CFD OPTIMIZATION VIA SENSITIVITY-BASED SHAPE MORPHING

[1]Carsten Othmer, [2]Evangelos M. Papoutsis-Kiachagias, [3]Konstantinos Haliskos
[1]Volkswagen AG, Germany, [2]National Technical University of Athens, Greece, [3]BETA CAE Systems S.A., Greece

KEYWORDS – CFD, Optimization, Sensitivities, Adjoint method, Morphing

ABSTRACT

Since the availability of professional adjoint CFD solvers, adjoint-based optimization methods are getting increasingly popular for industrial CFD – both for topology and shape optimization. A critical component of an adjoint-based shape optimization process chain is the translation of surface sensitivities into an improved shape. We present a novel implementation of such a translation process: Surface sensitivities, as obtained from the adjoint CFD code, are used to steer the shape morphing process in ANSA, thus giving rise to an efficient, goal-oriented shape optimization method. The mapping of surface sensitivities to morphing control points as realized within this study closes the gap persisting until now for an industrial process chain of sensitivity-based CFD optimization. The mapping technique and its technical realization within ANSA are described and demonstrated using an automotive example: the optimization of aerodynamic drag.

1. INTRODUCTION

Across all industries, optimization processes for fluid dynamic devices are getting increasingly important. Such processes have generally two major ingredients: a suitable parameterization of the geometry to be optimized and an optimization algorithm. Among the plethora of existing optimization algorithms, the industrial working horses for CFD optimization are evolutionary strategies, being applied to either CAD-based or morphing-based geometry parameterizations.

As an alternative to these black-box type optimization methods, Volkswagen investigates the industrial feasibility of optimization methods driven by sensitivities. With several partners, development work has gone into this kind of optimization methods over the last years. The main hurdle – the availability of reliable sensitivities for geometries of industrial complexity – has been overcome by programming an adjoint code into the CFD toolbox OpenFOAM® [1]. At the cost of only one primal and one adjoint run, this code delivers topological and surface sensitivities for a variety of objective functions [2,3,4].

On the basis of this code, the adjoint method is operational already for topological optimization [5,6,7]. Its application to shape optimization, however, still lacks a fundamental step: the translation of the surface sensitivity maps into a sufficiently smooth, ideally manufacturable new shape – the “reshape” step. Three different approaches to this translation of surface sensitivities are under investigation at Volkswagen: (1) direct mesh deformation, i.e. node-wise displacements with appropriate smoothing, (2) mapping to CAD parameters [8], and (3) mapping to morphing parameters. This paper documents the current status of our joint efforts on the third approach.

The morphing capabilities of ANSA have long been used for CFD optimization purposes (e.g. [9,10] and references therein). The new facet that this paper aims to present, is the control of the morphing process by surface sensitivities – more precisely, the mapping of element-wise sensitivities onto morphing control points in order to obtain an optimal shape deformation mode. The mapping procedure and its technical realization within ANSA will be described in detail in Sec. 2, while Sec. 3 demonstrates its application to the aerodynamic optimization of the Volkswagen low emission car L1.

2. THE MAPPING TECHNIQUE

Volkswagen in conjunction with partners has developed an adjoint CFD solver within the CFD toolbox OpenFOAM® [2,3,4]. For a given objective function J, e.g. aerodynamic drag or pressure drop, it allows the computation of the sensitivities $\partial J/\partial \alpha_k$ for each element centre of the surface mesh, with α_k being the outward normal displacement at element k. These sensitivity maps can be used for an efficient shape optimization. However, the sensitivities are rather noisy, and if they were directly translated into a surface mesh displacement, this would result in an unacceptably rough surface. To overcome this problem, it was suggested [11] to use the ANSA morphing capabilities, which, by definition, retain the smoothness of the surface. To that end, a mapping of the element-based sensitivities provided by the adjoint CFD code onto the ANSA morphing control points was required.

