Effect of gradient approximations on aero-structural gradient-based wing optimization

Mohammad Abu-Zurayk*
German Aereospace Center (DLR)
Institute of aerodynamics and Flow Technologies, Braunschweig,Germany
Email: mohammad.abu-zurayk@dlr.de

Časlav Ilić
German Aereospace Center (DLR)
Institute of aerodynamics and Flow Technologies, Braunschweig,Germany
Email: caslav.ilic@dlr.de

Andreas Schuster
German Aereospace Center (DLR)
Institute of Composite Structures and Adaptive Systems, Braunschweig, Germany
Email: andreas.schuster@dlr.de

René Liepelt
German Aereospace Center (DLR)
Institute of Aeroelasticity, Göttingen, Germany

Summary

This paper presents several studies in the context of gradient-based aerostructural optimization. The studies are motivated by the computational limitations and constraints industry has when attempting to run such optimization. The aero-structure wing optimization problem engages two parts; one that improves the aerodynamic performance and the other reduces the mass under the critical load cases. To setup a realistic test case, the aerodynamic critical load cases required for sizing the wing of the XRF1 configuration are identified, with the help of high-fidelity approaches and ROMs, and are afterwards employed to size the wing. Then the magnitudes of the disciplinary and the cross disciplinary gradients are investigated for the sized wing to check which of them drive the aerostructural optimization under the given parameterization. It was found that these gradients might differ in two to three orders of magnitude. At the end the configuration is aero-structurally optimized with different levels of gradient accuracy.

.***Keywords****: MDO, Aero-structural Optimization, Adjoint Approach, Gradient Approximation*

1 Introduction

Driven by the search for advanced designs, that consume less fuel and produce less emission, the aircraft industry is looking into aircraft optimization techniques using numerical methods. Multidisciplinary design optimization (MDO) techniques incorporate more than one discipline in the optimization, which yields realistic designs that fulfil the disciplinary constraints and reduce the development risks. However, the complexity and the computational cost of the problem in MDO are significantly higher than that of single disciplinary optimizations. Therefore, optimization algorithms that drive high-fidelity MDO need to be efficient, which motivates the use of gradient-based algorithms for such problems, provided that the computation of the required gradients itself is efficient.

Since gradient-based algorithms can efficiently find the nearest *local* optimum, we consider them more suitable for the final design stages where a generally good design is already available and is provided as a starting point to these algorithms in order to fine-tune it. Regarding the gradient-based algorithms as a fine-tuning tool is, however, not necessarily a common sense. Several studies used these algorithms to optimize the planform of the aircraft's wing.

The most visited MDO problem in aeronautical research is the aero-structural optimization problem. The importance of this problem lies in including two main disciplines in the aircraft design that directly affect the fuel consumption during a mission. In academic circles, this problem was tackled intensively by Martins and his MDO group[1] where they showed several high-fidelity aero-structural optimizations using gradient-based algorithms on the CRM[2] research model, among other configurations. In governmental aeronautics research labs, this problem was investigated by Ronzheimer[3] and Ilic[4] where it was applied to the Dornier728 and XRF1 research model configurations respectively, using gradient-free algorithms and hence for little number (5-10) of design parameters. Moreover, at ONERA, the XRF1 configuration was optimized using a bi-level optimization technique[5]. In industry, Piperni et al[6] showed a complex application of aero-structural optimization at BOMBARDIER where a business jet was optimized. The level of fidelity, however, was compromised in some aspects to suit their numerical environment limitations and constraints.

The aerostructural optimization problem contains two parts. The first aims at improving the aerodynamic performance for the wing's current flight shape. This can be done for the aircraft's cruise point only (in single point optimizations) or additionally for extra other points (in multipoint optimizations) simultaneously. The second part aims at reducing the mass of the structure while guaranteeing that the structure holds under the current critical loads, whether they were aerodynamic loads or other types of loads.

In aerostructural gradient-based optimizations, the full set of gradients are expected by the optimizer, including disciplinary as well as cross disciplinary gradients. Our definition, whether a gradient is disciplinary or cross-disciplinary, considers shape design variables to be aerodynamic parameters and the structural thicknesses to be structure design parameters. Then, an example on disciplinary gradients would be the sensitivity of wing’s mass to the change in structural material thicknesses or the gradient of aerodynamic drag with respect to shape design parameters. On the other hand, following our definition, the gradient of structural failure criteria with respect to shape design parameters is a cross disciplinary gradient.

