prediction and optimization [34, 57, 58]. The model was originally conceived as a purely aerodynamic benchmark and as such the geometry of the wing was designed to take the shape of the deflected wing at a 1.0 g flight condition. There has been growing interest in extending this model to aeroelastic studies as well. However, due to its predefined in-flight shape, the model is not suitable for aeroelastic analysis and design as is. To address this, we defined the undeflected Common Research Model (uCRM), which includes the outer mold line (OML) geometry of the undeflected wing and the corresponding internal wing box structure, with the goal of achieving the flying shape of the CRM under nominal flight conditions. The topology of the wing box is designed to be similar to that of a Boeing 777. The jig shape was generated through an inverse design procedure where the objective was to minimize the difference between the 1.0 g aerostructurally deflected shape and the CRM. Since the original CRM has an aspect ratio of 9 we refer to this model as the uCRM-9. Since modern transport aircraft are trending toward higher aspect ratio wing designs in an attempt to reduce induced drag, and therefore fuel burn, there was a need for a higher aspect ratio variant of this model to assess the design potential of new future technologies. This variant, uCRM-13.5, has a larger aspect ratio of 13.5 and is defined through an aerostructural optimization of the uCRM-9 with considerations for buffet. The goal for both of these models is to define a useful benchmark for aeroelastic wing analysis and design optimization studies. We will start by describing the uCRM-9 geometry and structural wing box model. We will then describe the procedure for developing the uCRM-13.5. A summary of the key geometric parameters of the uCRM-9 and uCRM-13.5 planforms is given in Table 3.

Table 3: uCRM specifications (metric)

Parameter	uCRM-9	uCRM-13.5
Aspect ratio	9.0	13.5
Span	58.76 m	72.00 m
Root chord	13.62 m	16.02 m
Side of body chord	11.92 m	13.85 m
Yehudi chord	7.26 m	5.36 m
MAC	7.01 m	5.36 m
Tip chord	2.73 m	1.915 m
Wimpress area	383.74 m^2	384.05 m^2
Gross area	412.10 m^2	465.22 m^2
Exposed area	337.05 m^2	377.45 m^2
1/4 chord sweep	35°	35°
Taper ratio	0.275	0.250
Gear post depth (777)	0.736 m	0.648 m

5.2 Definition for uCRM-9

The uCRM-9 wing model definition is based on the outer mold line (OML) of the CRM wing, which is the shape of the wing at 1 g at the CRM nominal flight condition ($M = 0.85$, $C_L = 0.5$ at 37 000 ft). No structural model or jig shape are provided in the original CRM model. Thus we create a wingbox structure and jig OML whose aerostructural solution at the nominal flight condition reproduces the original CRM OML. The CRM geometry does not provide any information
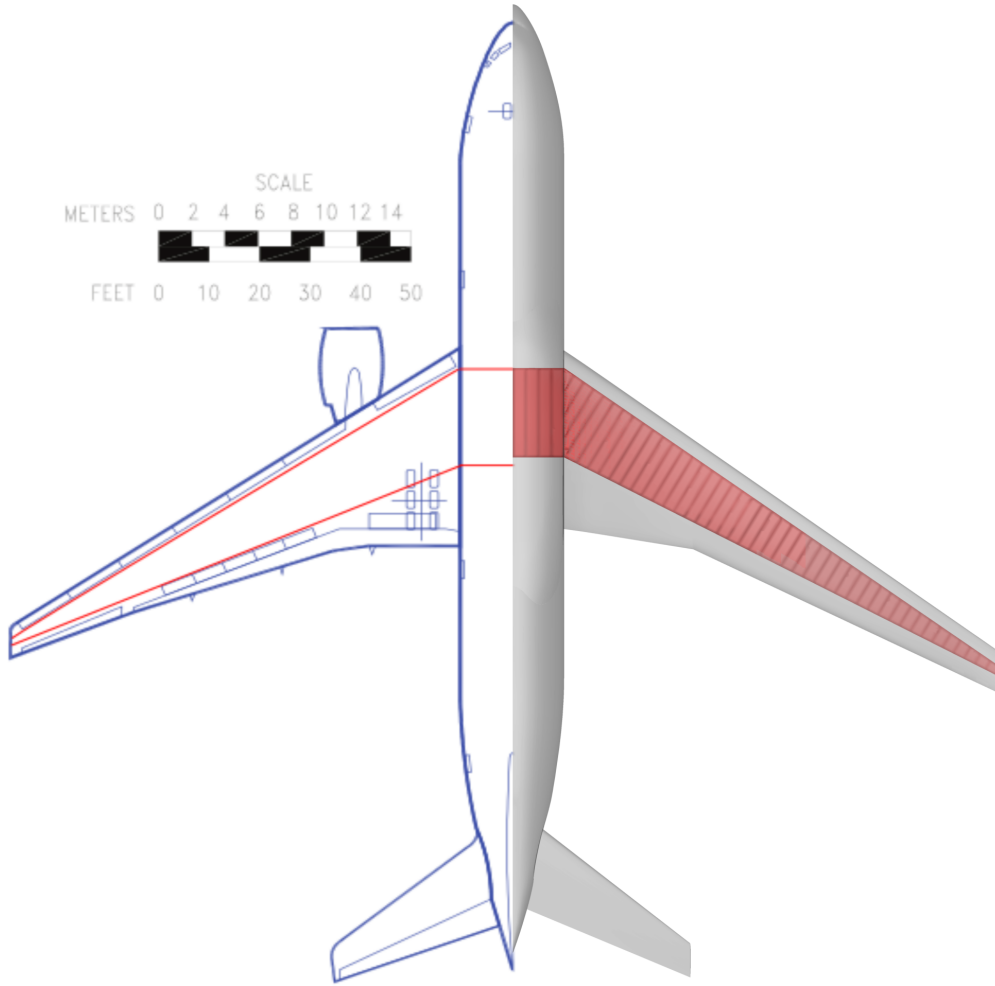
Figure 18: Boeing 777 (left) and CRM-9 (right). The CRM has a slightly smaller area wing and 3.5 deg more wing sweep than the Boeing 777.

with respect to the wing internal structure. To produce a structural model as representative as possible, we examined cutaway views of the Boeing 777 aircraft to determine the layout of the wingbox. The wingbox consists of two spars: a leading edge and trailing edge spar. Figure 18 shows the planform view of the 777-200ER extracted from the aircraft planning document [59] with a best-guess superimposed wingbox locations. This estimate of the geometry of the Boeing 777 wingbox could not be used directly in the CRM wing because the CRM wing has more sweep. Our approach was to generate a wingbox for the CRM using the same proportions (front and rear spar location, rib spacing, etc.) and structural layout of the Boeing 777 wingbox. Using digital versions of the Boeing 777 drawing, we measured the extents of the estimated wingbox at the root and tip locations as a fraction of local chord. We then rounded off these measurements to obtain the percentages of the locations with respect to local chord. This provided the required information to define the wingbox structure planform for the CRM geometry. The rib spacing for the uCRM wingbox is chosen to be approximately the same as for the Boeing 777, with a normal-distance spacing of 73.15 cm (28.8 in). Since the CRM wing has a slightly lower span than the Boeing 777, there are 44 ribs, not including the root close-out rib.

Having determined the planform of the wingbox, we used an in-house tool (`pyLayoutGeo`) to generate the geometry of the wingbox and its components. Based on the information given above,

Table 4: Material properties.

Parameter	Value
Density	2780 kg/m^3
Modulus	73.1×10^9 Pa
Poisson Ratio	0.33

we generated the wingbox structure shown in Figure 19. The meshing of this geometry was then performed by specialized grid generation software that can generate surface meshes.

A set of constant aerodynamic loads is required to find the undeflected shape of the wing. These aerodynamic loads were generated on the CRM wing-body-tail configuration using meshes from the 5th Drag Prediction Workshop, but were adapted for flight Reynolds number. We then determined the structural sizing for the wingbox geometry previously described using a structural optimization based on the aerodynamic loads. The objective of the optimization is to minimize the structural weight of the wingbox subject to material failure and buckling constraints. We assume that all structural panels are made of a 7000 series aluminum alloy with the properties listed in Table 4.

All main components of the wingbox are modeled using a blade stiffened panel approach. The remaining variables are free to change (within bounds) from one panel to the next. Figure 20 shows the panel definition distribution across the primary wing structure. The total mass of the wingbox is 13,416 kg. This initial optimized wingbox provides a guess for the structure of our jig wing; however, because it is based on the CRM OML, it also features the 1.0 g deflection in its geometry. This deflection will then be removed in the subsequent inverse design procedure such that the CRM wing geometry is recovered from the jig through structural deflection of the wingbox alone.

Given the 1 g OML of the original CRM, the wingbox structural layout and sizing, as well as the aerodynamic loads, we then determined the jig OML and wingbox. This is accomplished using the inverse design procedure described by Kenway et al. [24]. This is done by defining a least squares optimization problem to minimize the difference between points on the CRM geometry, X_T, and points on the deflected jig shape under the nominal 1.0 g loads X_{deflect}. The deflected points are found by taking the jig points and adding the structural deflections from the applied aerodynamic loads ($X_{\text{deflect}} = X_{\text{jig}} + u$). The wing outer mold line geometry is parameterized using free-form deformation (FFD) [60]. The optimizer changes the x, y, and z coordinates of the FFD control points such that the initial deflection of the CRM wing is achieved through structural deflection only. This inverse design procedure is not straightforward, since as the geometry of the structure changes, the structural deflections change as well. Once the FFD shape variables are found, they are then applied to the initial CRM geometry and wingbox to remove the deflection and achieve a new approximate jig geometry. The new structural geometry is then structurally resized using another structural optimization, and the procedure is repeated until convergence. This procedure required three inverse design cycles before a converged jig solution was achieved. Figure 21 shows the sequence of inverse design result leading to the finalized wing jig. The final wing jig and structural geometry were dubbed as the uCRM-9.

To verify the result, we compared an aerostructural analysis of the uCRM-9 to an aerodynamic analysis of the CRM. The shape and aerodynamic solution should match between these two cases. A comparison of the upper surface C_p contours is shown in Figure 22. Overall, the two sets of contours match well, although there are slight differences that can be attributed to the 0.04 deg angle of attack difference in the aerostructural solution. This difference in angle of attack was
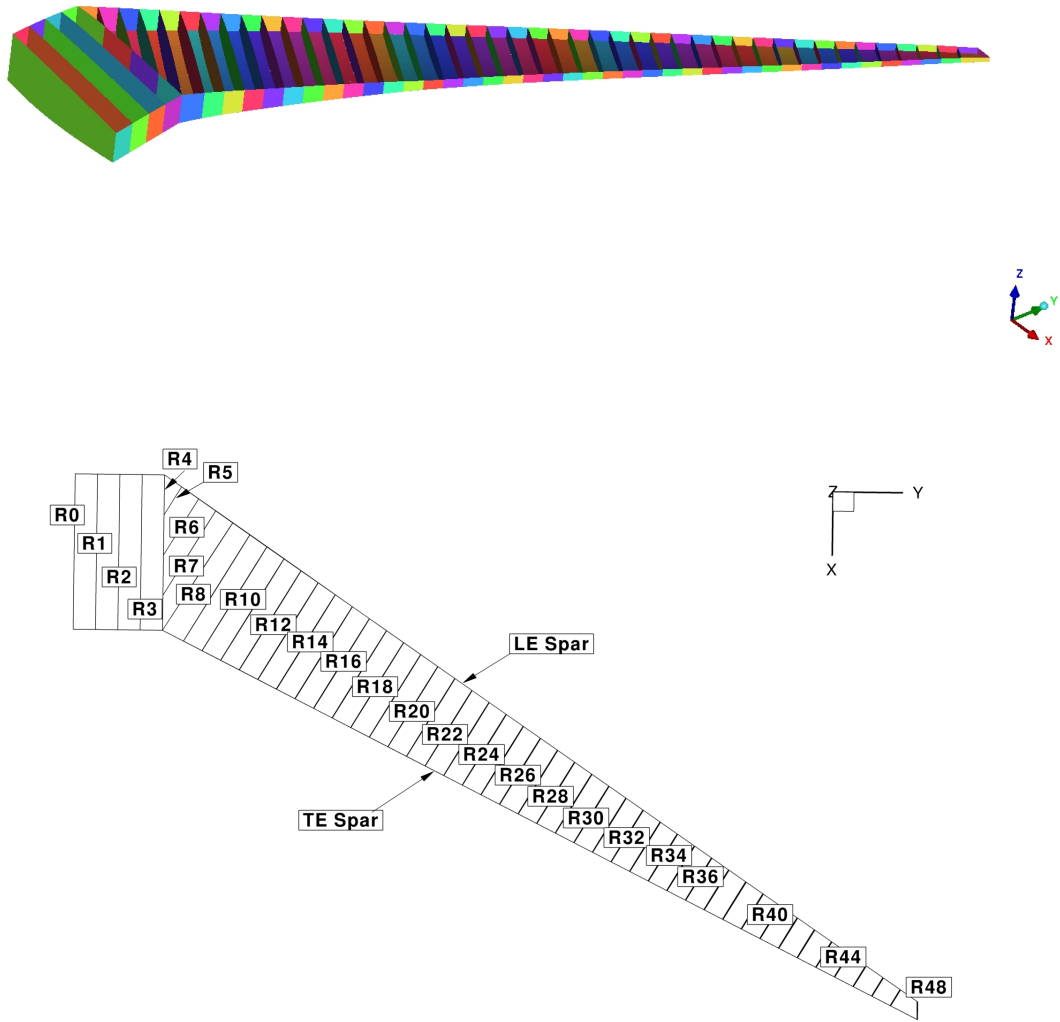
Figure 19: Overview of the uCRM-9 wingbox geometry.
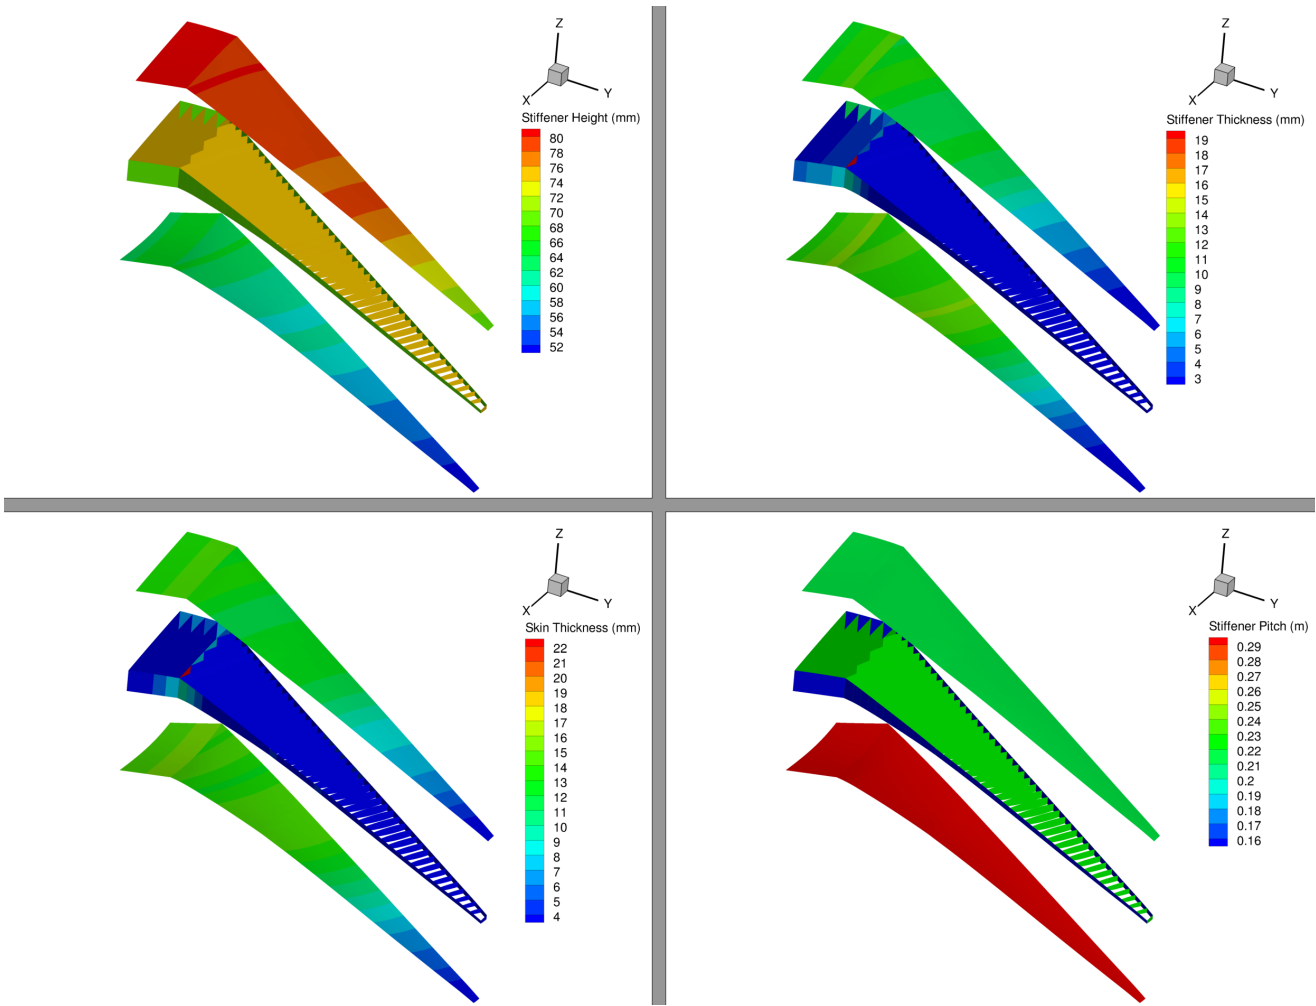
Figure 20: Distribution of stiffener height, stiffener thickness, skin thickness and stiffener pitch for the uCRM-9 structural design.
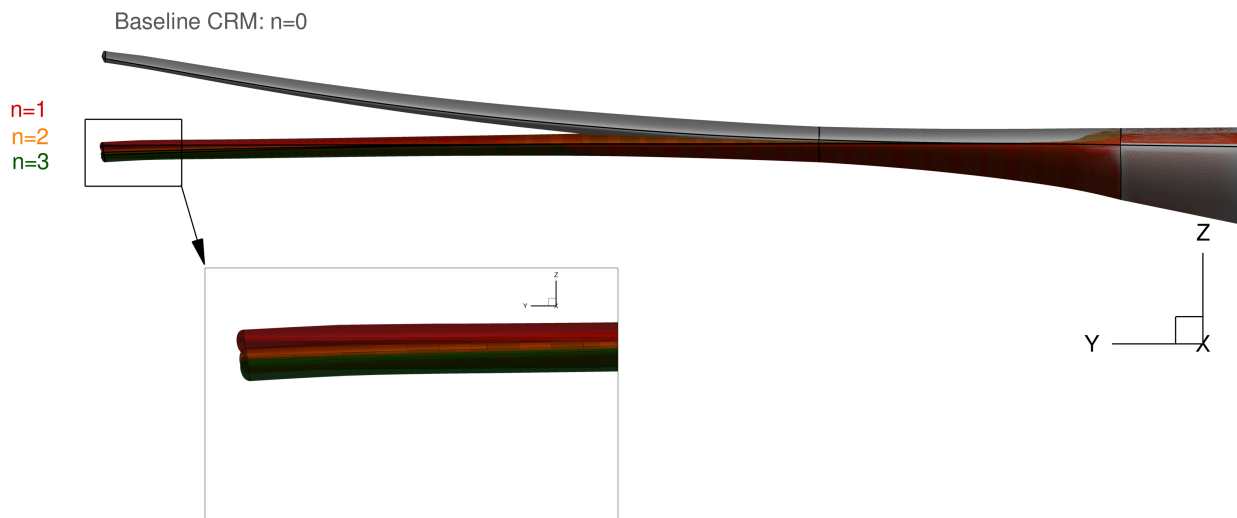
Figure 21: Iteration history of inverse design procedure
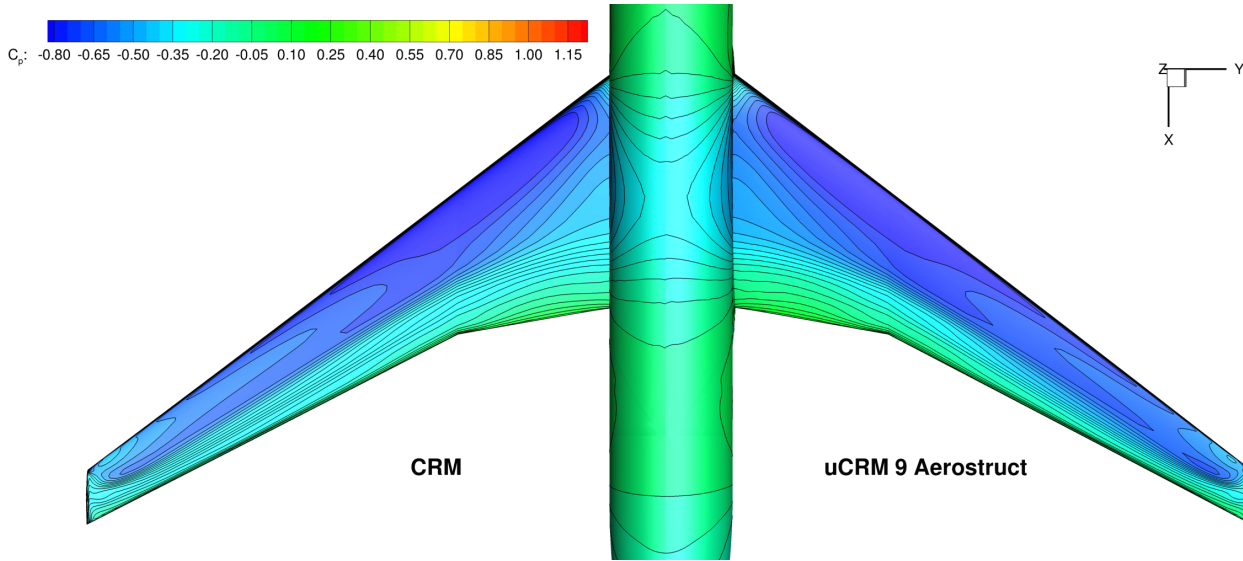
Figure 22: C_p contours comparison between original geometry and aerostructural solution.

Table 5: Aerodynamic coefficient comparison

Parameter	CRM wing-only	uCRM-9 aerostructural
C_L	0.5	0.5
C_D (counts)	200.75	199.03
α	2.208	2.168

necessary to match the lift coefficient of the original CRM case. A summary of the comparison is given in Table 5. The drag coefficient values differ by less than 1%, which we consider acceptable given the significant additional complexity of the aerostructural solution.

5.3 Definition for uCRM-13.5

Because there exists no standard transonic transport aircraft geometry with such a large aspect ratio, the structure and wing geometry of the uCRM-13.5 had to be defined through the process of a full aerostructural design optimization. This is in contrast with the uCRM-9 model, where the CRM OML is preserved, and only a structural sizing optimization is performed. The extension of the uCRM-9 to obtain the required 13.5 aspect ratio is complicated by the fact that we want to preserve the same wing loading and landing gear location, in order to avoid a cascade of design changes that would affect the preliminary sizing. The new 13.5 aspect ratio uCRM wing attempts to maintain overall compatibility with the remainder of the CRM aircraft. In this way, the resulting aircraft design may be considered a higher span derivative of the CRM. Keeping in mind the constraints imposed by a common fuselage, empennage, and propulsion system, we decided to keep the cruise altitude and planform area fixed. Fixing the cruise altitude avoids having to account for the engine performance variation, and keeps the original wing loading to avoid resizing the rest of the aircraft. To ensure sufficient space is available in the uCRM-13.5 planform, we first determined an appropriate value for this physical constraint. To determine this, we examined the planform
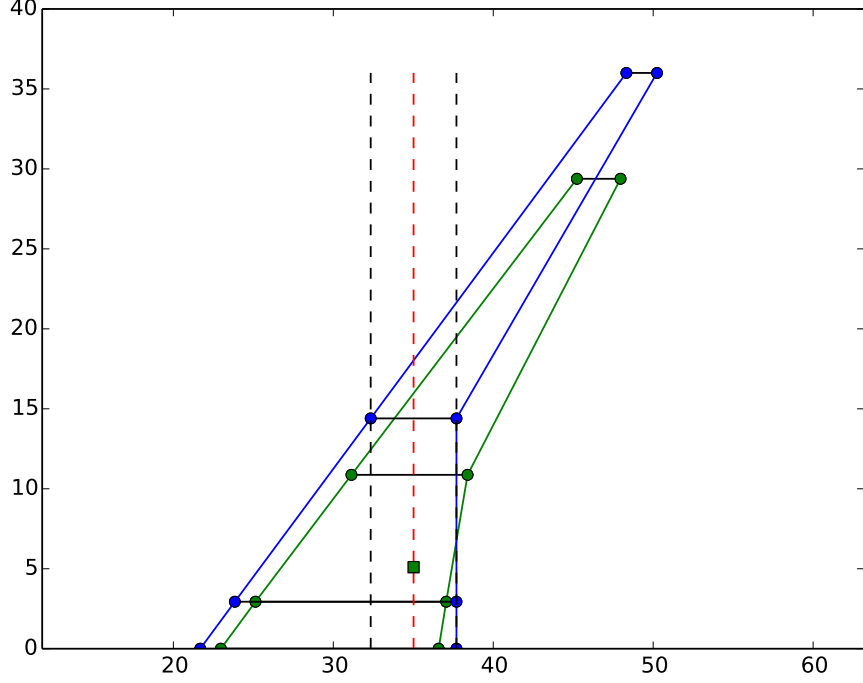
Figure 23: Proposed uCRM-13.5 wing planform (blue) and CRM-9 planform (green). The dotted lines show the 0, 50% and 100% MAC locations for the uCRM-13.5. The green dot on the 50%MAC location signifies the main gear post.

diagrams from the 777 Airport planning document. The new proposed uCRM-13.5 planform is shown in Figure 23.

A three-view rendering of the uCRM-13.5 planform is given in Figure 24. Note that the in-built dihedral of the CRM has been removed. As suggested by Boeing, this configuration is a straightforward modification of the well-tested CRM configuration with a clear connection to that configuration. We expect this geometry will also be made available for further research.

The structural wingbox was chosen to exhibit the same topology of the uCRM-9 wingbox, with the exception of the number of ribs. The absolute spacing was kept the same, and hence the uCRM-13.5 has 9 additional ribs for a total of 54.

Once the initial geometry of the uCRM-13.5 was defined an aerostructural optimization had to be performed to arrive at baseline design with realistic performance. The aerostructural design optimization was performed by the MACH framework.

Like the uCRM-9 design procedure, this optimization was based on a traditional aluminum design. However, unlike the uCRM-9 the uCRM-13.5 design procedure was based on an aerostructural optimization featuring a set of buffet onset constraint (explained in Section 7). The initial wing weight is estimated from a structural optimization preformed before the start of the aerostructural optimization. The flight conditions considered included a 5 point cruise stencil: 2 buffet constraint conditions and 3 maneuver conditions, shown in Table 6. The first point of the cruise stencil was defined as the center or nominal cruise condition (Mach = 0.85, $C_L = 0.5$). Points 2 and 3 of the stencil were defined by offsets of ± 0.025 in C_L relative to point 1. Cruise points 4 and 5 were defined as ± 0.01 offsets in Mach number relative to point 1, but restricted to have the same di-
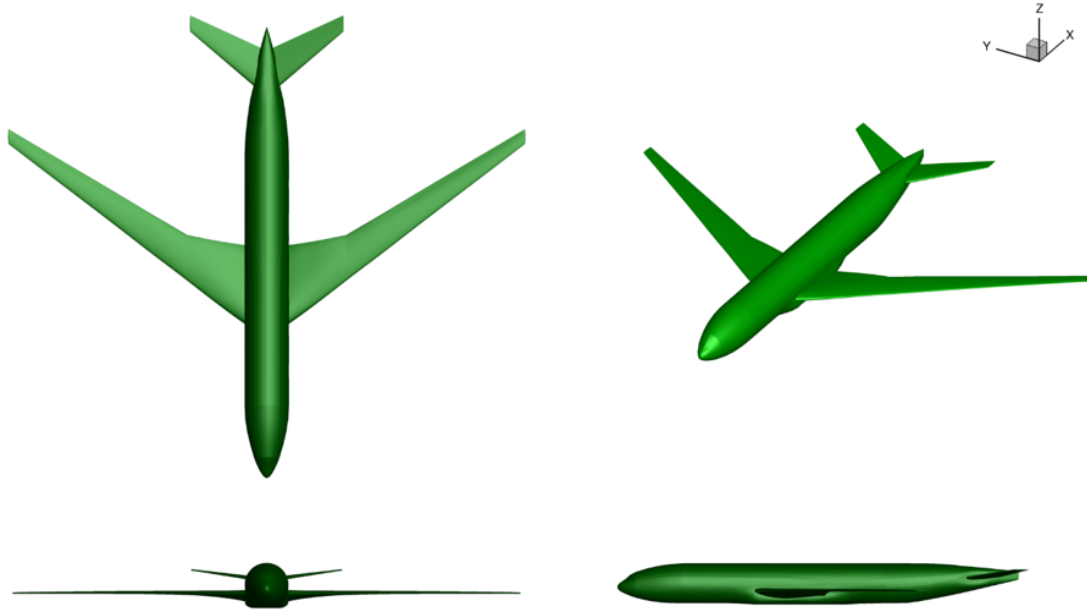
Figure 24: Three view of the 13.5 configuration.

mensional lift. All buffet and cruise conditions were evaluated at an altitude of 37,000 ft. The first buffet point was specified to be a 1.3 g margin on the highest lift of the cruise conditions, cruise condition 3. The second buffet condition was placed at a high Mach design point ($M = 0.89$) with the constraint that the lift match that of the nominal cruise condition, cruise condition 1. The 3 maneuver conditions included a -1.0 g, 2.5 g, and 1.0 g cruise with gust condition with yield and buckling constraints for structural sizing. The gust condition is estimated using a static 1.3 g load. The purpose of the gust condition is to ensure the operating stresses at normal operating conditions are not too high. It is meant to be a surrogate for a gust load that may be expected to load the wing sufficiently quickly that the passive aeroelastic tailoring cannot redistribute the load inboard fast enough. This maneuver condition is analyzed at a steady cruise condition near the Mach cross-over altitude of 28 000 ft, M = 0.85, TOGW and full fuel load and uses a 2.67 safety factor to ensure a sufficient structural margin for a gust encountered near this altitude.