As of today, there is, to our knowledge, no professional tool available for an efficient translation of element- or node-based CFD sensitivities into a new, smooth shape for geometries of industrial complexity. As will be shown in the following, the required mapping boils down to applying the chain rule, and all the "ingredients" were already available inside ANSA. They just had to be combined and made accessible to the user [11].

Required functionality

Let $\boldsymbol{R}_j = (X_j , Y_j , Z_j)$ be the coordinates of control point j, and $\boldsymbol{r}_i = (x_i , y_i , z_i)$ those of surface node i. The quantities we are interested in are $\partial J/\partial X_j$, $\partial J/\partial Y_j$ and $\partial J/\partial Z_j$ for all control points j. These will allow the efficient steering of the morphing process, e.g. towards the direction of steepest descent of the cost function. Applying the chain rule yields

$$\frac{\partial J}{\partial X_j} = \sum_i \left(\frac{\partial J}{\partial x_i}\frac{\partial x_i}{\partial X_j} + \frac{\partial J}{\partial y_i}\frac{\partial y_i}{\partial X_j} + \frac{\partial J}{\partial z_i}\frac{\partial z_i}{\partial X_j} \right) = \sum_i \frac{\partial J}{\partial \boldsymbol{r}_i}\frac{\partial \boldsymbol{r}_i}{\partial X_j} \tag{1}$$

and similarly for the other directions. The sum goes over all nodes, or, more precisely, over all nodes that are affected by the movement of control point j. While the last term of Eqn. (1), $\partial \boldsymbol{r}_i /\partial X_j$, represents the so-called DVGRID-card (or shape basis vector) for a unit X-displacement of control point j and is therefore available inside ANSA, the vectorial sensitivity $\partial J/\partial \boldsymbol{r}_i$ can be simply computed as:

$$\frac{\partial J}{\partial \boldsymbol{r}_i} = \frac{\partial J}{\partial \beta_i}\boldsymbol{n}_i\,. \tag{2}$$

Here, $\boldsymbol{n}_i$ is the outward facing normal vector at node i and β_i the nodal displacement in normal direction. $\partial J/\partial \beta_i$ is computed by mapping the element sensitivities $\partial J/\partial \alpha_k$ as delivered by the adjoint CFD solver onto the adjacent nodes. In this way, we finally arrive at the formula for the desired sensitivities of control point j:

$$\frac{\partial J}{\partial X_j} = \sum_i \frac{\partial J}{\partial \beta_i}\boldsymbol{n}_i \cdot \frac{\partial \boldsymbol{r}_i}{\partial X_j} \tag{3}$$

and similarly for Y_j and Z_j.

The update of the control point coordinates is performed with a simple steepest descent step according to

$$X_j^{new} = X_j^{old} - \lambda \frac{\partial J}{\partial X_j}, \tag{4}$$

where λ is a fixed step size – in the current implementation determined by defining an upper limit on the shape perturbation of the first optimization cycle.

Realization within ANSA

Given the fact that shape sensitivities, as computed by an adjoint solver, are usually quite noisy, the box approach to morphing, rather than the direct morphing, was chosen as the basis for the new method. By moving the edges or the control points defining a morphing box, the shape of the geometry and/or of mesh entities that are loaded to it, can be smoothly modified. In addition, combined movements of control points can be aggregated into morphing parameters.

For the computation of the parameter or control point sensitivities and the subsequent morphing, the equations outlined above were hard-coded into ANSA. The resulting process looks as follows:

1. The element-based sensitivities (generated by the adjoint solver) are imported into ANSA in OpenFOAM® format.
2. A maximum shape displacement value is specified by the user.
3. The morphing parameters to be updated by the sensitivity-based morphing algorithm are chosen.
4. The model is morphed automatically according to the given specifications and can be exported for further analysis.

This process is what we call the "single improvement step" mode of the new methodology: On the basis of the sensitivities, the model is updated once within the ANSA GUI. Of course, this update can be automated and included as part of an automatic optimization loop.

