

This is a preprint of the following article, which is available from mdolab.engin.umich.edu

Neil Wu, Charles A. Mader, and Joaquim R. R. A. Martins. Large-scale Multifidelity Aerostructural Optimization of a Transport Aircraft. *33rd Congress of the International Council of the Aeronautical Sciences*, Stockholm, Sweden. September 4, 2022.

The published article may differ from this preprint.

Large-scale Multifidelity Aerostructural Optimization of a Transport Aircraft

Neil Wu, Charles A. Mader, and Joaquim R. R. A. Martins
University of Michigan, Ann Arbor, Michigan, 48109

Abstract

High-fidelity aerostructural aircraft optimizations have become more widespread in recent years, especially when combined with the efficiencies offered by the adjoint method and gradient-based optimizers. However, it is still computationally expensive to perform such optimizations with hundreds of design variables and multiple flight conditions. In this work, we address some practical issues with large-scale aerostructural optimizations using a multifidelity approach. We perform an optimization of a notional transport aircraft analyzed over ten flight conditions, including close to 1000 design variables and 1000 constraints. We compare the result against a reference single-fidelity optimization, and show that by leveraging multiple lower-fidelity models, we can reduce the computational cost by 14%.

1 Introduction

Multidisciplinary design optimization (MDO) has emerged in recent years as an effective strategy for improving the design of complex aerospace systems. In particular, aircraft design has become a prominent area of focus, given the complex interactions and tradeoffs between aerodynamic and structural considerations. Leveraging the efficiency of the adjoint method, these aerostructural optimizations have gradually increased in complexity, both in terms of the analysis fidelity and the scope of the optimization problems considered.

Kenway *et al.* [1] first performed aerostructural optimizations of a wing-body-horizontal tail configuration, coupling a computational structural mechanics (CSM) solver and a computational fluid dynamics (CFD) solver together, and using the coupled adjoint method to compute the total derivatives efficiently. A large multipoint problem was considered, with close to 500 design variables and over 50 constraints. However, by solving the Euler equations, the missing viscous effects caused the airfoil cross-sections to become unrealistically thick. Kenway *et al.* remedied this by solving the Reynolds-averaged Navier–Stokes (RANS) equations and a more elaborate optimization problem, as well as higher-fidelity structural modelling [2]. The updated optimization problem considered close to 1000 design variables and 1000 constraints.

Since then, further improvements in modelling have been developed. Kenway *et al.* [3] later introduced a separation-based constraint for buffet, and showed that it was effective at preventing the optimal design from buffeting. Brooks *et al.* [4] added modelling for tow-steered composites and developed two manufacturing constraints [5] suitable for gradient-based optimization. More

recently, Gray *et al.* [6] added geometrically-nonlinear structural modelling in an aerostructural optimization framework.

With the introduction of further modelling, aerostructural optimizations have become an even more challenging topic. One prominent issue with these optimizations is that they are computationally expensive. For example, the tow-steered composite optimizations by Brooks *et al.* [4] took 14 400 core h, and the undeflected Common Research Model (uCRM) optimizations took approximately 48 000 core h [7]—about two days with 1000 processors. As the computational costs scale linearly with the number of flight conditions in the multipoint optimization problem, it can quickly become prohibitive to perform large-scale optimizations, especially considering that some lower-fidelity analysis methods consider over 20 000 load cases [8].

Multifidelity methods have emerged as a way to reduce computational costs. By combining models of different fidelities, we can still obtain the high-fidelity optimum, but at significantly reduced costs. These multifidelity methods typically fall into two categories: global and local methods [9, 10]. While global methods have been shown to be effective at obtaining the global optimum for a variety of problems, the use of surrogate models make the approach prohibitive for large-scale problems involving hundreds of design variables. In local methods, gradient-based approaches are used to efficiently converge to a local optimum. Such approaches have been demonstrated on large-scale aerodynamic [11] and aerostructural [12, 13] optimizations. However, There have been little exploration of multifidelity methods applied to multidisciplinary problems, where the fidelity of each discipline can be changed independently. Furthermore, these methods do not support explicit nonlinear constraints, often using either the CFD solver or the penalty method to convert the optimization problem to an unconstrained one.

To tackle these issues, we use a sequential gradient-based multifidelity approach for MDO. We first provide an overview of the aerostructural optimization framework in section 2. Then, we explain the multifidelity approach in section 3. The optimization problem formulation is explained in section 4. Finally, we discuss the results in section 5, and summarize the conclusion in section 6.

2 Aerostructural Optimization Framework

We use the MACH Framework (MDO of Aircraft Configurations with High-fidelity) to perform aerostructural optimizations [1, 14]. It uses the multidisciplinary-feasible (MDF) formulation of the optimization problem and gradient-based optimizers to perform aerostructural optimizations. We use the extended design structure matrix (XDSM) [15] to show the high-fidelity optimization process in fig. 1.

2.1 Geometric Parameterization and Mesh Deformation

We use pyGeo [16] for geometric parameterization using free-form deformation (FFD) volumes, which are used to deform both the aerodynamic surface mesh and the structural mesh. We then use IDWarp [17] to warp the aerodynamic volume mesh based on the surface mesh deformations.

2.2 Aerostructural Analysis and Adjoint

The aerodynamic analysis is performed using ADflow [18], a finite-volume CFD solver with an efficient adjoint implementation suitable for aerodynamic shape optimization. ADflow can solve both the Euler and RANS equations with a variety of turbulence models, and we use the Spalart–Allmaras (SA) turbulence model in this work. We use the approximate Newton–Krylov approach [19] as the globalization method to accelerate convergence, then switch to a full Newton–Krylov solver towards the end of the solution.