The objective of the optimization was to minimize the average fuel burn of the cruise conditions. The optimization design variables included angle of attack, tail trim angles, cross-sectional shape, twist and structural variables. In addition a target nominal lift coefficient variable was added to allow the nominal cruise condition and thereby the entire stencil to find the optimal C_L to cruise about. Only the cruise and buffet points were affected by the moving stencil, the maneuver conditions were constrained. The specifics of the optimization problem are given in Table 8.

The results of the optimization can be found in Figure 25. From the results we find that the optimizer was able to improve the L/D, but the structural weight of the wing box increased. Despite the increase in structural weight the improvement in L/D is enough to lead to an overall improvement in fuel burn performance of the optimized aircraft. If we look at the t/c distribution of the two designs we find that the reason for the increase in structural weight was due to the fact that the depth of the wing box decreased considerably during the optimization so it was necessary

38

Table 6: Initial flight condition stencil

Point	Condition	FB Weight (T_i)	Mach	C_L/Lift	Altitude (ft)
1	Cruise	1/5	0.85	0.5	37 000
2	Cruise	1/5	0.85	0.475	37 000
3	Cruise	1/5	0.85	0.525	37 000
4	Cruise	1/5	0.84	0.512	37 000
5	Cruise	1/5	0.86	0.488	37 000
6	Buffet	0	0.85	0.683	37 000
7	Buffet	0	0.89	0.456	37 000
8	2.5 g maneuver	0	0.64	$2.5 \cdot$ MTOW	0
9	-1.0 g maneuver	0	0.64	$-$MTOW	0
10	Cruise gust	0	0.86	MTOW	27 300

Table 7: Optimized flight condition stencil

Point	Condition	FB Weight (T_i)	Mach	C_L
1	Cruise	1/5	0.85	0.516
2	Cruise	1/5	0.85	0.491
3	Cruise	1/5	0.85	0.541
4	Cruise	1/5	0.84	0.528
5	Cruise	1/5	0.86	0.504
6	Buffet	0	0.85	0.703
7	Buffet	0	0.89	0.470

to reinforce the structure to account for the loss in stiffness. Figure 26 provides a fuel burn contour plot as a function of Mach and C_L for both designs. The benefit of this plot is it allows to visually see both the robustness of the design and the margin between buffet onset and the cruise flight envelope. The first thing to note on these plots is the line denoted 1.3 g buffet boundary. This line shows the boundary for all of the operational conditions that have at least a 1.3 g margin from buffet. Ideally all points in the cruise stencil should lie within this region. We can see that from the initial contour plot some of the cruise conditions are not within this feasible region and as a result the optimizer increases the feasibility region, ensuring that all points in the cruise stencil are now feasible. The decrease in t/c that we saw in Figure 25 is primarily due to the prevalence of the buffet constraint; the optimizer makes the airfoil thinner in order to delay flow separation and in effect satisfy the buffet constraint. The other point of interest is the box under the 1.3 g buffet boundary line which is designated the fuel burn performance integration region. In this region the fuel burn is averaged over the entire region to assess the robust performance of the design. Ideally this region should correspond with the cruise stencil, and both should lie very close to the global minimum fuel burn on the plot, however due to the buffet boundary being so restrictive in the case of the initial design the integration region is located further down in the Mach-C_L design space than it should be, causing a reduction in performance. We can see that the optimizer is able to raise the performance integration region by moving the stencil up in C_L which ultimately leads to a decrease from 3,814 kg (4%) in average fuel burn. Once this optimization was performed the converged wing shape and structurally sized result was dubbed as the uCRM-13.5 model.

Table 8: uCRM 13.5 multipoint optimization problem with buffet conditions

	Variable/function	Description	Quantity
Minimize	$\sum_{i=1}^{N} T_i \mathrm{FB}_i$	Average Fuel Burn	
with respect to	x_{α_i}	Angle of attack for each case	3
	x_{twist}	Wing twist	10
	x_{tail}	Tail trim angle for each case	3
	x_{airfoil}	FFD control points	192
	$x_{\mathrm{stiff\ thick}}$	Panel stiffener thickness skin/spars/ribs	266
	$x_{\mathrm{stiff\ height}}$	Panel stiffener height skin/spars/ribs	266
	$x_{\mathrm{panel\ length}}$	Panel length skin/spars/ribs	266
	$C_{L_0}^*$	Nominal cruise target lift coefficient	1
		Total design variables	**1007**
Subject to	$C_{L_i} = C_{L_i}^*$	Cruise and Buffet lift conditions	7
	$L = n_i W$	Maneuver lift conditions	3
	$c_{m_y}^i = 0$	Trimmed flight	10
	$t_{LE}/t_{LE_{init}} \geq 1.0$	Leading edge radius	20
	$t_{TE}/t_{TE_{init}} \geq 1.0$	Trailing edge thickness	20
	$(t/c)_{\mathrm{TE\ Spar}} \geq 0.80(t/c)_{\mathrm{TE\ spar}_{init}}$	Minimum trailing edge spar height	20
	$L_{\mathrm{panel}} - x_{\mathrm{panel\ length}} = 0$	Target panel length	266
	$\mathrm{KS_{stress}} < 1.0$	2.5 g Material Failure	8
	$\mathrm{KS_{buckling}} < 1.0$	2.5 g and -1.0 g Buckling	9
	$\left\lvert x_{\mathrm{panel\ thick}_i} - x_{\mathrm{panel\ thick}_{i+1}} \right\rvert \leq 0.0025$	Skin thickness adjacency	258
	$\left\lvert x_{\mathrm{stiff\ thick}_i} - x_{\mathrm{stiff\ thick}_{i+1}} \right\rvert \leq 0.0025$	Stiffener thickness adjacency	258
	$\left\lvert x_{\mathrm{stiff\ height}_i} - x_{\mathrm{stiff\ height}_{i+1}} \right\rvert$	Stiffener height adjacency	258
	$x_{\mathrm{stiff\ thick}} - x_{\mathrm{panel\ thick}} < 0.005$	Maximum stiffener-skin difference	172
	$\Delta z_{\mathrm{TE,upper}} = -\Delta z_{\mathrm{TE,lower}}$	Fixed trailing edge	8
	$\Delta z_{\mathrm{LE,upper}} = -\Delta z_{\mathrm{LE,lower}}$	Fixed leading edge	8
	$\mathrm{Sep}_i \leq 0.04$	Buffet separation constraints	2
		Total constraints	**1328**
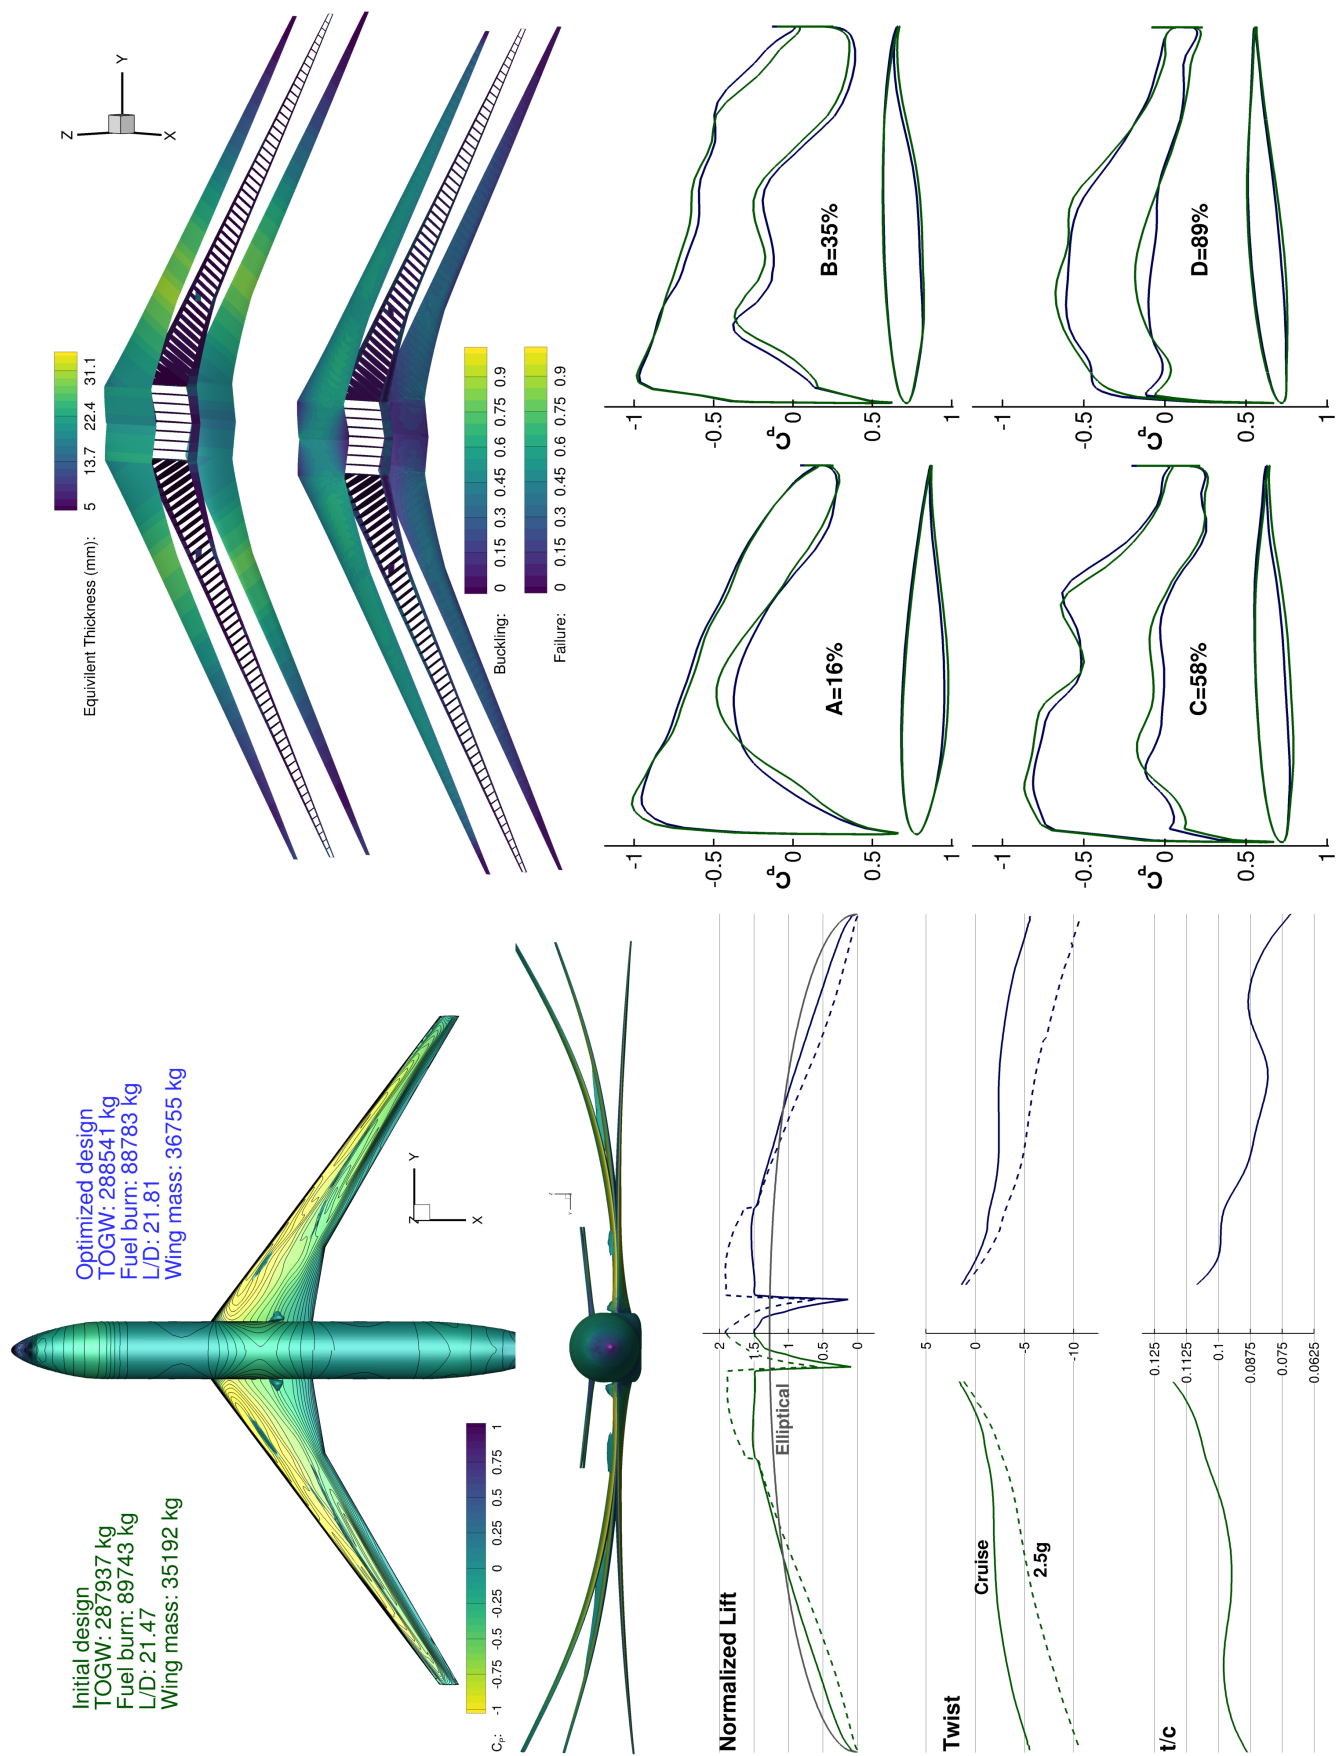
Figure 25: Comparison between initial (left) and optimized (right) aerostructural designs for the buffet-constrained multipoint optimization
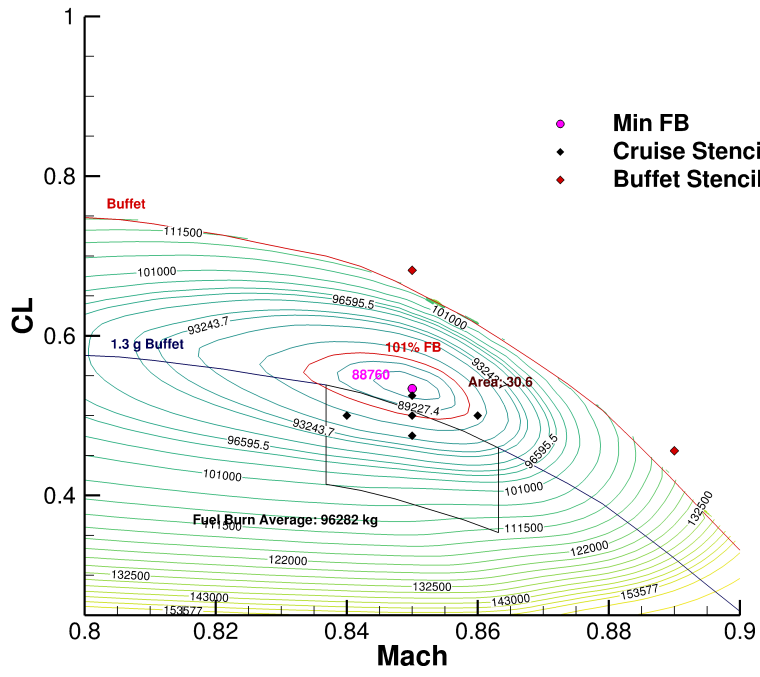
(a) Initial
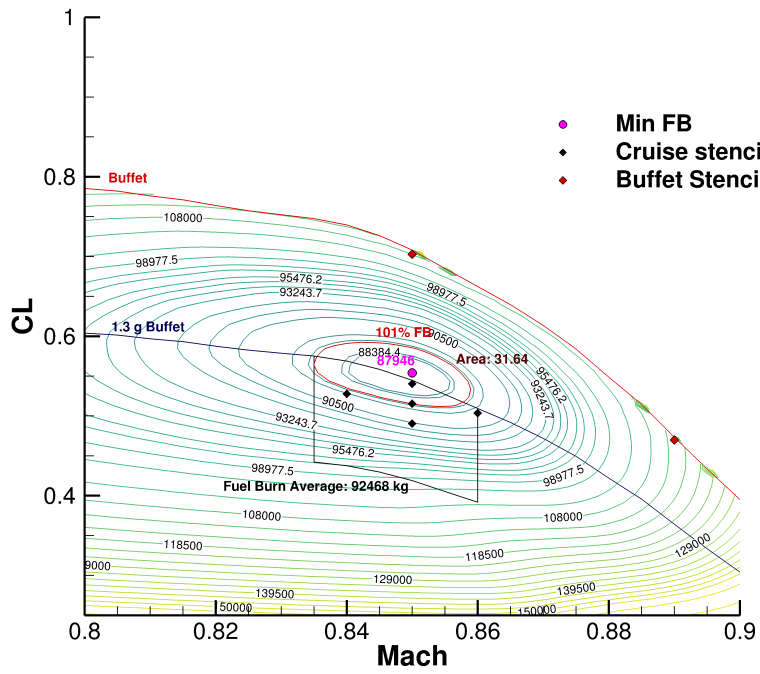
(b) Optimized

Figure 26: Fuel burn contours (initial vs optimized)

6   Continuous morphing trailing edge technology assessment

In an effort to improve fuel efficiency, aircraft designers are studying a number of next-generation configurations and technologies. One such technology is continuous morphing trailing edge devices. Designing morphing trailing edge devices is very challenging given the conflicting objectives involved. A morphing mechanism needs to manage the trade-offs between the mechanism weight, the ability of the mechanism to support aerodynamic loads, and the device's ability to morph.

Within literature and in industry there are a number of approaches currently in development to address these challenges. Vos et al. [61] considered the use of post-buckled pre-compressed piezoelectric actuators for UAV wing morphing. Pankonien [62] developed and studied a hybrid morphing concept utilizing shape memory alloys and macro fiber composites, again for use on UAVs. For commercial aircraft, a number of researchers have considered the Variable Camber Continuous Trailing Edge Flap (VCCTEF) approach, which uses a series of small flaps joined with an elastic material on the wing's trailing edge. Urnes et al. [63], Kaul and Nguyen [64], Ting et al. [65] have performed a number of studies using this configuration. Stanford [66] also considered this configuration, with the inclusion of open loop load cases in the measurement of a morphing device's effectiveness.

Commercially, FlexSys has produced the FlexFoil, a morphing design utilizing global compliant structures. This mechanism has the potential to produce a smooth continuous wing surface, with a wide variety of deflected shapes. The morphing devices considered herein are a simplification of this configuration. While these global compliant mechanisms offer a wide range of shapes, determining the optimal in flight shape and designing the set of deformed shapes that can be produced by the morphing structure remains a challenging task. To address these challenges, we perform a number of idealized high fidelity aerostructural optimizations of morphing trailing edge configurations.

Within our model the structures outside the wing box—like those within the morphing region—are not modeled explicitly, but rather idealized to remove many morphing shape restrictions, and allow a wider range of morphing versatility. While most of the morphing deformations in literature consist of a number of rigid rotations or a number of spanwise polynomial deformation profiles, the deformations in this work are much more general. We use Free Form Deformation (FFD) volumes to parameterize geometric shape changes, including those in the morphing section of the wing. This approach embeds the geometry in a volume covered in a number of control points. Moving those control points stretches the entire volume, and thus the geometry embedded within. This relatively unrestrictive parameterization permits a wide variety of morphing shapes, which when coupled with gradient based optimization allows us to explore the potential of morphing technology. Rather than starting with a mechanism and determining the best deformation within the mechanism's design space, this approach seeks a more mechanism-independent optimal morphing shape. This approach produces a more general set of optimal morphing shapes, which are valuable for measuring of the potential of morphing technology and for informing the design of future morphing mechanisms.

Our goals in this study were to accurately quantify the fuel efficiency improvements enabled by morphing trailing edge technology, and subsequently to gain a more detailed understanding about the mechanisms by which this improvement is achieved. Towards these ends, we performed a number of single- and multi-point aerostructural optimizations of the uCRM configuration with and without a morphing trailing edge. Detailed results and discussions of the various optimizations and comparisons herein can be found in the related papers [15, 18].

6.1   Single point optimization

The first study we completed consisted of a collection of single point optimizations. We compared the optimized results of a uCRM wing-body (with a correction for the trim contributions) with

no morphing, with a retrofit morphing trailing edge, and with a complete redesign including a morphing trailing edge. Detailed results on these three optimizations can be found in the associated paper [18].

Each optimization consisted of one 2.5-g sea level maneuver condition at Mach 0.64 and a single cruise condition at a Mach number of 0.85, a lift coefficient of 0.5, and an altitude of 37,000 feet. The cruise condition is used to quantify the objective function—fuel burn, calculated using the Breguet range equation—at a typical flight condition, while the maneuver provides limits for the structural design. Design variables include angle of attack, wing twist, structural sizing of members within the wingbox, and wing shape variables. The wing shape is manipulated using an FFD, with 192 control points distributed across the wing. Eighty of those control points are used on the aft 40% of the wing to control the morphing region of the wing. A number of geometric, aerodynamic, and structural constraints are used to ensure the optimized design is reasonable. The specific set of design variables and constraints used for each optimization are defined in the related publication.

The baseline aerostructural optimization included no active morphing. Leveraging improved aerodynamic cruise performance and aeroelastic tailoring produced by changes to the jig shape and wingbox structure, this baseline optimization reduced the fuel burn of the uCRM configuration by approximately 5.8%. The first morphing optimization represented a retrofit of the uCRM with a morphing trailing edge device. While this optimization provided no shape control on the front 60% of the wing, independent shape changes were possible in the morphing region at both cruise and maneuver. The final single point optimization represented a clean sheet redesign of the wing including a morphing trailing edge. This optimization is essentially the same as the baseline, non-morphing optimization, with the addition of a set of morphing design variables at the maneuver condition. The fuel burn and structural weight results of each optimization are shown in Table 9.

Table 9: Comparison of the fuel burn and wing mass of the three single point optimized wing designs.

Optimization	Fuel Burn [kg]	Wing mass [kg]
Baseline	105,993	33,839
Retrofit	105,959	31,957
Clean Sheet	105,613	33,131

Each of the single point optimizations produced very similar fuel burn values. It is notable that in this case, adding a retrofitted morphing trailing edge produced a fuel burn reduction essentially equal to that resulting from a redesign of the wing without morphing. The small 0.36% fuel burn difference between the baseline optimized wing and the clean sheet design is attributed to the specification of the optimization problems. Considering one cruise and one maneuver condition does not take advantage of the active shape changing capabilities of morphing technology. Without morphing, the baseline optimized wing uses aeroelastically tailored bend-twist coupling to achieve maneuver load alleviation. This aeroelastic mechanism provides a very good design, which offers limited potential for improvement with the addition of morphing capabilities. In order to more appropriately compare morphing and non-morphing wings, more flight conditions need to be considered, as is done in the next section.

6.2 Multipoint optimization

To perform a legitimate comparison of clean-sheet morphing and non-morphing wing designs, we ran a series of multipoint optimizations of both configurations. Optimizations with and without morphing are performed on both three- and seven-point stencils, with maneuver conditions at 2.5-g

and -1-g. More details about the optimization problem definitions and results can be found in the associated paper [15]. A comparison between the conventional and morphing optimization results for the seven point stencil is shown in Figure 27.

Figure 27 demonstrates that the 5% fuel burn improvement produced by morphing for a seven point stencil is achieved with aerodynamic improvements at cruise and lightening of the structure enabled by maneuver load alleviation. At maneuver the outboard morphing regions reflex, reducing the lift in that region, and producing the sizable inboard shift visible in the spanwise lift distribution. More details about these results are available in the reference alluded to previously.

Table 10: As the multipoint stencil size is increased from 3 to 7 points, the fuel burn savings increases from 2.53% to 5.04%, respectively.

Stencil	Morphing	Fuel Burn[kg]	Wing Mass [kg]
3 point	No	94,421	29,573
	Yes	92,034	22,938
7 point	No	98,627	30,060
	Yes	93,656	22,300

Table 10 summarizes the results of the 3 and 7 point optimizations. The maneuver conditions used to size the structures in both optimizations were the same, which leads to an interesting finding. Wing weight changes are subject to an aerostructural trade-off. As the weight of the wing is increased, there is a direct penalty to fuel burn resulting from the increased weight of the aircraft. Increasing wing weight also produces an aerodynamic benefit due to the improved consistency of the deflected shapes at various flight conditions. These results demonstrate that the addition of morphing changes the balance of these two effects. Without morphing, increasing stencil size encourages an increase in the structural weight, as there is a larger benefit from improved deflected shape consistency. With morphing however, the coupling between performance at various flight conditions is decreased, weakening the aerodynamic incentive to increase the structural weight. This is seen in the optimized results, as the weight of the morphing wing optimized at 7 points is less than that optimized at 3 points, while this trend is reversed for the non-morphing wings.

After completing these multipoint studies, we conducted two additional comparisons of optimized results. The first additional study investigated the importance of the size of the morphing region. A realistic wing design including a morphing trailing edge may not allow a device to span the entire chordwise region aft of the wingbox. To address this, the morphing region was reduced to the aft 30% of the wing. The three and seven point optimization comparisons were repeated, and fuel burn increases of 0.22% and 0.81%, respectively, were observed compared to the results with the larger morphing region. The second additional study considered the influence of the aspect ratio on the effectiveness of morphing devices. The three point comparison was repeated on the uCRM13.5, and showed promising results. Morphing improved the fuel efficiency of the uCRM and uCRM13.5 by 2.53% and 3.79%, respectively. This result clearly demonstrates the synergistic effect between increasing a wing's aspect ratio and adaptive morphing.

6.3 Morphing dexterity study

The morphing trailing edge optimizations up to this point have not included an explicit consideration of the morphing mechanism. Morphing shapes were were not limited based on any mechanism restrictions, and no design-dependent weight penalty was added for actuators. As such, the previous results can be interpreted as the maximum performance using an ideal morphing device. Based on discussions with collaborators at FlexSys, this problem formulation is reasonably representative

and realistic, however these assumptions prevent a study considering the dexterity of morphing devices. To address this issue, we developed a surrogate model representing the actuator weight as a function of aerodynamic loading, designed for use within our high fidelity framework.

There is a physical relationship between the morphing shape and the weight of the required actuator/mechanism(s). The idealistic assumptions used in the previous optimizations ignore this relationship. Without modelling this relationship, it is impossible to capture and study the tradeoff between a morphing configuration's weight and dexterity. As more spanwise actuators are added, the morphing configuration gains dexterity, and therefore can produce better aerodynamic performance; however, the addition of actuators also adds weight to the morphing configuration. To study this tradeoff we developed the necessary actuator surrogate models, and subsequently sought to perform a number of aerostructural optimizations.

We first consider the development of the actuator weight surrogate model. The model is built around a dataset for linear electromechanical actuator specifications, as shown in Figure 28. A large portion of the dataset for smaller actuators is based on off-the-shelf actuators available online [1, 2], and data for a larger actuator comes from a design outline produced by Fronista and Bradbury [3]. The data is fit to a power regression of the power density as a function of the power, as defined in Figure 28. This formulation, yielding power density as a function of power formulation is typical in industry, and with an R-squared value of 0.9464, showed strong performance compared with alternative forumlations. Within the framework, the actuator formulation needs to produce masses as a function of required power. Replacing the power density with power per kilogram mass and rearranging reformulates the surrogate as:

$$m = 0.5903 P^{0.2878} \tag{6.1}$$

,where m is the actuator mass and P is the required actuator power. This relationship provides a sizing function for the morphing actuators as a function of the power those actuators require. The power requirement is based on the aerodynamic loads produced on the panels of the morphing region. We assume a transition region is included between each panel to avoid nonlinear effects produced by the influence of deflections in adjacent actuators. The pressure and viscous forces on each panel are aggregated into an aerodynamic moment about the actuation axis. This calculation of a moment about an arbitrary axis produced by a subset of the surface was produced for this work. Adjoint-based gradients of this calculation were developed and validated as well. The aerodynamic moment is multiplied by a required actuation rate of 70 deg/s [67], which was selected as an upper bound of required actuation rates found in literature. The result is the required actuator power as a function of the aerodynamic loading. An additional surrogate was added to account for the requirements in deforming the morphing structure, however that model is based on work by FlexSys [68], and is therefore not described herein. The total required power for each actuator is then used in Eqn. (6.1) to produce the required actuator mass.

The key difference between these optimizations and those previously discussed is the addition of the morphing regions and actuator weights. Four different optimizations were developed in this study, with 0, 3, 6, and 12 actuators. The smallest division of the morphing region, into 12 panels, is shown in Figure 29. Figure 29 also shows the wingbox superimposed, showing the morphing panels locations aft of the wingbox. This results in a morphing region spanning a variable percentage of the chord, between 30 and 40%. These panels, defined as sub-families within the mesh, define the regions of integration used when calculating the aerodynamic moment used to size the actuators. For the case with 12 actuators, each panel is considered individually. In the optimizations with 3 and 6 actuators, each actuator is sized using a combination of the moments from multiple panels.

Defining the morphing region used to size the actuators generates the dependence of the actuator weight on the design, however it does not model the change in morphing dexterity associated with

adding actuators. To include the changes in dexterity, the design space of each optimization is defined differently. The baseline shape of the wing is defined the same way in each case, with the FFD shown in the upper-left corner of Figure 30. The morphing regions for each case are defined with a separate sub-FFD. Each of the sub-FFDs used for 3, 6, and 12 actuators are shown on the bottom of Figure 30. In the upper right of Figure 30 is a view with all of the FFDs included. Note that in this aggregated view, control points are sized such that those used with fewer actuators are larger. Therefore green control points from the sub-FFD for 3 actuators are always covering purple and red control points from the sub-FFDs for 6 and 12 actuators. In the same way, there is a red control point from the case with 12 actuators under each purple control point for the 6 actuator case.

