

AEROELASTIC ADJOINT-BASED OPTIMISATION OF AN HIGHLY FLEXIBLE AIRCRAFT WING CONFIGURATION

M. MEHEUT^{*}, M. CARINI^{*} AND C. BLONDEAU[†]

^{*}ONERA, Université Paris-Saclay, F-92190 Meudon, France
e-mail: michael.meheut@onera.fr, marco.carini@onera.fr

[†]ONERA, Université Paris-Saclay, F-92322 Châtillon, France
email: christophe.blondeau@onera.fr

Key words: Adjoint, MDO, aero-elastic optimisation.

Abstract. This paper investigates benefits resulting from the use of coupled aeroelastic analysis for aerodynamic shape optimisation of a highly flexible wing. The study is carried out on the eXternal Research Forum model (XRF-1) specified by Airbus Commercial Aircraft, representative of a long-range aircraft configuration. Improvements delivered by considering aeroelastic effects for the evaluation of both the aerodynamic performance and the associated gradients are assessed with respect to the results obtained by freezing the wing flexibility in both primal and adjoint computations. An analysis of the impact on the different drag components is also illustrated based on the far-field drag breakdown. Results show that for induced drag, engaging flexibility only at the primal level still allows to capture first-order gain on the final performance. However, engaging coupled-adjoint sensitivities is key to completely master wave drag reduction on the considered highly flexible wing. Performance improvement obtained by increasing the number of design parameters is also investigated.

1 INTRODUCTION

Today in industry, during aircraft or engine design phases, most optimisation studies (both gradient-free and gradient-based) performed using High-Fidelity (HiFi) tools focus on a single discipline (aerodynamics, structural analysis, acoustics, ...) while new emerging configurations feature tight interactions and critical trade-offs among disciplines only coupled multi-physics analyses can properly capture and identify. In order to demonstrate the benefits of engaging Multi-Disciplinary Optimisation (MDO) early in the industrial design campaign, the European H2020 MADELEINE project was launched in 2018 [1]. MADELEINE aimed at strengthening the capabilities and use of multi-physics adjoint solvers [1] to maximise the benefit obtained from computationally intensive simulations that are key enablers for future airframe and engine design.

In this context, one of the ONERA activities was focused on the aeroelastic adjoint capability applied to flexible wing design in collaboration with Airbus Commercial Aircraft and DLR. The Airbus XRF-1 research test case, representative of a long-range aircraft configuration, was shared among partners. In [4] ‘frozen-flexibility’ approaches have been

illustrated in order to consider structural deformation in the shape optimisation loop while avoiding the increase of both computational cost and code complexity associated with a fully-differentiated aero-elastic solver. The jig-to-flight shape displacements are pre-computed for each flight condition of interest by performing fully-coupled aeroelastic computations on the baseline geometry. Then, a classical multipoint “rigid” adjoint-based aerodynamic optimisation is performed on the flight shape resulting from the jig one augmented by the baseline structural displacements. Such weakly-coupled approaches, although computationally attractive and compliant with industrial legacy tools and company structure, are not able to capture the structural deformation feedback on aerodynamic shape changes during the optimisation process. This coupled effect requires a fully-differentiated aeroelastic formulation leading to the so-called coupled-adjoint capability [5][6] which have been successfully demonstrated in several studies to conduct aeroelastic and aero-structural optimisations [5][11], but often restricted to academic configurations.

In [3], ONERA in collaboration with DLR, presented a demonstration of coupled adjoint-based optimisations on the XRF-1 test case, with the integration of industrial tools in the complete optimisation chain. Starting from these preliminary achievements, the present paper aims at quantifying the impact of considering aeroelastic effects in the optimisation loop for aerodynamic performance evaluation on the one hand, and for gradient computation (through coupled-adjoint) on the other hand. The investigation is carried out for an increased flexibility of the wing structure with respect to that considered in [3] and based on the far-field drag analysis [12] of the resulting flow fields, with focus on the induced and wave drag contributions. Finally, the further performance improvement that can be achieved by increasing the number of shape design variables is also assessed.

2 METHODOLOGY

2.1 Aerodynamic Shape Parameterisation

The same computational chain described in [3] is employed here. In particular, the Airbus CAD modeller PADGE (Parametric And Differentiated Geometrical Engine) is used to represent the XRF-1 wing shape geometry and to parametrise it, enabling realistic industrial shape optimisations. Two parameterisation levels are considered, the *Medium* one and the *Fine* one; the corresponding main features are summarised in Table 1. Both parameterisation levels share the same 11 active control sections, which are illustrated in Figure 1 (left panel). For each section, 6 camber control points are used for the Medium level, which are increased to 10 in the Fine level. An example of the local camber variation associated to a single control point is illustrated in Figure 1, top-right panel. For the two shape parameterisations, the precise chordwise distributions of the camber control points (in terms of chord percentage) is reported in Table 1: a finer control is introduced both at leading and trailing edges in the Fine level compared to the Medium one. For both levels, the same twist law control is adopted with a total of 7 design variables. More precisely, the twist distribution is linearly interpolated between the crank and the tip sections (as illustrated in the middle-right panel of Figure 1) with an additive third-order spline correction for the inner and outer wing, separately (as illustrated in the bottom-right panel of Figure 1). The total number of shape design variables

is 73 and 117 for the Medium and Fine levels, respectively.