Computing the full set of gradients is not practical when the number of design parameters and the number of the constraints are high. The Adjoint approach is appropriate when the number of design variables is high and the number of constraints and objectives is little, since the adjoint equation needs to be solved once for each constraint or objective. The finite differences approach on the other hand is useful if the number of objectives and constraints is high but the number of design variables is little since it has to be solved once per design variable.

For a realistic aircraft *fine-tuning* optimization problem, the number of shape and structure design parameters is in the order of thousands, the number of objectives and constraints for the aerodynamic performance part of the problem is small whereas it is in the order of tens of thousands for the structure sizing part of the problem, mainly due to the failure criteria constraints of each finite element in the structure model and the high number of the critical load cases. Computing the full set of gradients for such a problem is evidently not practical, since the structure sizing part of the problem contains a huge number of constraints which makes the adjoint approach useless. Martins and his group avoided this problem[7] by aggregating the thousands of constraints using the Kreisselmeier-Steinhauser technique into several constraints where the adjoint approach becomes feasible to use for both parts of the aero-structural problem. This, however, is a conservative approach that can reduce the improvement one expects in wing's structural mass.

Additionally, going for the full aero-structure adjoint option requires access to the source code of the solvers in order to be able to implement the approach, and this is not always the case in industry, which relies heavily on commercial software that is not necessarily completely differentiated.

The aim of this work is to explore the effect of approximating the gradients in aero-structural optimizations by comparing optimizations driven by approximate gradients, to others driven by the full set of gradients, once for complete set of constraints and once for aggregated set of constraints. The optimizations are performed to a realistic aircraft configuration (the XRF1 model by AIRBUS). This study will not only help understanding the cross disciplinary effects at the fine-tuning level, but also shows the impact of approximation on the computational cost.

While this gradient approximation study is performed on the wing-body XRF1 configuration, the lessons learned are applied to optimize, at the end of the work, the full XRF1 configuration trimmed with tail and engine.

2 MDO Problem Definition

At the German Aerospace Center (DLR), within the MDO work packages under the cross institutes’ projects Digital-X (2012-2016) and VicToria (2016-2019), two main optimization directions are taking place for developing novel and reusable MDO processes. The first direction gathers three levels of fidelity, starting from preliminary design, passing by a dynamic level for the prediction of the critical load cases and ending with detailed high-fidelity aerostructural level. Due to the high complexity of this direction, it was decided there to employ a limited number of global design parameters, that control the planform of the wing in addition to some main section twists. The second direction concentrates on sharpening and exploring the gradient-based approach for high-fidelity MDO and then gradually adding other levels of fidelity. Whilst the first direction aims at finding a global optimum for the global design parameters and taking constraints from all levels of aircraft design, the second direction focuses at refining the optimum design produced by the global optimization and uses hundreds of local wing parameters on the aerodynamic and the structure sides to optimize the resulting aircraft using a gradient-based approach. This work describes the studies performed in the second stream that adopts the high-fidelity gradient based MDO.

The model employed in this study is the research configuration by Airbus (XRF1); which is a transonic wide body transport configuration. The twin engine model has design Mach number of 0.83, a design lift coefficient of 0.5 and a range of 8000 nm. Fig. 1 shows a general view of the complete model.

Two CFD models and one CSM model were used in this work. For the gradient approximation studies a wing-body configuration was employed where a structured grid with 1 million points (half model) was generated. For the final optimization however, an unstructured grid of the full configuration with 4 million points (half model) was generated, Fig. 2. Both grids were coupled to a CSM model for the wing-box only. The model has approximately 30000 degrees of freedom, Fig. 3. As shown the wing box is made of 3 spars, 45 ribs, upper skin and lower skin. Aluminum structure is assumed throughout this study.

