

RESEARCH

Open Access

Adjoint methods for car aerodynamics

Carsten Othmer*

*Correspondence:
carsten.othmer@volkswagen.de
Vehicle Technology, Group
Research, Volkswagen AG, 38436
Wolfsburg, Germany


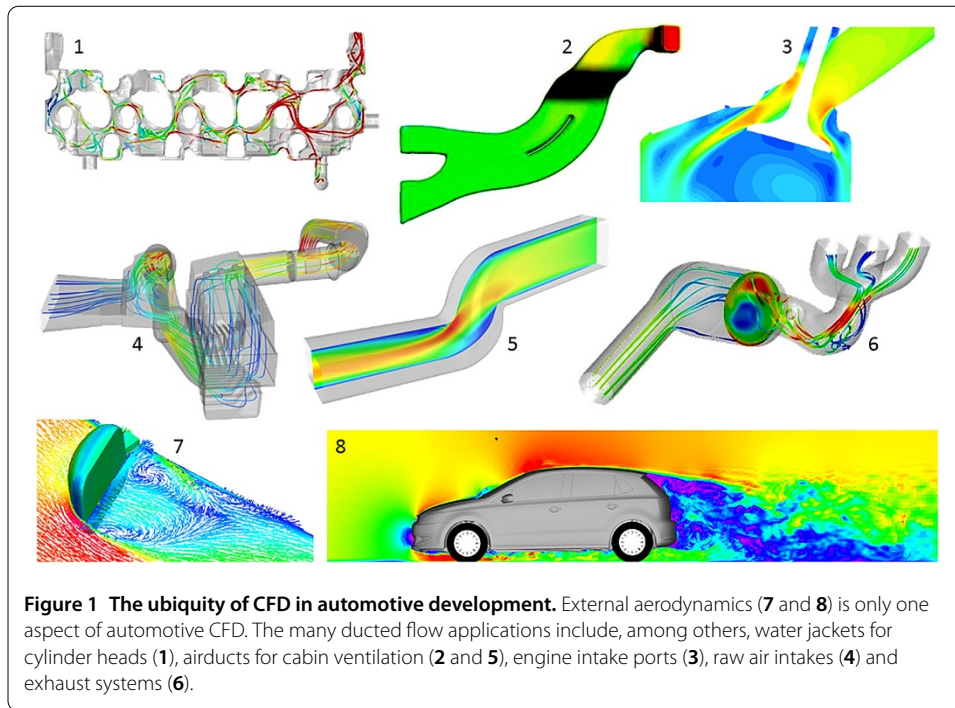
Abstract

The adjoint method has long been considered as the tool of choice for gradient-based optimisation in computational fluid dynamics (CFD). It is the independence of the computational cost from the number of design variables that makes it particularly attractive for problems with large design spaces. Originally developed by Lions and Pironneau in the 70's, the adjoint method has evolved towards a standard tool within the development processes of the aeronautical industries. Its uptake in the automotive industry, however, lags behind. The first systematic applications of adjoint methods in automotive CFD have interestingly not taken place in the classical shape design arena, but in a relatively young discipline of sensitivity-based optimisation: fluid dynamic topology optimisation. While being an established concept in structure mechanics for decades already, its transfer to fluid dynamics took place just ten years ago. We demonstrate that specifically for ducted flow applications, like airducts for cabin ventilation or engine intake ports, it constitutes a very powerful tool and has matured over the last years to a level that allows its systematic usage for various automotive applications. To drive adjoint-based shape optimisation to the same degree of maturity and robustness for car applications is the subject of ongoing research collaborations between academia and the car industry. Achievements and challenges encountered during these efforts are presented.