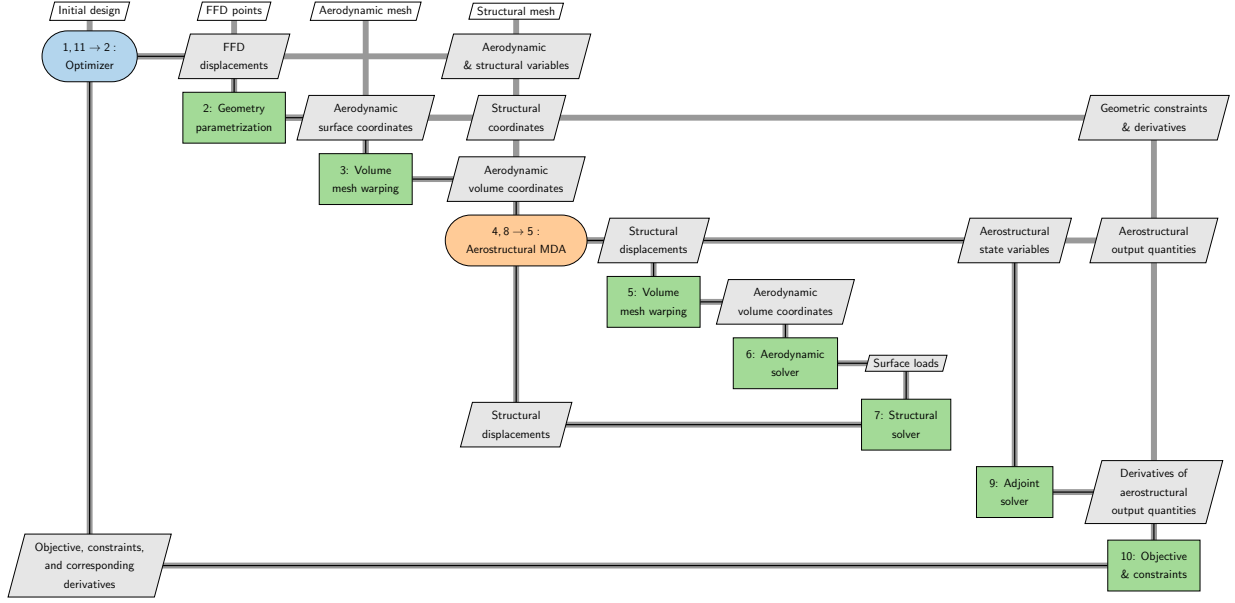


Figure 1: An XDSM diagram of the high-fidelity MDO process implemented in MACH. For simplicity, only a single MDA is shown.

The structure is analyzed using the Toolkit for Analysis of Composite Structures (TACS) [20], a CSM package designed for analysis and gradient computation of thin-walled structures. We use the approach outlined by Kenway *et al.* [14] to compute the aeroelastic deformation and the coupled sensitivities. The coupled aeroelastic system is converged using nonlinear block Gauss–Seidel with Aitken acceleration, and the coupled adjoint equation is solved using a coupled Krylov approach.

2.3 Optimization

We use pyOptSparse [21], an optimization framework tailored for large-scale gradient-based optimizations that provides wrappers for several popular optimizers. It has been shown in the past that to tackle large-scale optimizations involving hundreds or thousands of design variables, gradient-based optimization is the only viable option [22]. In this work, we use the optimizer SNOPT [23], which uses sequential quadratic programming (SQP) to solve general nonlinear constrained optimization problems.

3 Multifidelity Optimization Framework

We use a sequential multifidelity approach to accelerate optimization convergence [10], that leverages the existing high-fidelity approach for MDO. The approach consists of a sequence of single-fidelity sub-optimizations, starting from the low-fidelity model, and with the next fidelity chosen automatically via a selection algorithm. In this way, we preserve the benefits of the high-fidelity gradient-based approach: efficiency in tackling hundreds of design variables; direct handling of nonlinear constraints; and guaranteed convergence to the high-fidelity optimum. We restrict fidelity at the discipline level, such that we can update the fidelity of an analysis independently. This is particularly important in multipoint optimizations, where the appropriate fidelity could be different for the different flight conditions.

The process begins by quantifying the errors present in each fidelity, which are assumed to be

fixed throughout the optimization. We then start with the lowest-fidelity model, and perform one sub-optimization. At the end of the sub-optimization, we estimate the errors introduced into the objective and constraints due to the use of low-fidelity models. We then select the next fidelity combination as one that offers the most error reduction, normalized by the cost of each fidelity. A new sub-optimization is then performed with the updated model fidelity, and this process is continued until we select the high-fidelity model and reach the optimum.

During this process, we highlight some salient features. First, we do not fully converge each sub-optimization, since they use low-fidelity models that are inaccurate. Instead, we use error-based criteria to terminate the optimization at an appropriate moment. Furthermore, to preserve the efficiency of gradient-based optimizations, we hot-start each subsequent optimization from the previous one. By doing so, we can transfer optimizer state variables such as the approximate Hessian from one optimization to the next, such that we leverage low-fidelity models not only for an improved starting design, but also further information regarding the design space.

The XDSM diagram of the multifidelity approach is given in fig. 2.

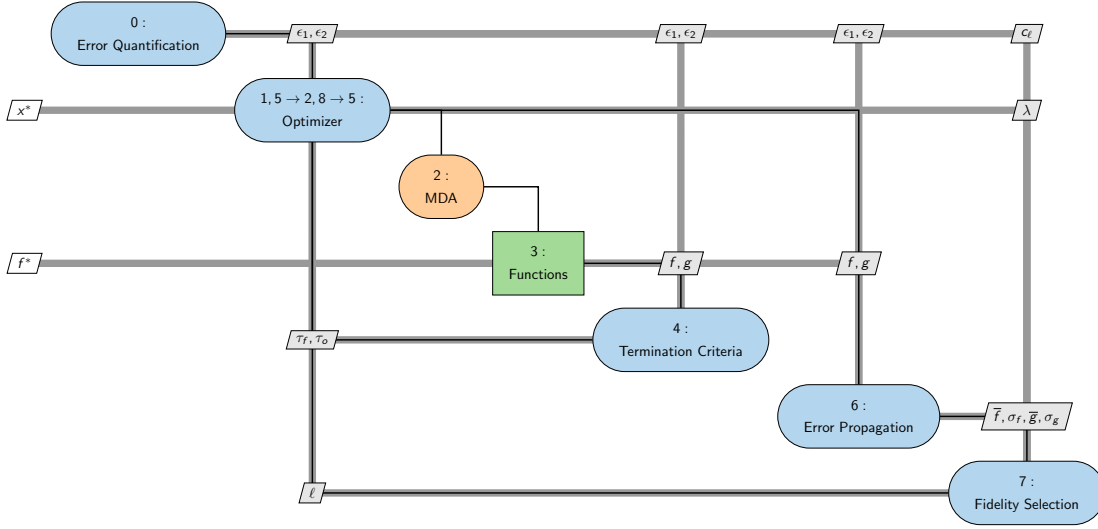


Figure 2: An XDSM diagram of the sequential multifidelity framework, again showing just a single MDA.