Alternatively, the whole metric information connecting the mesh nodes and the control points, i.e. the DVGRID cards, can be exported in order to run the mesh morphing and the optimization with suitable tools outside ANSA. For the current study, OpenFOAM® routines were used to aggregate the whole optimization process – solving the primal, solving the adjoint, applying sensitivity-based mesh-morphing – into a single OpenFOAM® application. After generating the DVGRID cards in ANSA, this application can run independently as a so-called "one-shot optimization".

3. APPLICATION

In the following, it will be shown how the newly developed method of sensitivity-based morphing was applied to an aerodynamic optimization problem: drag minimization of the Volkswagen low emission car L1 (see Fig. 1). With a drag coefficient of less than 0.2, this car is aerodynamically already nearly perfect – any further optimization is therefore a challenging task.

All computations were carried out on a half model of the fully symmetric L1. The mesh was created with ANSA by using the CFD variable size, curvature dependent surface meshing algorithm. Several refinement zones have been defined in order to increase precision in specific areas of the car model. A refined boundary of 20 prismatic layers on the car and on the street, which results in y+-values around 1, and hexahedral elements in the interior make up the volume mesh consisting in total of 18 Mio cells (Fig. 2).

Fig. 1: The Volkswagen low emission car L1.

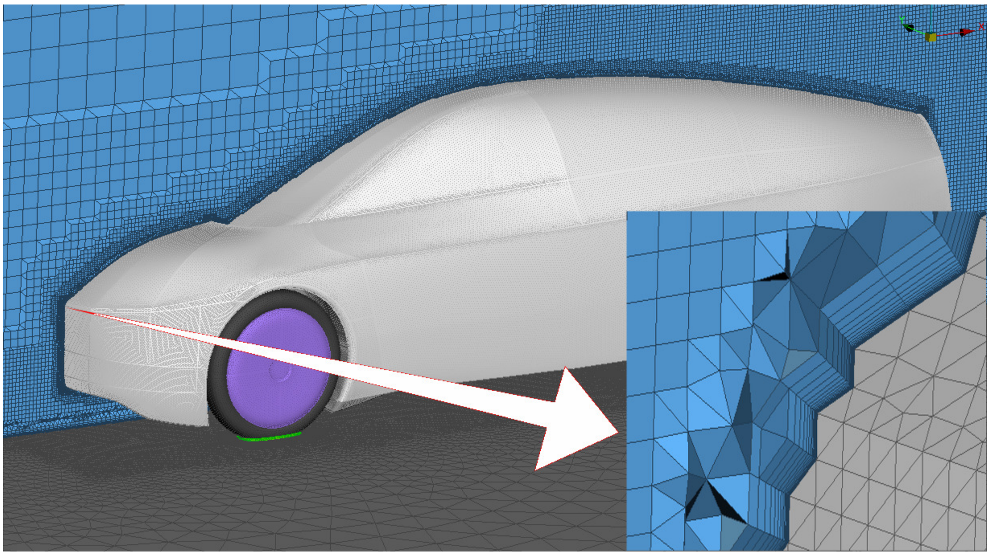

Fig. 2: The computational mesh generated by ANSA.

In order to get a detailed picture of the optimization potential, a drag sensitivity map was computed by running a RANS computation with the Spalart-Allmaras turbulence model and a subsequent adjoint computation including the adjoint turbulence model of Zymaris et al. [4]. According to Fig. 3, the L1 exhibits significant potential for drag optimization in several areas, among them the front wing, the front end of the bonnet and the car's back: interestingly, the sensitivity map "suggests" the formation of a rear spoiler here.