1 Background

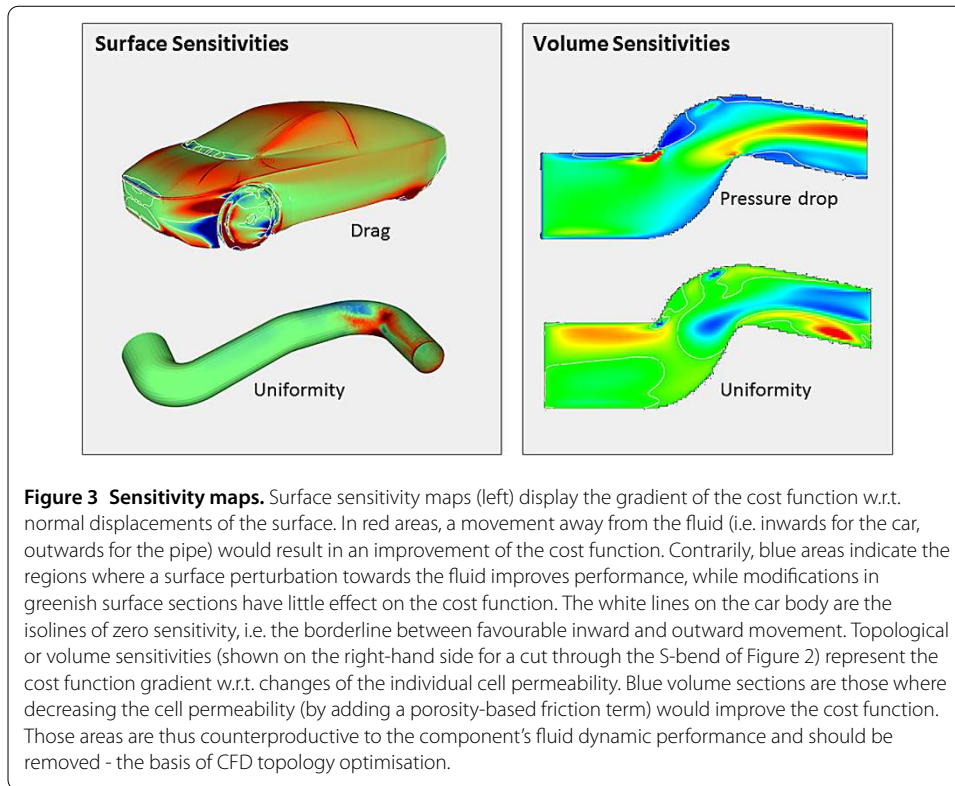
Computational Fluid Dynamics (CFD) is a central element of the automotive development process. Besides the classical external aerodynamics for the prediction of drag and lift coefficients, there is a whole plethora of applications for ducted flows: airducts for cabin ventilation, engine intake ports, exhaust systems including catalytic converters, air intakes, water jackets of cylinder heads, cooling plates for electric vehicle battery packs and many more (Figure 1). The standard procedure of developing these parts still consists of manual iteration loops between designer and computational engineer, and automatic optimisation methods are being used systematically only for some selected CFD applications (e.g. for airducts, Figure 2). Via an automatic process chain consisting of a CAD (Computer-Aided Design) system, meshing software, CFD solver and post-processor a black-box optimisation algorithm, typically of the evolutionary strategy type, is employed to drive a parameterised CAD-model into an optimal state.

Due to the high computational effort associated with black-box optimisation, where the number of CFD evaluations scales roughly linearly with the number of design variables, the explorable design space is very limited: In our current practice, these methods cannot afford more than ten design parameters. An additional obstacle that inhibits the further



adoption of black-box optimisation within the regular automotive development process is the necessity of a versatile but stable CAD parameterisation. For e.g. an engine intake, the setup of the parameterised model might take an experienced CAD engineer several weeks before it comes up to the stability requirements of an automatic optimisation process chain. It is these limitations that call for a radically different approach to automotive CFD optimisation - especially in view of applications like entire vehicle aerodynamics that entail shape optimisation of highly complex free-form geometries with a practically unlimited design space.