4 Problem Description

We perform aerostructural optimization of the XRF1 model [24], which is an aircraft research model developed by Airbus that contains both an aerodynamic and structural model. In this work, the wing-body-horizontal tail configuration is used. Figure 3 shows a sample aerostructural solution of the configuration, where the aeroelastic deflection of the wingtip due to the aerodynamic loads can be seen.

4.1 Objective Function

The objective function used for the optimization is to minimize the fuel burn at the cruise point. The fuel burn is computed based on a rearranged form of the Breguet range equation:

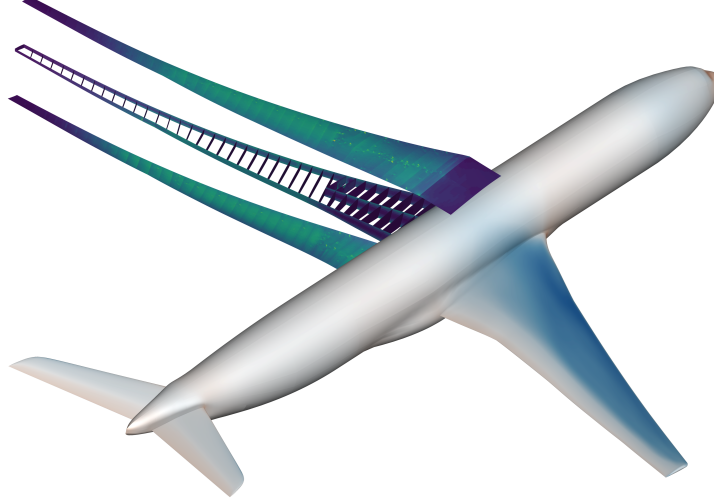


Figure 3: A sample aerostructural solution of the XRF1 wing-body-horizontal tail configuration.

$$\text{TOGW} = \text{LGW} \exp \left(\frac{R \text{ TSFC}}{V (L/D)} \right) \quad (1)$$

$$\text{FB} = \text{TOGW} - \text{LGW}, \quad (2)$$

where TOGW is the take-off gross weight, FB is the fuel burn, LGW is the aircraft landing 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 landing weight is computed using the following formula:

$$\text{LGW} = 1.25 \times W + \text{Area Weight} + \text{Fixed Weight} + \text{Payload} + \text{Reserve Fuel Weight}, \quad (3)$$

where W is the weight computed by the structural finite-element model. The factor of 1.25 used here accounts for additional weight associated with fasteners and other components not modeled in the idealized structural wingbox model. The “Area Weight” refers to the additional mass associated with the leading and trailing edges, and the necessary actuation equipment. The relevant values are listed in table 1.

Table 1: The reference values used for the optimization.

Parameter	Value
Area Weight	6000 kg
Fixed Weight	90 595 kg
Payload	23 950 kg
Reserve Fuel Weight	8000 kg
Range	8000 nm
TSFC	15.57 g/(kN s)

4.2 Cruise and Maneuver Flight Conditions

There are in total one cruise point and three maneuver points. The cruise point is used to compute the objective, while the maneuver points are used to compute the structural constraints at higher loads. The three maneuver points include symmetric 2.5 g pull-up and -1.0 g push-over maneuvers, as well as a 1.0 g quasi-static gust load following FAA Part 23.343/Part 25.341 gust criteria. Contrary to most optimizations, we do not fix the cruise altitude but set it as a design variable h . The full set of flight conditions are listed in table 2.

Table 2: Flight conditions for the multipoint optimization.

Flight condition	Mach	Altitude	Angle of attack	Fuel fraction	Load factor
1	0.83	h	α_1	0.50	1.0
1	0.78	16 000 ft	α_2	1.0	2.5
2	0.78	22 000 ft	α_3	1.0	-1.0
3	0.84	27 000 ft	α_4	1.0	1.0

4.3 Design Variables

There are over 900 design variables for this optimization problem. First are the geometric design variables, which directly impact both the aerodynamic and structural analyses. Those include local shape variables that modify the sectional airfoil shape at each spanwise station, sectional twist for the wing, and tail rotation angles for each flight condition to trim the aircraft.

Next, there are aerodynamic design variables. For this optimization, they are the individual angles of attack for the different flight conditions, and the altitude h .

Lastly are the structural design variables. We use a smeared stiffness approach to model the wingbox [25], where the effects of the stiffeners are accounted for by changing the material properties of the panel rather than explicitly modelled. As a result, instead of a single panel thickness variable for each patch of the wingbox, we have three more: stiffener pitch, thickness, and height. However, we do not allow each panel to control these variables independently. For example, we keep a single stiffener pitch variable for each of the skins, ribs, and spars. The full breakdown of the structural design variables are shown in table 4. In this work, we model the wingbox using 7000 series aluminum alloy with a yield stress of 420 MPa.

Figure 4 shows the geometric and structural design variables on the XRF1 model.

4.4 Constraints

There are also many constraints present in the optimization problem. First are the aerodynamic constraints, consisting of lift and moment constraints for each flight condition. We also enforce a separation constraint for the three maneuver cases.

Next are the geometric constraints. We impose limits on the leading edge radii and trailing edge thickness constraints on the wing, such that their values cannot go below their initial values. We also impose a fuel volume constraint, and linear leading and trailing-edge constraints to prevent the shape design variables from emulating a twist deformation.

Finally, there are structural constraints on the wingbox. These include yield and buckling constraints computed on the maneuver flight conditions, and aggregated using the Kreisselmeier–Steinhauser function [26]. We also add linear adjacency constraints to prevent significant differences in the structural design between adjacent components.