As actuators are added, design variables are also added, giving a larger design space, and providing the aerodynamic benefit produced by improved dexterity. In order to maintain $c1$ continuity on the OML, only the aft two design variables chord-wise are given freedom on each sub-FFD. Additionally, thickness constraints like those used in previous optimizations are not given to maintain the wing thickness in the morphing region. Instead, control points on the morphing sub-FFDs are constrained to move together with the point above or below them.

A three point multipoint stencil is used for each of these optimizations. The nominal point of the stencil is defined at $L = W$, and the off-design points of the stencil are at a load factor of 0.9 and 1.1. The nominal Mach number is once again 0.85. Two maneuver conditions are used to constrain the structure, a 2.5-g pull up and a -1.0g push over.

The remainder of the optimization problem definition closely mirrors the problem definitions from previous studies. At each flight condition the angle of attack and tail rotation variables are provided, and the lift and pitching moment are constrained. Along with 240 local shape variables, the baseline shape of the wing is optimized with 8 twist variables. There are 1029 structural variables, including panel length, skin thickness, stiffener height, stiffener pitch, and stiffener thickness. The volume of the wing is constrained not to decrease, to maintain enough space for fuel. The thickness cannot decrease at the leading edge, trailing edge, or aft spar. As in previous optimizations, these constraints provide better low speed performance, assure manufacturability, and leave space for actuation mechanisms and other avionics etc. near the aft spar. The leading edge control points are always constrained to move in equal and opposite directions. When there is no morphing, the trailing edge is constrained to prevent shear twist. A subset of the design variables—like panel lengths, actuator weights, and some shape variables—are specified in multiple locations for implementation reasons, and require associated consistency constraints. The structural members are constrained not to fail at 2.5g, not to buckle at 2.5g, and the lower skin, ribs, and spars are constrained not to buckle at -1.0g. An overview of the four optimization problems is given in Table 11.

Table 11: Optimization problem overview for the high aspect ratio uCRM with 0, 3, 6, or 12 actuators.

	Function/variable	Description	0 Actuators	3 Actuators	6 Actuators	12 Actuators
minimize		Fuel burn				
w.r.t.	x_{α_c}	Cruise AoA	3	3	3	3
	x_{α_m}	Maneuver AoA	2	2	2	2
	x_{shape}	Wing shape (FFD)	240	240	240	240
	x_{morph}	Morphing shape (FFD)	0	16	28	52
	x_{twist}	Wing twist	8	8	8	8
	x_{tail}	Tail rotation angle	5	5	5	5
	x_{struct}	Structural sizing	1029	1029	1029	1029
	x_{act}	Actuator weight	0	3	6	12
	Total DVs		**1287**	**1306**	**1321**	**1351**
subject to	$L = n_i W$	Lift	5	5	5	5
	$M = 0$	Pitching moment	5	5	5	5
	$V/V_{init} \geq 1$	Fuel volume	1	1	1	1
	$t/t_{init}\vert_{LE} \geq 1$	Leading edge thickness	20	20	20	20
	$t/t_{init}\vert_{TE} \geq 1$	Trailing edge thickness	20	20	20	20
	$t/t_{init}\vert_{spar} \geq 1$	Aft spar thickness	20	20	20	20
	$\Delta z_{LE_u} = -\Delta z_{LE_l}$	Fixed leading edge	9	9	9	9
	$\Delta z_{TE_u} = -\Delta z_{TE_l}$	Fixed trailing edge	9	0	0	0
	$L_{panel} - x_{panel} = 0$	Panel consistency	335	335	335	335
	$KS_{stress} \leq 1$	Maneuver stress	3	3	3	3
	$KS_{buckling} \leq 1$	Maneuver buckling	5	5	5	5
	$\vert x_{s_i} - x_{s_{i+1}} \vert \leq 5mm$	Adjacency constraints	811	811	811	811
	$x_{act} - m_{act}$	Actuator mass consistency	0	3	6	12
	$x_{shape}[i,j,0] - x_{shape}[i,j,1] = 0$	Linear shape constraints	24	40	52	76
	Total constraints		**1267**	**1277**	**1292**	**1322**
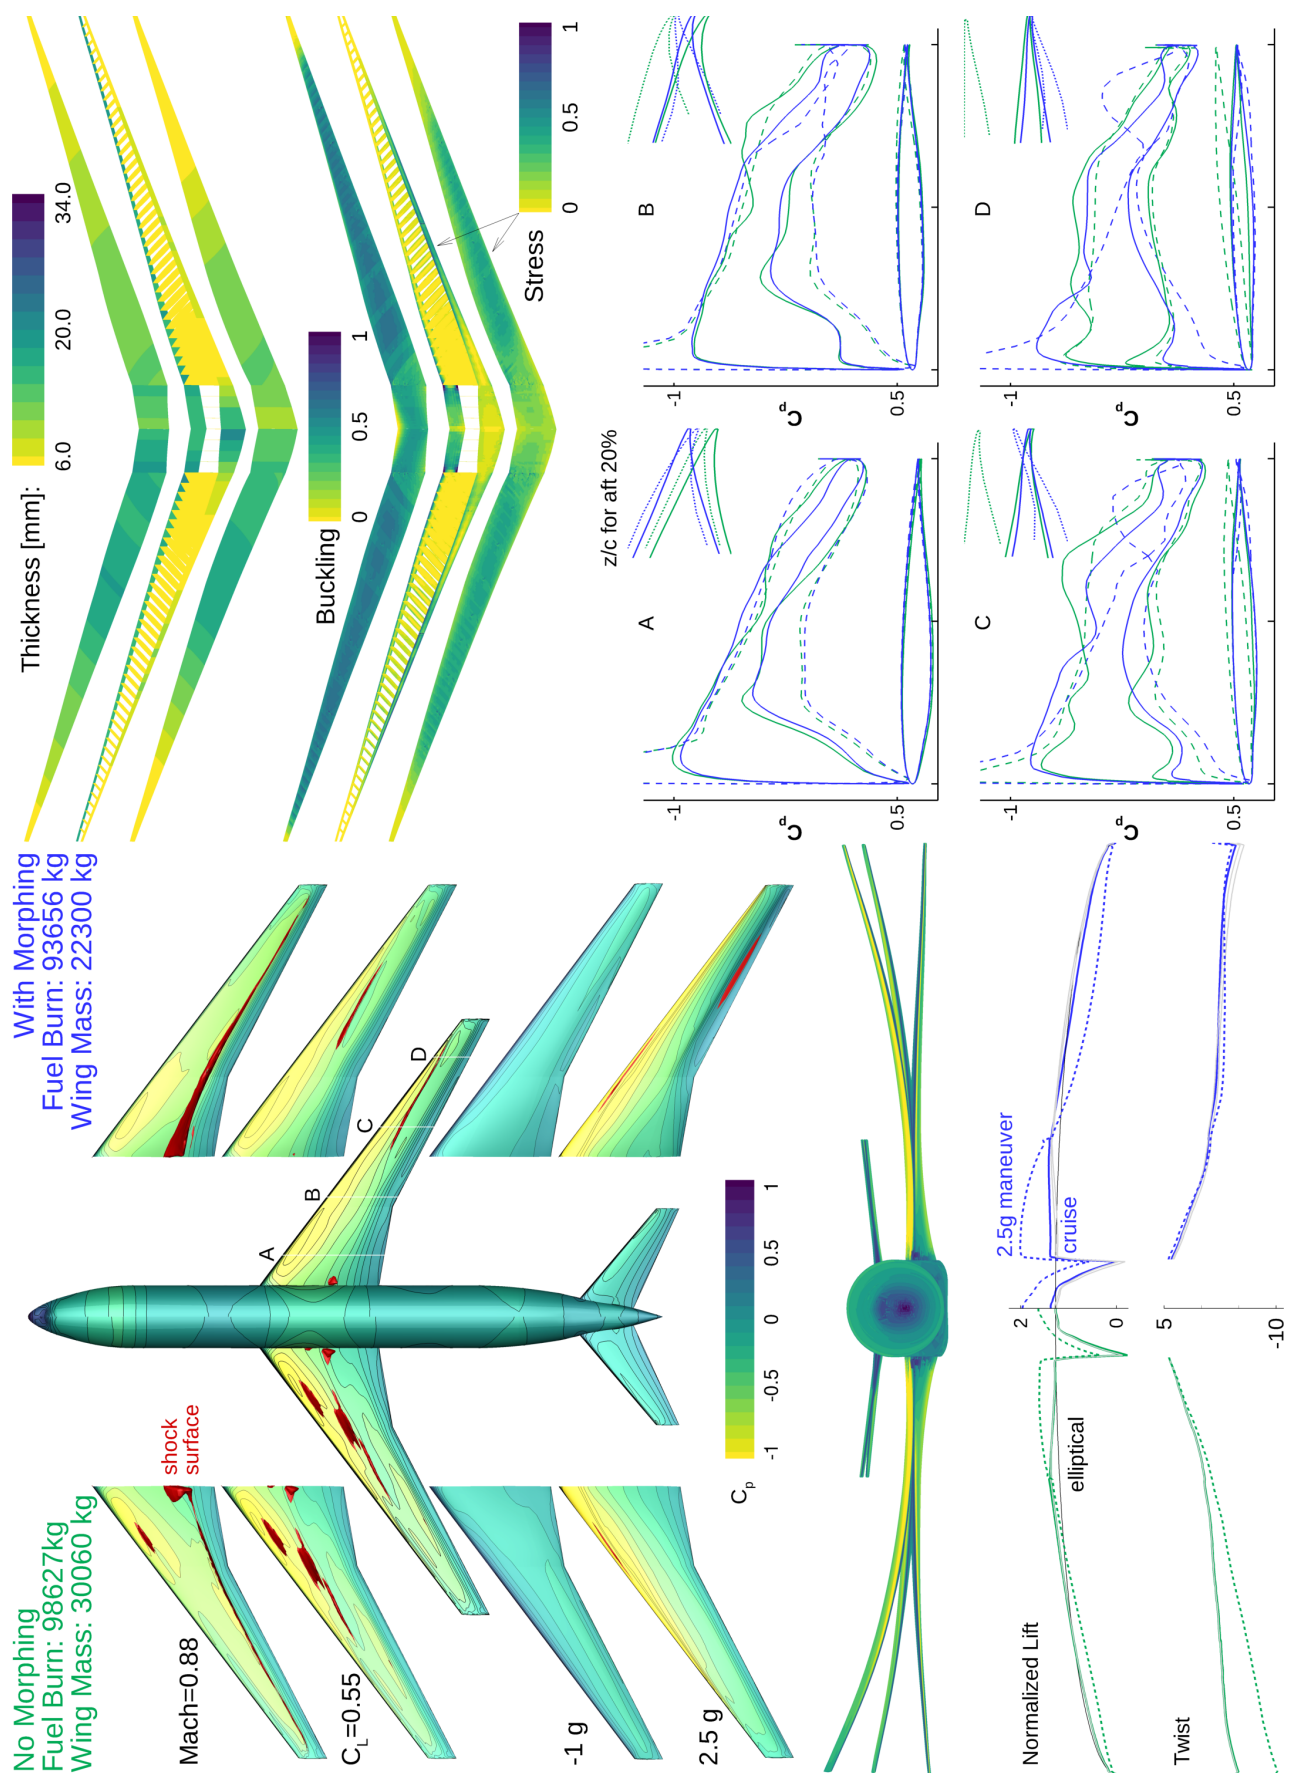
Figure 27: The morphing wing optimized for performance on a seven point stencil reduced the fuel burn by 5.04%, and the structural weight by 25.8% compared to the optimized non-morphing wing.

$\rho = 1.6941*P^{0.7122}$
$R^2 = 0.9464$
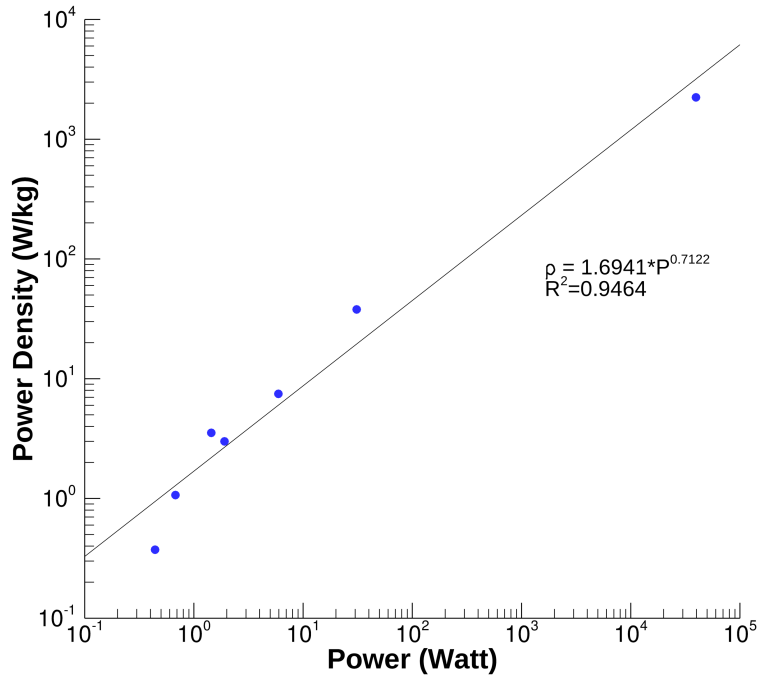
Figure 28: The actuator sizing surrogate is based on a fit of the relationship of power (P) to power density (ρ) for a number of electromechanical actuators [1–3].
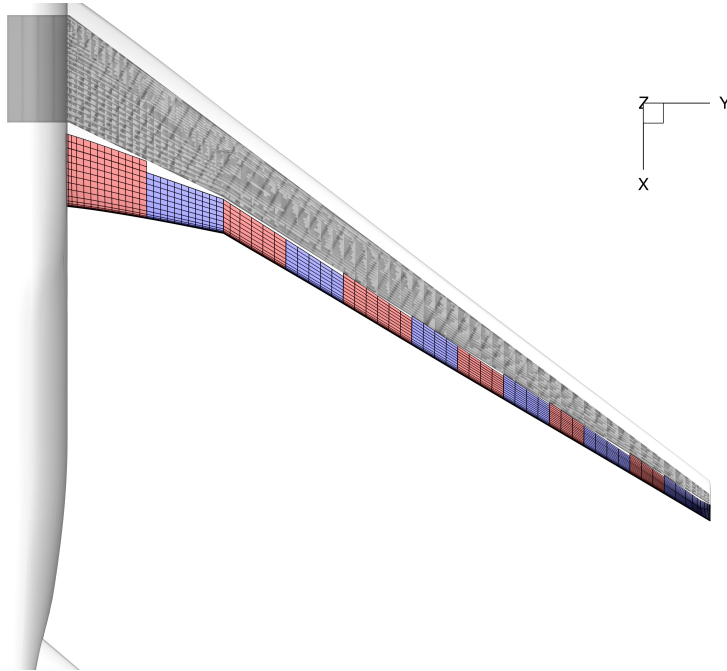
Figure 29: The morphing region is divided into 12 spanwise panels, each of which is aft of the wingbox.
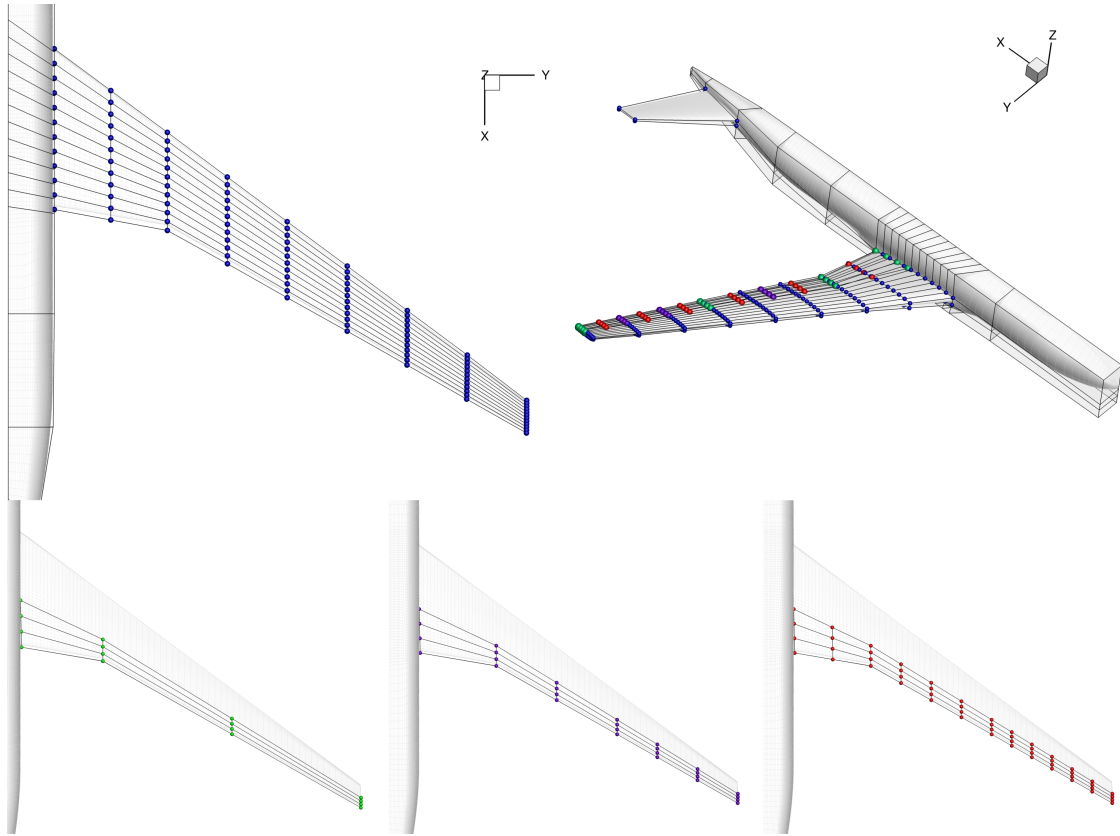
Figure 30: The baseline FFD used to define the undeformed shape of the wing shown in blue, while the sub-FFDs defining the shape changes of the morphing region for 3, 6, and 12 actuators are shown in green, purple, and red, respectively.

To complete this study we developed a surrogate model relating actuator weights to aerodynamic moments on regions of the wing. We also added functionality in our flow solver to calculate those aerodynamic moments produced by viscous and pressure forces on a region of the wing around arbitrary axes. This calculation included gradient computation with an adjoint approach, which was verified. The additions to the framework necessary for the optimizations and the required inputs were developed, but we have not yet completed the full optimization study. These high fidelity optimizations typically require a number of manual iterations to adjust parameters, problem definitions, etc., and finding the correct problem setup proved difficult for this problem. A particular obstacle in this case was the inability to use multigrid capabilities in our flow solver. Given the discretization of the mesh for the definition of the morphing panels, the indexing required for the use of multigrid capabilities was lost. This resulted in substantially slower analyses, and subsequently a bottleneck in the development of a properly formulated optimization problem.

6.4   Conclusions

In this assessment, we used high fidelity adjoint-based optimization techniques to quantify the fuel efficiency improvement potential of morphing trailing edge technology. The morphing trailing edge produced fuel burn savings of over 5% for a seven point multipoint optimization. Morphing technology was shown to improve aircraft efficiency using a combination of improvements to cruise aerodynamic performance and maneuver load alleviation. The importance and influence that morphing has on the trade-off governing the optimal structural weight was also identified and studied. We additionally addressed outstanding questions associated with the technology. We demonstrated that the chordwise size of the morphing device could be reduced without severe performance penalties. We also demonstrated the direct relationship between a wing's aspect ratio and the effectiveness of morphing technology. As aerospace materials and composites continue to improve and enable more flexible, higher aspect ratio wings, the incentive to include morphing technology will increase.

The use of our high fidelity optimization framework enabled us to produce an accurate quantification of the efficiency improvement potential of morphing technology that was previously unavailable in literature. Our framework additionally enabled an accurate and detailed investigation of the subtle effects that the inclusion of morphing technology will have on the trade-offs that engineers navigate during aircraft design.

7   Buffet prediction and constraint

Although high-fidelity aerostructural optimization was essential for the wing material study (Section 4) and the morphing technology assessment (Section 6), we found that the designs are susceptible to buffet, and are not guaranteed to satisfy the safety requirements. These requirements stipulate that commercial transport aircraft maintain at least a 30% margin from the cruise operating condition to buffet onset. To address this issue, we developed a separation-based constraint formulation that constrains buffet onset in an aerodynamic shape optimization. The separation metric developed in this work was verified against a common buffet prediction method and validated against experimental wind tunnel data. We performed a series of optimizations based on the AIAA Aerodynamic Design Optimization Discussion Group wing-body-tail case to show that buffet-onset constraints are required and to demonstrate the effectiveness of the proposed approach. Here, we present a summary of the separation constraint formulation and some selected results. Much more detail is included in Kenway and Martins [6, 17].

One classic way of buffet prediction is to use the lift curve slope of the aircraft configuration—the $\Delta\alpha = 0.1$ method [69]. However, this approach is not ideal for optimization and does not utilize the rich flow field information that is available from RANS solutions. To develop a more direct way of constraining buffet onset, we focus on the physical mechanism of shock-induced flow separation. This approach is also suitable for constraining separation in general. An example showing the typical progression of this type of separation with increasing angle of attack is shown in Fig. 31. The first row in Fig. 31 shows the friction lines and pressure coefficient, as well as a shock sensor (in orange) [70].

To determine if the flow is separated at a given location on the surface, we check if the surface flow velocity has a component in the negative freestream direction (which is approximately the negative x axis) direction, i.e., if

$$\cos\theta = \frac{\vec{V} \cdot \vec{V}_\infty}{|\vec{V}||\vec{V}_\infty|} < 0, \tag{7.1}$$

where θ is the angle between the local surface velocity and the freestream. The surface velocity is taken from the center of the cells right above the solid surface. We can then define a *separation sensor* as

$$\chi = \begin{cases} 1 & \text{if} \quad \cos\theta \leq 0 \\ 0 & \text{if} \quad \cos\theta > 0. \end{cases} \tag{7.2}$$

Thus, χ is specific to each surface location and is a Heaviside function: It is equal to one when the flow is separated, and equal to zero when the flow is attached. The blue areas on the surface for $\alpha = 3.00°$ and $\alpha = 3.29°$ in the bottom row of Fig. 31 show the regions where $\chi = 1$, which approximately coincide with the regions where the flow is separated.

Since we need to use this function as a constraint in a gradient-based optimization, we would like this function to be smooth. However, because this integral will be discretized based on a CFD surface mesh, and χ is either zero or one for a given cell, the value of this area does not change continuously with the design variables. To address this issue, we use a smooth Heaviside function to blend the discontinuity as follows:

$$\bar{\chi} = \frac{1}{1 + e^{2k(\cos\theta + \lambda)}}. \tag{7.3}$$

In this equation, k and λ are free parameters, where k determines the sharpness of the transition, and λ is a parameter that can be used to shift the smoothing function to the left or right as a function of the angle.

Next, we can integrate the smooth separation sensor (7.3) over the surface and normalize it by the aircraft reference area to obtain the proposed *separation metric*:

$$S_{\text{sep}} = \frac{1}{S_{\text{ref}}} \int_S \bar{\chi} \, \mathrm{d}S. \tag{7.4}$$

This is equivalent to performing a weighted area integration of the sensor value shown in the bottom row of Fig. 31.

To find out if the separation metric (7.4) is correlated to buffet onset, we used the $\Delta\alpha = 0.1$ method as a reference. We start by using the $\Delta\alpha = 0.1$ method to compute the buffet boundary for the baseline CRM configuration [57] at a flight altitude of $37\,000\,\text{ft}$ and for Mach numbers ranging from 0.8 to 0.9. The resulting reference buffet boundary is shown as the orange line in Fig. 32. We then plot the lines corresponding to various values of the separation metric, and determine that a cutoff value of $\chi = 4\%$ yields the best agreement when compared with the $\Delta\alpha = 0.1$ method.

The overall shape of the buffet onset is consistent with the separation-based criteria. Because the separation-based approach is more representative of the actual physics, we argue that it is the more accurate of the two methods [6].

To validate the proposed separation metric as a way to constrain buffet-onset, we compare the results obtained by using this approach with experimental results by Balakrishna and Acheson [4], who tested the CRM wind tunnel model. They estimated the buffet onset by making high-speed measurements of the strain at the wing root. Buffet onset can be identified by the increase in the strain gauge signal amplitude, which is caused by the shock oscillations interacting with the separated flow. Based on this increase in signal amplitude, they define the *buffet coefficient*, C_B.

Figure 33, which is reproduced from Fig. 4 in Balakrishna and Acheson [4], shows the evolution of the buffet coefficient for two Mach numbers: a high subsonic Mach number ($M = 0.70$) and a transonic Mach number ($M = 0.85$). We overlay lines at α values for which our method yields a separation sensor value of 4%, with $\lambda = -0.1$. We see that the results of the separation sensor method correlate well with the increase in C_B, providing more evidence that the separation-metric approach correctly predicts buffet onset.

To understand the need for the buffet constraint, and to understand the consequences of enforcing the constraint, we solve a sequence of seven design optimizations. These cases—numbered 5.1 through 5.7—are summarized in Table 3. Cases 5.1 and 5.2 are currently specified by the AIAA Aerodynamic Design Optimization Discussion Group (ADODG) [71]. We added the other cases (cases 5.3 through 5.7) to further study the effects of including buffet-onset conditions on the optimized geometries. The objective of all optimizations (except for Case 5.7, which is discussed separately) is to reduce the weighted drag coefficient at the N operating conditions. The optimization problem statement can be written as

$$
\begin{array}{lll}
\text{minimize} & \sum_{i=1}^{N} \mathcal{W}_i C_{D_i} & \text{Quantity} \\
\text{with respect to} & \text{Airfoil shape variables} & 240 \\
& \text{Wing twist} & 9 \\
& \text{Angle of attack } (\alpha_i) & N \\
& \text{Tail rotation angle } (\eta_i) & N \\
\text{subject to} & C_{L_i} - C_{L_i}^* = 0.0 & N \\
& C_{M_{y_i}} = 0.0 & N \\
& t_j \geq t_{j_{\text{CRM}}} & 750 \\
& S_{\text{sep}_i} \leq 0.04 & N
\end{array} \tag{7.5}
$$

Each flight condition i is assigned a weight $\mathcal{W}_i$ that specifies to what extent the drag of the given flight condition influences the objective function. The lift and moment coefficient constraints

ensure that the aircraft is trimmed at each flight condition, which can be achieved by the appropriate combination of angle of attack and tail rotation angle. The thickness t_j is computed at 750 points arranged in a 25×30 regular grid in the chordwise and spanwise directions, respectively. These thicknesses are constrained to be greater than or equal to the original thicknesses of the CRM geometry at the corresponding points. Because making the wing as thin as possible is desirable in transonic flow [10], these constraints ensure that the wing does not become too thin, which would result in a significant increase in structural weight. Imposing thickness constraints means that only changes in the wing camber are available to the optimizer.

Only cases 5.4, 5.5, and 5.7 use the separation constraint to satisfy the buffet margin. In these cases, the separation constraint is only enforced for the last two flight conditions, and the drag coefficient for these conditions does not contribute to the objective function (i.e., $\mathcal{W}_i = 0$). Therefore, the adjoint for C_D is not evaluated for the buffet-onset conditions. Conversely, the separation metric adjoint is not evaluated for the conditions where the drag coefficient weights are nonzero. This results in a total of three adjoint solutions being required for both the cruise and buffet flight conditions, which is desirable from a computational load-balancing perspective.

The ADODG specification for Case 5 disallows changes to the wing planform, and any shape modification may only be made in the vertical direction. Additionally, twist rotation is permitted for the wing, as is a solid-body rotation of the horizontal tail for trimming the aircraft.

Although drag-coefficient divergence curves yield useful insights into optimized designs, examining the performance in the full M-C_L space is particularly instructive. In the context of transonic transport wing design, ML/D is a better measure of performance because it includes the benefit imparted on overall aircraft efficiency by a higher cruise speed. This overall performance can be approximated by the Breguet range equation
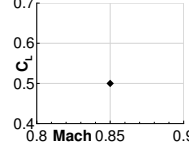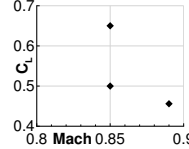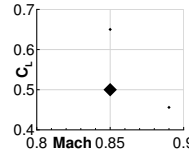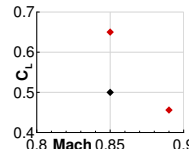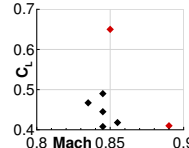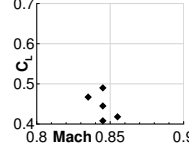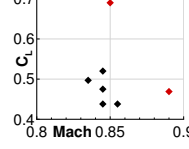
$$R = \frac{Ma}{c} \frac{L}{D} \ln\left(\frac{W_1}{W_2}\right), \tag{7.6}$$

where L/D is the lift-to-drag ratio, a is the speed of sound, c is the thrust-specific fuel consumption, and W_1 and W_2 are the initial and final cruise weights, respectively. For a purely aerodynamic optimization at fixed Mach number, only L/D varies if we assume a constant c and weight ratio W_1/W_2, so we are left with ML/D.

The procedure for generating contour plots is detailed by Kenway and Martins [6]. Figure 34 shows contour plots for all seven optimizations and the baseline design. The contours in each figure extend up to the predicted buffet-onset curve shown in red. The orange curve shows the buffet onset predicted by using the $\Delta\alpha = 0.1$ method. Several regions appear where the orange curves are missing data, which we attribute to the failure to find an intersection between the two lift curves. Overall, the separation-metric method continues to produce results that are close to those of the $\Delta\alpha = 0.1$ method, despite the large changes in the buffet-onset boundary.