Fig. 1: the XRF1 model

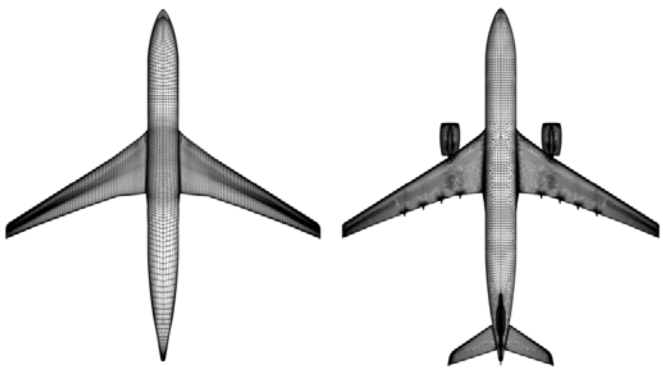

Fig. 2 the CFD models employed for the gradients study (left) and the final optimization (right)

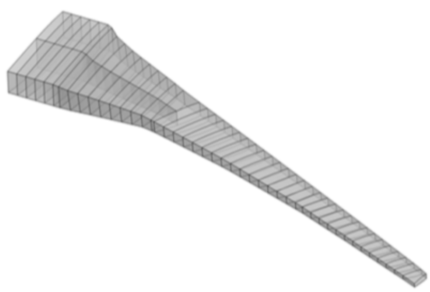

Fig. 3 wing box

To couple aerodynamics with structure, a loose coupling is used between the RANS unstructured solver DLRTAU and the commercial structure solver ANSYS Mechanical8. It is worth mentioning here that both tools are used in industry. In this work, RANS equations with the 1-equation Spalart-Allmaras turbulence model are used to compute the flow, and linear elasticity is used to compute the structure deformations. S_BOT, DLR's structure sizing robot, is employed afterwards to compute the structure failure criteria or to completely size the structure. To interpolate the pressure distribution from the CFD model to the CSM model, a linear interpolation algorithm or a nearest neighbor algorithm is used[9]. To interpolate the deformations the other way around, a radial-basis function approach is employed.

Free-form deformation (FFD)[10] nodes control the shape of the CFD mesh directly where the sectional wing profiles can be perturbed for a fixed wing planform. This perturbation is interpolated to the CSM model as well, which is additionally altered by changing the structural thicknesses. Sets of finite elements are grouped into the so called optimization regions, or structural design variables, where this grouping follows the manufacturability constraints. One material thickness parameter is associated with each optimization region.

The objective function in the optimizations is inspired by the Breguet range equation for continuous cruise climb. It is made of the structure mass and the aerodynamic lift and drag coefficients as given in equation (1)

$$objective = \left(\frac{m}{m_{ref}}\right)\left(\frac{C_D}{C_L}\right) \tag{1}$$

Where m and m_{ref} are the current and baseline structural wing masses, and C_D and C_L are the aerodynamic coefficients of drag and lift, respectively. The baseline configuration is sized with the sizing tool S_BOT before the optimization starts using a fully stressed design algorithm. This allows us to judge fairly how beneficial the gradient-based optimization is since the starting point represents the best point reachable without optimization algorithms.

A pyOpt environment connected to a feasible sequential quadratic programming (FSQP) algorithm is employed to drive the optimizations. This algorithm computes the gradients once per optimization cycle where an optimization cycle is made of one gradient computation and several design iterations in a non-gradient line search.

Each design iteration includes the computation of the performance points and the computation of the critical load cases that size the wing's structure. The performance points are usually set as a target of the design and are consequently known in advance. The critical load cases are, however, not known and theoretically need to be identified for each new design. In the following section, the identification of these loads is discussed.

3 Identification and Computation of Sizing Loads

The loads process in aircraft design is the process responsible for identifying and computing the critical load cases that an aircraft might experience during the mission. Hundreds of thousands of load cases are computed within this process where hundreds of them might be critical for some parts of the aircrafts structure. These loads are usually computed in a conservative manner (overestimated) which increases the safety factor in aircraft design. The loads are applied to the structure of the aircraft in order to size it (adapt the structure material thickness) by making sure that it can withstand these loads without breaking.