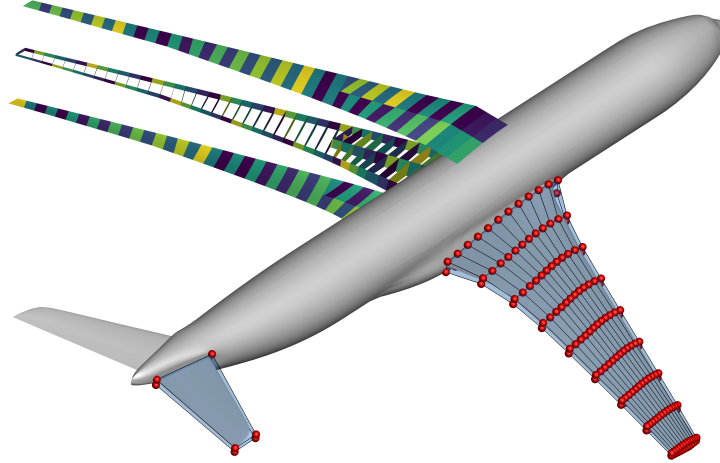


Figure 4: The geometric and structural parameterization for the XRF1 model, showing the component breakdown of the wingbox and the FFD volume used for shape control.

4.5 Available Fidelities

There are a number of fidelities available for the multifidelity approach. For aerodynamics, there are three different RANS grids available, which we label L2.5, L2, and L1.5. For structures, there are a total of six grids of varying solution order and mesh level. These fidelities are listed in table 3.

Table 3: The different fidelities available for both aerodynamics and structures.

Fidelity	# of DOF
L2.5	282 408
L2	908 832
L1.5	2 259 264
O2L3	6254
O2L2	21 937
O3L3	26 300
O2L1	85 814
O3L2	90 336
O3L1	348 470

In this work, we only consider two aerodynamic fidelities in each analysis. Fidelities L2 and L2.5 are used for the maneuver cases, while L1.5 and L2 are used for the cruise case. In total, there are 12 different fidelity combinations for each MDA, and with four different flight conditions, the number of total possible combinations is $(2 \times 6)^4 = 20\,736$, although the actual number is slightly lower due to the initial Pareto filter applied to the discipline outputs.

The full optimization problem is given in table 4.

5 Results

We perform two aerostructural optimizations. The reference optimization uses just the high-fidelity model, while the multifidelity optimization will use a number of different fidelities sequentially. Since the two approaches solve the same optimization problem, we should obtain the same optimum.

Table 4: The aerostructural optimization problem.

	Function/variable	Description	Quantity
minimize with respect to	FB	Fuel burn	1
	x_{twist}	Wing twist	9
	x_{shape}	FFD control points	216
	α_i	Angle of attack for each flight condition	4
	h	Cruise altitude	1
	x_{η_i}	Tail rotation angle for each flight condition	4
	$x_{\text{panel thick}}$	Panel thickness for skins/spars/ribs	273
	$x_{\text{skin pitch}}$	Stiffener pitch for skins	2
	$x_{\text{spar pitch}}$	Stiffener pitch for LE/TE/mid spar	3
	$x_{\text{rib pitch}}$	Stiffener pitch for ribs	1
	$x_{\text{skin height}}$	Stiffener height for skins	116
	$x_{\text{spar height}}$	Stiffener height for LE/TE/mid spar	3
	$x_{\text{rib height}}$	Stiffener height for ribs	1
	$x_{\text{skin thick}}$	Stiffener thickness for skins/spars/ribs	273
	Total		906
subject to	$L = n_i W$	Lift constraint	4
	$C_{M_y} = 0$	Trim constraint	4
	$A_{\text{sep}} \leq 4\% A_{\text{ref}}$	Amount of separation at maneuver	3
	$t_{\text{LE}}/t_{\text{LEinit}} \geq 1$	Leading edge radius	20
	$t_{\text{TE}}/t_{\text{TEinit}} \geq 1$	Trailing edge thickness	20
	$V_{\text{wing}} > V_{\text{fuel}}$	Minimum fuel volume	1
	$\Delta z_{\text{TE,upper}} = -\Delta z_{\text{TE,lower}}$	Fixed trailing edge (L)	8
	$\Delta z_{\text{LE,upper}} = -\Delta z_{\text{LE,lower}}$	Fixed leading edge (L)	8
	$\text{KS}_{\text{buckling}} \leq 1$	2.5 g buckling	3
	$\text{KS}_{\text{yield}} \leq 1$	2.5 g yield stress	4
	$\text{KS}_{\text{buckling}} \leq 1$	1.0 g gust buckling	3
	$\text{KS}_{\text{yield}} \leq 1$	1.0 g gust yield stress	4
	$\text{KS}_{\text{buckling}} \leq 1$	-1.0 g buckling	3
	$ x_{\text{panel thick}_i} - x_{\text{panel thick}_{i+1}} \leq 0.0005$	Skin thickness adjacency (L)	217
	$ x_{\text{skin thick}_i} - x_{\text{skin thick}_{i+1}} \leq 0.0005$	Stiffener thickness adjacency (L)	217
	$ x_{\text{skin height}_i} - x_{\text{skin height}_{i+1}} \leq 0.0005$	Stiffener height adjacency (L)	217
	$x_{\text{skin thick}} - x_{\text{panel thick}} < 0.0025$	Maximum stiffener-skin difference (L)	186
	Total		922

Table 5 shows a summary of the two optimizations. Compared to the single-fidelity optimization, the multifidelity approach took seven sub-optimizations and 14% less computational resources in total, as measured in processor-hours. Although it took more major iterations in total, many of the earlier iterations were substantially cheaper since they were computed using lower-fidelity models.

Upon closer examination, the fidelity choices presented in table 5 are quite intuitive. The fidelity selection algorithm preferred improving the structural fidelity of the maneuver points first, as the structural constraints had a significant impact on the overall optimization. Furthermore, the order of improvement was correlated with the amount of constraint violation at each maneuver points, as the 2.5 g point was most structurally constrained, and the -1.0 g point least so. On the other hand, since the cruise structural fidelity was only responsible for computing the structural mass, it had much less impact on the optimization and was therefore only improved at the last sub-optimization. In terms of the aerodynamic fidelities, the 2.5 g point was improved at the same time as the cruise point, showing that the algorithm allowed for simultaneous improvement in the fidelity of multiple analyses. The cruise aerodynamic fidelity was also improved earlier than most of the structural

Table 5: The sequence of fidelities for the aerostructural optimization, showing the fidelity used, the computational costs, and the number of major iterations taken. Here the fidelity combination is represented by four pairs of columns, corresponding to the four analysis points. The cost is given in proc-hours, and the relative cost is normalized by the total cost of the multifidelity approach.