The blue curve represents the 30% margin to buffet-onset boundary and is computed directly from the red buffet-onset curve. For normal operation, only operating conditions below the buffet-margin curve can be considered. The absolute maximum ML/D value for each configuration is shown in pink. Two specific contours for the optimization configuration (one for the baseline configuration) are highlighted: the contour of 99% $(ML/D)_{\max}$ for the particular design is shown in blue, and the contour of 99% $(ML/D)_{\max}$ for the baseline configuration is shown in red. The motivation for plotting these 99% contours is that airliners typically fly between the Mach number yielding maximum range (approximated by the maximum ML/D value in the figures) and a higher Mach number that yields a 1% fuel-burn penalty but decreases in the flight time. The area enclosed by both of these contours is used to quantify the robustness of the design in these figures. The areas are scaled by a factor of 100^2 so that the area of the rectangle measuring 0.01 in M and 0.01 in C_L has unit area.

Table 12: Operating conditions for each optimization. Red diamonds denote separation-constrained points. Operating conditions for Case 5.7 are determined by the optimization process itself. Zero weight means that only the flight condition is considered for the constraints.

Case	Point	Weights ($\mathcal{W}_i$)	Mach	C_L	Re	M–C_L plot
5.1	1	1	0.85	0.500	43.00×10^6	
5.2	1	1/3	0.85	0.500	43.00×10^6	
	2	1/3	0.85	0.650	43.00×10^6	
	3	1/3	0.89	0.456	45.00×10^6	
5.3	1	2/3	0.85	0.500	43.00×10^6	
	2	1/6	0.85	0.650	43.00×10^6	
	3	1/6	0.89	0.456	45.00×10^6	
5.4	1	1	0.85	0.500	43.00×10^6	
	2	0	0.85	0.650	43.00×10^6	
	3	0	0.89	0.456	45.00×10^6	
5.5	1	1/5	0.845	0.490	42.75×10^6	
	2	1/5	0.845	0.445	42.75×10^6	
	3	1/5	0.845	0.408	42.75×10^6	
	4	1/5	0.835	0.467	42.24×10^6	
	5	1/5	0.855	0.418	43.25×10^6	
	6	0	0.85	0.650	43.00×10^6	
	7	0	0.89	0.456	45.00×10^6	
5.6	1	1/5	0.845	0.490	42.75×10^6	
	2	1/5	0.845	0.445	42.75×10^6	
	3	1/5	0.845	0.408	42.75×10^6	
	4	1/5	0.835	0.467	42.24×10^6	
	5	1/5	0.855	0.418	43.25×10^6	
5.7	1	1/5	0.845	0.520	42.75×10^6	
	2	1/5	0.845	0.475	42.75×10^6	
	3	1/5	0.845	0.438	42.75×10^6	
	4	1/5	0.835	0.497	42.24×10^6	
	5	1/5	0.855	0.448	43.25×10^6	
	6	0	0.85	0.690	43.00×10^6	
	7	0	0.89	0.469	45.00×10^6	

The design operating conditions listed in Table 12 are shown as diamonds. The operating conditions considered for the objective function are shown in black, whereas the buffet-onset constraint

conditions are shown in red. The first buffet point ($M = 0.85$, $C_L = 0.65$) is at the nominal cruise Mach number and the C_L value corresponding to a $1.3g$ maneuver. The second buffet point ($M = 0.89, C_L = 0.456$) is 0.04 higher in Mach number, which is a typical margin between a nominal cruise Mach number and the maximum Mach number (M_{MO}) condition. The lift coefficient for this condition is adjusted to give the same dimensional lift as the nominal cruise condition at the same altitude.

Two additional regions are highlighted in black and orange, which we refer to as *integration regions*. They are constructed as follows: The Mach range is from 0.83 to 0.86, which corresponds to the typical range of operating Mach numbers for an aircraft such as the CRM. The upper line corresponds to the buffet-margin boundary, which is equivalent to specifying the maximum altitude the aircraft can fly for a particular weight. The bottom line corresponds to the reduced C_L for a 4000 ft decrease in altitude. To put it in another way, the integration region contains all operating conditions within 4000 ft of the buffet-constrained ceiling and for all normal operating Mach numbers. The aircraft spends the vast majority of cruising flight in this region. The black integration region corresponds to the baseline design, whereas the orange regions are adjusted to reflect the actual buffet-margin boundary for each design. In addition, the upper edge of the black region indicates how the buffet-onset boundary changes for each design relative to the baseline configuration for the specific Mach range of integration.

Given the results above (and the more extensive results in Kenway and Martins [6]), we recommend that physics-based buffet-onset constraints be enforced for aerodynamic and aerostructural shape optimization of transonic transports, as we do in this work. Although one must be cautious when using RANS to model separated flow, our results are encouraging because the approach was verified against another established numerical approach, and also validated against experimental data. The separation metric we develop herein is easily implemented and yields robust results, so it provides a much needed constraint formulation for the CFD-based wing design community not only for constraining buffet, but flow separation in general.
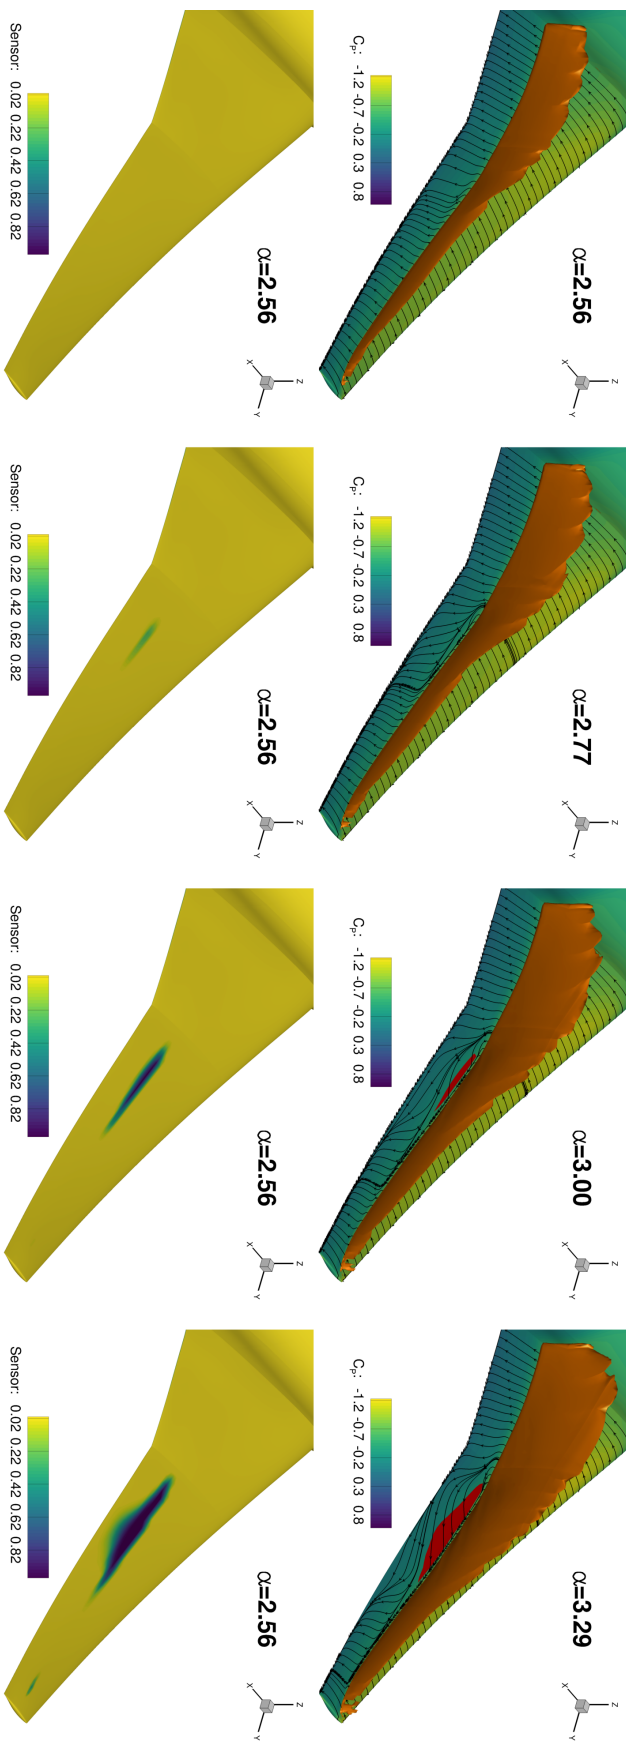
Figure 31: Progression of separated flow for the CRM configuration at $M = 0.85$ with increasing angle of attack. Top row shows the surface streamlines and pressure coefficient, as well as the reversed flow (red) and the shock (orange). Bottom row shows the value of the separation area integrand from Eq. (7.4).
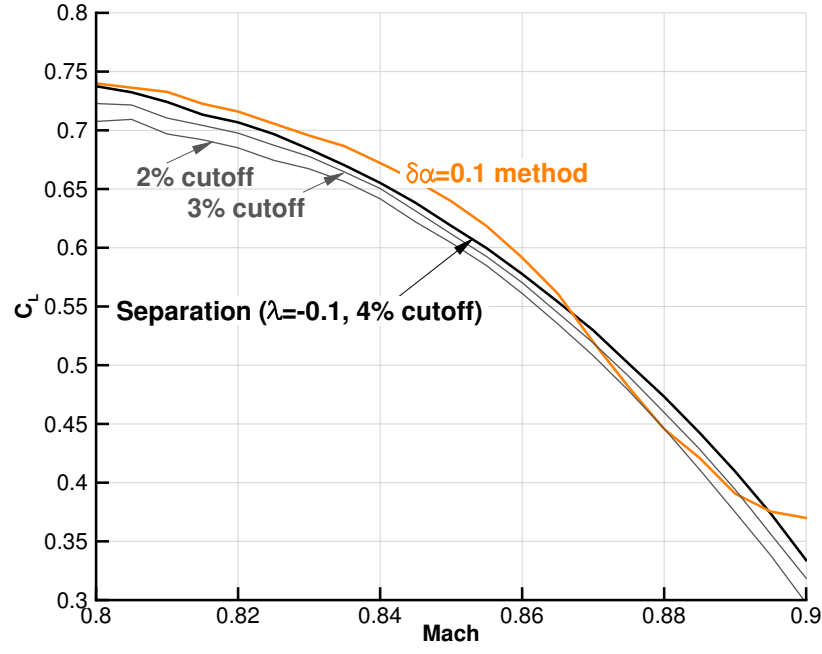
Figure 32: Results of the $\Delta\alpha = 0.1$ method compared with those of the proposed separation metric method.
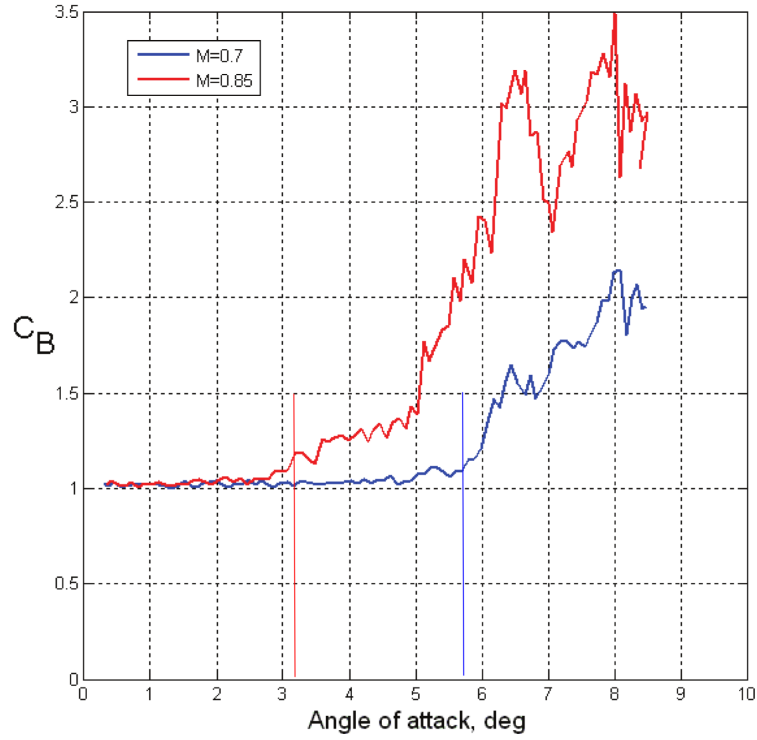
Figure 33: Buffet coefficient C_B obtained from wind tunnel data [4]. Vertical lines are the buffet-onset locations predicted by the separation sensor.
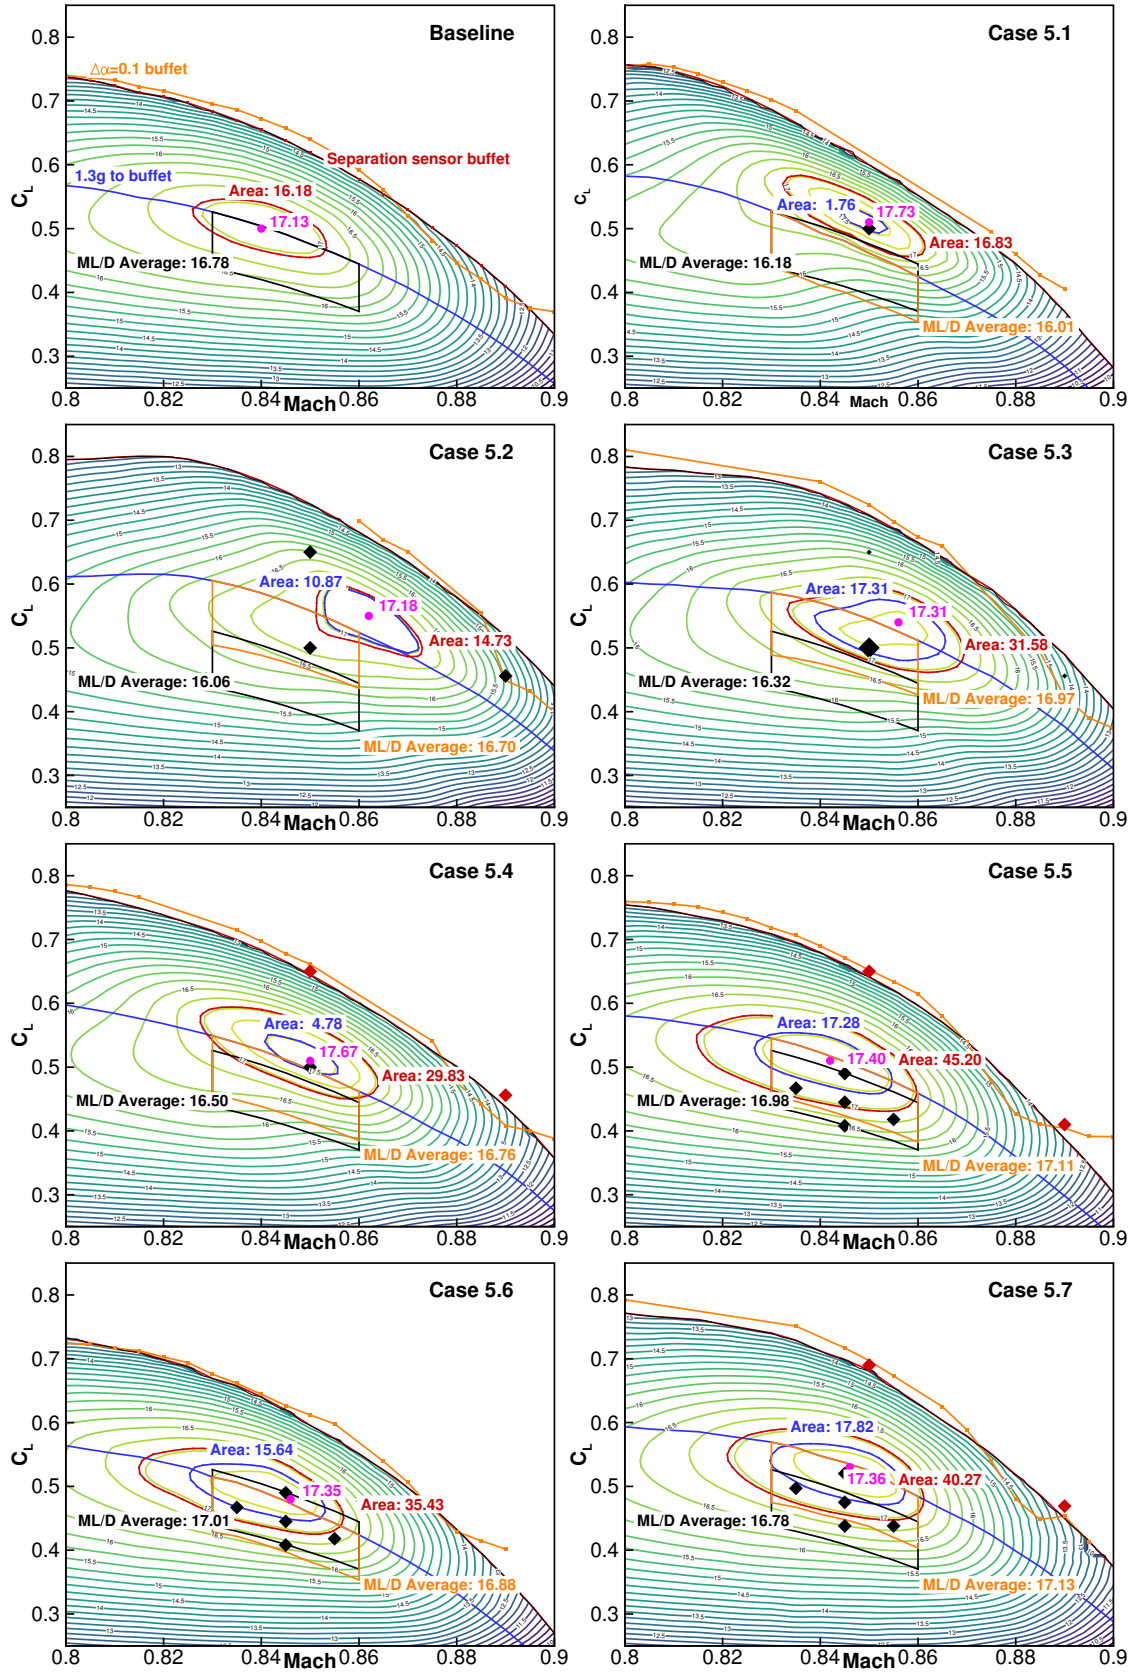
Figure 34: Contours of ML/D for the baseline and for each optimized configuration.

8 Flutter boundary constraint

8.1 Introduction

One of the issues that arise from the high-fidelity aerostructural optimizations performed in Section 4, especially in the case of fuel burn minimization, is that the resulting long span wings are probably flutter critical. To address this issue, we decided to develop a flutter constraint formulation suitable for gradient-based optimization. Several methods have been developed for computing flutter boundaries, but the critical requirements in the context of our optimization that have not been addressed so far are: 1) Efficient computation of the flutter boundary derivatives for optimization, and 2) Computation of *coupled* derivatives with respect to not only structural sizing variables, but with respect to wing shape as well.

In the research community, a wing's flutter boundary is often analyzed with a time-accurate coupled CFD-CSD solver similar to that proposed by Liu et al. [72]. However, this method incurs a high computational cost since hundreds if not thousands of time steps are required to simulate the flutter motion making it ill suited for optimization.

The current standard for flutter prediction in industry are methods based on panel codes, Doublet-Lattice Methods (DLM), and Transonic Small Disturbance (TSD) equations [73]. In the transonic region there is a significant reduction in the flutter speed, called the transonic dip or *flutter bucket*. Corrections using wind tunnel experimental data can be applied to panel method aerodynamic influence coefficients (AIC). For aerostructural design optimization this data is unavailable. Although other methods, such as full potential flow or the Euler models, are able to predict shock waves in the flow, they are limited in certain cases. As shown by several authors [74–76] viscous effects are necessary to predict flutter onset accurately.

Despite the aforementioned shortcomings of linear methods, the DLM method can be used to efficiently calculate flutter onset and thus can be incorporated as a flutter constraint in an otherwise high-fidelity aerostructural optimization. Using this approach, the high-fidelity CFD solver solving the RANS equations will account for nonlinearities in the flow solution, whereas the optimizer tries to smooth out any shock that may appear on the wing surface.

Stanford and Dunning [5], Stanford et al. [77] implemented a flutter constraint for several mass minimizations of a transport wing. To retain some of the nonlinearities in the flow, a transonic AIC matrix was constructed from an Euler CFD code. This AIC matrix is however not updated during the optimization process. Further, the derivatives of the mode shapes with respect to design variables are neglected during the optimization. However, the mode shapes are updated at every design iteration. For a mass minimization problem these approximations may be acceptable [78], but for a change in planform and aspect ratio, these approximations may give unrealistic results. For the flutter computation in this work, sensitivities of the mode shapes from the reduced system with respect to the design variables, including planform, will be considered.

We implement the DLM method coupled to an FEM solver TACS. The entire code base is differentiated to give sensitivities of an aggregated damping value (flutter eigenvalue) with respect to the design variables. To compute the sensitivities accurately and efficiently we employ an Automatic Differentiation (AD) adjoint method. Sensitivities are verified using a complex-step method. We plan to apply the newly developed flutter constraint to two geometries of interest, a flat plate verification problem and a transport wing. Further we perform optimization on both geometries where in both cases we consider planform design variables.

8.2 Methodology

There are several techniques and components necessary to enable flutter computations. In the following section we will outline the major characteristics of these components. Further, we outline the theory necessary for predicting flutter.

The doublet-lattice method (DLM) by Albano and Rodden [79] consists of a lifting surface method that is formulated in the frequency domain. The DLM is based on the vortex-lattice method (VLM) where the VLM is extended to harmonically oscillating surfaces where a flat wake is assumed. A substantial body of literature exists on the DLM. An excellent reference worth mentioning is the work done by Blair [80]. The DLM is widely adopted in the aeroelastic community and has been a valuable tool in flutter analysis of subsonic aircrafts. Commercial software tools such as MSC/NASTRAN have adopted the DLM [81]. In this work the implementation is based in part on the method of Albano and Rodden [79], and the extension by Rodden et al. [82].

In this work, we use a flutter analysis that takes the following form:

$$\mathbf{F}(s) = \left[s^2\mathbf{M}(\mathbf{x}) + \mathbf{K}(\mathbf{x}) - q_\infty\mathbf{A}(s)\right]\mathbf{u} = 0 \tag{8.1}$$

where $\mathbf{F}(s)$ is the overall flutter system, $\mathbf{M}(\mathbf{x})$ and $\mathbf{K}(\mathbf{x})$ are the mass and stiffness matrices from the finite-element equations, which are functions of the design variables $\mathbf{x}$. Furthermore, q_∞ is the dynamic pressure and $\mathbf{A}(s) = \mathbf{T}^T\mathbf{A}_{\mathrm{IC}}(s)\mathbf{T}$, where $\mathbf{A}_{\mathrm{IC}}$ is the aerodynamic influence coefficient matrix, and $\mathbf{T}$ is the load and displacement transfer interpolation matrix. The Laplacian parameter (eigenvalue) $s = \omega(\gamma + i)$ is complex and gives the frequency and damping of the motion. Note that the influence coefficient matrix is complex and is a nonlinear function of s. Therefore, the flutter equation (8.1) is a generalized nonlinear eigenvalue problem with a solution given by the triplet $(s, \mathbf{v}, \mathbf{u})$, where s is the complex eigenvalue and $\mathbf{v}$ and $\mathbf{u}$ are the left and right eigenvectors, respectively. The triplet satisfies the following equations:

$$\left[s^2\mathbf{M}(\mathbf{x}) + \mathbf{K}(\mathbf{x}) - q_\infty\mathbf{A}(s)\right]\mathbf{u} = 0$$
$$\mathbf{v}^H\left[s^2\mathbf{M}(\mathbf{x}) + \mathbf{K}(\mathbf{x}) - q_\infty\mathbf{A}(s)\right] = 0$$

Instead of using the full eigenvalue problem (8.1), flutter analysis techniques often employ a reduced eigenproblem using a small number of natural frequencies. The reduced modes are the eigenvectors of the problem:

$$\left[\mathbf{K}(\mathbf{x}) - \omega_i^2\mathbf{M}(\mathbf{x})\right]\mathbf{u} = 0,$$

where ω_i is the natural frequency. The eigenvectors $\mathbf{u}$ for $i = 1, \ldots, r$ are collected in the matrix $\mathbf{Q}_r \in \mathbb{R}^{n \times r}$ where n is the size of the square mass and stiffness matrices. These eigenvectors are $\mathbf{M}$-orthonormal, such that $\mathbf{Q}_r^T\mathbf{M}\mathbf{Q}_r = \mathbf{I}_r$. The reduced eigenproblem can now be written as follows:

$$\mathbf{F}_r(\tilde{s}) = \left[\tilde{s}^2\mathbf{M}_r + \mathbf{K}_r - q_\infty\mathbf{A}_r(\tilde{s})\right]\mathbf{u}_r = 0, \tag{8.2}$$

with the solution $(\tilde{s}, \mathbf{u}_r, \mathbf{v}_r)$. The reduced matrices take the form:

$$\mathbf{M}_r = \mathbf{Q}_r^T\mathbf{M}\mathbf{Q}_r = \mathbf{I}_r \in \mathbb{R}^{r \times r},$$
$$\mathbf{K}_r = \mathbf{Q}_r^T\mathbf{K}\mathbf{Q}_r = \mathrm{diag}\{\omega_i^2\} \in \mathbb{R}^{r \times r},$$
$$\mathbf{A}_r(\tilde{s}) = \mathbf{Q}_r^T\mathbf{A}(\tilde{s})\mathbf{Q}_r \in \mathbb{C}^{r \times r}.$$

Note that $\mathbf{A}_r$ has no sparsity structure and is a dense matrix in general, while $\mathbf{M}_r$ and $\mathbf{K}_r$ are sparse in general.

In this work, we use a shifted Lanczos method, which is a generalized eigenvalue solution algorithm. The Lanczos algorithm extracts eigenvalues for symmetric generalized eigenvalue problems.

Here, we use this algorithm to solve for the natural frequencies of the structural problem without aerodynamic loads:

$$\mathbf{Ku} = \lambda \mathbf{Mu}.$$

Instead of solving this problem directly, we use a shift and invert strategy to zero-in on the desired spectrum to reduce the number of iterations required. This shift and invert technique produces the following eigenproblem that has the same eigenvectors but different eigenvalues [83, 84]

$$\mathbf{M}(\mathbf{K} - \sigma \mathbf{M})^{-1} \mathbf{Mu} = \mu \mathbf{Mu},$$

where the transformed eigenvalue μ is related to the original eigenvalue λ through the relationship:

$$\mu = \frac{1}{\lambda - \sigma}.$$

When σ is chosen such that it lies close to the desired λ, the corresponding transformed eigenvalues, μ, become well separated, making the Lanczos algorithm more efficient.

The Lanczos algorithm uses an $\mathbf{M}$-orthonormal subspace, written as $\mathbf{V}_m \in \mathbb{R}^{n \times m}$, such that $\mathbf{V}_m^T \mathbf{MV}_m = \mathbf{I}_m$. In exact arithmetic, this subspace can be formed directly from the Lanczos three-term recurrence. However, the resulting subspace loses orthogonality as the algorithm converges to an eigenvalue due to numerical truncation errors. Instead, we use an expensive, but effective, full-orthonormalization procedure (Gram–Schmidt) that enforces $\mathbf{M}$-orthonormality.

The Lanczos method can be easily extended to find multiple eigenpairs $(\lambda_i, \mathbf{u})$.

Lanczos method for computing eigenvalues and eigenvectors of $\mathbf{Ku} = \lambda \mathbf{Mu}$
Given: m, $\hat{\mathbf{v}}_1$, σ, ϵ_{tol}
Factor the matrix $(\mathbf{K} - \sigma \mathbf{M})$
Set $i = 1$
while $i \leq m$ **do**
    $\hat{\mathbf{v}}_{i+1} = (\mathbf{K} - \sigma \mathbf{M})^{-1} \mathbf{Mv}_i$
    Set $j = 1$
    while $j \leq i$ **do**   ▷ Full $\mathbf{M}$-orthonormalization
        $h_{ji} = \mathbf{v}_j^T \mathbf{M} \hat{\mathbf{v}}_{i+1}$
        $\hat{\mathbf{v}}_{i+1} \leftarrow \hat{\mathbf{v}}_{i+1} - h_{ji} \mathbf{v}_j$
        $j \leftarrow j + 1$
    end while
    $\alpha_i \leftarrow h_{ii}$
    $\beta_i = \sqrt{\hat{\mathbf{v}}_{i+1}^T \mathbf{M} \hat{\mathbf{v}}_{i+1}}$
    $\mathbf{v}_{i+1} = \hat{\mathbf{v}}_{i+1} / \beta_i$
    $\mathbf{T}_i = \text{tridiag}_k \{\beta_k, \alpha_k, \beta_{k-1}\}$ ▷ Solve the reduced eigenproblem
    Solve $\mathbf{T}_i \mathbf{y}_i = \theta \mathbf{y}_i$ for $(\theta, \mathbf{y}_i)$
    if $\beta_i \mathbf{y}_i^T \mathbf{e}_i < \epsilon_{tol}$ **then**   ▷ Test for convergence
        $\mathbf{u} = \mathbf{V}_i \mathbf{y}_i$
        $\lambda = \frac{1}{\theta} + \sigma$
        break
    end if
    $i \leftarrow i + 1$
end while

The flutter solution algorithm that we use is given by [85], which solves the reduced nonlinear eigenvalue problem Eq. (8.2). This method is a secant method applied to the determinant equation:

$$\Delta(s) = \det \mathbf{Q}_r^T \mathbf{F}(s) \mathbf{Q}_r.$$

Note that the columns of $\mathbf{Q}_r \in \mathbb{R}^{n \times r}$ are the eigenvectors from the natural frequency eigenproblem. Given initial guesses s_1, and s_2, the method computes s_{k+2} as follows:

$$s_{k+2} = \frac{s_{k+1}\Delta(s_k) - s_k\Delta(s_{k+1})}{\Delta(s_k) - \Delta(s_{k+1})},$$

the iteration is continued until $|\Delta(s_{k+2})| \leq \epsilon_{tol}$ for some specified tolerance. In this implementation all flutter eigenvalues s, at a velocity (or dynamic pressure) range of interest, are found for mode i first before advancing to the next mode.