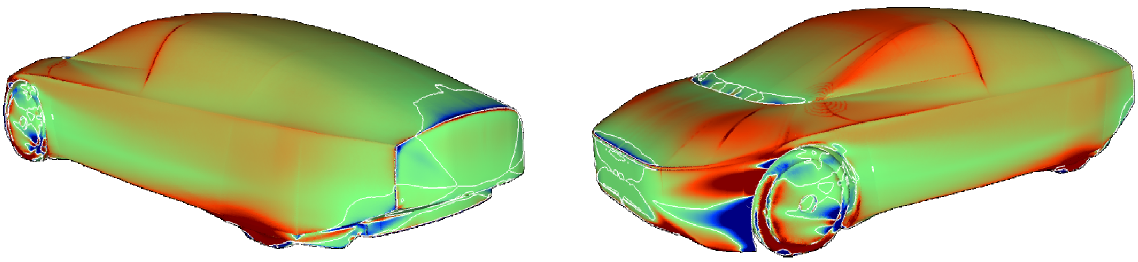

Fig. 3: Drag sensitivity maps for the L1. Strong red and blue areas pinpoint those regions on the car surface, where shape changes have a significant impact on the aerodynamic drag, whereas greenish areas are rather insensitive. The sensitivity map also provides information on the direction of the favourable shape change: red areas have to be moved inwards and blue areas outwards in order to reduce drag.

As presented previously [5], a number of different morphing boxes were defined in order to modify the sensitive areas in the "single improvement step" mode, i.e. reading the sensitivities into ANSA, computing the relevant parameter sensitivities and automatically morphing the model accordingly. Following this process for each of those areas, it was successfully shown that the updated car shape resulted in a reduced drag. The current study, however, focuses on the "one shot optimization" mode of the sensitivity-based morphing method: The aim was to specifically exploit the sensitivity information on the car's back, in order to find the optimal shape of the rear spoiler w.r.t. drag.

Several 3D morphing boxes have been generated around the car's back (Fig. 4) and, where required, have been snapped onto the rear edge in order to add more accuracy to the movement response. In order to allow for sufficient design freedom for the rear spoiler to be built up, five control points are used per half model to control the shape of the rear edge. The movement of the control points was restricted to the z-direction, thus keeping the length of the car unchanged.

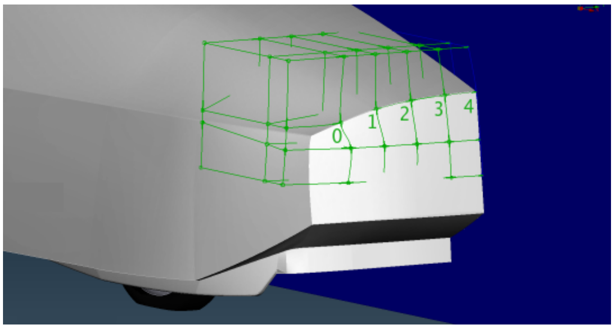

Fig. 4: Morphing boxes and control points for the rear spoiler area.

The corresponding DVGRID information for the morphing boxes was extracted and provided to our adjoint-based steepest descent optimization algorithm in OpenFOAM®. A fixed step size corresponding to a maximum shape displacement of 5mm in the first iteration was chosen, and 10 optimization cycles, amounting to a total computational effort of about 3.6 equivalent flow solutions, were performed (see Fig. 5). During this first part of the optimization, the drag decreased by roughly 1.5%. Subsequently, the step size was reduced to one half and 9 further iteration steps were carried out, as a result of which the drag continued to decrease and reached a total reduction of more than 2%. As a side effect, also the total lift decreased – by as much as 30%. Fig. 6 depicts how the reduction of drag and lift arises: the strong pressure increase upstream of the automatically formed rear spoiler pushes the car forward and downward.