#	Cruise		Maneuver 1		Maneuver 2		Maneuver 3		Cost	% Cost	# Iter
	Aero	Struct	Aero	Struct	Aero	Struct	Aero	Struct			
1	L2	O2L1	L2.5	O2L1	L2.5	O2L1	L2.5	O2L1	1661	7	95
2	L2	O2L1	L2.5	O3L0	L2.5	O2L1	L2.5	O2L1	703	3	32
3	L2	O2L1	L2.5	O3L0	L2.5	O2L1	L2.5	O3L0	911	4	47
4	L1.5	O2L1	L2	O3L0	L2.5	O2L1	L2.5	O3L0	5802	23	79
5	L1.5	O2L1	L2	O3L0	L2	O2L0	L2	O3L0	1134	4	19
6	L1.5	O2L1	L2	O3L0	L2	O3L0	L2	O3L0	876	3	16
7	L1.5	O3L0	L2	O3L0	L2	O3L0	L2	O3L0	14 278	56	248
Total									25 365	100	536
HF	L1.5	O3L0	L2	O3L0	L2	O3L0	L2	O3L0	28 891	114	454

fidelities, as it is responsible for computing both lift and drag, while only lift is needed at the maneuver points. Therefore, the total combined benefit of improving the cruise aerodynamic fidelity was likely higher, and the fidelity selection algorithm therefore favoured such a choice.

Figure 5 shows the progress of the optimization, measured in terms of the normalized distance as computed by:

$$d(\mathbf{x}) = \frac{\|\mathbf{x} - \mathbf{x}_{\text{final}}\|_2}{\|\mathbf{x}_{\text{initial}} - \mathbf{x}_{\text{final}}\|_2}, \quad (4)$$

where $\mathbf{x}_{\text{initial}}$ is the initial design and $\mathbf{x}_{\text{final}}$ the final optimum as computed by the high-fidelity optimization.

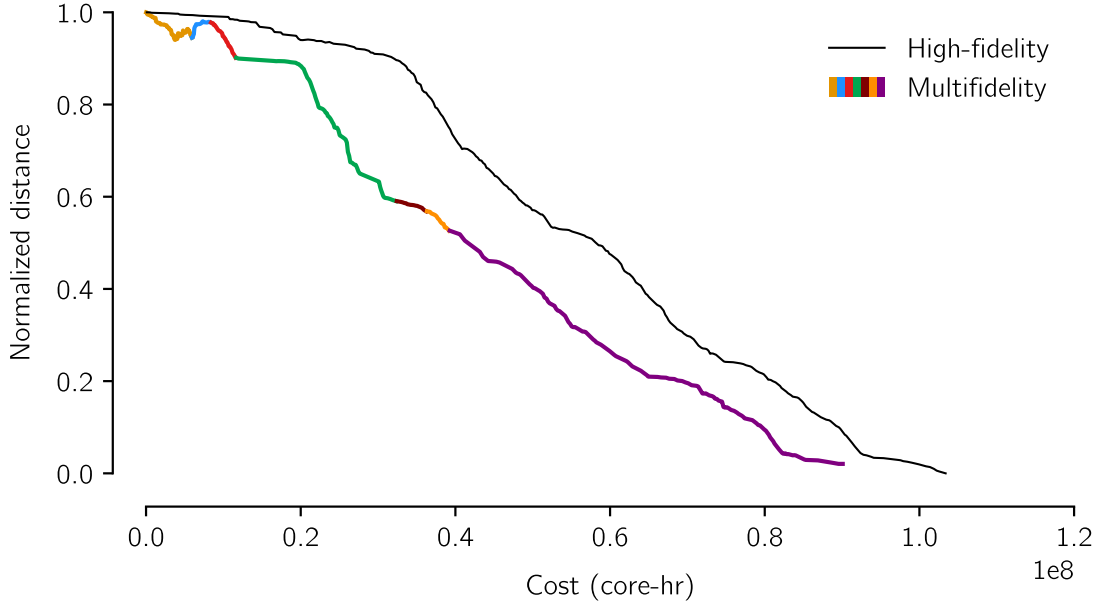


Figure 5: Normalized distance from the initial to the final optimum.

At each optimization iteration, the distance is normalized such that the initial design is one unit

away from the optimum. Since both optimizations solve the same problem, if the design space is unimodal then we expect both approaches to obtain the same numerical optimum. While there are subtle differences in the final designs, they are expected given the level of convergence for the optimizations. In terms of the objective, the difference between the fuel burn of the two optimal designs is 0.54 kg, for a relative difference of 5×10^{-6} .

Next we examine the intermediate solutions in more detail. Figure 6 shows three airfoil profiles at different spanwise stations along the wing. At the tip, the effect of structural deformation is evident. Due to changes in both the design and the analysis fidelity, the tip deflection changes significantly between intermediate designs. Based on the pressure distribution and airfoil profiles, the optimum found by the fourth sub-optimization (labelled “Opt 4”) is already close to the final optimum, despite only taking about 30% of the high-fidelity cost.