When evaluating the flutter determinant, the Aerodynamic Influence Coefficient (AIC) matrix is not evaluated explicitly but rather interpolated using 4^{th} order b-splines from a precomputed AIC matrices at a range of reduced frequencies. These AIC matrices are precomputed at startup for a specified range of reduced frequencies where the upper limit is the maximum reduced frequency k_{max} of the reduced problem. The number of evaluation points for this range is chosen a priori and is equally spaced over the selected range. This procedure is commonly applied for an AIC method like the DLM and will improve the overall efficiency of the code as the evaluation of the AIC matrix is generally expensive operation. Although the AIC matrix is independent of the structures a new set of precomputed AIC matrices are constructed for each optimization iteration as the planform of the geometry, hence the aerodynamic mesh is changed or updated.

We use a Kreisselmeier–Steinhauser (KS) function [52, 86] to aggregate the real parts of all the flutter eigenvalues $\Re(s_i)$ for each mode, into a single value. Note that s_i is the flutter eigenvalue at the i^{th} dynamic pressure for that mode. The KS function gives a differentiable maximum and can be defined as

$$KS(\mathbf{x}) = \frac{1}{\rho} \ln \left(\sum_{i=1}^{N} e^{\rho x_i} \right) \tag{8.3}$$

where ρ is the KS parameter. This parameter can be used as a buffer or tolerance. As $\rho \to \infty$ the KS function approaches the maximum, but too large a value can cause sharp changes in the gradient. To avoid numerical difficulties due to overflow we use an alternate form of the KS function

$$KS(\mathbf{x}) = a + \frac{1}{\rho} \ln \left(\sum_{i=1}^{N} e^{\rho(x_i - a)} \right) \tag{8.4}$$

where $a = \max(\mathbf{x})$.

Here we present a brief verification study of the DLM code and the methods used in our flutter analysis. Natural frequencies of the flat plate problem presented in Section 8.6 are obtained using the Lanczos method. The subspace size was 10, and 4 modes are used. A comparison with MSC/NASTRAN is shown in Table 13, where the overall agreement is good.

Flutter analysis is obtained using the uCRM wingbox presented in Section 8.8. Specifically, in Fig. 35, we compare the flutter eigenvalues for the first eight modes of the structure with MSC/NASTRAN. Overall trends of our implementation are good although some discrepancy is present. This is likely due to several reasons such as: (1) The precomputed AIC matrix could have different resolution when constructed (evaluated a different number of reduced frequencies), resulting in a discrepancy when the AIC is interpolated for a specific reduced frequency. (2) The

use of a parabolic approximation kernel for the AIC matrix rather than the quartic approximation [82] found in MSC/NASTRAN and different flutter equation formulation used in MSC/NASTRAN [87].

Table 13: Flat plate natural frequencies agree well with NASTRAN. Units presented here are radians.

Mode number	NASTRAN	TACS	Error (%)
1	7.130	7.134	0.058
2	44.575	44.693	0.263
3	57.704	58.268	0.969
4	125.070	125.838	0.610
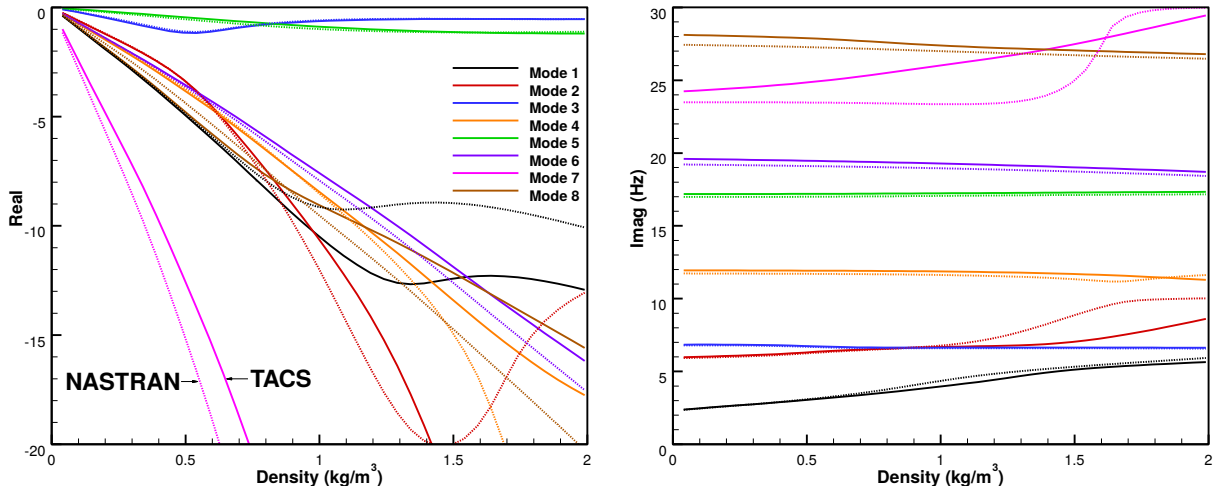
Figure 35: Real and imaginary parts shown for the flutter eigenvalue of a CRM wing-box. Solid lines show results from this proposed work, the DLM with TACS flutter implementation, where the dashed lines are from MSC/NASTRAN aeroelastic module.

As in the other work described in this report, we use FFD for the wing shape parametrization [60]. The FFD for the plate used in this section is shown in Fig. 36, where the red points represent the control points and the black lines connecting them denotes the outer edges of the volume.

8.3 Adjoint implementation

To produce accurate gradient information, the code is analytically as well as automatically differentiated (AD). We now describe the automatic differentiation adjoint approach that we used in this work. We write the quantity of interest as $\mathbf{I}$, which depends implicitly on the independent variables of interest $\mathbf{x}$. In the case of aerodynamics this quantity could for example be the lift or the drag coefficients. State variables $\mathbf{y}(\mathbf{x})$ implicitly depend on the independent variables.

$$\mathbf{I} = \mathbf{I}(\mathbf{x}, \mathbf{y}(\mathbf{x}))$$

The residual equation can be written as

$$\mathbf{r} = \mathbf{R}(\mathbf{x}, \mathbf{y}(\mathbf{x})) = 0$$

For the derivation of the adjoint we start with the total derivative of the value of interest

$$\frac{d\mathbf{I}}{d\mathbf{x}} = \frac{\partial \mathbf{I}}{\partial \mathbf{x}} + \frac{\partial \mathbf{I}}{\partial \mathbf{y}} \frac{d\mathbf{y}}{d\mathbf{x}} \tag{8.5}$$

The total derivative of the residual equation can also be obtained as

$$\frac{d\mathbf{r}}{d\mathbf{x}} = \frac{\partial \mathbf{R}}{\partial \mathbf{x}} + \frac{\partial \mathbf{R}}{\partial \mathbf{y}} \frac{d\mathbf{y}}{d\mathbf{x}} = 0$$

Rewriting the above equation we get

$$\frac{\partial \mathbf{R}}{\partial \mathbf{y}} \frac{d\mathbf{y}}{d\mathbf{x}} = -\frac{\partial \mathbf{R}}{\partial \mathbf{x}} \tag{8.6}$$

$$\Rightarrow \frac{d\mathbf{y}}{d\mathbf{x}} = -\left[\frac{\partial \mathbf{R}}{\partial \mathbf{y}}\right]^{-1} \frac{\partial \mathbf{R}}{\partial \mathbf{x}} \tag{8.7}$$

Note that the inverse of the square Jacobian $\partial \mathbf{R}/\partial \mathbf{y}$ is not calculated explicitly but rather used to indicate that this matrix is solved as a linear system along with some right hand side vector. Combining Eq. (8.7) and Eq. (8.5) we obtain

$$\frac{d\mathbf{I}}{d\mathbf{x}} = \frac{\partial \mathbf{I}}{\partial \mathbf{x}} - \frac{\partial \mathbf{I}}{\partial \mathbf{y}} \left[\frac{\partial \mathbf{R}}{\partial \mathbf{y}}\right]^{-1} \frac{\partial \mathbf{R}}{\partial \mathbf{x}} \tag{8.8}$$

The *adjoint equations* are written as

$$\left[\frac{\partial \mathbf{R}}{\partial \mathbf{y}}\right]^{T} \psi = -\left[\frac{\partial \mathbf{I}}{\partial \mathbf{y}}\right]^{T}$$

where ψ is the adjoint matrix and is of size $n_x \times n_I$. This linear system needs to be solved n_I times or for each column of $[\partial I/\partial y]^{T}$. Substituting in to Eq. (8.8) we get

$$\frac{d\mathbf{I}}{d\mathbf{x}} = \frac{\partial \mathbf{I}}{\partial \mathbf{x}} - \psi^{T} \frac{\partial \mathbf{R}}{\partial \mathbf{x}} \tag{8.9}$$

We could also solve the system with the *direct method* where we would solve

$$\frac{\partial \mathbf{R}}{\partial \mathbf{y}} \frac{d\mathbf{y}}{d\mathbf{x}} = -\frac{\partial \mathbf{R}}{\partial \mathbf{x}}$$

This linear system needs to be solved n_x times to get the full $d\mathbf{y}/d\mathbf{x}$. Since many aerodynamic optimization formulations contain many more design variables than outputs of interest (or $n_x \gg n_I$) we solve the adjoint equations in order to obtain the sensitivities as efficiently as possible.

Automatic differentiation, also know as algorithmic differentiation, is a well established method that systematically applies the differentiation chain rule to source code. The method uses source transformation tools that takes in the original computer program, augments it, and generates a new code, such that it computes the analytical derivatives along with the original program [88, 89].

Two modes exist, the forward mode and the reverse mode. For a generic system with scalar input x and output y we can write it as:

$$\text{System} \quad x \rightarrow \boxed{F(x)} \rightarrow y$$

$$\text{Forward AD} \quad \dot{x} \rightarrow \boxed{F'(x)} \rightarrow \dot{y}$$

$$\text{Reverse AD} \quad \bar{x} \leftarrow \boxed{F'^*(x)} \leftarrow \bar{y}$$

where the arrows represent the flow of information, the box represents the system or a function. The forward mode, know as tangent, is denoted with a dot (˙) over the variable. Given some small variations on the input (independent) variables $\mathbf{x}$ we can compute the resulting variations of the dependent variables $\mathbf{y}$. The Jacobian matrix $\mathbf{J}$ contains the partial derivatives of each output (dependent) variable y_j with respect to each independent variable x_i. The forward mode thus computes $d\mathbf{y} = \mathbf{J} \times d\mathbf{x}$ for each given $d\mathbf{x}$ or

$$d\mathbf{y} = \mathbf{J} \times d\mathbf{x}$$

$$
\begin{bmatrix} \dot{y}_1 \\ \dot{y}_2 \\ \vdots \\ \dot{y}_n \end{bmatrix}
=
\begin{bmatrix}
\frac{\partial y_1}{\partial x_1} & \frac{\partial y_1}{\partial x_2} & \cdots & \frac{\partial y_1}{\partial x_m} \\
\frac{\partial y_2}{\partial x_1} & \frac{\partial y_2}{\partial x_2} & \cdots & \frac{\partial y_2}{\partial x_m} \\
\vdots & \vdots & \ddots & \vdots \\
\frac{\partial y_n}{\partial x_1} & \frac{\partial y_n}{\partial x_2} & \cdots & \frac{\partial y_n}{\partial x_m}
\end{bmatrix}
\times
\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \vdots \\ \dot{x}_m \end{bmatrix}
$$

Conversely the reverse mode, known as the adjoint, is denoted by a bar (¯) over the variable. The order of operations reverses and we compute the transposed Jacobian product $d\mathbf{x} = \mathbf{J}^* \times d\mathbf{y}$ for each given $d\mathbf{y}$ or

$$d\mathbf{x} = \mathbf{J}^* \times d\mathbf{y}$$

$$
\begin{bmatrix} \bar{x}_1 \\ \bar{x}_2 \\ \vdots \\ \bar{x}_n \end{bmatrix}
=
\begin{bmatrix}
\frac{\partial y_1}{\partial x_1} & \frac{\partial y_2}{\partial x_1} & \cdots & \frac{\partial y_n}{\partial x_1} \\
\frac{\partial y_1}{\partial x_2} & \frac{\partial y_2}{\partial x_2} & \cdots & \frac{\partial y_n}{\partial x_2} \\
\vdots & \vdots & \ddots & \vdots \\
\frac{\partial y_1}{\partial x_m} & \frac{\partial y_2}{\partial x_m} & \cdots & \frac{\partial y_n}{\partial x_m}
\end{bmatrix}
\times
\begin{bmatrix} \bar{y}_1 \\ \bar{y}_2 \\ \vdots \\ \bar{y}_m \end{bmatrix}
$$

In other words, the gradient of the independent variable is a linear combination of the variation in the dependent variable. This is a very important observation particularly in the case where we have many fewer output variables than input variables. In this work we used the AD source-transformation tool, Tapenade [90, 91] to obtain part of the derivatives used in this work.

8.4   Flutter derivatives

A simplified flow of information for the flutter calculation, obtaining derivatives in forward and reverse mode can be written as:

$$\mathbf{x} \rightarrow \mathbf{K}, \mathbf{M} \rightarrow \mathbf{Q}_r, \mathbf{K}_r, \mathbf{M}_r \rightarrow s_i \rightarrow KS_i$$

$$\dot{\mathbf{x}} \rightarrow \dot{\mathbf{K}}, \dot{\mathbf{M}} \rightarrow \dot{\mathbf{Q}}_r, \dot{\mathbf{K}}_r, \dot{\mathbf{M}}_r \rightarrow \dot{s}_i \rightarrow \dot{K}S_i$$

$$\bar{\mathbf{x}} \leftarrow \overline{\mathbf{K}}, \overline{\mathbf{M}} \leftarrow \overline{\mathbf{Q}}_r, \overline{\mathbf{K}}_r, \overline{\mathbf{M}}_r \leftarrow \bar{s}_i \leftarrow \overline{KS}_i$$

where $\mathbf{K}, \mathbf{M}$ are the siffness and mass matrices and $\mathbf{Q}_r$ the reduced natural modes shapes. The KS_i represents the real part of flutter eigenvalue s_i for mode i aggregated over the velocity range interest using the KS function defined in Eq. (8.4).

Total derivatives of the flutter constraints $d\mathbf{KS}/d\mathbf{x}$ are generated with a combination of AD and analytically differentiated code. Routines such as LAPACK's [92] complex and real linear solutions are not automatically differentiated, but instead must be differentiated analytically and implemented as such.

The adjoint method was also implemented to obtain total derivatives for the flutter constraint where the partial derivatives are computed mainly with AD. However, likely due to bugs in the implementation, accuracy was lacking compared to full differentiation of the code or complex step. More work on the adjoint implementation is therefore needed as that will improve efficiency over automatically differentiating the full code.

8.5 Derivative verification

To demonstrate correct and accurate derivatives we go through a rigorous verification. Each function in the code is unit tested where sensitivities are computed using a second order central finite-difference stencil, a complex-step method [93] as well as forward and reverse mode AD. The finite-difference stencil used here is

$$\frac{dI}{dx} = \frac{I(x+h) - I(x-h)}{2h} + O(h^2),$$

with a step size ranging from $h = 10^{-3}$ to $h = 10^{-6}$ depending on the function under consideration. It has to be mentioned that in order to get a reasonable prediction with finite-difference, a step size study was performed to get the most accurate gradient possible. For the complex-step method the sensitivity of a function is computed as

$$\frac{dI}{dx} = \frac{\text{Im}[I(x+ih)]}{h} + O(h^2),$$

where $i = \sqrt{-1}$ with a step size of $h = 10^{-40}$. The complex-step method does not suffer from subtractive cancellation errors as does the finite-difference method. The step size can be made very small, hence the $O(h^2)$ truncation error becomes negligible, thus giving same precision as the function $I(x)$.

One drawback of the complex-step however is that it cannot be used by default in programs that already contain complex numbers and do complex arithmetic. Such programs need to be modified such that the complex number is represented by two real numbers, one for the real part and one for the imaginary part. Complex arithmetic thus need to be done by hand and cannot be done using intrinsic functions. In this work, all code has been modified such that it utilizes only real numbers for complex calculations and is complex-step safe.

Since we have many more design variables than functions of interest $I(x)$, we want to employ the adjoint or the reverse mode AD in the optimization. A common technique to verify the reverse mode is to implement the forward mode and perform a *dot product test*. The forward mode and the complex-step should also match close to machine precision so implementing the forward mode as well is an important step. We can then write the *dot product test* [94]

$$\begin{aligned}
\bar{\mathbf{x}}^* \dot{\mathbf{x}} &= (\mathbf{J}^* \bar{\mathbf{y}})^* \dot{\mathbf{x}} \\
&= \bar{\mathbf{y}}^* (\mathbf{J} \dot{\mathbf{x}}) \\
&= \bar{\mathbf{y}}^* \dot{\mathbf{y}} \qquad set \quad \bar{\mathbf{y}} = \dot{\mathbf{y}} \\
&= \dot{\mathbf{y}}^* \dot{\mathbf{y}}
\end{aligned}$$

This equality should be exact to machine precision.

8.5.1 Intermediate derivatives

We now present derivative results of the flutter constraint with respect to the reduced stiffness matrix $\mathbf{K}_r$, the aerodynamic mesh nodes $\mathbf{X}_{aero}$ and $\mathbf{Q}_m$ which are the reduced mode shapes $\mathbf{Q}_r$ transferred on to the structure. As shown in Table 14 sensitivities are very accurate as the complex-step and the reverse AD agree very well. As expected, the second order finite difference method does however not perform as well, and was sensitive to variation in step size. In order to get the best finite difference derivative the step size was varied from $h = 10^{-3}$ to $h = 10^{-6}$ depending on which derivative was being calculated.

8.5.2 Flat plate derivatives

Sensitivities for the flat plate, geometry presented in Section 8.6 are considered. Sensitivities for the design variable of interest, chord, span and thickness are shown in Table 15. The thickness variable is for the entire plate as opposed to one thickness variable per element. As before finite difference offers less accuracy compared to AD or complex-step. Its noted that number of matching digits for the full chain is somewhat less than what is presented in Table 14 which is coming from the Fortran code. Once passed through the reverse implementation of TACS and the Lanczos method some accuracy is lost. The reason why this is, are not obvious at this point and need further evaluation.

Table 14: Intermediate sensitivities of aggregated eigenvalue mode 1 with respect to a single value in the reduced stiffness $\mathbf{K}_r$ matrix, aerodynamic mesh points $\mathbf{X}_{aero}$ matrix and the reduced transferred mode shapes $\mathbf{Q}_m$. Reverse AD compares very well to complex-step and is close to matching machine precision. Finite difference however lacks accuracy compared to the other methods.

	$\frac{\partial KS_1}{\partial \mathbf{K}_r}$	$\frac{\partial KS_1}{\partial \mathbf{X}_{aero}}$	$\frac{\partial KS_1}{\partial \mathbf{Q}_m}$
Finite Difference	0.0043664132309829	0.1058895990890818	-0.0230063250672430
Complex-step	0.0043664135454205	0.1058895982342961	-0.0230063273513534
AD (Reverse)	0.0043664135454142	0.1058895982342963	-0.0230063273513523

Table 15: Sensitivities of the of aggregated eigenvalue mode 1 with respect to the design variables, chord, span and plate thickness.

	$\frac{dKS_1}{dx_{\text{chord}}}$	$\frac{dKS_1}{dx_{\text{span}}}$	$\frac{dKS_1}{dx_{\text{thickness}}}$
Finite Difference	1.59234199	-2.55801140	3930.43055292
Complex-step	1.5903748	-2.56075992	3929.75519514
AD (Reverse)	1.5981803	-2.56035533	3929.74618753

Table 16: Derivatives for flutter constraint on mode 1 versus the design variables for 8 spars, 8 ribs and 8 skin elements of the structure. RAD is reverse AD, FD is finite differencing, and CS is the complex step.

Variable	i	FD	CS	RAD	Relative error FD-CS	Relative Error RAD-CS
Spar Elements	1	-2.014985468	-1.999974758	-1.999965907	7.51E-003	4.43E-006
	2	-0.088696547	-0.090194566	-0.090198427	1.66E-002	4.28E-005
	3	-0.095666766	-0.085257553	-0.085258866	1.22E-001	1.54E-005
	4	0.203193368	0.210156947	0.210212359	3.31E-002	2.64E-004
	5	0.263933667	0.259278622	0.259260509	1.80E-002	6.99E-005
	6	-0.071987279	-0.032305006	-0.032318974	1.23E+000	4.32E-004
	7	0.006071139	-0.000334429	-0.000327057	1.92E+001	2.25E-002
	8	-0.020179093	0.000032238	0.000026615	6.27E+002	2.11E-001
Rib Elements	1	0.253309222	0.224356643	0.224359214	1.29E-001	1.15E-005
	2	0.191743890	0.175807433	0.175809351	9.06E-002	1.09E-005
	3	0.207341327	0.230292459	0.230290811	9.97E-002	7.16E-006
	4	0.185916602	0.181973005	0.181975321	2.17E-002	1.27E-005
	5	0.237979611	0.233055462	0.233046967	2.11E-002	3.65E-005
	6	0.181381180	0.185529844	0.185526430	2.24E-002	1.84E-005
	7	0.231692141	0.231781620	0.231775388	3.86E-004	2.69E-005
	8	0.150392537	0.184741162	0.184729822	1.86E-001	6.14E-005
Skin Elements	1	0.441049044	0.434946729	0.434919882	1.40E-002	6.17E-005
	2	0.276180216	0.287471575	0.287450964	3.93E-002	7.17E-005
	3	0.247139578	0.285488076	0.285466885	1.34E-001	7.42E-005
	4	0.470635907	0.436378482	0.436469380	7.85E-002	2.08E-004
	5	0.328335009	0.287666266	0.287733113	1.41E-001	2.32E-004
	6	0.291624919	0.285354574	0.285422324	2.20E-002	2.37E-004
	7	0.485841301	0.477884284	0.477930152	1.67E-002	9.60E-005
	8	0.306040695	0.312406435	0.312433467	2.04E-002	8.65E-005

8.5.3 uCRM derivatives

Here we present sensitivities for the uCRM geometry presented in Section 8.8. Sensitivities for the KS function with respect to selected variables are listed in Table 16. As before the finite difference values are obtained using a second order stencil. A step size study would be desirable for each design variable or element, but this was not done in this case. Instead, the step size that gave the best overall results was chosen, which was $h = 10^{-6}$. This step size resulted in more elements having sensitivities that where closer to the complex-step. The last two columns give the relative error between the finite difference and complex-step and reverse AD and complex-step. We note that the reverse AD and complex-step compare very well as the relative error is three to six orders of magnitude smaller than compared to finite difference.
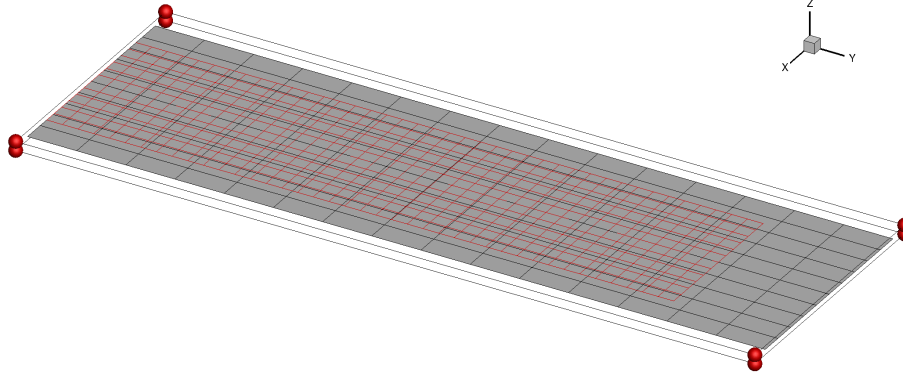
Figure 36: Flat plate structural and aerodynamic mesh shown in red and black respectively. The plate is cantilevered at the left edge. A Free-Form-Deformation (FFD) volume is also shown with 8 control points which are depicted as red spheres connected by solid black lines.

8.6 Flat plate analysis and optimization

8.6.1 Model description

The flat plate geometry shown in Fig. 36 is chosen for verification purposes due to its simplicity and short optimization turnaround. The flat plate structure, shown in red is embedded within a larger flat aerodynamic mesh, shown in gray. The red circles are control points of the FFD volume, which both meshes have been embedded in. The structural model consists of 12 elements in streamwise direction and 40 elements in the spanwise direction with a total of 480 finite elements. The finite element used here is the MITC4 shell element. The aerodynamic model consists of 10 elements in streamwise direction and 15 elements in the spanwise direction. Material properties, dimensions and discretization for the baseline flat plate are summarized in Table 17.

For the flat plate analysis and optimization the air density is kept fixed and the Mach number is set to zero resulting in incompressible flow. The velocity range is varied in 40 increments from 2 to 15 m/s for each design. The flight conditions are summarized in Table 18. The flutter problem is solved at each velocity (dynamic pressure effectively) for each mode. In this work we constrain the first 4 modes of the plate. The first 4 natural modes and mode shapes are calculated using the Lanczos algorithm with a subspace of size 10.

8.6.2 Problem statements

For this simple flat plate problem design variables are chosen as thickness of the entire plate $x_{\text{thickness}}$, its span x_{span} and chord length x_{chord}. No sweep, taper or dihedral is applied here. The objective is to maximize its range using the Breguet range equation,

$$R = \frac{V}{c_T} \frac{C_L}{C_D} \ln \left(\frac{W_{\text{init}}}{W_{\text{final}}} \right) \tag{8.10}$$

where R is range, V/c_T is the flight speed to thrust-specific fuel consumption ratio, C_L/C_D is the lift to drag ratio, and $W_{\text{init}}/W_{\text{final}}$ is the initial to final cruise weight ratio. For simplicity we assume that $V/c_T = 1$, and we define the cruise weights as

$$\begin{aligned} W_{\text{final}} &= W_{\text{fixed}} + W_{\text{plate}} \\ W_{\text{init}} &= W_{\text{final}} + W_{\text{fuel}} \end{aligned} \tag{8.11}$$

71

Table 17: Flat plate mechanical properties, dimensions and discretization of the structure and the aerodynamic surface.

Variable	Symbol	Value
Density	ρ_s	2800 kg/m^3
Modulus of Elasticity	E	70 GPa
Poisson ratio	ν	0.3
Yield stress	σ_y	400 KPa
Thickness	t	0.002 m
Structure Span	b_s	0.85 m
Structure Chord	c_s	0.21 m
Number of Finite elements, streamwise	n_x	12
Number of Finite elements, spanwise	n_y	40
Span	b	1.0 m
Chord	c	0.3 m
Number of DLM elements, streamwise	n_x	10
Number of DLM elements, spanwise	n_y	15
Planform area	A_{init}	0.3 m^2

where the W_{fixed} is a fixed weight and W_{fuel} is the fuel weight. The lift coefficient is fixed at $C_L = 0.5$ and the drag coefficient is calculated assuming it consists only of the lift induced drag.

$$C_D = \frac{C_L^2}{\pi e AR} \tag{8.12}$$

where the wing span efficiency factor is set to $e = 1$ for simplicity. The AR is the aspect ratio defined as $AR = b^2/S$ where b is the reference span and the reference area S is the planform area. The range can for example be increased by reducing the drag coefficient. The drag coefficient is reduced by increasing the aspect ratio as other parameters are fixed.

The chord and span directly affect the aspect ratio so we want to formulate the problem in terms of the aspect ratio rather than directly the chord and the span. This allows us to make contour plots that can give valuable insight into the design space of such a small problem. By adding an area equality constraint we ensure that there is a link between the chord and span. This prevents the chord and span from shrinking or growing as the optimizer will otherwise try to increase the aspect ratio by modifying these variables independently. This will reduced the number of design variables making this in effect a problem with two design variables, thickness and aspect ratio. Flutter constraints are added on the first 4 modes such that $KS_i \leq 0$ for $i = 1, \ldots, 4$ and the planform area must be kept constant. Due to the primitive nature of the flutter constraint applied here no minimum flutter speed or flutter point is defined explicitly. The constraint forces the geometry to be flutter free for the entire flight envelope. For this particular problem no flutter must occur for the velocity range of 2-15 m/s. One could then think of the minimum flutter point to be 15 m/s as no constraints are enforced for higher velocities. The initial area is given in Table 17. The side constraints are as follows, chord and span are specified such that the aspect ratio is allowed to vary from $1 \leq AR \leq 6$, and the material thickness of the plate is allowed to vary

Table 18: Flat plate operating conditions under investigation. The airspeed is incremented from 2 to 15 m/s in 40 increments.

Variable	Symbol	Value
Mach	M_∞	0
Lift Coefficient	C_L	0.5
Air Density	ρ_∞	1.225 kg/m^3
Air Speed Range	U_∞	2-15 m/s
# Speed Increments	n_U	40
Flight speed (range eqn.)	V	1.0 m/s
Thrust specific fuel consumption	c_T	1.0 lb/(lbf $\cdot$ h)
Fixed weight	W_{fixed}	1.0 kg
Fuel weight	W_{fuel}	0.25 kg

from $0.0012 \leq t \leq 0.0025$ m. The optimization problem is summarized in Table 19.

8.6.3 Design space analysis

Before running an optimization the design space is analyzed since we have the luxury of plotting the design space. A contour plot is generated by sweeping over both the aspect ratio range $1 \leq AR \leq 6$ and the thickness range $0.0012 \leq t \leq 0.0025$ m. Due to a relatively cheap cost of each analysis a grid of 32 steps in each variable is done giving a total of 1024 points. Figure 37 shows the contour plot of the objective function where the four constraints have been applied to the contour one at a time. The objective function appears smooth with a clear maximum at $AR = 6, t = 0.0012$ m as expected. For each flutter constraint where $KS_i \geq 0$ the objective function value has been blanked out as this part of the design space is infeasible. Constraints boundaries on modes 1,2,3,4 are shown with black, blue, red, and purple curves, respectively. For mode 1 there appear small pockets of infeasible region in the design space. This anomaly will be addressed later. For the other modes 2-4 we see that they are represented as continuous lines in the design space.