Figure 1: XRF-1 wing-shape PADGE parametrisation. Left: wing control sections. Right: example of camber and twist modifications.

Table 1: Summary of the main features of the employed wing-shape PADGE parametrisations.

Shape Parametrisation Features	Medium	Fine
<i>N. control sections</i>	11	11
<i>N. twist variables</i>	7	7
<i>N. camber control point per section</i>	6	10
<i>Camber control point $x/c\%$ distribution</i>		5
		10
	15	15
	30	30
	45	45
	60	60
	80	80
		85
	90	90
		95
<i>N. design variables</i>	73	117

2.2 Aeroelastic analysis

Both CFD and coupled CFD-CSM analyses are carried out using the structured *elsA* solver (ONERA-Airbus-Safran property) [14]. The CFD mesh of the XRF-1 wing-body configuration (jig shape) consists of 5.3M points and the associated surface mesh is illustrated in Figure 2 (left panel). The Spalart-Allmaras turbulence model is adopted for RANS modelling and the Jameson scheme is used for the inviscid flux discretization, with $k_2=0.5$ and $k_4=0.016$. The resulting discretised equations are solved using pseudo-time iterations (with maximum CFL=100) and multigrid acceleration. For aeroelastic analyses, a segregated two-way coupled CFD-CSM approach is implemented in *elsA* and the mesh deformation is handled by classical Inverse Distance Weighting (IDW) technique combined with Trans-Finite Interpolation (TFI) to reduce the associated computational cost. An example of residual convergence for primal aeroelastic computations at cruise conditions is reported in Figure 2 (right panel), showing a decrease of four orders of magnitude in the density residual. Advanced Krylov subspace methods based on flexible inner-outer GMRES solver combined with ad-hoc preconditioning strategies and deflation techniques [13] are employed for the numerical solution of the adjoint problem.

The employed wing structure model is derived from the FS2 finite element model generated by DLR through the CPACS-MoNa sizing process and employed in the aeroelastic optimisation studies presented in [3]. The FS2 structure already featured an increased flexibility due to the reduced number of load cases considered in the sizing process with respect to those associated with the datum XRF-1 wing structure. The resulting flight shape at $CL=0.5$ (and cruise conditions) is compared to the jig shape in Figure 3 (left panel). In the aeroelastic analyses and optimisation studies illustrated in the following, the FS2 flexibility matrix has been scaled by a factor of 1.75 to artificially increase the wing deflection, thus further challenging the aeroelastic optimisation. The resulting flight shape at $CL=0.5$ is also illustrated in Figure 3 (left panel), clearly showing a remarkably larger displacement of the wing tip compared to the FS2 flight shape. In Figure 3 (right panel) the computed pressure coefficient distribution is also illustrated.

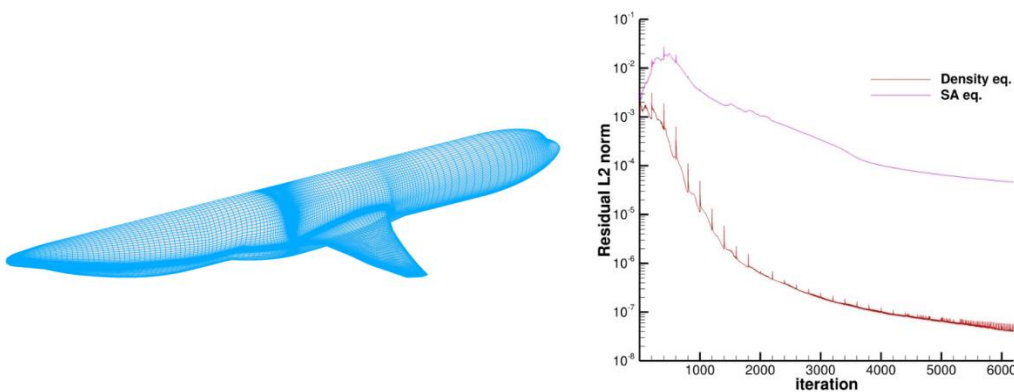


Figure 2: Aeroelastic analysis using the *elsA* structured solver. Left: surface mesh of the XRF-1 wing_body jig shape. Right: example of residual convergence for aeroelastic computations at cruise conditions.




Figure 3: Aeroelastic analysis using the elsA structured solver. Left: comparison of jig shape in input (grey shaded) and resulting flight shapes for FS2 DLR's wing structure model (blue shaded) and the scaled one used in the present study (orange shaded). Right: pressure coefficient distribution on suction (upper half model) and pressure side (lower half model) on the computed flight shape at cruise conditions ($CL=0.5$).