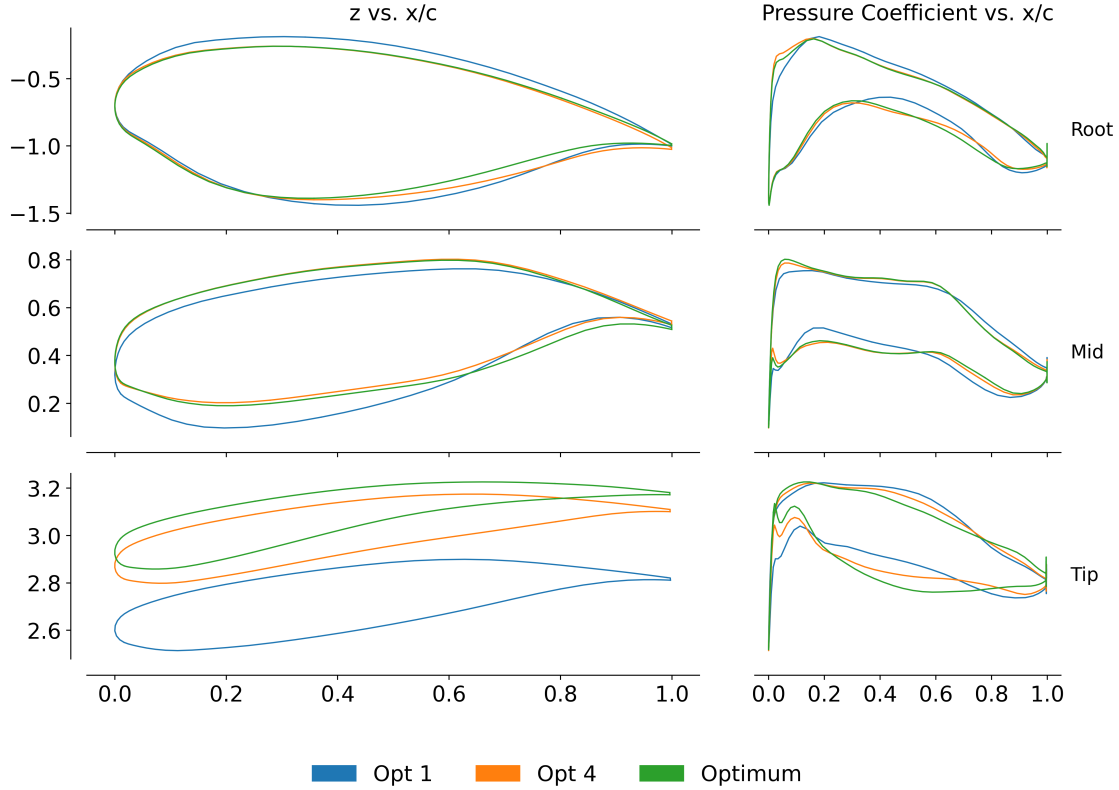


Figure 6: Airfoil profiles and pressure distributions along the wing.

Next we look at the structural design. Figure 7 shows the initial, intermediate, and final structural solutions, both in terms of the mass-equivalent structural thicknesses and the stress distribution. The structural design variables converge relatively quickly, and there are little discernable differences between the final few solutions. The stresses are normalized such that a value of 1 corresponds to the yield limit. We see that the initial design was infeasible, as there are large stresses near the tip of the mid-spar. However, by the end of the first optimization, these stress concentrations have been removed and the design has become nearly feasible. Regardless of the analysis fidelity, the optimizer keeps the overall stress distribution relatively similar.