All constraints have been applied to the objective function contour plot in Fig. 38. The constraint pockets appearing on mode 1 will never be active as the these scattered regions are blanked out by mode 2 and 3. Similarly constraint boundaries for mode 2 and 3 appearing in the low aspect ratio region will never be active as mode 4 will blank them out. Hence modes 2,3,4 are only needed as constraints for this optimization as mode 1 will never be active. After having applied the constraints to the design space the optimum now appears to be close to $AR^* \approx 6, t^* \approx 0.0015$.

8.6.4 Optimization results

The flat plate is optimized using flight conditions in Table 18. Multiple random initial designs within and outside the feasible design space where chosen to test the robustness of the code and optimizer. The optimizer is able to find or get very close to the global optimum, but experiences numerical difficulties close to the global optimum in almost all cases.

To illustrate the issue we give the following example. One initial design is chosen ($AR = 4.38, t = 0.002$), which in this case is within the feasible design space. The optimizer finishes
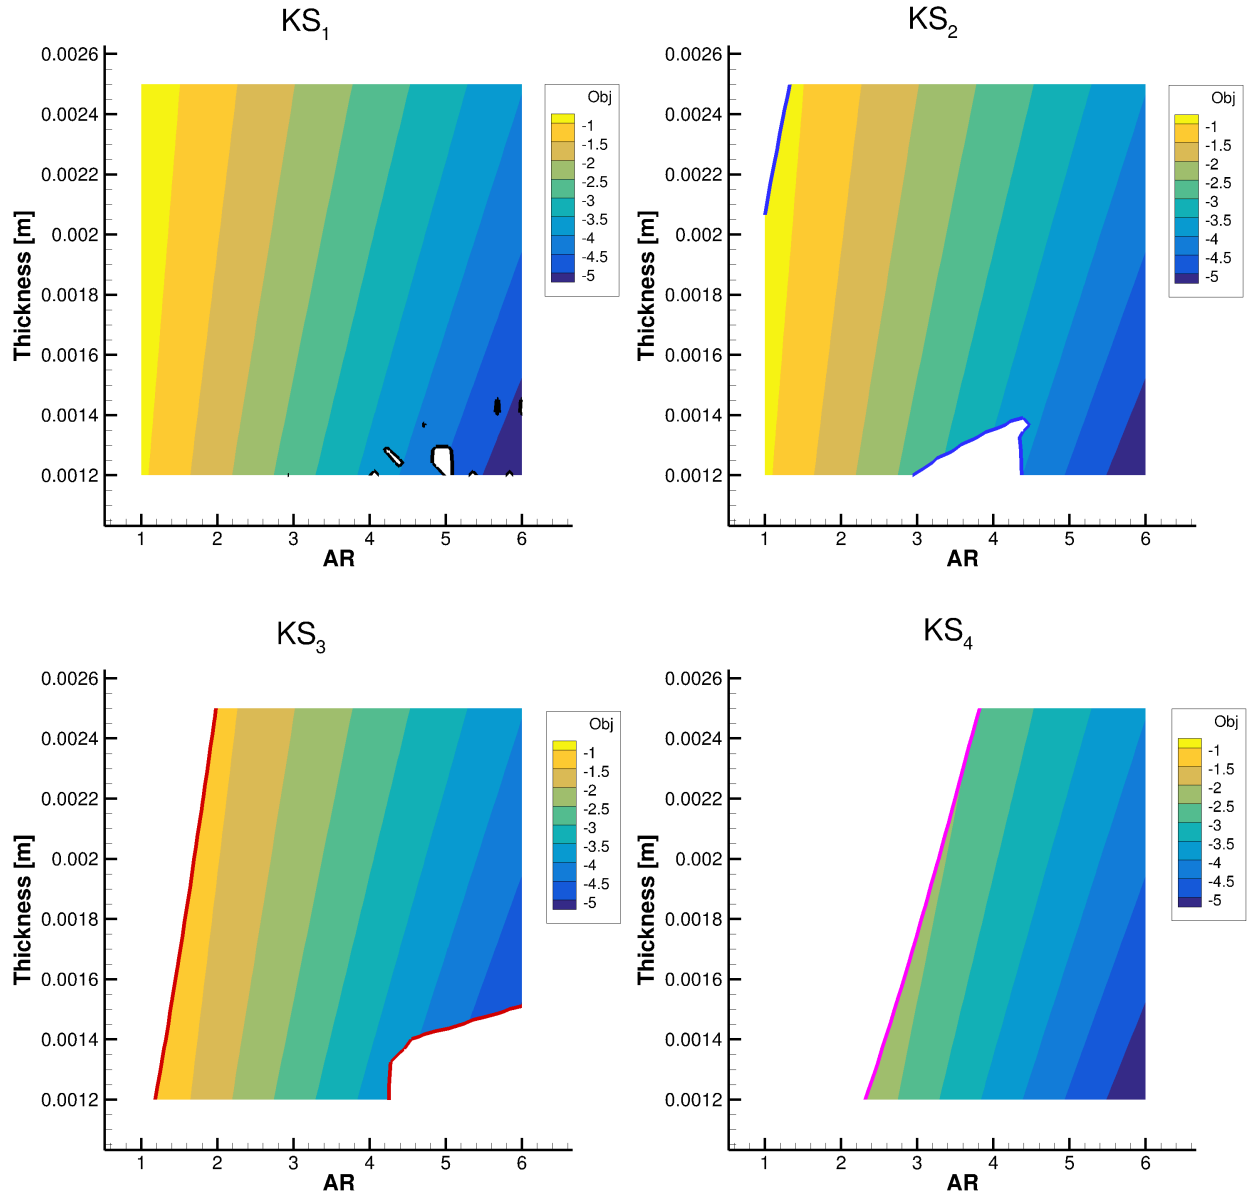
Figure 37: Contour plot of the objective function shown with the four flutter constraints applied to the contour plot. To generate the contour the design space is sampled using 32 points in both variables.

Table 19: Optimization formulation of the flat plate problem

	Function/variable	Description	Quantity
maximize	Range	Breguet equation	
with respect to	x_{span}	Plate span	1
	x_{chord}	Plate chord	1
	$x_{\text{thickness}}$	Plate thickness	1
		Total design variables	**3**
subject to	$A - A_{\text{init}} = 0.0$	Fixed plate area	1
	$\text{KS}_i \leq 0$	Aggregate damping for velocity range	4
		Total constraints	**5**
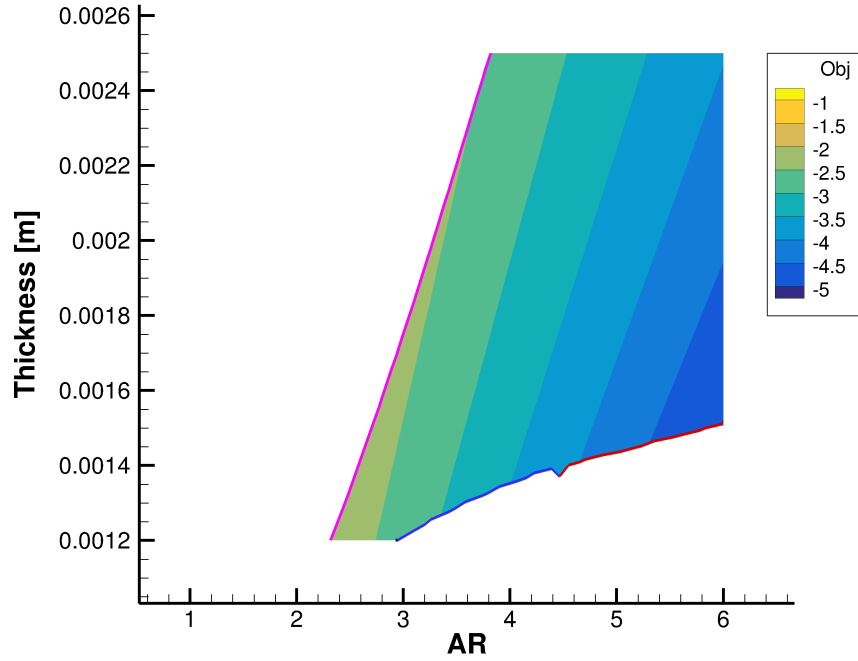
Figure 38: Objective function shown where the infeasible design space has been blanked out by the four flutter constraints.

but exits indicating numerical difficulties after 6 major iterations with the objective function and design variables equal to -5.024517 and $AR^* = 5.98, t^* = 0.001497$, respectively. Note that we are maximizing the range, so the value is in fact positive (5.024517). Although the result is labeled as the global optimum (with an asterisk) it is clearly not. The aspect ratio, for instance, is not at its maximum of 6, which indicates that the optimum found is not the global optimum. The initial and the optimized geometry along with their damping plot is shown in Fig. 39.

By inspecting Fig. 39, we see that the initial geometry (dashed lines) is a feasible design as it does not have any modes with positive damping values for the velocity range of interest 2-15 m/s. As stated in Table 19 the KS aggregate for each mode must be less than zero. For the optimized geometry (solid lines) we see that mode 3, once aggregated, is the one that makes KS_3 constraint active. Other constraints are inactive. This had already been revealed in our previous design space study. The point on mode 3 that causes the constraint to be active is at the end of the flight envelope, 15 m/s, where the value is zero, hence causing the constraint to be active. Note that the constraints KS_i are not shown explicitly in Fig. 39. All constraints would form a horizontal line at 0 rad/s spanning the entire velocity range (2-15 m/s). There are no constraints before or after 2 m/s and 15 m/s, respectively.

Numerical difficulties close to the optimum could indicate incorrect or noisy gradients, or issues in the design space and problem formulation, such as discontinuous constraints. Further investigation of the design space around the optimum is needed; this is addressed in the next section.

8.7 Further analysis of flutter constraints

To analyze the numerical difficulties experienced in the optimization we turn our attention to the smoothness of the constraints. Figure 40 shows each of the constraints plotted over the entire design space. Each row of plots represents the constraints where the first row of figures are the constraint on mode 1 and the last constraint on mode 4. The left column shows each constraint for the full design space while the right column show a zoomed in region $5.8 \leq AR \leq 6.2$ and $0.00148 \leq t \leq 0.00152$ m where the optimum is found. The zoomed region is sampled with 16 extra points in each variable, a total of 256 points. The zoom region is highlighted with a red square on the full design space in the left column.

Investigating the left column in Fig. 40 we find that constraint mode 1 has discontinuous regions close to the optimum in the lower right corner of the design space. Looking at the zoomed region where the optimum is found we see that the mode is not smooth which is certainly contributing to the numerical difficulties that the optimizer is experiencing. By dropping out mode 1 as a constraint from the optimization problem will likely improve the behavior of the optimization but would be a temporary fix and not a good solution. By investigating the remaining mode constraints 2, 3, and 4 we can see that there is a discontinuity appearing between aspect ratio of 2–3 for all thicknesses. Constraints 2, 3, and 4 are however smooth close to the optimum.

This is further illustrated in Fig. 41. Thickness is fixed as $t = 0.001619$ m and the aspect ratio range $1 \leq AR \leq 6$ is sampled with 256 points. The constraints are shown in the top figure in two dimension. The damping and frequency plots are shown at three aspect ratio of interest labeled from one to three. The numerical value of the aspect ratio is not important here as it is only chosen to explain the discontinuity issue. Aspect ratios chosen are indicated by a gray vertical lines as well as labeled from left to right using the letters A, B, and C.

By inspecting the top figure we immediately notice the discontinuity in constraint 4 and constraint 1 as well as the mode hopping or mode swapping where constraint 2 and constraint 3 switch places. By comparing the damping plots at aspect ratio A and B we see that mode 4 is picking up another flutter root that is close to the root that it is starting at. By comparing the frequency plots for aspect ratios A and B, we notice that as the frequency of modes 2 and 3 become close to
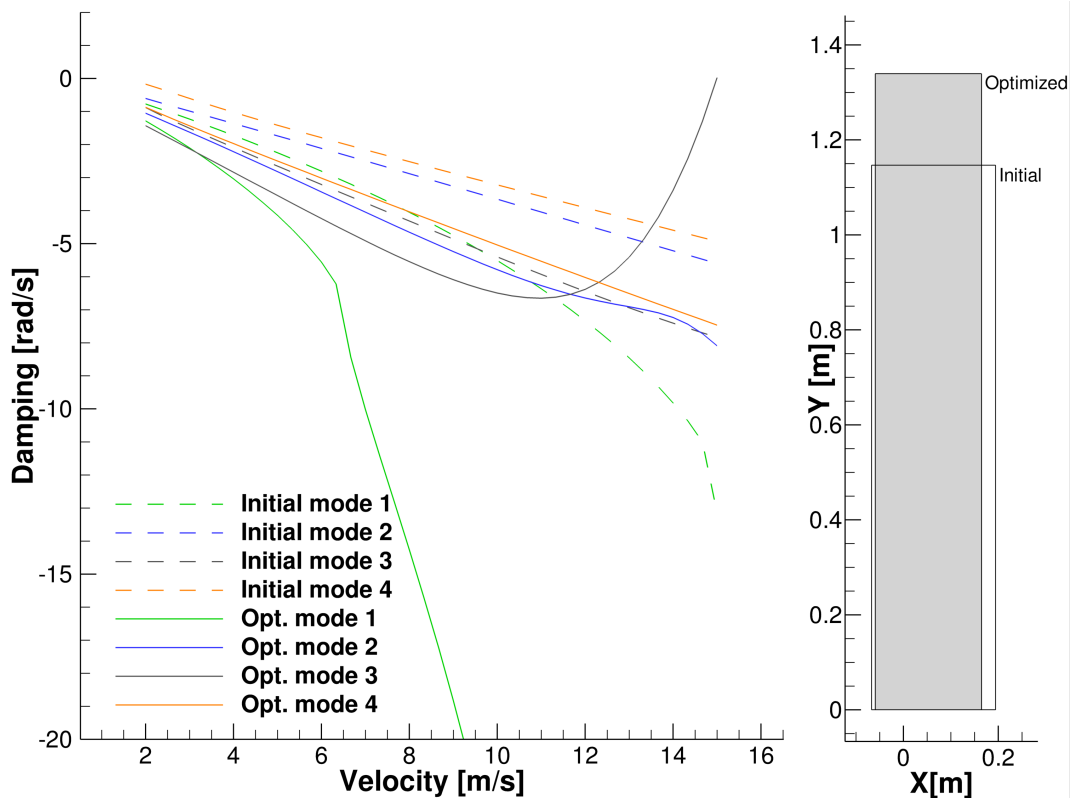
Figure 39: Damping and planform view of the initial ($AR = 4.38, t = 0.002$) and the optimized ($AR^* = 5.98, t^* = 0.001497$) design. As expected mode 3 is causing the constraint KS_3 to be active. It can be seen that the maximum for mode 3 is at 15 m/s and it is very close to zero. Other constraints are inactive. Constraints are not shown explicitly in the figure. Initial geometry shown with outlines and the optimized geometry shown in gray solid color.
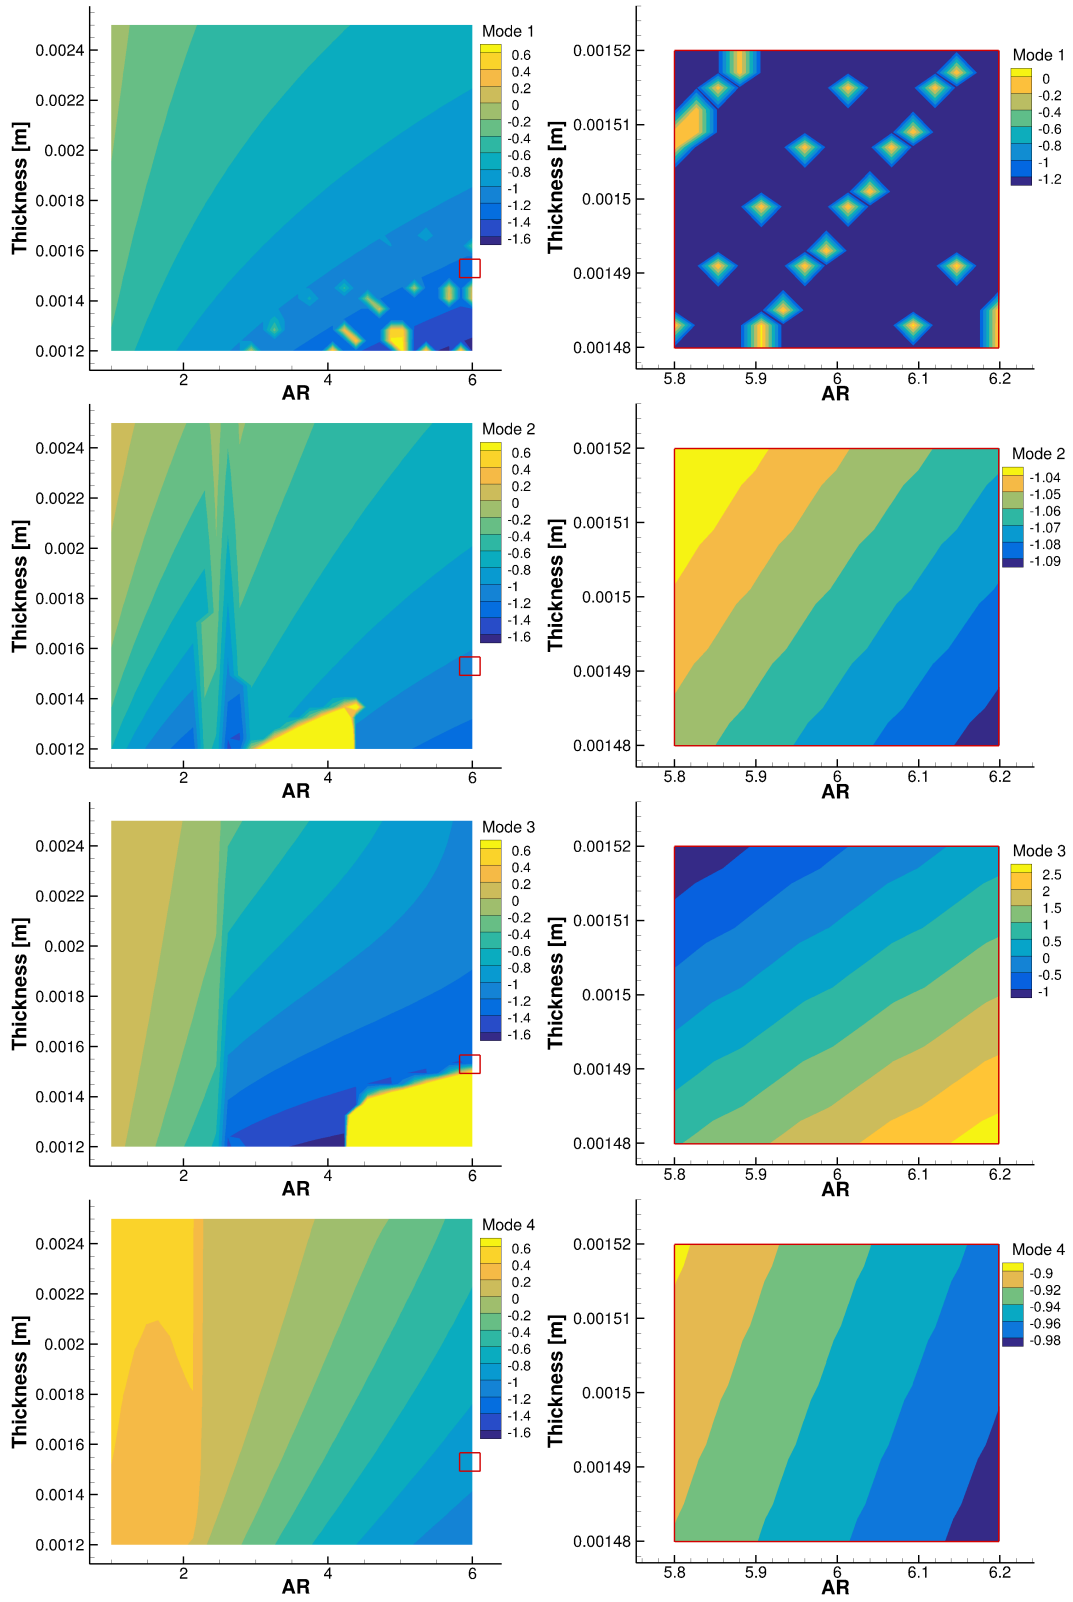
Figure 40: KS values for mode 1 to 4 shown in the left column for the entire design space. The right column shows the region $5.8 \leq AR \leq 6.2$ and $0.00148 \leq t \leq 0.00152$ m sampled using 16×16 stencil. The zoom region is highlighted with a red square on the full design space in the left column.

78
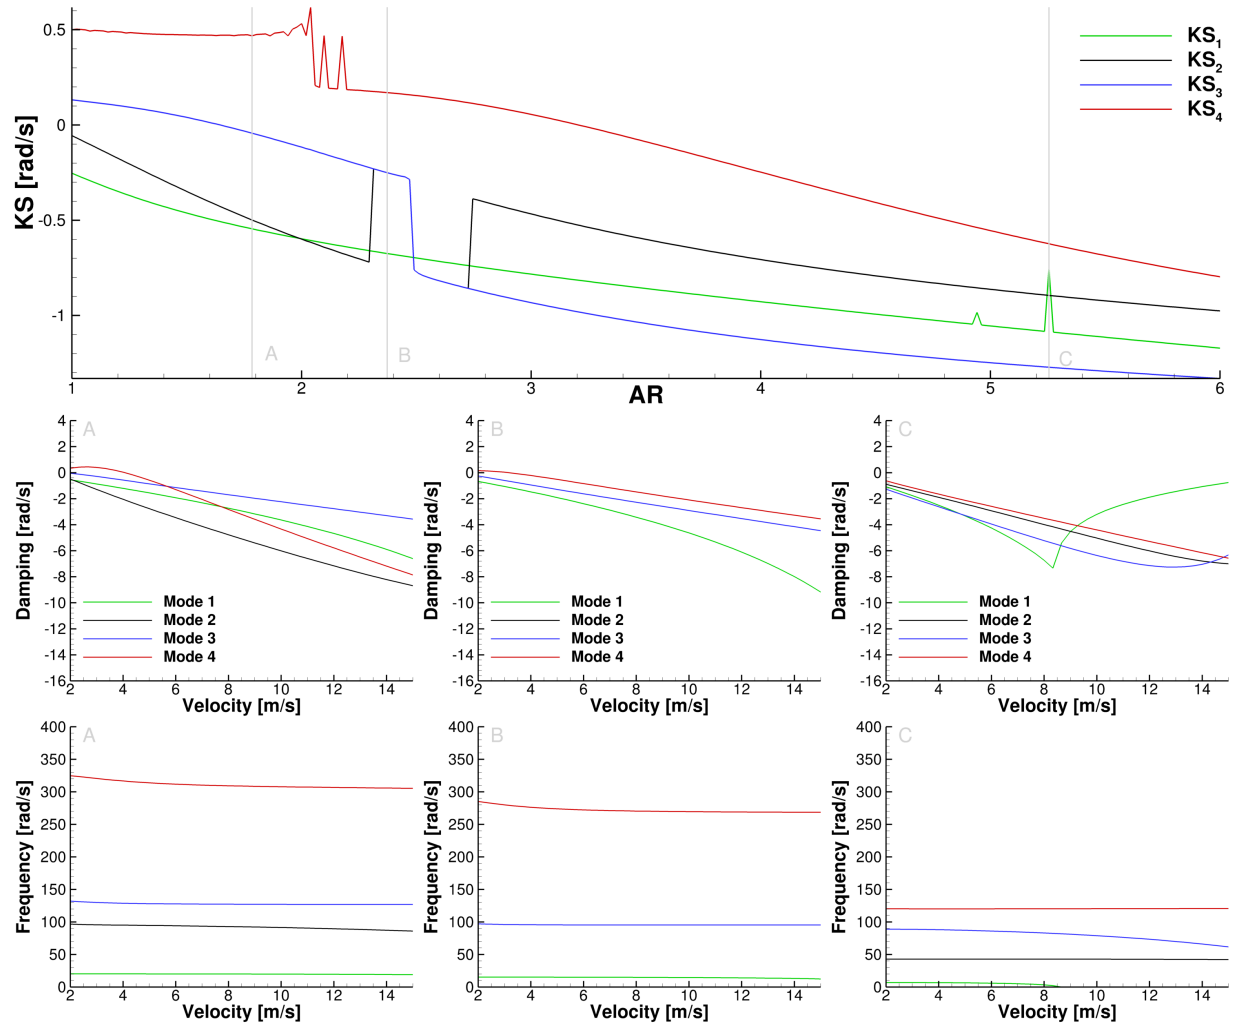
Figure 41: Thickness is fixed as $t = 0.001619$ m and the aspect ratio range $1 \leq AR \leq 6$ is sampled with 256 points. Damping and frequency plots are shown for three different aspect ratios labeled A to C.

each other. The eigenvalue solution method (determinant iteration) can not distinguish between the modes, and both of them pick up the same root. When increasing the aspect ratio further, we see that modes 2 and mode 3 switch, and are therefore not representing the correct mode. Inspecting the frequency plot at aspect ratio C, we notice that the frequency for mode 1 becomes zero. When this happens two real roots emerge as the complex root (and its conjugate) will have zero imaginary part. This is indicated as a bifurcation on the damping plot. In this case however we are only tracking one of the real roots as current algorithm is only capable of tracking one for each mode. This explains the discontinuity in mode 1 at aspect ratio C. The bifurcated real root starts to appear between $B \leq AR \leq C$ but depending on which real root the algorithm "locks" onto, changes between different AR causing these discontinuities to show up. This can be prevented by using an algorithm that can track both real roots. To avoid mode swapping and track the real roots properly, a more sophisticated algorithm than presented here is needed.
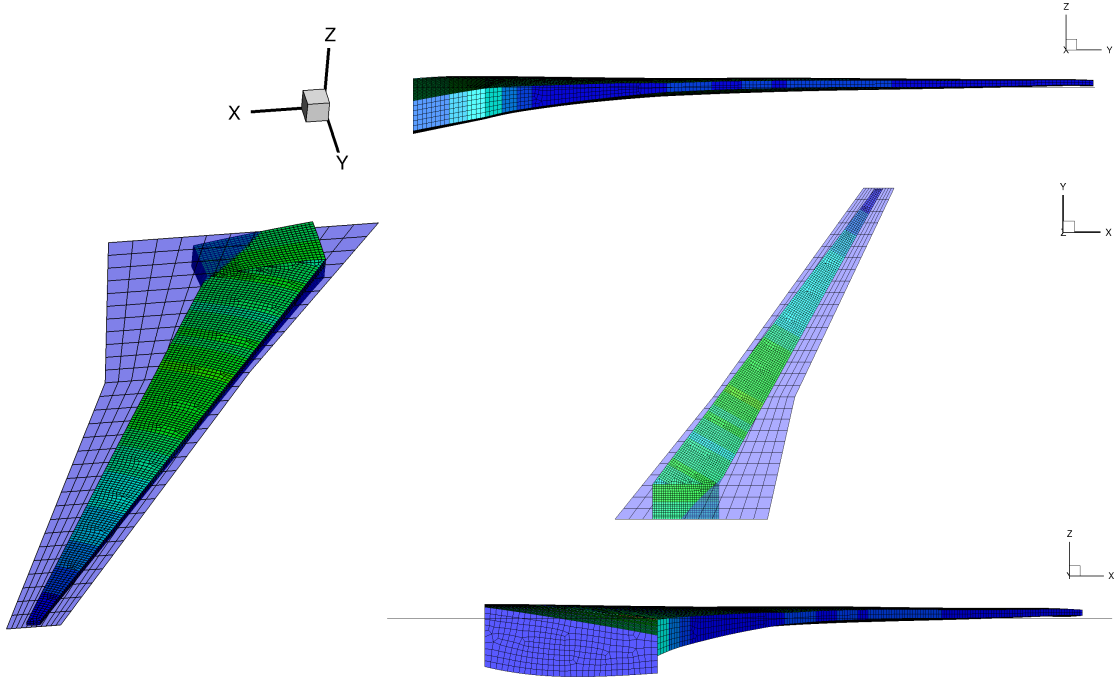
Figure 42: The uCRM wingbox shown in green along with the aerodynamic mesh shown in purple. The FFD volume is not shown here.

8.8   uCRM flutter analysis

8.8.1   Model description

The geometry under consideration is the undeflected Common Research Model (uCRM) developed by Kenway et al. [24] for aerostructural optimization. This geometry is based off the Common Research Model which has been studied extensively in multiple Drag Prediction Workshops (DPW) as well as by the Aerodynamic Design Optimization Discussion Group test cases [26, 57, 95]. The purpose of the uCRM is to provide the undeformed jig geometry and structural layout for aerostructural optimization, which reflects the overall size and shape of a typical transport aircraft such as the Boeing 777-200ER.

The aerodynamic surface of the wing is discretized using 11 elements in both the streamwise and spanwise directions for the inboard section but 11 and 18 elements in streamwise and spanwise direction for the outboard section of the wing. The wingbox is made up of 10584 MITC4 shell finite elements. The aerodynamic mesh and wingbox are shown in Fig. 42 where the aerodynamic mesh is shown in purple, and the wingbox in green. The FFD used here encapsulates both meshes but is not shown here. The material properties for the uCRM along with discretization is summarized in Table 20. Skin, spar and rib thickness and stiffener height, thickness and pitch are not presented here. The load transfer scheme between the aerodynamic and structure meshes is rigid links approach [48] and is discussed in [96].

For the analysis of the uCRM the Mach number is kept fixed at $M = 0.85$. The air speed and the air density are varied in 79 increments as a pair ranging from 253 - 305 m/s and 0.04 - 1.99 kg/m^3 receptively. The operating conditions are summarized in Table 21. We use the Lanczos method to calculate the 20 first natural modes and mode shapes using a subspace size of 50.

Table 20: uCRM mechanical properties, dimensions and discretization of the structure and the aerodynamic surface.