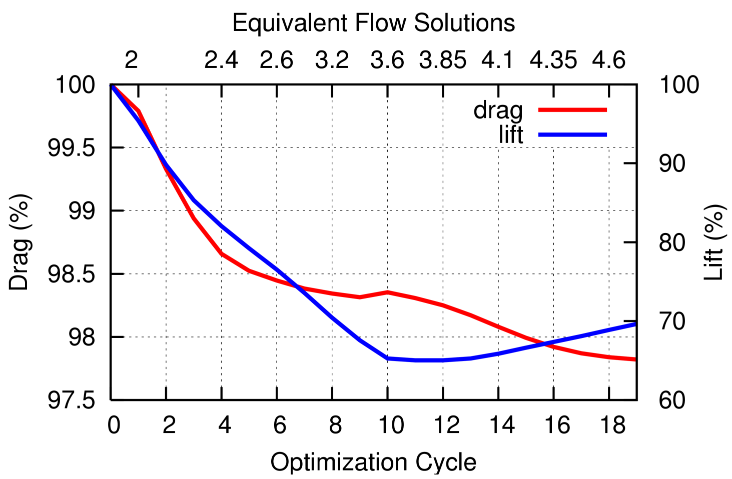

Fig. 5: Drag and lift evolution during the one-shot optimization. After 10 cycles, the step size was reduced by one half. Note that the whole optimization required a computational effort of less than 5 equivalent flow solutions (see upper x-axis).

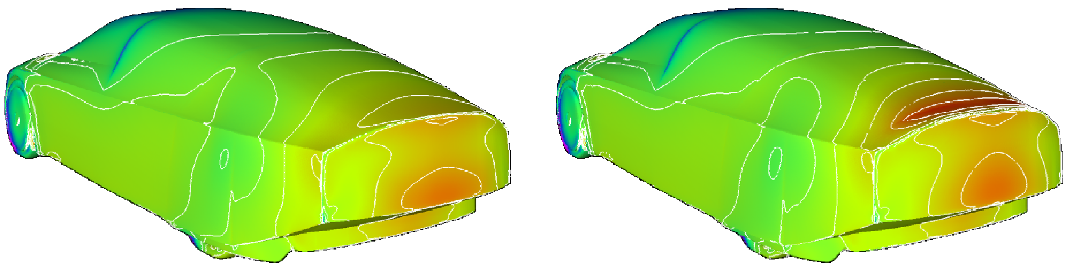

Fig. 6: Pressure contours on the car surface before (left) and after the optimization (right). Note the large footprint of the rear spoiler on the right picture.

This kind of vehicle drag improvement after almost 20 iteration cycles has been achieved in the past with other optimization methods as well – the innovative aspect of the new adjoint-based methodology lies, however, in its efficiency: Note from Fig. 5 that the overall optimization process requires a total effort of less than 5 equivalent flow solutions – two for running the primal and the adjoint for the initial computation of the sensitivities plus an effort of about three flow solutions for the actual one-shot optimization comprising iterations of the primal, the adjoint and the shape update.

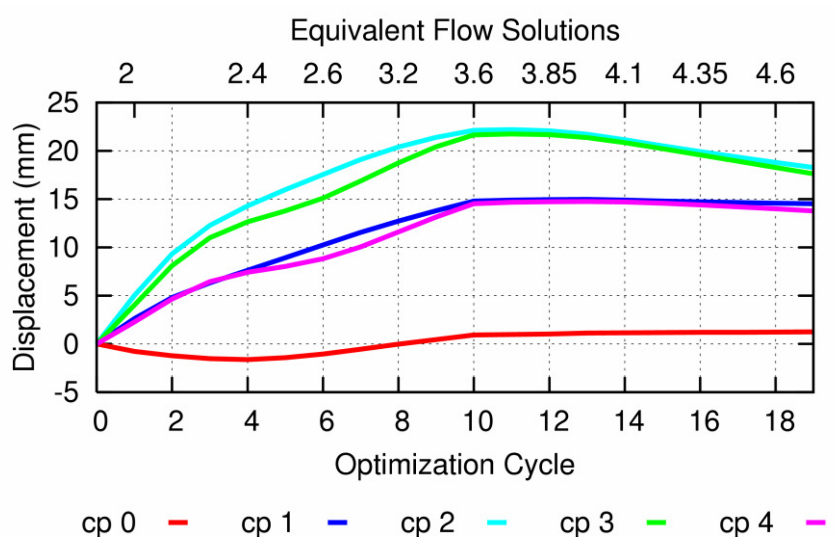

Fig. 7: The evolution of the control point coordinates during the one-shot optimization (cf. Fig. 4 for the respective position of the control points cp0 to cp4).