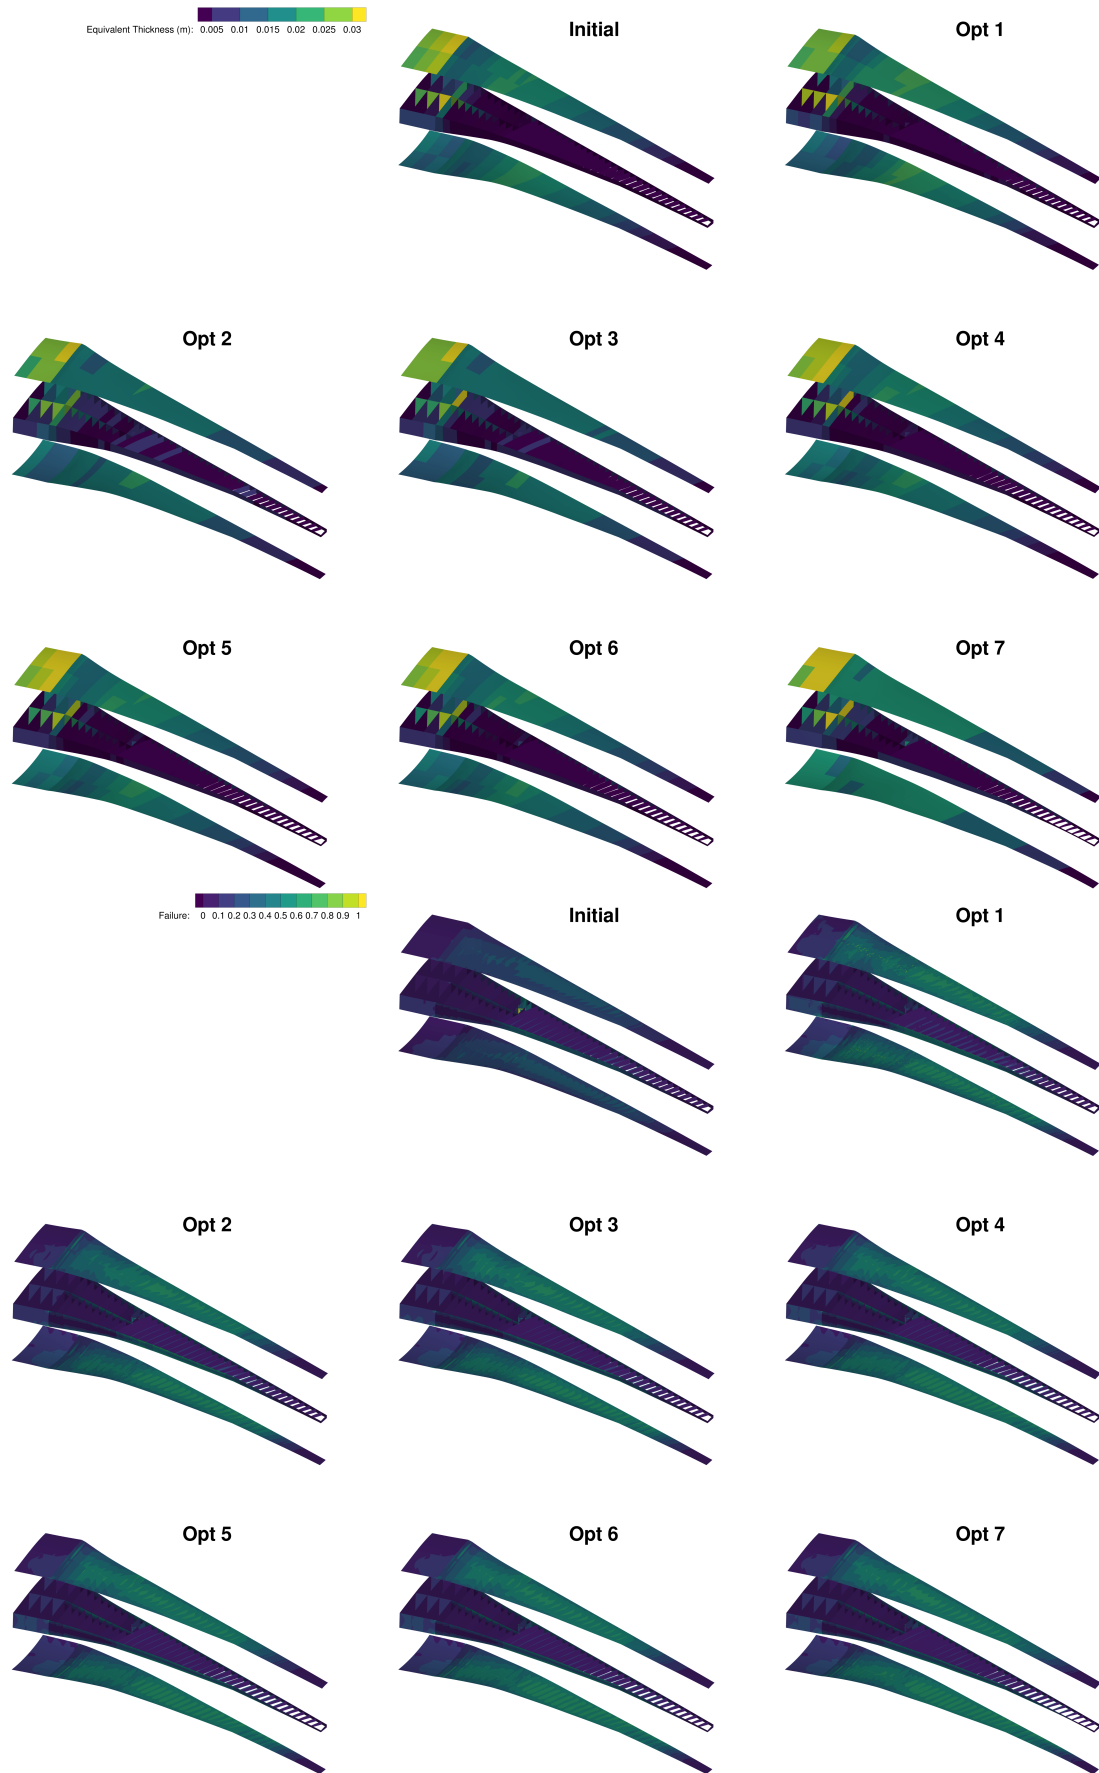


Figure 7: Initial, intermediate, and final structural designs.

While the cost reduction offered by the multifidelity method is notable, it is substantially less than the savings previously reported [10]. We believe this can be partially attributed to the difference in the cost ratios between the low and high-fidelity models. In the previous work, the computational cost of the high-fidelity model was a few hundred times higher than the low-fidelity model, which contributed to the overall effectiveness. In this case, the cost ratios were only around 4. Without the availability of substantially cheaper models, the multifidelity approach was unable to replicate the gains that were seen earlier. This highlights the importance of obtaining appropriate model fidelities, and that despite the lack of cheap models, moderate gains still could be obtained.

6 Conclusion

In this work, we performed aerostructural optimizations of the wing-body-horizontal tail configuration of the XRF1 design using a multifidelity MDO framework. The multipoint optimization contains over 900 design variables and a similar number of constraints, analyzed at four different operating points. Through the use of the sequential multifidelity approach and gradient-based optimization techniques, we were able to obtain the same numerical optimum with 14% less computational cost.

Acknowledgments

The material is based upon work supported by Airbus in the frame of the Airbus / Michigan Center for Aero-Servo-Elasticity of Very Flexible Aircraft. Special thanks to Tom Gibson and Joël Brezillon for their expert advice and review of this paper.

References

- [1] Kenway, G. K. W., and Martins, J. R. R. A., “Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration,” *Journal of Aircraft*, Vol. 51, No. 1, 2014, pp. 144–160. doi:[10.2514/1.C032150](https://doi.org/10.2514/1.C032150).
- [2] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Aerostructural Optimization of the Common Research Model Configuration,” *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, 2014. doi:[10.2514/6.2014-3274](https://doi.org/10.2514/6.2014-3274), aIAA 2014-3274.
- [3] Kenway, G. K. W., and Martins, J. R. R. A., “Buffet-Onset Constraint Formulation for Aerodynamic Shape Optimization,” *AIAA Journal*, Vol. 55, No. 6, 2017, pp. 1930–1947. doi:[10.2514/1.J055172](https://doi.org/10.2514/1.J055172).
- [4] Brooks, T. R., Martins, J. R. R. A., and Kennedy, G. J., “High-fidelity Aerostructural Optimization of Tow-steered Composite Wings,” *Journal of Fluids and Structures*, Vol. 88, 2019, pp. 122–147. doi:[10.1016/j.jfluidstructs.2019.04.005](https://doi.org/10.1016/j.jfluidstructs.2019.04.005).
- [5] Brooks, T. R., and Martins, J. R. R. A., “On Manufacturing Constraints for Tow-steered Composite Design Optimization,” *Composite Structures*, Vol. 204, 2018, pp. 548–559. doi:[10.1016/j.compstruct.2018.07.100](https://doi.org/10.1016/j.compstruct.2018.07.100).
- [6] Gray, A. C., and Martins, J. R. R. A., “Geometrically Nonlinear High-fidelity Aerostructural Optimization for Highly Flexible Wings,” *Proceedings of the AIAA SciTech Forum*, 2021. doi:[10.2514/6.2021-0283](https://doi.org/10.2514/6.2021-0283).