Variable	Symbol	Value
Density	ρ_s	2780 kg/m^3
Modulus of Elasticity	E	70 GPa
Poisson ratio	ν	0.3
Yield stress	σ_y	420 MPa
Span	b	58.6 m
Aspect ratio	AR	9.0
Reference Area	S	383.7 m^2
Sweep (leading edge)	Λ	37.4 deg
Number of finite elements	n_{FE}	10584
Number of DLM elements, inboard streamwise	n_x^i	11
Number of DLM elements, inboard spanwise	n_y^i	11
Number of DLM elements, outboard streamwise	n_x^o	11
Number of DLM elements, outboard spanwise	n_y^o	18

8.8.2 Problem description

In order to gain insight into the uCRM flutter characteristics we formulate the problem similarly as the flat problem where we use two effective design variables, the aspect ratio and thickness scaling variable. Skin, spar and rib thickness and stiffener height, thickness and pitch are modeled as individual variables but are then scaled using a single scaling factor t_{scaling} that is allowed to vary from $0.8 \leq t_{\text{scaling}} \leq 1.75$. When $t_{\text{scaling}} = 1.0$ all structural variables are at their initial value. When $t_{\text{scaling}} < 1.0$, this will decrease their values and vise versa. This approach is only used here in order to be able to visualize the design space in a similar manner as for the flat plate example.

The objective is to maximize the range using Eq. (8.10) but here the range is in nautical miles. As before the lift coefficient is fixed at $C_L = 0.5$. The drag coefficient is modeled using the lift induced drag and the zero lift drag coefficient,

$$C_D = \frac{C_L^2}{\pi e AR} + C_{D_0} \tag{8.13}$$

To simplify we set $e = 1$ and for the zero lift drag coefficient we use $C_{D_0} = 0.0162$. The flight speed in the range equation is fixed to $V = 250$ m/s and the thrust-specific fuel consumption also kept fixed $c_T = 0.53$ lb/(lbf $\cdot$ h). The initial and final cruise weights are defined as

$$\begin{aligned} W_{\text{final}} &= W_{\text{fixed}} + 2.5 W_{\text{wing}} \\ W_{\text{init}} &= W_{\text{final}} + W_{\text{fuel}} \end{aligned} \tag{8.14}$$

All parameters used there are summarized in Table 21. Flutter constraints are added to the first 10 modes or $KS_i \leq 0$ where $i = 1, \ldots, 10$ and as before area is kept constant. As for the plate problem the flutter constraints is enforced over the entire flight envelope (density and velocity pairs) and thus there is no flutter point specified explicitly. Optimization problem is summarized in Table 22.

Table 21: The uCRM operating conditions under investigation. The airspeed and density are incremented together in 79 steps.

Variable	Symbol	Value
Cruise Mach	M_∞	0.85
Cruise lift coefficient	C_L	0.5
Cruise air density	ρ_∞	0.04–1.99 kg/m^3
Cruise air speed range	U_∞	253–305 m/s
Number of speed/density increments	n_U	79
Zero lift drag coefficient	C_{D_0}	0.0169
Flight speed (range eqn.)	V	250 m/s
Thrust specific fuel consumption	c_T	0.53 lb/(lbf · h)
Fixed weight	W_fixed	197000 kg
Fuel weight	W_fuel	90000 kg

The aspect ratio is allowed to vary from $8.18 \leq AR \leq 16.43$ and as mentioned before the thickness scaling factor can vary from $0.8 \leq t_\text{scaling} \leq 1.75$.

8.8.3 Design space analysis

As for the flat plate problem the design space is analyzed. To generate the contour plot the design space is sampled using 32 samples in each variable, the aspect ratio and the thickness scaling giving a total of 1024 samples. Figure 43 shows the first 9 modes, the KS aggregated values, plotted as contour plots. None of the modes, excluding mode 9 which shows discontinuities, violate the constraint or where $KS_i > 0$. Mode 2 and 5 have however regions where they are very close to zero or practically zero. Mode 9, shows similar discontinuities as the flat plate geometry where $KS >> 0$. Mode 10, which is not shown here, shows similar behavior as the other.

Figure 44 highlights this further where a slice was done at fixed thickness scaling of $t_\text{scaling} = 1.1$ over all aspect ratios. Inspecting the figure we see that mode 1 experiences a discontinuity at lower aspect ratios. Similar behavior was seen with the flat plate case where in this case the flutter root finding method is locking onto a different eigenvalue. Mode 9 show similar behavior and experiences discontinuity around $AR = 12.7$. Here this mode has likely bifurcated and split into two real roots where at this aspect ratio the root finding method is locking onto which ever root it finds first. As noted before mode 5 is also very close to zero and is practically zero at the higher aspect ratios. This design space appears to be very uninteresting as none of the constraints become active although mode 2 and 5 come close to it. Optimization of this simplified problem does thus offer limited insight into the flutter characteristics of the uCRM. No optimization is thus done at this point.

A possible explanation why none of the modes are active could be due to the scaling of the entire structure using only one thickness scaling variables. By lowering the bounds on the thickness scaling parameter we might able to produce points where the structure flutters. Another explanation why the structure does not flutter is the omission of the masses due to the engine as well as leading edge and trailing edge devices such as slats, flaps and ailerons. One possible way of modeling these

Table 22: Overview of the simplified optimization problem for the uCRM wing. All structural variables are scaled using one scaling parameter.

	Function/variable	Description	Quantity
maximize	Breguet range equation		
with respect to	x_{planform}	Planform variables	3
	x_{struct}	Structural sizing variables	582
		Total design variables	**585**
subject to	$A - A_{\mathrm{init}} = 0.0$	Fixed area	1
	$\mathrm{KS}_i \leq 0$	Aggregate damping for velocity range for mode $i = 1, \ldots, 10$	10
		Total constraints	**11**

masses is to lump them at as a point mass at discrete locations and connect them to the structure using rigid body elements (RBE3) elements such that they do not add stiffness to the structure. Their effect is thus purely inertial. Figure 45 is a sketch of such a solution where the rigid links are shown as lines extending out from the structure. These lumped masses have been shown to be very important [5] as they reduce the torsional frequency substantially resulting in a lower flutter speed of the structure.

8.9 Conclusions

Sensitivities of the implemented flutter constraints in this work are both accurate and obtained efficiently. This allows us to perform optimization of the uCRM wing with tens or hundreds of design variables. As illustrated with the flat plate problem and the uCRM geometry the current root finding method, the determinant iteration, is not adequate or robust for even the simplest problems. Issues encountered in this work that relate to the method can be described as follows.

1. If an eigenvalue changes rapidly with increasing velocity the method might converge to an incorrect value due to a large step size in velocity. This issue can be mitigated by increasing the number of velocities analyzed substantially.

2. If two eigenvalues have similar or same imaginary parts this method can lock onto an incorrect eigenvalue. This issue as illustrated above was frequently observed in this work. Increasing the number of velocity points analyzed substantially may sometimes alleviate this issue, although it might become a computational burden if many modes and velocities are needed for the flutter analysis.

3. This method offers no direct way of tracking two real roots that emerge from a bifurcated root where the imaginary part vanishes.

To determine the flutter speed accurately and efficiently, and to implement it as a constrain in an optimization, these issues have to be addressed. Using more sophisticated root finding methods or reformulating the problem are options that we are currently looking into. The chosen method must also be differentiated such that accurate and efficient sensitivities are obtained.
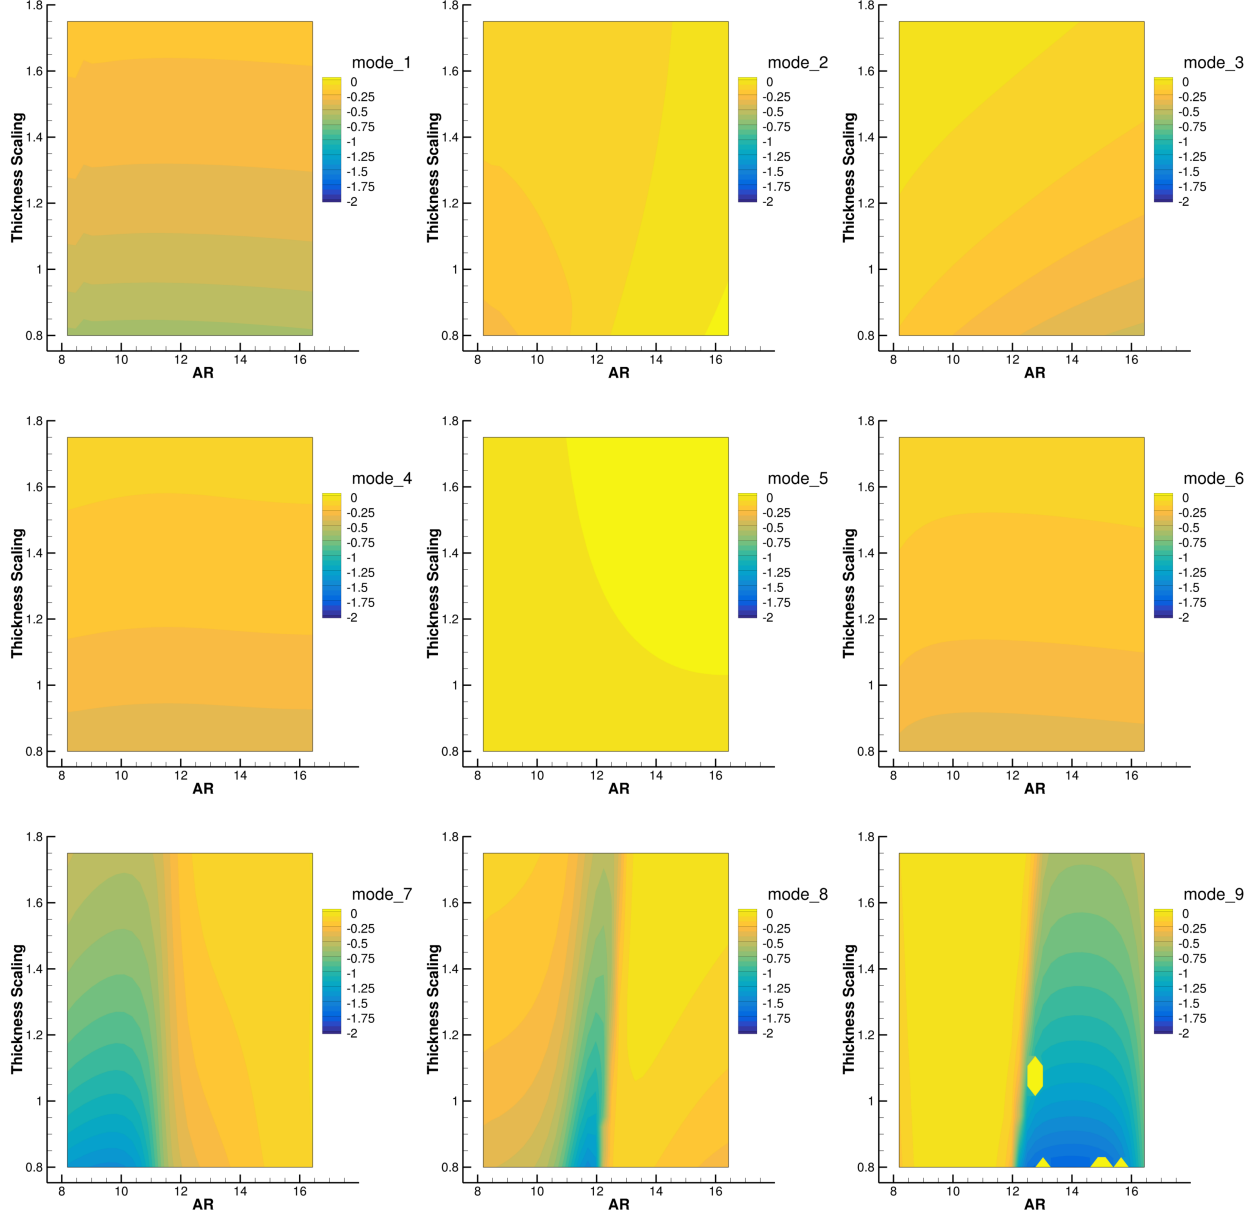
Figure 43: KS values for mode 1 to 9 shown for the entire design space. Design space was sampled with 32×32 stencil
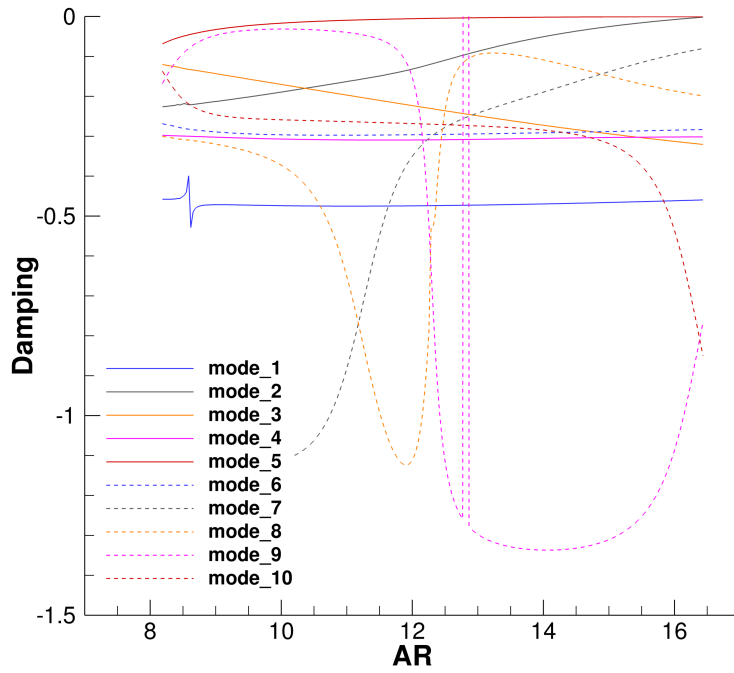
Figure 44: KS values for mode 1 to 10 shown as a slice at thickness scaling $t_{\text{scaling}} = 1.1$ for all aspect ratios. Notice the discontinuity in mode 9.
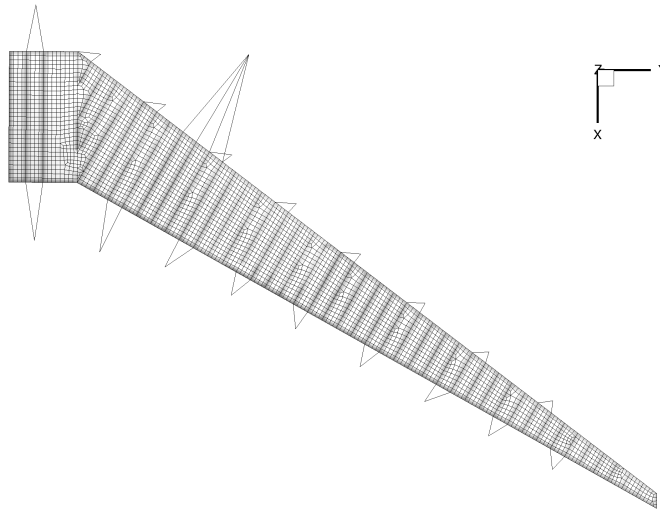
Figure 45: uCRM wingbox, where the engine mass as well as leading and trailing edge devices are modeled as lumped point masses connected to the structure using rigid body elements (RBE3); adapted from Stanford and Dunning [5].

Necessary capabilities in TACS need to be developed to model the rigid body elements (RBE3), or something similar, to support lumped discrete masses that do not add stiffness to the structure. These point masses are lumped from the engine, leading and trailing edge devices and have been shown to be very important as they decrease the predicted flutter speed [5].

8.10   Ongoing and Future Work

The flutter equation (8.1) can be written in a state space form and solved as a generalized eigenvalue problem. Introducing the trivial expression

$$\mathbf{I}\dot{\mathbf{u}} - \mathbf{I}\dot{\mathbf{u}} = 0 \tag{8.15}$$

where $\mathbf{I}$ is the identity matrix and is a $N_m \times N_m$ matrix, where N_m is the number of modes. We now can write it on a state space form substituting $\mathbf{u} = \bar{\mathbf{u}}e^{pt}$

$$p \begin{bmatrix} \mathbf{I} & 0 \\ 0 & \left(\frac{U}{b}\right)^2 \mathbf{M} \end{bmatrix} \begin{Bmatrix} \bar{\mathbf{u}} \\ p\bar{\mathbf{u}} \end{Bmatrix} - \begin{bmatrix} 0 & \mathbf{I} \\ -(\mathbf{K} - q_\infty \mathbf{A}^R) & -\frac{U}{b}\mathbf{C} + \frac{q_\infty}{k}\mathbf{A}^I \end{bmatrix} \begin{Bmatrix} \bar{\mathbf{u}} \\ p\bar{\mathbf{u}} \end{Bmatrix} = 0 \tag{8.16}$$

which is $2N_m \times 2N_m$ of size. The non-dimensional eigenvalue $p = g + ik$ is obtained by multiplying the Laplacian eigenvalue s by b/U i.e. $p = s(b/U)$, and $\mathbf{u}$ is the generalized coordinate vector.

To solve Eq. (8.1) iterative procedures are commonly applied since the aerodynamic matrix depends on the reduced frequency k, the imaginary part of p. One such procedure is the determinant iteration [85]. Another popular method is found in commercial software [81, 87], where an eigenvalue problem is solved based on an assumed reduced frequency k. The resulting eigenvalue for the mode under study is identified and compared to the assumed reduced frequency k. If the difference exceeds some predefined tolerance the iteration continues with the imaginary part of the new eigenvalue used as the reduced frequency $k = \text{Im}(p)$.

Both methods however have drawbacks or can experience convergence issues or converge to incorrect values. For example if two eigenvalues are close to each other, meaning when the imaginary component (frequency) of the modes are close to each other or identical, an incorrect mode may be picked up. This will result in mode hopping. Another potential issue is when the eigenvalue changes rapidly with dynamic pressure (velocity). This can result in convergence issues or in the incorrect eigenvalue being picked up.

Further these methods do not distinguish aerodynamic lag roots from structural modes as they are fixed to find as many roots as there are modes. If aerodynamic lag roots become unstable these methods may converge to the lag root over the structural mode. A more sophisticated root finding method is thus needed.

In this work we implement a non-iterative proposed by van Zyl [97]. A similar method is also proposed by Chen [98]. The method is outlined here:

1. At each dynamic pressure q_∞ increment, an eigenvalue problem is solved for a range of reduced frequencies k.

2. Eigenvalues are valid roots of the flutter equation if the imaginary part of the eigenvalue equals the assumed k value i.e. a matched point solution $\text{Im}(p) = k$.

3. A change in sign of the difference $\text{Im}(p) - k$ signifies the existence of a root.

4. The root is then determined by a linear interpolation.

Figure 46 shows the reduced frequency sweep for a single dynamic pressure q_i for a hypothetical system with two modes shown in red and blue. The black dots represent an intersection of a mode with the black diagonal line, $\text{Im}(p) = k$. There are 5 valid roots for this particular system, 4 from the red mode and 1 from the blue mode. For the red mode, there are two real roots at $k = k_0 = 0$ and two complex roots at $k = k_1$ and $k = k_2$. The blue mode is complex throughout the reduced frequency sweep and has only one root at $k = k_3$. An iterative method such as the determinant iteration would have issues converging to $k = k_1$ and would converge to $k = k_2$. A similar behavior is noted in [99] where the iterative method converges to a real root when it should have converged to oscillatory complex root. Further the determinant iterative method is only be able to find as many roots as there are oscillatory modes but the non-iterative places no limit on the number of roots that can be found.

The modes must be tracked at two stages. During the reduced frequency sweep, which finds all the roots for a given dynamic pressure, and during the migration of the modes between dynamic pressure increments. The modes are tracked using the mode tracking method of Van Zyl [100]. The tracking of the modes during the mode migration is more challenging and requires extra care since new modes can show up as well as modes can disappear. A correlation metric is used to determine if the dynamic pressure increment is too big. If the increment is too large, the step is halved until its either accepted (moving on to the next dynamic pressure), or fails due to being halved too often where a minimum dynamic pressure increment has been reached. If it fails and the minimum step size is reached, the roots are accepted and the dynamic pressure is incremented using a normal step size.
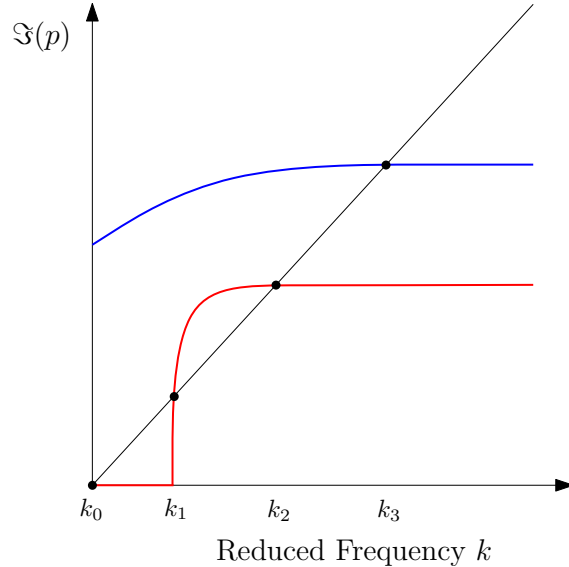
Figure 46: Hypothetical system with two modes shown in blue and red. Black dots represent a match point solution, i.e., where the modes intersect the diagonal line, $\text{Im}(p) = k$, depicted in black.

The method has already been implemented in Python. A complex-step-able version in Fortran is currently being developed. A preliminary result from the python version is shown in Fig. 47

There are several challenges in adding efficient flutter constraint in eigenvalue problems. A naive approach would be to specify the flutter point which is defined as the speed or dynamic pressure which the structure flutters at or equivalently when one of the modes becomes unstable [101]. For optimization this approach may introduce discontinuities in the constraints due to mode switching or a hump mode so additional steps are required to get continuous constraints.
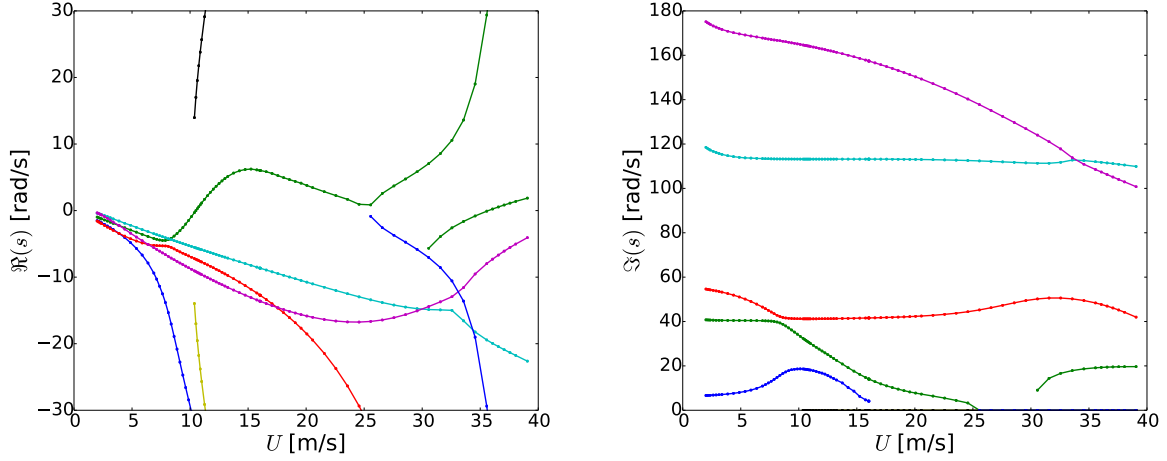
Figure 47: The modes found using the van Zyl root finding method (python implementation) for a swept plate (not shown here).

In the case of mode crossing, as shown in Fig. 48, two modes switch (real part only shown), but the flutter point is very similar. This can cause discontinuity in the predicted flutter point, and results in issues for a gradient based optimizer. The imaginary part of the eigenvalue (frequency) will also switch or in many cases the frequencies will coalescence causing a mode to become unstable.
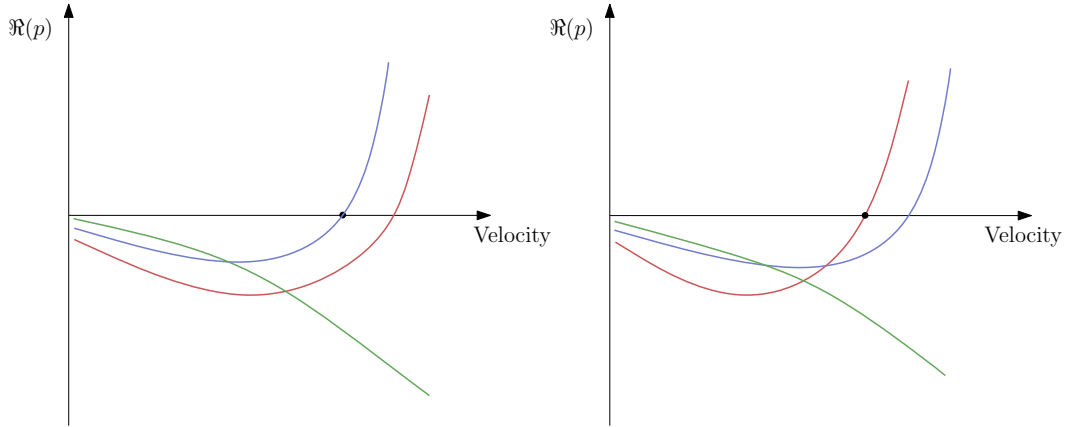
Figure 48: Mode crossing. Here two modes change role of being the fluttering mode.

A more serious problem is when a hump mode is present and it becomes the unstable mode. This is shown in Fig. 49. This is C^0 discontinuity and for a gradient based optimizer is a nonstarter.

Techniques exist to handle these problems and are summarized by Stanford et al. [102]. One method is to force the unstable mode to remain the unstable mode. Series of frequency separation constraints can be used to prevent the mode switching which will force the critical mode to remain the same thus mitigating the discontinuity issue. For a hypothetical system this is shown in the left figure of Fig. 50. The drawback of this method is that for a system with N_m modes, $N_m - 2$ constraints are needed (here we expect two modes to coalesce hence no constraint is needed for two modes). Other limitations are that always the same mode will become unstable which might over-constrain the optimization problem. Another more serious flaw is that the possibility of a hump mode becoming active is still present.

88
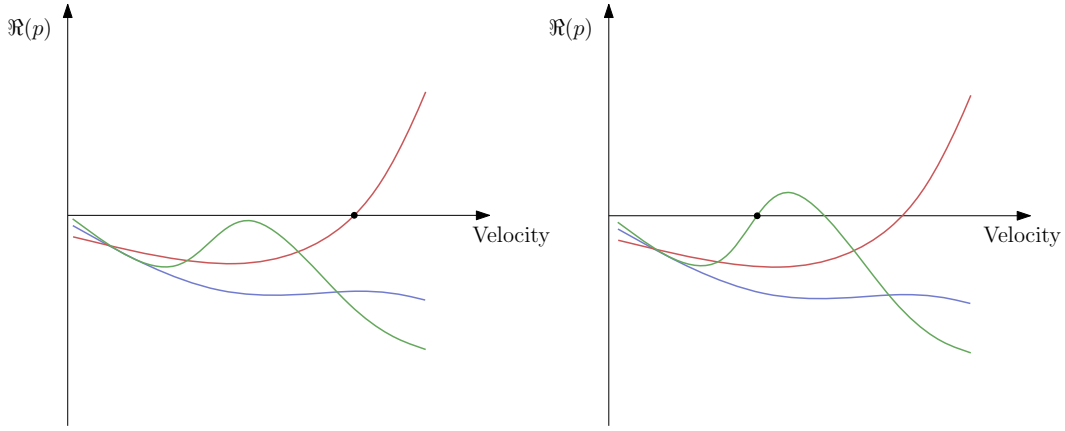
Figure 49: Illustration of a hump mode present where it becomes active with small changes in design.
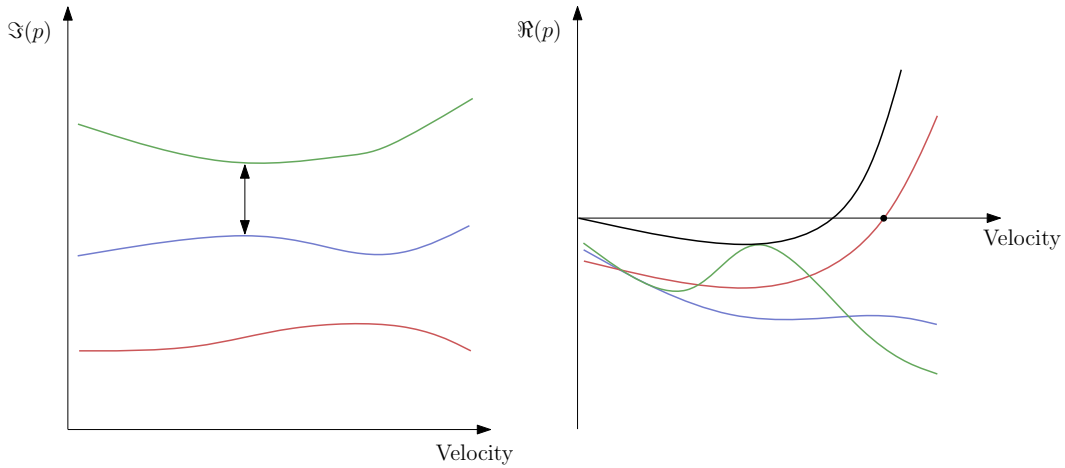
Figure 50: Two possible methods to prevent discontinuity in the flutter constraint. To the left is frequency separation method and to the right is the damping boundary shown in black.

More elaborate constraints are thus necessary to both handle the switch as well as the hump mode. One way is to add a damping curve or a boundary such that we force the real part of the eigenvalue to be below certain boundary. This boundary spans the operating conditions of interest from wind off to some high value of interest. Such boundary is shown in the right figure of Fig. 50 where the black line represents that boundary. This approach mitigates the aforementioned issues and has further other benefits. No constraint is placed on the flutter point itself and thus no need to calculate ii specifically. A mode that becomes unstable abruptly and has steeper slope than the boundary, will be handled as well.