Loads are mainly grouped in equilibrium manoeuvers, dynamic manoeuvers, ground manoeuvers and gust/turbulence encounters[11]. The first group concern steady manoeuvers, the second group involve the dynamic response to transient inputs, the ground manoeuver concern all load conditions on the ground such as taxiing and landing, and the final group involve responding to discrete gusts. According to "Luftfahrttechniches Handbuch"[12] equilibrium and dynamic manoeuvers size around 38% of the wing, gusts size around 36% of the wing whereas ground loads size 28% of the wing.

Within DLR, a process that includes load case definition, parametric model generation, load analysis and structural sizing[13] was developed. Following industry approaches, the loads in this automated process are computed over two levels of fidelity. The first is a conceptual design level that produces a pre-sized FEM. The second level uses higher-fidelity aerodynamics loads by employing the doublet-lattice method implemented in NASTRAN[14] to apply them on the FEM model produced earlier. After that the FEM model is sized. The tools can also identify the critical load cases and forward their definition to high-fidelity load computation tools like CFD (whether viscous or inviscid) to be computed. In the meantime these tools are being further developed to accept correction of loads from high-fidelity flow solvers during the sizing process.

In this work, it was decided to compute only equilibrium manoeuvers using the high-fidelity tools at DLR, and neglect the rest of the loads sizing the wing. The reason of this decision is to understand the steady flight envelope and to test the ability of the CFD tools to go beyond cruise points, in order to be able, in the future, to estimate the boundaries that the tools can deal with and to provide the lower fidelity load computing tools with high-fidelity corrections. It is worth mentioning here that a parallel study at DLR is performing the same type of optimization using the NASTRAN loads process from the aeroelastic institute[15], where more type of loads conditions are considered.

To compute the high-fidelity loads, the configuration's CL polars (**Fig. 4**) were computed and provided to the aeroelastic institute which accordingly constructed the V-n envelopes to define the load cases including different mass conditions. A mass case defines payload, fuel load, and center of gravity. Five mass cases were provided, for each of which a Mach-altitude envelope was computed using the high-fidelity (RANS - Linear Elasticity) loads with load factors of -1g and 2.5g. The mass cases include the maximum take-off weight, once for maximum payload and once for maximum fuel load, a mass case with maximum payload and zero fuel, a mass case with maximum fuel and zero payload, and the mass case of operating empty weight.

Since this process can compute the loads only at discrete points for pre-given intervals in Mach number and altitude, reduced order models (ROMs)[16] are afterwards generated and used to help identifying if there are other sizing loads that were not computed between these intervals. After getting the new ROMs candidates they are computed via high-fidelity aeroelastic coupling and it is checked if these loads are sizing loads or not. The process here (ROMs, CFD-CSM coupling, Sizing) which is described in a paper[17] also keeps running until convergence, see Fig. 5.

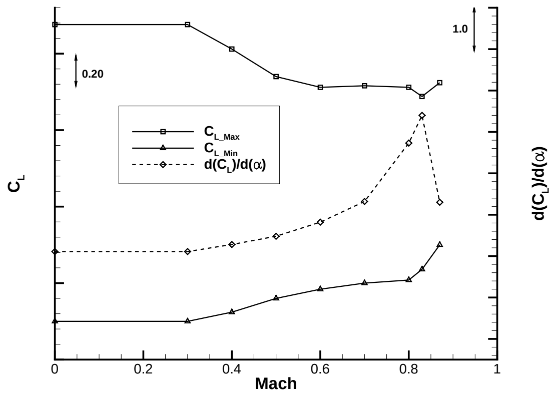

Fig. 4 lift polars for the configuration

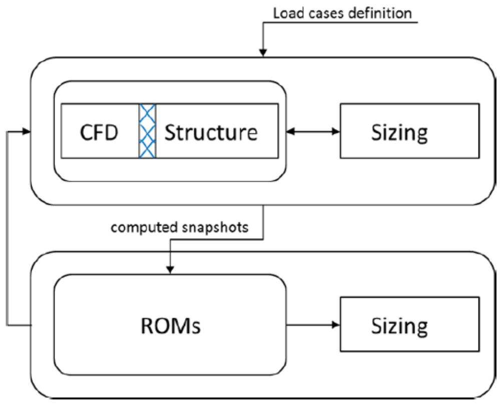

Fig. 5 using ROMs for sizing load identification

It is worth mentioning here that the ROMs are very beneficial here since the computation of the snapshots is required anyway for the loads selection process, and no evaluations need to be performed specifically and only for the purpose of building the ROMs.