- [7] Brooks, T. R., Kenway, G. K. W., and Martins, J. R. R. A., “Benchmark Aerostructural Models for the Study of Transonic Aircraft Wings,” *AIAA Journal*, Vol. 56, No. 7, 2018, pp. 2840–2855. doi:[10.2514/1.J056603](https://doi.org/10.2514/1.J056603).
- [8] Rajpal, D., Gillebaart, E., and De Breuker, R., “Preliminary aeroelastic design of composite wings subjected to critical gust loads,” *Aerospace Science and Technology*, Vol. 85, 2019, pp. 96–112.
- [9] Peherstorfer, B., Willcox, K., and Gunzburger, M., “Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization,” *SIAM Review*, Vol. 60, No. 3, 2018, pp. 550–591. doi:[10.1137/16M1082469](https://doi.org/10.1137/16M1082469).
- [10] Wu, N., Mader, C. A., and Martins, J. R. R. A., “A Gradient-based Sequential Multifidelity Approach to Multidisciplinary Design Optimization,” *Structural and Multidisciplinary Optimization*, Vol. 65, 2022, pp. 131–151. doi:[10.1007/s00158-022-03204-1](https://doi.org/10.1007/s00158-022-03204-1).
- [11] Olivanti, R., Gallard, F., Brézillon, J., and Gourdain, N., “Comparison of Generic Multi-Fidelity Approaches for Bound-Constrained Nonlinear Optimization Applied to Adjoint-Based CFD Applications,” *AIAA Aviation 2019 Forum*, 2019. doi:[10.2514/6.2019-3102](https://doi.org/10.2514/6.2019-3102).
- [12] Bryson, D. E., and Rumpfkeil, M. P., “Aerostructural Design Optimization Using a Multifidelity Quasi-Newton Method,” *Journal of Aircraft*, Vol. 56, No. 5, 2019, pp. 2019–2031. doi:[10.2514/1.C035152](https://doi.org/10.2514/1.C035152).
- [13] Thelen, A. S., Bryson, D. E., Stanford, B. K., and Beran, P. S., “Multi-Fidelity Gradient-Based Optimization for High-Dimensional Aeroelastic Configurations,” *Algorithms*, Vol. 15, No. 4, 2022, p. 131. doi:[10.3390/a15040131](https://doi.org/10.3390/a15040131).
- [14] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations,” *AIAA Journal*, Vol. 52, No. 5, 2014, pp. 935–951. doi:[10.2514/1.J052255](https://doi.org/10.2514/1.J052255).
- [15] Lambe, A. B., and Martins, J. R. R. A., “Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes,” *Structural and Multidisciplinary Optimization*, Vol. 46, 2012, pp. 273–284. doi:[10.1007/s00158-012-0763-y](https://doi.org/10.1007/s00158-012-0763-y).
- [16] Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A., “A CAD-Free Approach to High-Fidelity Aerostructural Optimization,” *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, 2010. doi:[10.2514/6.2010-9231](https://doi.org/10.2514/6.2010-9231).
- [17] Secco, N., Kenway, G. K. W., He, P., Mader, C. A., and Martins, J. R. R. A., “Efficient Mesh Generation and Deformation for Aerodynamic Shape Optimization,” *AIAA Journal*, Vol. 59, No. 4, 2021, pp. 1151–1168. doi:[10.2514/1.J059491](https://doi.org/10.2514/1.J059491).
- [18] Mader, C. A., Kenway, G. K. W., Yildirim, A., and Martins, J. R. R. A., “ADflow: An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization,” *Journal of Aerospace Information Systems*, Vol. 17, No. 9, 2020, pp. 508–527. doi:[10.2514/1.I010796](https://doi.org/10.2514/1.I010796).
- [19] Yildirim, A., Kenway, G. K. W., Mader, C. A., and Martins, J. R. R. A., “A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations,” *Journal of Computational Physics*, Vol. 397, 2019, p. 108741. doi:[10.1016/j.jcp.2019.06.018](https://doi.org/10.1016/j.jcp.2019.06.018).

- [20] Kennedy, G. J., and Martins, J. R. R. A., “A Parallel Finite-Element Framework for Large-Scale Gradient-Based Design Optimization of High-Performance Structures,” *Finite Elements in Analysis and Design*, Vol. 87, 2014, pp. 56–73. doi:[10.1016/j.finel.2014.04.011](https://doi.org/10.1016/j.finel.2014.04.011).
- [21] Wu, N., Kenway, G., Mader, C. A., Jasa, J., and Martins, J. R. R. A., “pyOptSparse: A Python framework for large-scale constrained nonlinear optimization of sparse systems,” *Journal of Open Source Software*, Vol. 5, No. 54, 2020, p. 2564. doi:[10.21105/joss.02564](https://doi.org/10.21105/joss.02564).
- [22] Yu, Y., Lyu, Z., Xu, Z., and Martins, J. R. R. A., “On the Influence of Optimization Algorithm and Starting Design on Wing Aerodynamic Shape Optimization,” *Aerospace Science and Technology*, Vol. 75, 2018, pp. 183–199. doi:[10.1016/j.ast.2018.01.016](https://doi.org/10.1016/j.ast.2018.01.016).
- [23] Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131. doi:[10.1137/S0036144504446096](https://doi.org/10.1137/S0036144504446096).
- [24] Pattinson, J., and Herring, M., “High Fidelity Simulation of Wing Loads with an Active Winglet Control Surface,” *IFASD 2013, International Forum on Aeroelasticity and Structural Dynamics, 24-26 June 2013, Bristol (GB)*, Royal Aeronautical Society, Bristol, UK, 2013.
- [25] 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,” *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, National Harbor, MD, 2014. doi:[10.2514/6.2014-0596](https://doi.org/10.2514/6.2014-0596).
- [26] Kreisselmeier, G., and Steinhauser, R., “Systematic Control Design by Optimizing a Vector Performance Index,” *International Federation of Active Controls Symposium on Computer-Aided Design of Control Systems, Zurich, Switzerland*, 1979. doi:[10.1016/S1474-6670\(17\)65584-8](https://doi.org/10.1016/S1474-6670(17)65584-8).