2.3 Optimisation problem formulation and approaches

The optimisation of the wing shape parameters introduced in §2.1 seeks the constraint minimisation of the sum of the aerodynamic drag over three flight conditions for the same cruise altitude and speed. The values of $CL=0.45$, $CL=0.5$ and $CL=0.55$ are selected, where the lower and higher lift conditions are essentially introduced to ensure a smooth performance in the neighbourhood of the target cruise point $CL=0.5$. The expression of the objective function J is simply given by:

$$J = CD_{ff}|_{CL=0.45} + CD_{ff}|_{CL=0.5} + CD_{ff}|_{CL=0.55}, \quad (1)$$

with CD_{ff} denoting the far-field drag coefficient [12], extracted from the post-processing of the whole flow field through the in-house ONERA code FFD72. Moreover, thanks to such analysis the drag coefficient CD_{ff} is decomposed into its different physical contributions: friction drag (CD_f), viscous-pressure drag (CD_{vp}), wave drag (CD_w) and induced drag (CD_i). In particular, the impact of the optimisation process on CD_w and CD_i will be analysed in the following.

In order to assess the benefit of engaging flexibility in primal and adjoint computations in the optimisation process, three different approaches are considered:




- a pure *rigid* aerodynamic optimisation of the flight shape at $CL=0.5$. In this approach the structural displacement computed from the aeroelastic analysis of the baseline configuration at $CL=0.5$ is frozen, neglecting flexibility both in primal and adjoint computations. Differently from the weakly-coupled approach presented in [4], the same rigid flight shape is employed to evaluate the aerodynamic performance and the associated rigid gradients for all flight points. Although, this approach is less consistent than storing the baseline structural displacement for the different lift conditions, the performance evaluation at higher and lower CL are still included in the

cost function to affect the robustness of the optimisation result, although not completely meaningful from physical point of view.

- A *hybrid* approach where flexibility is considered in primal computations only. In this case the optimisation is carried out on the jig shape and aeroelastic analyses are integrated in the optimisation loop to evaluate lift and drag coefficients. However, wing flexibility is frozen in the adjoint computations, and rigid (pure aerodynamic) gradients are computed around each of the three flight shapes resulting from primal aeroelastic analyses.
- A *flexible* approach where a fully consistent aeroelastic optimisation of the jig shape is carried out by engaging flexibility both in primal and adjoint computations. In this case, aeroelastic analyses are performed to evaluate lift and drag coefficients and the aeroelastic coupled-adjoint capability is used to compute the shape (flexible) gradients.

These distinct optimisation strategies are summarised in Table 2. For each flight condition and consistently with the associated lift constraint, the corresponding value of the angle of attack is introduced in addition to the shape parameters in the optimisation process. The total number of optimisation variables is thus 76 and 120 when using the Medium and Fine parametrisations, respectively. The pipeline of the complete workflow is shown in Figure 4: for given admissible values of the shape parameters α , the wing CAD geometry is first updated through the PADGE tool and then the CFD mesh X is deformed accordingly, by means of the combined IDW and TFI techniques. If sensitivity information are required by the optimiser, the mesh gradient $dX/d\alpha$ is also computed starting from the analytical geometrical gradients provided by PADGE through the differentiated CAD representation of the wing surface. Once X is available, CFD or CFD-CSM computations are carried out for each considered flight condition, depending on the adopted optimisation approach (see Table 2). The resulting conservative flow fields W are then analysed through the far-field drag post-processing to accurately estimate the aerodynamic drag and its associated physical components. The same analysis can also provides the gradient of the aerodynamic coefficients (and hence of the objective function J) with respect to both W , to feed the adjoint problem, and to X , to finally assemble the total derivative $dJ/d\alpha$. Both J and $dJ/d\alpha$ are provided as input to the optimiser: for the present studies, the SLSQP algorithm available in the open-source pyOpt library [15] is employed.

Table 2: Summary of optimisation approaches.

	Rigid	Hybrid	Flexible
<i>Wing Shape</i>	Flight shape at CL=0.5	JIG shape	JIG shape
			
<i>CL & CD_{ff} computation</i>	CFD	CFD-CSM	CFD-CSM

CL & CD_{ff} gradient computation	CFD	CFD (aeroelastic flight shape)	CFD-CSM
--	-----	-----------------------------------	---------