Following this costly but necessary process, the sizing loads were identified for the baseline configuration. Since the gradient-based optimization expects always the same number of constraints, the same number of load cases should be considered throughout the optimization. This assumption can be excused by the fact that this is a fine-tuning optimization which is not expected to change the geometry drastically and hence the same load cases are expected to be *the* sizing load cases. If we were interested in gradient-based optimizations that drastically change the geometry (eg. variable planform), the set of sizing load cases would have been necessary to identify at each optimization cycle and the optimizer would need to be recalled for one cycle each time. This means that information like Hessian that is transferred from an optimization cycle to another would need to be neglected.

For this configuration, it was found that 11 aerodynamic load cases size the structure model, 7 of them bring the mass to around 95% of the full mass without changing the thickness distribution. These 7 will be considered in the final optimization.

4 Aerostructural Gradients Study

This section is dedicated to studying the magnitudes of the gradients that are required for an aerostructural optimization and to investigating the gradient components that can be neglected during the optimization, especially if that saves considerable amount of computational power but have little effect on the optimization results. Two parts are presented here; the first part compares the magnitudes of the different gradients and the computational power required to compute them. The second part presents and compares several optimizations, which are performed at different levels of gradient accuracies. The compromises on gradients accuracies are driven by the constraints and limitations industry has because of the tools employed there or the computational power available. The presentation of the first point in this section (gradients' magnitudes) is necessary since gradient-based algorithms search for nearest local optimums and hence presenting the second point only (the optimizations) can be misleading, since it is possible to have accurate gradients leading to a worse optimum (less improvement). Of course analysing the optimization results should be accompanied with examining the KKT conditions to understand which optimization came closer to an optimum.

4.1. Gradients' Magnitudes

The test case considered for this investigation is the XRF1 wing-body configuration. The structure model is sized before the study started using the sizing aerodynamic loads that were identified and computed in the last section. As mentioned earlier, the wing is parameterized with thousands of design parameters. Since this study is prohibitively expensive for such a large design space, a representative set of design parameters is considered that includes 24 shape design parameters and 20 structural material thicknesses. The shape design parameters engaged in this work are FFD parameters that are applied directly on the CFD computational grid. They deform locally the wing's airfoils at different spanwise stations. On the structure side, 20 optimization regions, taken from different positions on the wing, are considered.

Following the previously mentioned objective, the gradients of interest here are the drag and pitching moment gradients at constant lift, the gradients of the structural mass and the structural failure criteria constraints.

Fig. 6 presents the gradients of the drag coefficient, at constant C_L, with respect to all design parameters. The first 24 design parameters are the FFD parameters and the rest 20 are the structure material thicknesses. The figure shows the gradients for two computations; one that includes the elasticity effect on the gradients and one that neglects them (in blue). Two main observations can be concluded here, the first is that the gradients of drag with respect to the shape design parameters are generally around one order higher than those with respect to the structure material thicknesses. Secondly, not including the elasticity effects in the gradient computation (computing aerodynamic gradients at the flight shape), results generally in the same trend as in that where the elasticity effects are included, at least for the gradients with high magnitude, that, actually, drive the optimization.

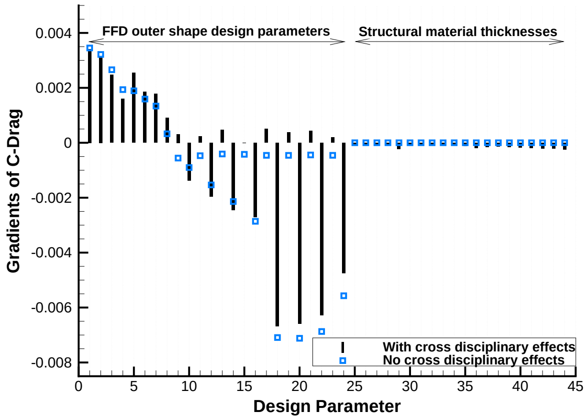

Fig. 6 Gradients of drag coefficient

Fig. 7 presents the gradients of pitching moment coefficient. Unlike C_D, the gradients here are almost of the same order, and the effect of including elasticity in the gradient computation is evident. The reason is that C_{MY} is more sensitive to the wing's elasticity than C_D. If the disciplinary C_{MY} gradients are used in an optimization, it is expected that the pitching moment constraint gets violated as soon as the direction of reducing the objective is not in favor of maintaining the constraint (if it becomes active), unless the optimizer itself guarantees feasibility as in the feasible sequential programming algorithm (FSQP) which is made to be used when the gradients are not fully accurate.