Here we use a boundary shape proposed by Stanford et al. [77] which is a modified version of a boundary proposed by Ringertz [103]. The boundary is defined as

$$
G(U) = \begin{cases} g^* \left(3U^2 U^* - 2U^3\right)/(U^*)^3 & 0 \leq U < U^* \\ \beta(U - U^*)^2 + g^* & U \geq U^* \end{cases} \tag{8.17}
$$

where $g^* < 0$ rad/s, $U^* > 0$ m/s, and $\beta > 0$ rad $\cdot$ s/m^2 are chosen based on the criteria of the problem at hand. Figure 51 shows the boundary where we have chosen $g^* = -2$ rad/s, $U^* = 20$

m/s, and $\beta = 3$. The constraint is self is written as

$$g_i \le G(U) \quad i = 1, \dots, (N_r N_U) \tag{8.18}$$

where the damping value g_i has to be less than $G(U)$ at every root N_r at every speed of interest N_U.
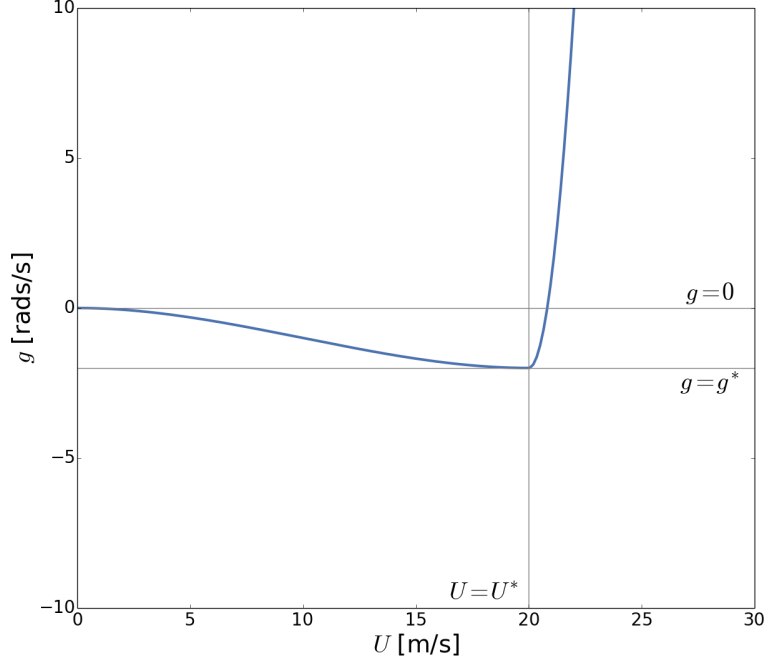
Figure 51: Example of a constraint curve that can be applied.

Although solving the issues with discontinuity in the flutter constraint the number of constraints have increased dramatically. To reduce the number of constraints we use a method proposed by Haftka [104] where a parametric constraint, such as the one introduced above, is replaced by the minimum value constraint. This reduces the number of constraints to the total number of modes found in the analysis process for the entire flight envelope. This reduced set of constraints can then be aggregated into a single constraint by using a KS function. Using this approach we can reduce the number of constraints from $N_r N_U$ to a single constraint. Although this hides information from the optimizer this allows for any number of modes to show up in the analysis process. As before the KS aggregated value has to be less than zero for the constraint to be inactive.

9 Summary of conclusions and future work

The overall goal of this NRA was to advance the sate-of-the art in high-fidelity aerostructural optimization of aircraft configurations, and to use high-fidelity aerostructural optimization tools to explore the design of specific configurations.

To address the bottleneck of geometry generation, we developed GeoMACH a new geometry parametrization tool that is able to quickly create conventional and unconventional aircraft geometries that are appropriate for high-fidelity MDO. The distinguishing feature of this tool is that it is able to compute derivatives efficiently using an adjoint method. Thus, using this in conjunction with an adjoint-based aerostructural tool enables a high-fidelity optimization cycle for an aircraft of two days or less using one thousand cores in a high-performance computing cluster. Some challenges remain in the treatment of intersections when two intersecting bodies move too much relative to each other, and this remains an area of future work. There also needs to be a better interoperability with CAD, which is often necessary in practice.

The development of the overset capability in ADflow addressed another major bottleneck in the workflow for high-fidelity aircraft MDO: mesh generation. We developed not only the capability for ADflow to solver overset meshes, but we also automated many tasks in the generation of single and overset meshes. A few steps are still not automated, and these might be tackled in the future. When we compared overset versus multiblock mesh solution convergence, we observed that the overset meshes exhibit better convergence. This was due to the fact that our overset meshes consist of multiple high quality meshes generated using an hyperbolic mesh generator for each component. The quality of the single meshes more than compensated for the interpolation errors that are inherent in the overset approach. The ability to mesh complex aircraft configurations allows us to perform design optimization studies of practically any full aircraft configuration. It has is also now enabling us to study propulsion-airframe integration problems.

One of the applications of aerostructural optimization was the study of optimal high aspect ratio wings. Because the true objective function in aircraft design is dependent on the trade-off between acquisition cost and operating cost, we generated Pareto fronts with respect to TOGW and fuel burn. We compared optimal wings for three different materials: aluminum, conventional carbon reinforced composites, and an hypothetical carbon nanotube material. In the optimization, we simultaneously optimized the wing planform and structural sizing. For all materials, the minimum fuel burn wings were found to be longer, heavier, thinner, more flexible, and more lightly loaded than their minimum TOGW counterparts.

Another application of high-fidelity aerostructural optimization was the assessment of the continuous morphing trailing edge technology. We demonstrated that the chordwise size of the morphing device could be reduced without severe performance penalties. We also demonstrated the direct relationship between a wings aspect ratio and the effectiveness of morphing technology. As aerospace materials and composites continue to improve and enable more flexible, higher aspect ratio wings, the incentive to include morphing technology will increase. The use of our high fidelity optimization framework enabled us to produce an accurate quantification of the efficiency improvement potential of morphing technology that was previously unavailable in literature. Our framework also enabled us to estimate the gains morphing technology will have once the trade-offs between weight and drag are considered with high-fidelity tools.

In our high-fidelity aerostructural optimization studies of high aspect ratio wings, we observed that buffet margin was being violated. This motivated the development of a method for constraining buffet based on a general physics-based formulation that can be used for any phenomena involving flow separation. Although this is a very nonlinear phenomenon, we derived a smooth formulation that is amenable to gradient-based optimization. This enabled us to perform high-fidelity aerostructural optimizations that satisfied buffet margins with little additional computational cost.

In our results, we showed the need for constraining buffet in a typical commercial airliner, and quantified the effect that this has on the performance in the flight condition space.

In the process of studying optimal high aspect ratio wings in collaboration with NASA, we identified the need for a common model for high-fidelity aerostructural design. Building on the CRM configuration, which was developed exclusively for aerodynamic studies, we developed two uCRM (undeflected CRM) configurations: one with the original aspect ratio of the CRM—the uCRM-9—and a high aspect ratio version—the uCRM-13.5. We have since used these models in our studies, and it has been used by a few other groups already.

Finally, we started the development of a flutter constraint, which is usually a critical design consideration in wing design, and is expected to be especially important for flexible high aspect ratio wings. We implemented a DLM-based flutter computation with adjoint derivative computation that included coupled derivatives with respect to both structural sizing (which had been done before by other researchers) and wing shape (which was a new development in the field). The derivatives of the implemented flutter constraints in this work are both accurate and obtained efficiently. This allows us to perform optimization of the uCRM wing with tens or hundreds of design variables. As illustrated with the flat plate problem and the uCRM geometry the current root finding method is not robust enough in its current form, and we plan to address this in future work.

In the course of the five years of this project, we advanced the state-of-the art in both methods and design knowledge related to high-fidelity MDO of aircraft configurations. MDO still has not reached its full potential and much remains to be done, but we have demonstrated that it is now possible to use high-fidelity aerostructural design optimization routinely in the study of new configurations and airframe technologies. This is made possible by improvements in the model preparation workflow, reductions in wall time by tackling computational bottlenecks, developing practical design constraint formulations, and creating effective visualizations of the results.

References

1. Aero Space Controls Corporation:, Aircraft actuators. `http://www.aerospace-controls.com/actuators.html`, February 2016.

2. CEF Industries:, Actuators. `http://www.cefindustries.com/pages/actuators.cfm`, February 2016.

3. Fronista, G. L.; and Bradbury, G.:, An electromechanical actuator for a transport aircraft spoiler surface. In *Energy Conversion Engineering Conference, 1997. IECEC-97., Proceedings of the 32nd Intersociety*, volume 1, 1997, pages 694–698.

4. Balakrishna, S.; and Acheson, M.:, Analysis of NASA Common Research Model dynamic data. In *49th Aersopace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2011. doi:10.2514/6.2011-1127.

5. Stanford, B. K.; and Dunning, P. D.:, Optimal topology of aircraft rib and spar structures under aeroelastic loads. *Journal of Aircraft*, 52(4):1298–1311, September 2014. doi:10.2514/1.C032913.

6. Kenway, G. K. W.; and Martins, J. R. R. A.:, Buffet onset constraint formulation for aerodynamic shape optimization. *AIAA Journal*, 55(6):1930–1947, June 2017. doi:10.2514/1.J055172.

7. Jonsson, E.; Kenway, G. K. W.; Kennedy, G. J.; and Martins, J. R. R. A.:, Development of flutter constraints for high-fidelity aerostructural optimization. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2017.

8. Kenway, G. K. W.; Secco, N.; Martins, J. R. R. A.; Mishra, A.; and Duraisamy, K.:, An efficient parallel overset method for aerodynamic shape optimization. In *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, January 2017. doi:10.2514/6.2017-0357.

9. Martins, J. R. R. A.:, *Encyclopedia of Aerospace Engineering*, volume Green Aviation, chapter Fuel burn reduction through wing morphing, pages 75–79. October 2016, pages 75–79. doi:10.1002/9780470686652.eae1007.

10. Kenway, G. K. W.; and Martins, J. R. R. A.:, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. *AIAA Journal*, 54(1):113–128, January 2016. doi:10.2514/1.J054154.

11. Hwang, J. T.; and Martins, J. R. R. A.:, An unstructured quadrilateral mesh generation algorithm for aircraft structures. *Aerospace Science and Technology*, 59:172–182, December 2016. doi:10.1016/j.ast.2016.10.010.

12. Lyu, Z.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of an adaptive morphing trailing edge wing. *Journal of Aircraft*, 52(6):1951–1970, November 2015. doi:10.2514/1.C033116.

13. Lyu, Z.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Aerodynamic shape optimization investigations of the Common Research Model wing benchmark. *AIAA Journal*, 53(4):968–985, April 2015. doi:10.2514/1.J053318.

14. Hwang, J. T.; Lee, D. Y.; Cutler, J. W.; and Martins, J. R. R. A.:, Large-scale multidisciplinary optimization of a small satellite's design and operation. *Journal of Spacecraft and Rockets*, 51 (5):1648–1663, September 2014. doi:10.2514/1.A32751.

15. Burdette, D. A.; Kenway, G. K.; and Martins, J. R. R. A.:, Performance evaluation of a morphing trailing edge using multipoint aerostructural design optimization. In *57th AIAA Structures, Structural Dynamics, and Materials Conference*, January 2016. doi:10.2514/6.2016-0159.

16. Burdette, D. A.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Aerostructural design optimization of a continuous morphing trailing edge aircraft for improved mission performance. In *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2016. doi:10.2514/6.2016-3209.

17. Kenway, G. K.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of the CRM configuration including buffet-onset conditions. In *54th AIAA Aerospace Sciences Meeting*, January 2016. doi:10.2514/6.2016-1294.

18. Burdette, D.; Kenway, G. K. W.; Lyu, Z.; and Martins, J. R. R. A.:, Aerostructural design optimization of an adaptive morphing trailing edge wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2015. doi:10.2514/6.2016-1294.

19. Kenway, G. K. W.; and Martins, J. R. R. A.:, Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2015. doi:10.2514/6.2015-0264.

20. Kenway, G. W. K.; and Martins, J. R. R. A.:, High-fidelity aerostructural optimization considering buffet onset. In *Proceedings of the 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2015.

21. Gray, J.; Hearn, T.; Moore, K.; Hwang, J. T.; Martins, J. R. R. A.; and Ning, A.:, Automatic evaluation of multidisciplinary derivatives using a graph-based problem formulation in Open-MDAO. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-2042.

22. Hwang, J. T.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Geometry and structural modeling for high-fidelity aircraft conceptual design optimization. In *Proceedings of the 15th AIAA Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-2041.

23. Kennedy, G. J.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Towards gradient-based design optimization of flexible transport aircraft with flutter constraints. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-2726.

24. Kenway, G. K. W.; Kennedy, G. J.; and Martins, J. R. R. A.:, Aerostructural optimization of the Common Research Model configuration. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-3274.

25. Lyu, Z.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of an adaptive morphing trailing edge wing. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2014. doi:10.2514/6.2014-3275.

26. Lyu, Z.; Kenway, G. K. W.; and Martins, J. R. R. A.:, RANS-based aerodynamic shape optimization investigations of the common research model wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2014. doi:10.2514/6.2014-0567.

27. Kennedy, G. J.; Kenway, G. K. W.; and Martins, J. R. R. A.:, High aspect ratio wing design: Optimal aerostructural tradeoffs for the next generation of materials. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, January 2014. doi:10.2514/6.2014-0596.

28. Hwang, J. T.; Lee, D. Y.; Cutler, J. W.; and Martins, J. R. R. A.:, Large-scale MDO of a small satellite using a novel framework for the solution of coupled systems and their derivatives. In *Proceedings of the 54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, April 2013. doi:10.2514/6.2013-1599.

29. Kennedy, G. J.; and Martins, J. R. R. A.:, An adjoint-based derivative evaluation method for time-dependent aeroelastic optimization of flexible aircraft. In *Proceedings of the 54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, April 2013.

30. Hwang, J. T.; and Martins, J. R. R. A.:, GeoMACH: Geometry-centric MDAO of aircraft configurations with high fidelity. In *Proceedings of the 14th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, September 2012. doi:10.2514/6.2012-5605.

31. Ivaldi, D.; Secco, N. R.; Chen, S.; Hwang, J. T.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of a truss-braced-wing aircraft. In *Proceedings of the 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2015. doi:10.2514/6.2015-3436.

32. Hwang, J. T.; and Martins, J. R. R. A.:, A computational architecture for coupling heterogeneous numerical models and computing coupled derivatives. *ACM Transactions on Mathematical Software*, 2017.

33. Lyu, Z.; and Martins, J. R. R. A.:, Aerodynamic design optimization studies of a blended-wing-body aircraft. *Journal of Aircraft*, 51(5):1604–1617, September 2014. doi:10.2514/1.C032491.

34. Chen, S.; Lyu, Z.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Aerodynamic shape optimization of the Common Research Model wing-body-tail configuration. *Journal of Aircraft*, 53(1):276–293, January 2016. doi:10.2514/1.C033328.

35. Liem, R. P.; Kenway, G. K. W.; and Martins, J. R. R. A.:, Multimission aircraft fuel burn minimization via multipoint aerostructural optimization. *AIAA Journal*, 53(1):104–122, January 2015. doi:10.2514/1.J052940.

36. Brooks, T. R.; Kennedy, G. J.; and Martins, J. R. R. A.:, High-fidelity aerostructural optimization of a high aspect ratio tow-steered wing. In *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, January 2016. doi:10.2514/6.2016-1179.

37. Lee, B.; and Kim, C.:, Aerodynamic shape optimization using discrete adjoint formulation based on overset mesh technique. In *European Conference on Computational Fluid Dynamics*, 2006.

38. Liao, W.; and Tsai, H. M.:, Aerodynamic shape optimization on overset grids using the adjoint method. *International journal for numerical methods in fluids*, 62(12):1332–1356, 2010.

39. Chan, W. M.; and Steger, J. L.:, Enhancements of a three-dimensional hyperbolic grid generation scheme. *Applied Mathematics and Computation*, 51(2):181–205, 1992.

40. Chan, W. M.; and Buning, P. G.:, Surface grid generation methods for overset grids. *Computers & fluids*, 24(5):509–522, 1995.

41. Lee, Y.; and Baeder, J.:, Implicit hole cutting-a new approach to overset grid connectivity. In *16th AIAA Computational Fluid Dynamics Conference*, 2003, page 4128.

42. Landmann, B.; and Montagnac, M.:, A highly automated parallel chimera method for overset grids based on the implicit hole cutting technique. *International Journal for Numerical Methods in Fluids*, 66(6):778–804, 2011.

43. Chan, W.:, Enhancements to the hybrid mesh approach to surface loads integration on overset structured grids. In *19th AIAA Computational Fluid Dynamics*, page 3990, 2009, page 3990.

44. Luke, E.; Collins, E.; and Blades, E.:, A fast mesh deformation method using explicit interpolation. *Journal of Computational Physics*, 231(2):586–601, January 2012. doi:10.1016/j.jcp.2011.09.021.

45. Kennedy, G. J.; Kenway, G. K.; and Martins, J. R. R. A.:, A comparison of metallic, composite and nanocomposite optimal transonic transport wings. Technical report, NASA, March 2014.

46. Kennedy, G. J.; and Martins, J. R. R. A.:, Parallel solution methods for aerostructural analysis and design optimization. In *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, September 2010.

47. Lyu, Z.; Kenway, G. K.; Paige, C.; and Martins, J. R. R. A.:, Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model. In *21st AIAA Computational Fluid Dynamics Conference*, Jul. 2013. doi:10.2514/6.2013-2581.

48. Brown, S. A.:, Displacement extrapolation for CFD+CSM aeroelastic analysis. In *Proceedings of the 35th AIAA Aerospace Sciences Meeting*, 1997.

49. Kenway, G. K. W.; Kennedy, G. J.; and Martins, J. R. R. A.:, Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and derivative computations. *AIAA Journal*, 52 (5):935–951, May 2014. doi:10.2514/1.J052255.

50. Kenway, G. K. W.; and Martins, J. R. R. A.:, Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration. *Journal of Aircraft*, 51(1):144–160, January 2014. doi:10.2514/1.C032150.

51. Vassberg, J.:, A unified baseline grid about the Common Research Model wing/body for the Fifth AIAA CFD Drag Prediction Workshop (invited). In *29th AIAA Applied Aerodynamics Conference*, Jul 2011. doi:10.2514/6.2011-3508.

52. Kreisselmeier, G.; and Steinhauser, R.:, Systematic control design by optimizing a vector performance index. In *International Federation of Active Controls Syposium on Computer-Aided Design of Control Systems, Zurich, Switzerland*, 1979.

53. Grossman, B.; Haftka, R. T.; Kao, P.-J.; Polen, D. M.; and Rais-Rohani, M.:, Integrated aerodynamic-structural design of a transport wing. *Journal of Aircraft*, 27(12):1050–1056, 1990. doi:10.2514/3.45980.

54. Chittick, I. R.; and Martins, J. R. R. A.:, An asymmetric suboptimization approach to aerostructural optimization. *Optimization and Engineering*, 10(1):133–152, March 2009. doi:10.1007/s11081-008-9046-2.

55. Brooks, T. R.; Kennedy, G. J.; and Martins, J. R. R. A.:, High-fidelity multipoint aerostructural optimization of a high aspect ratio tow-steered composite wing. In *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, January 2017. doi:10.2514/6.2017-1350.

56. Liem, R. P.; Martins, J. R. R. A.; and Kenway, G. K.:, Expected drag minimization for aerodynamic design optimization based on aircraft operational data. *Aerospace Science and Technology*, 63:344–362, April 2017. doi:10.1016/j.ast.2017.01.006.

57. Vassberg, J. C.; DeHaan, M. A.; Rivers, S. M.; and Wahls, R. A.:, Development of a common research model for applied CFD validation studies, 2008.

58. Vassberg, J.:, Introduction: Drag prediction workshop. *Journal of Aircraft*, 45(3):737–737, Jun 2008. doi:10.2514/1.37761.

59. Anonymous:, 777-200/300 airplane characteristics for airport planning. Technical Report D6-58329, Boeing Commercial Airplanes, July 1998.

60. Kenway, G. K.; Kennedy, G. J.; and Martins, J. R. R. A.:, A CAD-free approach to high-fidelity aerostructural optimization. In *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, number AIAA 2010-9231, September 2010. doi:10.2514/6.2010-9231.

61. Vos, R.; Barrett, R.; de Breuker, R.; and Tiso, P.:, Post-buckled precompressed elements: a new class of control actuators for morphing wing uavs. *Smart materials and structures*, 16(3): 919, 2007.

62. Pankonien, A. M.:, Smart material wing morphing for unmanned aerial vehicles, University of Michigan, 2015.

63. Urnes, J.; Nguyen, N.; Ippolito, C.; Totah, J.; Trinh, K.; and Ting, E.:, A mission adaptive variable camber flap control system to optimize high lift and cruise lift to drag ratios of future n+ 3 transport aircraft. In *51st AIAA Aerospace Sciences Meeting, Grapevine, TX*, 2013.

64. Kaul, U. K.; and Nguyen, N. T.:, Drag optimization study of variable camber continuous trailing edge flap (vcctef) using overflow. *AIAA Paper*, 2444, 2014.

65. Ting, E.; Dao, T.; and Nguyen, N.:, Aerodynamic load analysis of a variable camber continuous trailing edge flap system on a flexible wing aircraft. In *AIAA SciTech Conference*, 2015, pages 5–9.

66. Stanford, B. K.:, Static and dynamic aeroelastic tailoring with variable-camber control. *Journal of Guidance, Control, and Dynamics*, pages 2522–2534, 2016.

67. Yildiz, Y.; and Kolmanovsky, I. V.:, A control allocation technique to recover from pilot-induced oscillations (capio) due to actuator rate limiting. In *American Control Conference (ACC), 2010*, 2010, pages 516–523.

68. FlexSys:. Personal communication, February 2016.

69. Clark, R.; and Pelkman, R.:, High reynolds number testing of advanced transport aircraft wings in the national transonic facility. In *39th Aerospace Sciences Meeting and Exhibit*, January 2001. doi:10.2514/6.2001-910.

70. Lovely, D.; and Haimes, R.:, Shock detection from computational fluid dynamics results. In *Proceedings of the 14th Computational Fluid Dynamics Conference*, January 1999. doi:doi:10.2514/6.1999-3285.

71. Kenway, G. K. W.; and Martins, J. R. R. A.:, AIAA ADODG Case 5: CRM wing-body-tail optimization at flight reynolds number. Technical report, AIAA, May 2015.

72. Liu, F.; Cai, J.; Zhu, Y.; Tsai, H.; and F. Wong, A.:, Calculation of wing flutter by a coupled fluid-structure method. *Journal of Aircraft*, 38(2):334–342, 2001.

73. Yurkovich, R.:, Status of unsteady aerodynamic prediction for flutter of high-performance aircraft. *Journal of Aircraft*, 40(5):832–842, 2003.

74. Schuster, D. M.; Liu, D. D.; and Huttsell, L. J.:, Computational aeroelasticity: success, progress, challenge. *Journal of Aircraft*, 40(5):843–856, 2003.

75. Bennett, R. M.; and Edwards, J. W.:, An overview of recent developments in computational aeroelasticity. *AIAA paper*, (98-2421), 1998.

76. Gibbons, M. D.:, Aeroelastic calculations using cfd for a typical business jet model. 1996.

77. Stanford, B.; Wieseman, C. D.; and Jutte, C.:, Aeroelastic tailoring of transport wings including transonic flutter constraints. In *AIAA SciTech*, pages –, January 2015, pages –. doi:10.2514/6.2015-1127.

78. Haftka, R. T.; and Gürdal, Z.:, *Elements of structural optimization*, volume 11, 2012.

79. Albano, E.; and Rodden, W. P.:, A doublet-lattice method for calculating lift distributions on oscillating surfaces in subsonic flows. *AIAA Journal*, 7(2):279–285, February 1969. doi:10.2514/3.5086.

80. Blair, M.:, A compilation of the mathematics leading to the doublet lattice method. Technical report, DTIC Document, 1992.

81. MSC Software Corp:, *MSC Nastran 2012 Release Guide*.

82. Rodden, W. P.; Taylor, P. F.; and McIntosh, S. C.:, Further refinement of the subsonic doublet-lattice method. *Journal of Aircraft*, 35(5):720–727, September 1998. doi:10.2514/2.2382.

83. Trefethen, L. N.; and Bau III, D.:, *Numerical linear algebra*, volume 50, 1997.

84. Bathe, K.-J.:, *Finite element procedures*, 2006.

85. Hassig, H. J.:, An approximate true damping solution of the flutter equation by determinant iteration. *Journal of Aircraft*, 8(11):885–889, 1971.

86. Poon, N. M. K.; and Martins, J. R. R. A.:, An adaptive approach to constraint aggregation using adjoint sensitivity analysis. *Structural and Multidisciplinary Optimization*, 34(1):61–73, July 2007. doi:10.1007/s00158-006-0061-7.

87. Rodden, W. P.; Bellinger, E. D.; Harder, R. L.; and Center., L. R.:, *Aeroelastic addition to NASTRAN*, 1979.

88. Griewank, A.; and Walther, A.:, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, second edition, 2008.

89. Mader, C. A.; and Martins, J. R. R. A.:, Derivatives for time-spectral computational fluid dynamics using an automatic differentiation adjoint. *AIAA Journal*, 50(12):2809–2819, December 2012. doi:10.2514/1.J051658.

90. Hascoët, L.; and Pascual, V.:, Tapenade 2.1 user's guide. Technical report 300, INRIA, 2004.

91. Pascual, V.; and Hascoët, L.:, Extension of TAPENADE towards Fortran 95. In Bücker, H. M.; Corliss, G.; Hovland, P.; Naumann, U.; and Norris, B., editors, *Automatic Differentiation: Applications, Theory, and Tools*, Lecture Notes in Computational Science and Engineering, 2005.

92. Anderson, E.; Bai, Z.; Dongarra, J.; Greenbaum, A.; McKenney, A.; Du Croz, J.; Hammarling, S.; Demmel, J.; Bischof, C.; and Sorensen, D.:, LAPACK: A portable linear algebra library for high-performance computers. In *Proceedings of the 1990 ACM/IEEE Conference on Supercomputing*, Supercomputing '90, 1990, pages 2–11.

93. Martins, J. R. R. A.; Sturdza, P.; and Alonso, J. J.:, The complex-step derivative approximation. *ACM Transactions on Mathematical Software*, 29(3):245–262, September 2003. doi:10.1145/838250.838251.

94. Hascoet, L.; and Pascual, V.:, The tapenade automatic differentiation tool: Principles, model, and specification. *ACM Trans. Math. Softw.*, 39(3):20:1–20:43, May 2013. doi:http://dx.doi.org/10.1145/2450153.2450158.

95. Vassberg, J.; and Jameson, A.:, Influence of shape parameterization on aerodynamic shape optimization. Technical report, Von Karman Institute, Brussels, Belgium, April 2014.

96. Kennedy, G. J.; and Martins, J. R. R. A.:, Parallel solution methods for aerostructural analysis and design optimization. In *Proceedings of the 13^{t}h AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, September 2010.

97. van Zyl, L. H.:, Aeroelastic divergence and aerodynamic lag roots. *Journal of Aircraft*, 38(3): 586–588, May 2001. doi:10.2514/2.2806.

98. Chen, P.:, Damping perturbation method for flutter solution: the g-method. *AIAA journal*, 38(9):1519–1524, 2000.

99. Rodden, W. P.; and Bellinger, E. D.:, Aerodynamic lag functions, divergence, and the british flutter method. *Journal of Aircraft*, 19(7):596–598, July 1982. doi:10.2514/3.44772.

100. Van Zyl, L. H.:, Use of eigenvectors in the solution of the flutter equation. *Journal of Aircraft*, 30(4):553–554, July 1993. doi:10.2514/3.46380.

101. Bisplinghoff, R. L.; Ashley, H.; and Halfman, R. L.:, *Aeroelasticity*, 2013.

102. Stanford, B.; Beran, P.; and Bhatia, M.:, Aeroelastic topology optimization of blade-stiffened panels. *Journal of Aircraft*, 51(3):938–944, March 2014. doi:10.2514/1.C032500.

103. Ringertz, U.:, On structural optimization with aeroelasticity constraints. *Structural optimization*, 8(1):16–23, 1994.

104. Haftka, R. T.:, Parametric constraints with application to optimization for flutter using a continuous flutter constraint. *AIAA Journal*, 13(4):471–475, April 1975. doi:10.2514/3.49733.

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
01- 08- 2017	Contractor Report	06/2011 - 04/2017

4. TITLE AND SUBTITLE

High-fidelity Multidisciplinary Design Optimization of Aircraft Configurations

5a. CONTRACT NUMBER

NNX11AI19A

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Martins, Joaquim R.; Kenway, Gaetan K.; Burdette, David; Jonsson, Eirikur; Kennedy, Graeme J.

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

081876.02.07.02.01.03

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Michigan, Ann Arbor, Michigan 48109

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSOR/MONITOR'S ACRONYM(S)

NASA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

NASA-CR-2017-219647

12. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified
Subject Category 05
Availability: NASA STI Program (757) 864-9658

13. SUPPLEMENTARY NOTES Langley Technical Monitor: Karen M. Taminger

14. ABSTRACT
To evaluate new airframe technologies we need design tools based on high-fidelity models that consider multidisciplinary interactions early in the design process. The overarching goal of this NRA is to develop tools that enable high-fidelity multidisciplinary design optimization of aircraft configurations, and to apply these tools to the design of high aspect ratio exible wings. We develop a geometry engine that is capable of quickly generating conventional and unconventional aircraft configurations including the internal structure. This geometry engine features adjoint derivative computation for efficient gradient-based optimization. We also added overset capability to a computational fluid dynamics solver, complete with an adjoint implementation and semiautomatic mesh generation. We also developed an approach to constraining buffet and started the development of an approach for constraining flutter. On the applications side, we developed a new common high-fidelity model for aeroelastic studies of high aspect ratio wings. We performed optimal design trade-offs between fuel burn and aircraft weight for metal, conventional composite, and carbon nanotube composite wings. We also assessed a continuous morphing trailing edge technology applied to high aspect ratio wings.

15. SUBJECT TERMS

Adjoint methods; Aerostructural optimization; Buffet; CFD; Flutter; High aspect wings; MDO; Overset meshes; Subsonic transport aircraft

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			STI Help Desk (email: help@sti.nasa.gov)
U	U	U	UU	106	19b. TELEPHONE NUMBER *(Include area code)* (757) 864-9658