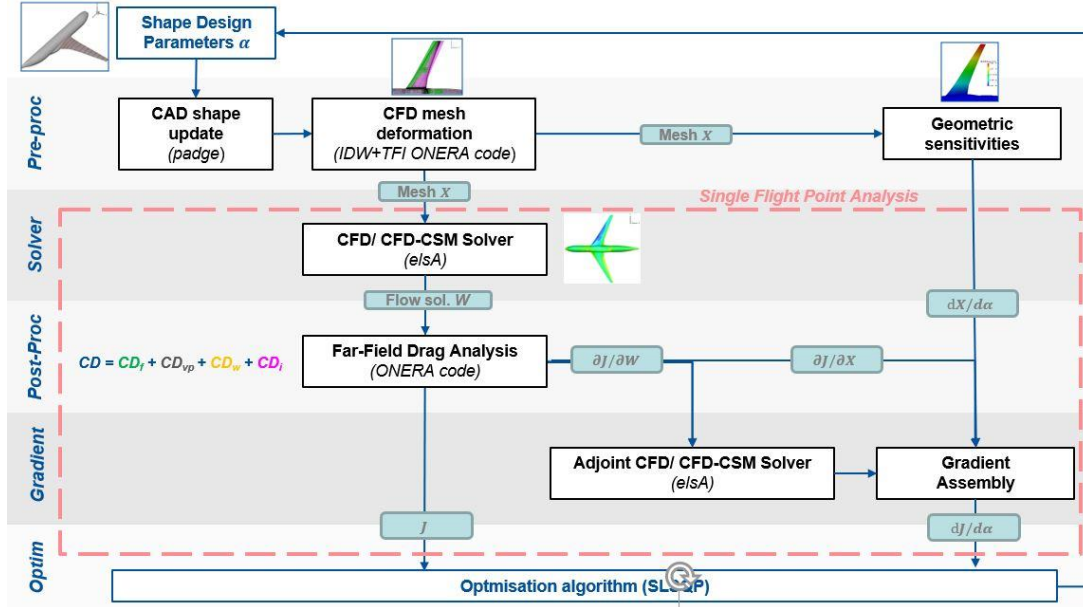


Figure 4: Schematic of the optimisation workflow.

3 OPTIMISATION RESULTS

3.1 Comparison of rigid, hybrid and flexible results

The optimisation results obtained using the different optimisation approaches summarised in Table 2 are compared in Figure 5 in terms of far-field drag polar curves: all these results are based on the Medium shape parametrisation (see Table 2). Drag values have been made non-dimensional with respect to the baseline aeroelastic performance and represented in terms of percentage variation. In the same plot, the rigid baseline curve obtained by taking the computed baseline flight shape at $CL=0.5$ and evaluating its pure aerodynamic performance (as a rigid shape) is also depicted: this curve corresponds to the initial performance level of the rigid optimisation. Such approach seemingly leads to a huge gain with respect to both rigid and aeroelastic baselines, by achieving $\sim 11\%$ and $\sim 13\%$ of $CD_{ff}|_{CL=0.55}$ reduction, respectively. However, a large fraction of this gain is lost when an aeroelastic assessment of the rigid optimised shape is performed, thus highlighting the relevant impact of flexibility and the lack of reliability of the estimated rigid performance. Conversely, the hybrid approach does not suffer from this drawback, since flexibility is considered at least in the evaluation of J . Moreover, the aeroelastic performance is improved by $\sim 0.5-1.0\%$ with respect to the rigid one undergoing structural flexibility. Finally, a further substantial improvement is achieved when a fully flexible optimisation approach is adopted, by introducing coupled-adjoint

capabilities. As expected, the impact of using flexible gradients becomes more important at higher load conditions: the difference between the flexible and the hybrid approach is quite small at $CL=0.45$ and increases up to $\sim 2\%$ at $CL=0.55$.

The analysis of induced and wave drag components is detailed in Figure 6 (left panel) and Figure 6 (right panel), respectively. For the induced drag, engaging flexibility essentially introduces an offset on the resulting performance. The rigid CD_i gain is indeed overestimated, being nearly double its effective value as computed from *a posteriori* aeroelastic analysis. The same performance is obtained by the hybrid approach, while it is reduced by $\sim 0.5\%$ when using the flexible one. These observations suggest that when engaging flexibility, the control on induced drag, essentially provided by the twist design variables, becomes less effective and the optimisation algorithm is not able to achieve the same performance level and the same wave-induced drag compromise, as in the rigid case. Indeed, with reference to Figure 6 (right panel), wave drag is progressively reduced by the hybrid and flexible optimisation processes, similarly to what observed for the far-field drag in Figure 5. Moreover, in the flexible case, the wave drag reduction is also slightly improved over the pure rigid performance, especially at $CL=0.5$. This is also confirmed by inspecting the corresponding load and wave drag spanwise distributions, which are illustrated in Figure 7 in terms of variation with respect to the aeroelastic baseline results. For the spanwise load distribution, see Figure 7 (left panel), the major difference is located in the outboard area. All the approaches (rigid, hybrid and flexible) tend to increase the load in this region and to reduce it on the inboard one, as usually observed in similar optimisation studies. However, the outboard load increment towards the wing tip is more pronounced in the rigid case compared to the hybrid and flexible ones, due to the lack of downwash effect caused by wing flexibility. This is also suggested by the inspection of the spanwise load distribution resulting from the aeroelastic analysis of the rigid optimised shape, which is found very close to that obtained by the hybrid approach. For the wave drag distribution, Figure 7 (right panel), all the approaches exhibit a similar pattern characterised by two main peaks in the spanwise distribution, around $y/b \approx 0.2$ and $y/b \approx 0.55$. Consistently with the results reported in Figure 6 (right panel), compressibility effects are reduced everywhere along the span by the hybrid optimisation and further improved by the flexible one.