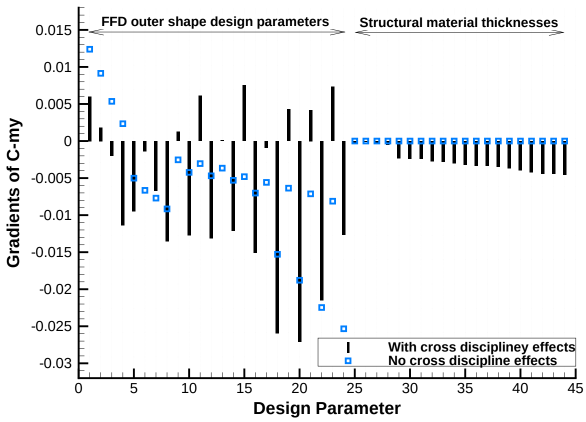

Fig. 7 Gradients of pitching moment coefficient

The sensitivity of wing's mass to the full set of design parameters is shown in Fig. 8. Mass depends linearly on the thicknesses for constant optimization regions area. The figure shows that the mass gradients with respect to structural thicknesses are 2-3 orders of magnitude higher than those with respect to shape design parameters.

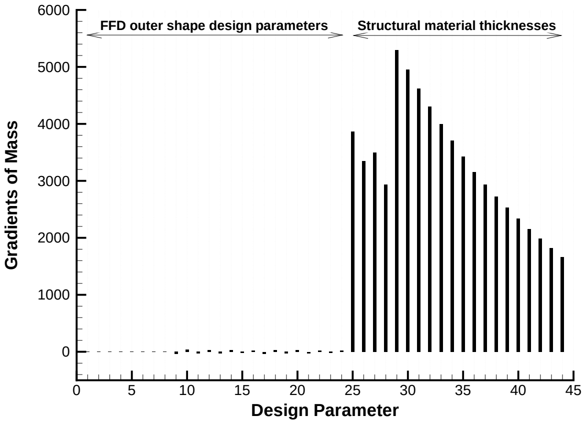

Fig. 8 Gradients of structural mass

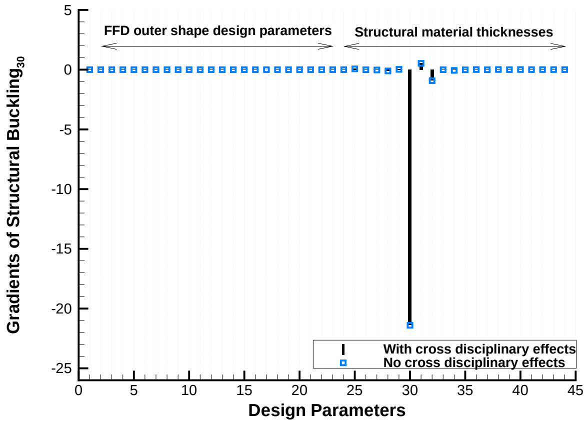

Fig. 10 Gradients of structure buckling

The last group of gradients to be checked are the gradients of the structure failure criteria, which are strength and buckling in this study. Unlike the previous set of gradients, this gradient is a matrix since the perturbation of one optimization region has effect on the failure criteria of all optimization regions. If only the structure problem is considered, the diagonal of this matrix represents the effect of perturbing an optimization region on its own failure criteria. Fig. 9 and Fig. 10 present the gradients of strength and buckling for one optimization region, respectively. The same behavior was observed when the failure criteria were plotted for all other optimization regions, and all load cases.

As depicted in the figures, the gradients of strength and buckling with respect to structural material thicknesses are 2-3 orders of magnitude higher than those with respect to shape design parameters. Additionally, the gradients for frozen loads are very similar to those for updated loads. Here we see a different behavior than that governing the aerodynamic performance gradients. This behavior allows us to save a lot of computational cost if we chose to run for frozen loads when computing the gradients.