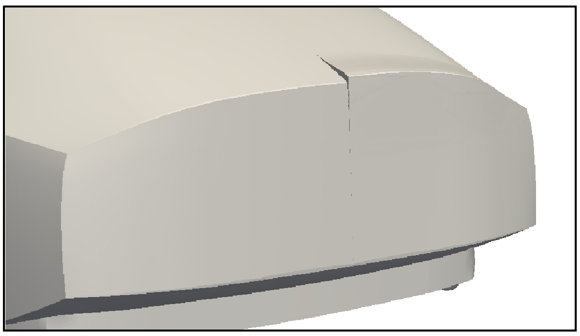

Fig. 8: The respective shapes of baseline (left) and optimum (right). The maximum distance between them is less than 20mm.

With the total displacements of the morphing control points being less than 20mm (Fig. 7), the overall shape change required to achieve the final drag improvement is quite small (Fig. 8). This is another characteristic of the presented method: it is "minimal invasive". Since the shape deformation is directly controlled by the sensitivities, it favours the morphing of sensitive areas over insensitive and can therefore translate a given maximum shape displacement into the biggest possible impact onto the cost function.

4. SUMMARY AND OUTLOOK

A sensitivity-based shape morphing method was developed, which closes an important gap towards an automatic adjoint-based shape optimization process. It can be applied in two ways: (1) for performing a single improvement step within the ANSA GUI, and (2) for running a one-shot optimization in OpenFOAM®, based on morphing information extracted from ANSA. With the latter, the optimization of shape details is now feasible at a cost of around 5 flow solutions, as demonstrated by the application to the external aerodynamics of the Volkswagen L1. Despite having delivered convincing results for the given example, this morphing methodology is still in its infancy and exhibits some shortcomings, like the lack of an automatic step-size control within the one-shot optimization, which are the focus of subsequent work.

REFERENCES

(1) http://www.openfoam.com

(2) C. Othmer, E. de Villiers and H.G. Weller, Implementation of a continuous adjoint for topology optimization of ducted flows, AIAA-2007-3947, 2007

(3) C. Othmer, A continuous adjoint formulation for the computation of topological and surface sensitivities of ducted flows, Int. J. Num. Meth. Fluids, 58: 862-877, 2008, DOI: 10.1002/fld.1770

(4) A.S. Zymaris, D.I. Papadimitriou, K.C. Giannakoglou and C. Othmer, Continuous Adjoint Approach to the Spalart Allmaras Turbulence Model for Incompressible Flows, Computers & Fluids, 38, pp. 1528-1538, 2009

(5) C. Othmer, K.C. Giannakoglou, A.S. Zymaris, D.I. Papadimitriou and E.M. Papoutsis-Kiachagias, Adjoint Methods for Automotive Optimisation: Achievements and Challenges, FLOWHEAD Workshop, Varna, 2010

(6) E. de Villiers, Multi-Objective Adjoint Optimization of Flows in Duct and Pipe Networks, Flow Problem Analysis in Oil & Gas Industry, Rotterdam, The Netherlands, 2011

(7) C. Hinterberger and M. Olesen, Application of a Continuous Adjoint Flow Solver for Geometry Optimisation of Automotive Exhaust Systems, FLOWHEAD Workshop, Varna, 2010

(8) T.T. Robinson, C.G. Armstrong, H.S. Chua, C. Othmer and Th. Grahs, Sensitivity-based optimization of parameterized CAD models, 8th World Congress on Structural and Multidisciplinary Optimization, June 2009, Lisbon, Portugal

(9) E. Ribaldone and G. Scantamburlo, Application of multi-objective optimization to the aerodynamic development of passenger cars at FIAT, 8th MIRA International Vehicle Aerodynamics Conference, October 2010, Oxford, UK

(10) F. Campos, P. Geremia and E. Skaperdas, Automatic optimization of automotive designs using mesh morphing and CFD, EACC 2007, Frankfurt, Germany

(11) C. Othmer, Sensitivity-based shape optimization with the ANSA Morphing Tool: Mapping of sensitivities from nodes to control points, Report, Volkswagen AG and BETA CAE Systems S.A., November 2006