Table 3 compares the total aerodynamic drag performance (the objective function J) resulting from rigid, hybrid and flexible optimisations for the Medium parameterisation and two levels of wing flexibility: FS2 (see [3]) and the increased one considered in the present paper. These results highlights that the discrepancy between the different optimisation processes increase with the wing flexibility. First, when the rigid optimal shape is re-evaluated by aero-elastic analysis (purple and dark green columns), the averaged drag increases of $\sim 2\%$ for FS2 but $\sim 5\%$ for the more flexible configuration. The use of aero-elastic simulations in the optimisation process (hybrid optimisation) for the primal analysis enables a drag reduction of $\sim 0.6\%$ for FS2 and of $\sim 0.8\%$ for the present configuration. Finally, when exploiting coupled aero-elastic sensitivities in the optimisation process (fully flexible optimisations) total drag is further reduced by 0.4% for FS2 and by $\sim 1.1\%$ for the present more flexible wing. This clearly underlines that the use of the coupled aero-elastic adjoint is mandatory for highly flexible wing configuration to design the most efficient aerodynamic wing shape.

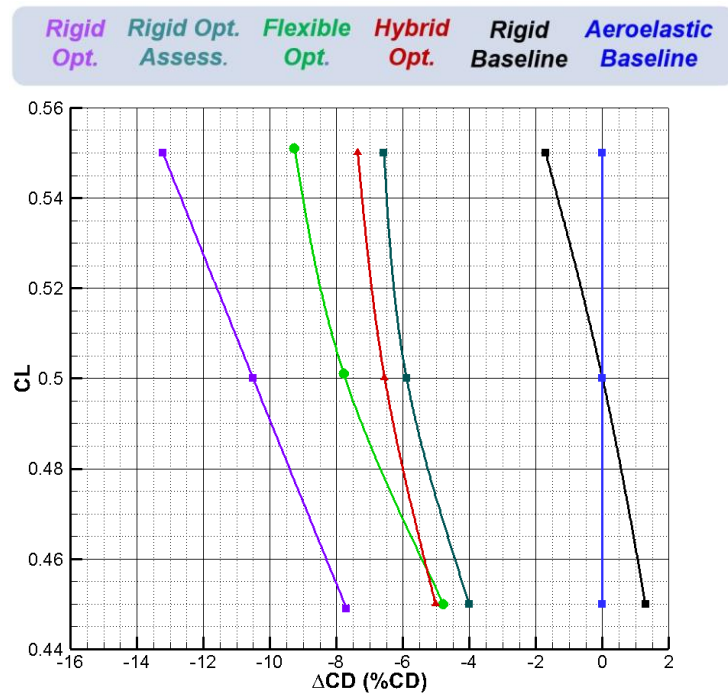


Figure 5: Comparison of aerodynamic drag performance resulting from rigid (violet line), hybrid (red line) and flexible (green line) optimisations of the XRF-1 wing shape using the Medium parametrisation (see

Table 1). Drag results are illustrated in terms of percentage variation with respect to aeroelastic baseline results (blue line). The black line denotes the rigid performance associated to the baseline flight shape at $CL=0.5$. The dark green curve corresponds to the aeroelastic performance re-evaluation of the rigidly optimised flight shape (violet line).

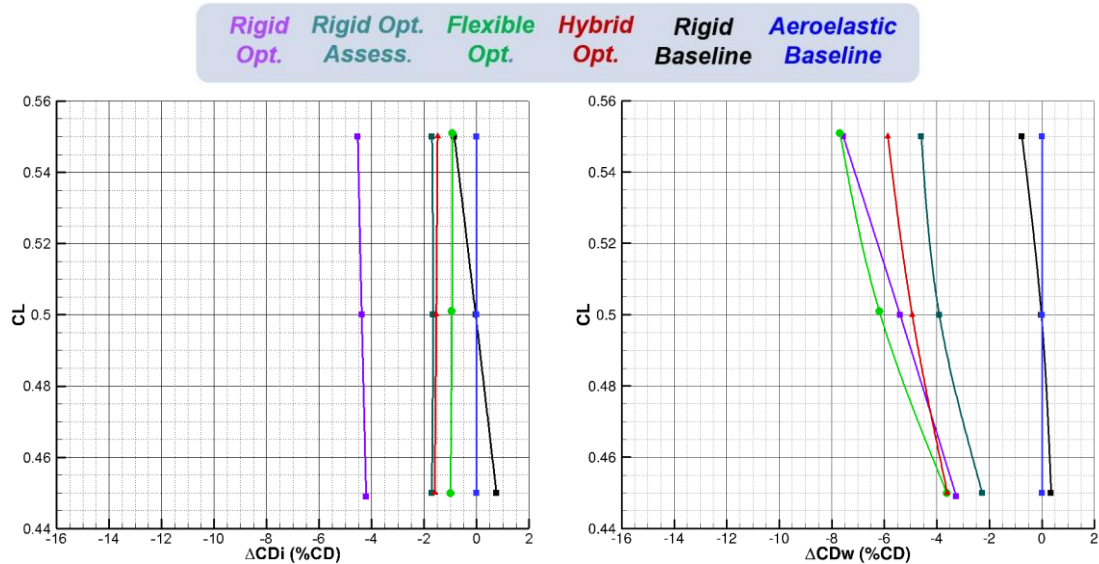


Figure 6: Comparison of aerodynamic drag performance resulting from rigid (violet line), hybrid (red line) and flexible (green line) optimisations of the XRF-1 wing shape using the Medium parametrisation (see

Table 1). Left: Induced drag component. Right: wave drag component. Percentage drag variations are referred to the value of the considered drag component resulting from aeroelastic baseline results.