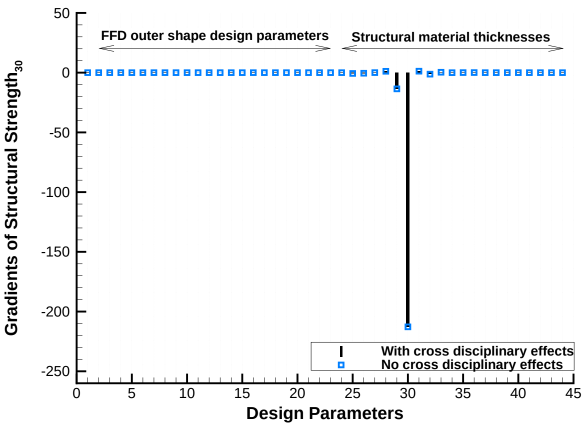

Fig. 9 Gradients of structural strength

In conclusion, it is remarked that the gradients of aerodynamics performance are more influenced by the aeroelastic effects than the gradients associated with the structure sizing process. This observation, which should not be generalized to all types of parameterizations or sizing techniques, suggests the use of the coupled aero-elastic adjoint approach for aerodynamic performance part of the problem, and the use of the parallelized finite differences for the structure sizing part of the problem to compute the required gradients. To test this observation, several optimizations are started, including this option.

4.2. Optimizing at different levels of accuracies

In the second part of this section several optimizations are performed with varying level of gradients accuracy, where the accuracy here is associated with including the cross disciplinary effects. The first optimization that serves as a reference for all following ones is an optimization with full set of gradients and full set of structural constraints. This optimization provides the highest possible coupling level between aerodynamics and structure. The second optimization has the full set of gradients, but aggregated set of structural constraints. This optimization will give us a feeling of what we get beside the ability to fully engage the efficient adjoint approach for the aero performance side as well as for the structure sizing side.

The third optimization tests the conclusions extracted from the gradients' magnitude study, and runs with the full set of gradients on the aerodynamic performance side where theoretically the coupled aeroelastic adjoint can compute the gradients, and employs the finite differences approach to compute the gradients on the structure sizing side for frozen loads. In other words, this optimization uses the full set of gradients on the performance side and only the disciplinary gradients on the structure sizing side, but still considers all structural constraints.

The fourth optimization is driven by only disciplinary gradients on both the performance and the structure sizing sides, while it considers all structural constraints. This means the aerodynamic adjoint runs on the flight shape to get the required gradients on the performance side, and as in the previous optimization scenario, finite differences are employed to get the gradients on the structure sizing side.

Finally, the fifth optimization runs sequential disciplinary aerodynamic and structure optimizations where the other discipline is kept frozen until convergence. This should provide the lowest level of coupling possible.

Since the computation of these gradients can be prohibitive for a high number of design variables, only the 44 design variables, shown in the gradient's study, are used for the optimizations, where a careful use of finite differences and the adjoint approach are employed to compute the gradients in these studies.

First results of this study can be shown in Fig. 11, where this study did not completely converge and is still running. The conclusion will be drawn when the all optimizations converge.

The fourth optimization and the fifth optimization (which is not shown in the plot yet) have the same quality of gradients, where in the fourth they are given at once to optimizer and in the fifth over steps, once aerodynamic once structure. The reason we test both options is to check what happens when we decrease the load on the optimizer as in the fifth so that it has to deal each step with less number of design variables or constraints. Here using different optimization algorithms that suit the problem (if more parameters or more constraints) is an advantage which's effect should not be underestimated.

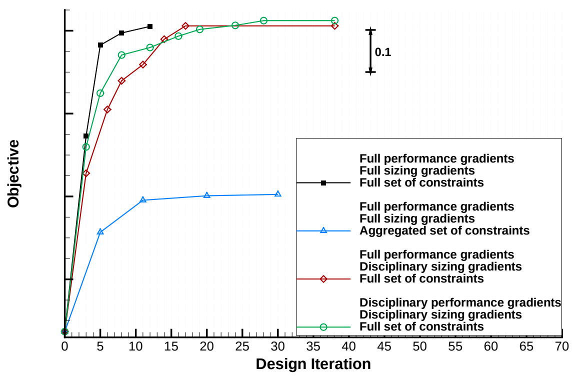

Fig. 11 optimizations convergences

The final paper will include a thorough discussion on the optimizations results are their final computational costs.