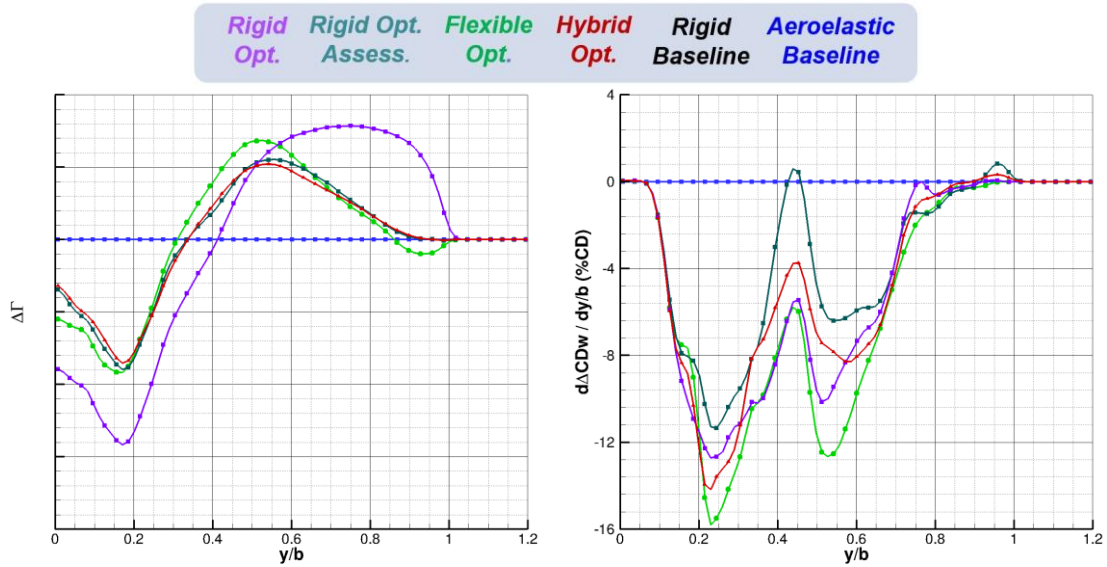


Figure 7: Spanwise analysis of rigid (violet line), hybrid (red line) and flexible (green line) optimisation results using the Medium parametrisation (see

Table 1). All results refer to $CL=0.5$. Left: spanwise load variation with respect to the load distribution of the aeroelastic baseline. Right: spanwise wave drag percentage variation with respect to the corresponding wave drag distribution evaluated for aeroelastic results.

Table 3: Comparison of total aerodynamic drag performance (J function) in terms of percentage variation resulting from rigid (violet line), hybrid (red line) and flexible (green line) optimisations of the XRF-1 for two levels of wing flexibility (FS2 and the increased one studied in the present paper) .

	Aeroelastic Baseline	Rigid Opt.	Rigid Opt. Assess.	Hybrid Opt.	Flexible Opt.
FS2	0.00%	-6.73%	-4.90%	-5.52%	-5.92%
		+1.8%	-0.6%	-0.4%	
		+5.1%	-0.8%	-1.1%	
Increased wing flexibility	0.00%	-10.72%	-5.60%	-6.41%	-7.47%

3.2 Refined parametrisation results

Flexible optimisation results obtained using the Fine parametrisation are illustrated in Figure 8 and compared to Medium ones. Thanks to the finer camber control, the final performance is further improved for all the three considered flight conditions, with a drag reduction increasing to $\sim 1\%$ at $CL=0.45$ and to $\sim 3\%$ at $CL=0.55$, close to the performance obtained by pure aerodynamic optimisation in Figure 5. The associated induced and wave drag breakdown is analysed in Figure 9: the induced drag is only slightly improved by the parametrisation refinement; this is partially expected since both parametrisations share the same twist law and design variables. However, the comparison of the spanwise load distributions at $CL=0.5$ in Figure 10 (left panel), shows that maximum positive and minimum negative load variations are less pronounced in the Fine case. For the wave drag, see Figure 10 (right panel), although the spanwise reduction pattern is still similar, results obtained using the Fine parametrisation show a wider inboard peak and a slightly reduced outboard one. This could be associated with the reduction of the aerodynamic load around the same spanwise position $y/b \approx 0.55$, which also contributes to mitigate compressibility effects.

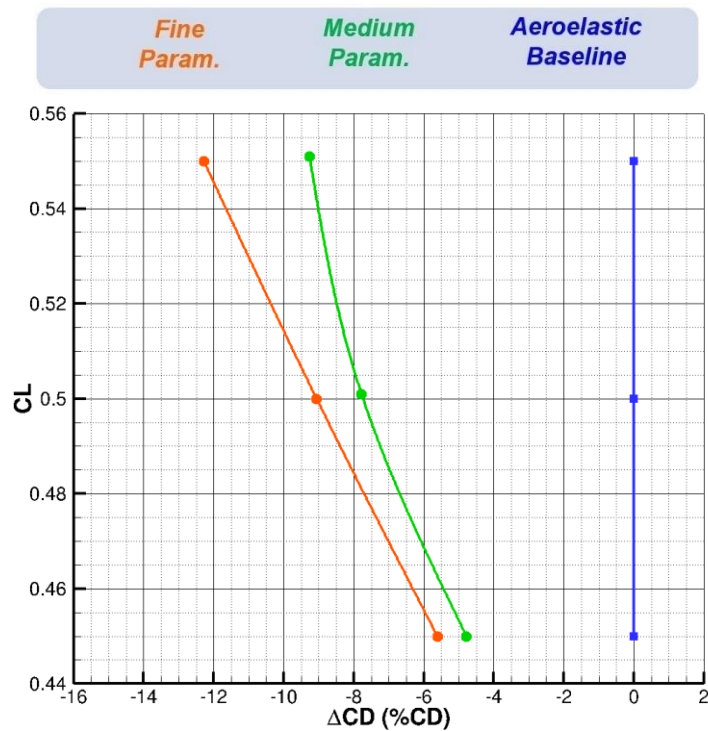


Figure 8: Comparison of aerodynamic drag performance resulting from flexible optimisation of the XRF-1 wing shape using the Medium (green line) and Fine (orange line) parametrisations; see also

Table 1. Drag results are illustrated in terms of percentage variation with respect to aeroelastic baseline results (blue line).

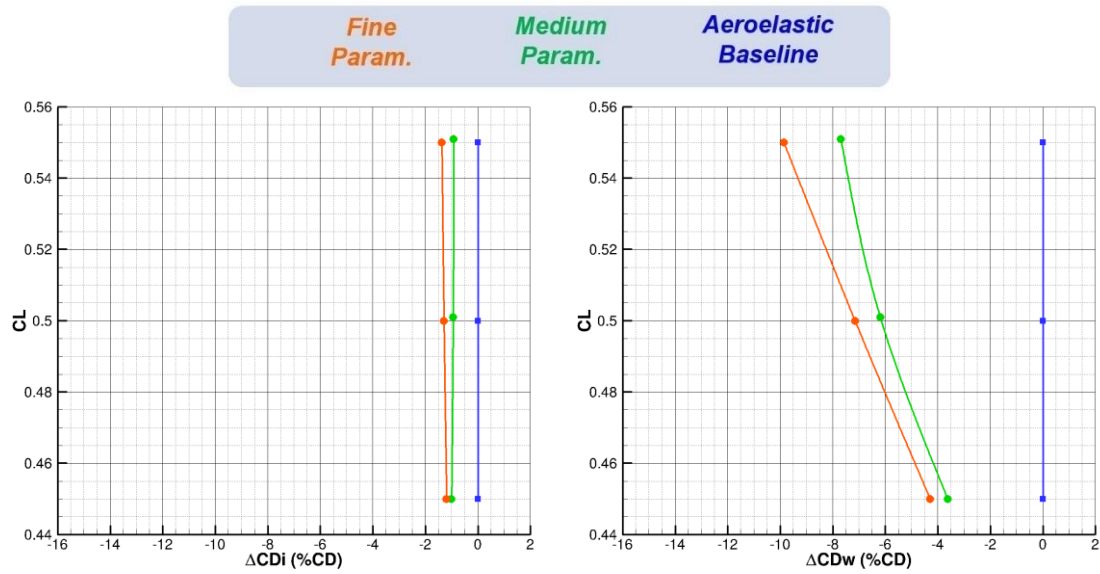


Figure 9: Comparison of aerodynamic drag performance resulting from flexible optimisation of the XRF-1 wing shape using the Medium (green line) and Fine (orange line) parametrisations; see also

Table 1. Left: Induced drag component. Right: wave drag component. Percentage drag variations are referred to the value of the considered drag component resulting from aeroelastic baseline results.




Figure 10: Spanwise analysis of flexible optimisation of the XRF-1 wing shape using the Medium (green line) and Fine (orange line) parameterisations; see also

Table 1. All results are referred to the $C_L=0.5$ condition. Left: spanwise load variation with respect to the load distribution of the aeroelastic baseline. Right: spanwise wave drag percentage variation with respect to the corresponding wave drag distribution evaluated for the aeroelastic baseline.