5 Aero-Structural Optimization of Full Configuration

The full XRF1 configuration shown in Fig. 2 will be aero-structurally optimized for its trimmed state. Since the aircraft needs to be trimmed, the gradients need to be corrected to include the effect of trimming. Especially, that the trim parameters will not be given as design variables to the optimizer. The final paper will include the optimization results.

6 References

[1] G. Kenway, G. Kennedy, J. Martins, "A scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations", 2012

[2] National Aeronautics and Space Administration (NASA), “Common Research Model,” [online web site], http://commonresearchmodel.larc.nasa.gov/, 2015

[3] J. Brezillon, M. Abu-Zurayk, A.Ronzheimer, D. Haar, W. Krüger, M. Lummer, "Development and application of multi-disciplinary optimization capabilities based on high-fidelity methods" AIAA Multidisciplinary Design Optimization Specialist Conference 2012

[4] C. Ilic, T. Führer, N. Banavara, M. Abu-Zurayk, G. Einarsson, M. Kruse, J. Himisch, D. Seider, R. Becker, "Comparison of Breguet and ODE Evaluation of the Cruise Mission Segment in the Context of High-Fidelity Aircraft MDO". New Results in Numerical and Experimental Fluid Mechanics. Notes on Numerical Fluid Mechanics and Multidisciplinary Design, 132. Springer Verlag. pp. 87-97. ISBN 978-3-319-27278-8. ISSN 1612-2909

[5] Blondeau C., Irisarri F.X., Leroy F.H., Salah El Din I., "A Bi level high fidelity Aero structural integrated design Methodology", RaeS 3rd Aircraft Structural Design Conference, October 2012

[6] P. Piperni, M. Abdo, F. Kafyeke, "Preliminary Aerostructural Optimization of a Large Business Jet " Journal of Aircraft , Vo. 44, Nr. 5, 2007

[7] N. Poon, J. Martins, "Adaptive Constraint Aggregation for Structural Optimization Using Adjoint Sensitivities", Canadian Aeronautics and Space Institute Annual General Meeting, 2005

[8] http://www.ansys.com/products/Structures/ANSYS-Mechanical-Enterprise

[9] R. Heinrich, “Development and Application of TAU–ANSYS Coupling Procedure”, Results of the Closing Symposium of the MEGADESIGN and MegaOpt Projects, Springer–Verlag, New York, 2009, pp. 151–168.

[10] A. Ronzheimer, Shape parameterization Using Freeform Deformation. In: MEGAFLOW- Numerical Flow Simulation for Aircraft Design, ISBN 3-540- 24383-6, 2005

[11] J. Wright, J. Cooper, “Introduction to Aircraft Aeroelasticity and Loads” AIAA Educational Series.

[12] http://www.lth-online.de

[13] R. Liepelt, V. Handojo, T. Klimmek "Aeroelastic Analysis Modelling Process to Predict the Critical Loads in an MDO Environment" IFASD 2015

[14] http://www.mscsoftware.com/product/msc-nastran

[15] S. Keye, T. Klimmek, M.Abu-Zurayk, M. Schulze and C. Ilic "Aero-Structural Optimization of the NASA Common Research Model" submitted to: AIAAAviation and Aeronautics Forum and Exposition (AVIATION 2017), Jun. 2017, Denver, CO.

[16] T. Franz, R. Zimmermann, S. Görtz, N. Karcher, „Interpolation-based Reduced-order Modeling for Steady Transonic Flows via Manifold Learning," in Special Issue: Reduced order modelling: the road towards real-time simulation of complex physics, International Journal of Computational Fluid Dynamics, Volume 28, Issue 3-4, 2014

[17] M. Verveld, T. Kier, N. Karcher, T. Franz, M. Abu-Zurayk, M. Ripepi and S. Görtz., "Reduced Order Models for Aerodynamic Applications, Loads and MDO", DLRK 2016, 13-15 Sep 2016