4 CONCLUSIONS

An effective demonstration of new capabilities in terms of multi-point aeroelastic

optimisation of industrial-aircraft wing-shape has been presented, with focus on the benefit of engaging flexibility in the optimisation approach. As expected, discrepancies between rigid and flexible optimisation scenarios are magnified when increasing wing structure flexibility as considered in the present study through a scaling factor applied to the wing stiffness matrix. Results show that the rigid approach, i.e., the pure aerodynamic shape optimisation, greatly overestimates the achievable performance gain. Therefore, a larger number of disciplinary iterations (between aerodynamic and structure) in the MDO loop can be necessary despite the reduced cost and complexity of the required computational tools. This drawback can be circumvented by engaging flexibility at least in the objective function evaluation through aeroelastic analyses of the considered flight points, while still freezing the structure flexibility in the gradient computation. Such hybrid approach shows an interesting compromise to cope with legacy tools while achieving a reliable and appreciable gain. Moreover, the far-field drag analyses of the obtained results suggest that the hybrid approach is satisfactory when seeking induced drag reduction. However, the same analysis also shows that only flexible gradient computations enable the largest reduction of the wave drag component, similarly to what observed for the total drag. These results confirm that coupled adjoint capabilities are a key-enabler for the design of future configurations featuring increased wing flexibility and deflection, such as those characterised by high aspect-ratio wing.

5 ACKNOWLEDGEMENTS

We would like to thank Airbus Operations SAS for providing us with the XRF-1 test case and the PADGE modeller, and the DLR Institutes of Aerodynamics and Flow Technology and of Aeroelasticity to provide us with the wing structure model. We would also thank them for fruitful discussions on the analysis of the results. The MADELEINE project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 769025.

REFERENCES

- [1] Méheut, M. on behalf of the MADELEINE consortium. Multidisciplinary Adjoint-based Optimizations in the MADELEINE Project: Overview and Main Results, *AIAA AVIATION 2021 FORUM* (2021) AIAA 2021-3052.
- [2] Jameson, A. Aerodynamic design via control theory, *Journal of Scientific Computing* (1988) **3**:233-260.
- [3] Carini M., Blondeau, C., Fabbiane, N., Méheut, M., Abu-Zurayk, M., Feldwisch, J., Ilic, C. and Merle, A., Towards industrial aero-structural aircraft optimization via coupled-adjoint derivatives, *AIAA AVIATION 2021 FORUM* (2021) AIAA 2021-3074.
- [4] Olivanti R. and Brézillon J., On the Benefits of Engaging Coupled-Adjoint to Perform High-Fidelity Multipoint Aircraft Shape Optimization, *AIAA AVIATION 2021 FORUM* (2021) AIAA 2021-3072.
- [5] Martins, J., Alonso, J., Reuther, J., and Reuther, J., A Coupled-Adjoint Sensitivity Analysis Method for High-Fidelity Aero-Structural Design, *Optimization and Engineering*, (2005) **6**:33–62.

- [6] Achard, T., Blondeau, C., and Ohayon, R., High-Fidelity Aerostructural Gradient Computation Techniques with Application to a Realistic Wing Sizing, *AIAA Journal* (2018) **56**:4487-4499.
- [7] Kenway, G. K. W., and Martins, J. R. R. A., Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration, *Journal of Aircraft* (2014) **51**:144-160.
- [8] Liem, R. P., Kenway, G. K. W., and Martins, J. R. R. A., Multimission Aircraft Fuel-Burn Minimization via Multipoint Aerostructural Optimization, *AIAA Journal* (2015) **53**:104-122.
- [9] Brooks, T. R., Martins, J. R. R. A., and Kennedy, G. J., Aerostructural Tradeoffs for Tow-Steered Composite Wings, *Journal of Aircraft* (2020) **57**:787-799.
- [10] Merle, A., Ilic, C., Abu-Zurayk, M., Häßy, J., Becker, R.-G., Schulze, M., and Klimmek, T., High-Fidelity Adjoint-based Aircraft Shape Optimization with Aeroelastic Trimming and Engine Coupling, EUROGEN 2019 (2019), Guimaraes, Portugal.
- [11] Abu-Zurayk, M., Merle, A., Ilic, C., Keye, S., Goertz, S., Schulze, M., Klimmek, T., Kaiser, C., Quero, D., Häßy, J., Becker, R., Fröhler, B., and Hartmann, J., Sensitivity based Multifidelity Multidisciplinary Optimization of a Powered Aircraft Subject to a Comprehensive Set of Loads,” *AIAA AVIATION 2020 FORUM* (2020) AIAA 2020- 3168.
- [12] Destarac, D., and van der Vooren, J., Drag/thrust analysis of jet-propelled transonic transport aircraft; Definition of physical drag components, *Aerospace Science and Technology* (2004) **8**:545-556.
- [13] Jadoui, M., Blondeau, C., Martin, E., Renac, F. and Roux, F., Comparative Study of Inner-Outer Krylov Solvers for Linear Systems in Structured and High-Order Unstructured CFD Problems, in WCCM-ECCOMAS2020 (2020).
https://www.scipedia.com/public/Jadoui_et_al_2021a
- [14] Cambier, L., Heib, S., and Plot, S., The Onera elsA CFD Software: Input from Research and Feedback from Industry, *Mechanics and Industry* (2013) **14**:159-174.
doi:10.1051/meca/2013056.
- [15] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., pyOpt: A Python-Based Object Oriented Framework for Nonlinear Constrained Optimization, *Structures and Multidisciplinary Optimization* (2012), **45**:101-118.