Adjoint-based optimisation [1–4] has long been recognized to be efficient for tasks with large numbers of design variables. Its underlying principle, the adjoint method, allows for a very elegant computation of sensitivities, i.e. of the derivative of the objective function



w.r.t. the design variables. The computational effort is thereby independent of the number of design variables and just involves one solve of the adjoint counterparts of the governing equation system, in our case the Navier-Stokes equations. When applied to a surface mesh representation or a volume mesh of the part to be optimised, information like the one depicted in Figure 3, for which we coined the term 'sensitivity maps,' can be generated: A surface sensitivity map tells us for each and every surface node how the objective function changes w.r.t. an infinitesimally small normal displacement of this surface node. On the other hand, a volume sensitivity map represents for each volume cell the sensitivity of the objective function w.r.t. a change of permeability of this volume cell. Both sensitivity maps give precise indications on where and how to change the geometry - perturb the surface inwards or outwards, and remove counterproductive cells from the flow domain, respectively - in order to improve the objective function. Obviously, this wealth of information would be inaccessible without the adjoint method, and, as will be shown in the following, allows to develop very efficient optimisation methods for automotive part design.

Despite being already a long-established tool in the aerospace sector, the adjoint method has only recently started to also enter the automotive development processes. This backlog is mainly due to two obstacles specific to the automotive industry.

Obstacle #1 *Rather than on in-house codes, the car industry relies almost exclusively on commercial CFD software.*

The implementation of the adjoint method requires access to the source code - either to implement the analytically derived adjoint equations into the code (so-called continuous

adjoint) or to perform a manual or automatic differentiation of the primal code (discrete adjoint, see [5] for a comparison). The sources of commercial CFD codes are, however, inaccessible for the user, and by the time we started our project to industrialise the adjoint method, the software vendors of the CFD solvers that are operational at Volkswagen declined to implement it into their codes. As a remedy, we selected the well-validated and versatile CFD toolbox OpenFOAM® [6] as our platform to develop adjoint-based optimisation methods for car applications. In addition to its open source character, OpenFOAM comes with a high-level symbolic programming style, which makes it ideally suited to implement additional partial differential equations - like the adjoint Navier-Stokes equations.

The first continuous adjoint implementation in OpenFOAM [7] is based on the derivation outlined in [8]. It focused on steady-state, incompressible ducted flows as described by the RANS (Reynolds-averaged Navier-Stokes) equations:

$$(\mathbf{v} \cdot \nabla)\mathbf{v} = -\nabla p + \nabla \cdot (2\nu\mathbf{D}(\mathbf{v})) - \alpha\mathbf{v}, \quad (1)$$

$$\nabla \cdot \mathbf{v} = 0, \quad (2)$$

where \mathbf{v} and p stand for velocity and pressure, respectively, and α for the porosity distribution. The effective kinematic viscosity ν is the sum of molecular and turbulent viscosity, and the rate of strain tensor $\mathbf{D}(\mathbf{v})$ is defined as $\mathbf{D}(\mathbf{v}) = \frac{1}{2}(\nabla\mathbf{v} + (\nabla\mathbf{v})^T)$. For a given cost function J composed of contributions from the design domain boundary $\Gamma = \partial\Omega$ and from the interior of Ω ,

$$J = \int_{\Gamma} J_{\Gamma} d\Gamma + \int_{\Omega} J_{\Omega} d\Omega, \quad (3)$$

the implemented adjoint field equations are

$$-(\mathbf{v} \cdot \nabla)\mathbf{u} - \nabla\mathbf{u} \cdot \mathbf{v} = -\nabla q + \nabla \cdot (2\nu\mathbf{D}(\mathbf{u})) - \alpha\mathbf{u} - \frac{\partial J_{\Omega}}{\partial \mathbf{v}}, \quad (4)$$

$$\nabla \cdot \mathbf{u} = \frac{\partial J_{\Omega}}{\partial p}, \quad (5)$$

with \mathbf{u} and q denoting adjoint velocity and adjoint pressure, respectively. While the cost function contribution from the domain interior J_{Ω} gives rise to source terms in the adjoint field equations, the surface contribution J_{Γ} enters the adjoint boundary conditions:

$$\mathbf{n} \cdot \nabla q = 0, \quad \mathbf{u}_t = 0, \quad u_n = -\frac{\partial J_{\Gamma}}{\partial p} \quad (6)$$

for boundaries with Neumann conditions for the primal velocity and a Dirichlet condition for the primal pressure (usually inlet and wall boundaries), and

$$q = \mathbf{u} \cdot \mathbf{v} + u_n v_n + \nu(\mathbf{n} \cdot \nabla)u_n + \frac{\partial J_{\Gamma}}{\partial v_n}, \quad (7)$$

$$0 = v_n \mathbf{u}_t + \nu(\mathbf{n} \cdot \nabla)\mathbf{u}_t + \frac{\partial J_{\Gamma}}{\partial \mathbf{v}_t} \quad (8)$$

for boundaries with Dirichlet conditions for the primal velocity and a Neumann condition for the primal pressure (e.g. outlet boundaries, for details we refer again to [8]). By solving

both primal and adjoint equation system, the desired (volume-normalised) topological and (surface-normalised) shape sensitivities can be computed as

$$\frac{\partial J}{\partial \alpha} = \mathbf{u} \cdot \mathbf{v} \quad \text{and} \quad \frac{\partial J}{\partial \beta} = -\nu \partial_n \mathbf{u} \cdot \partial_n \mathbf{v}, \quad (9)$$

where β denotes the normal node displacement.

Within the still ongoing collaboration with Prof. Giannakoglou's group at the National Technical University of Athens (NTUA) and Dr. de Villiers at Engys[®], this basic implementation was subsequently enhanced towards an inclusion of an adjoint turbulence model [9], optimisation of external flows [10], sensitivities for flow control [11], adjoint wall functions [12], the inclusion of heat conduction and constraints [13] and mildly compressible flows [14]. In addition, the uptake and further development of the basic adjoint code by other researchers lead to remarkable advances [15–18], especially in the area of topology optimisation for exhaust systems [19, 20]. As a result of these development efforts, a versatile continuous adjoint code in OpenFOAM for steady-state RANS is now available.

Obstacle #2 *The geometries of entire vehicles and even of single automotive parts are highly complex.*

The complexity of typical automotive geometries, especially of entire vehicles, which consist of an assembly of several hundred parts, necessitates the use of an automatic process for the generation of the computational mesh. Even though the sophistication of automatic meshing software is continuously increasing, the achievable mesh quality is still inferior to a hand-made mesh, especially concerning the creation of layers along the surface of the flow domain. While this quality suffices for obtaining reasonable results for primal computations, it has been found that the adjoint equation system reacts much more delicately to imperfections of the computational mesh. This is particularly the case in the presence of flow separation, which is commonly encountered in flows in and around automotive geometries.

This lack of robustness of the adjoint when running on typical industrial meshes is a problem that has been accompanying the adjoint development in OpenFOAM from the very beginning - and for certain applications it is still inhibiting the systematic use of the adjoint method in the regular computational processes. The source of this instability is the so-called 'adjoint transpose convection' $\nabla \mathbf{u} \cdot \mathbf{v}$ - the second convection term on the left-hand side of Eqn. (4). By developing several formulations for this term, as well as explicit and implicit treatments, major steps towards increased stability were made. As it introduces a high degree of cross-coupling between the three components of the velocity, recent efforts concentrate on solving the adjoint equation system as a single block rather than in a segregated manner [21].

While those developments are very promising, the current practice of running the adjoint - at least on meshes of the typically encountered quality in the car industry environment - still requires to trade robustness against accuracy: by using limiters, relaxed adjoint boundary conditions or by damping the adjoint transpose convection in an intelligent way that does not harm the *qualitative* correctness of the sensitivities. To report the achievements as well as the challenges of this pragmatic approach to solving adjoints for car applications is the intention of this article. It is organized as follows: The concepts of

topology optimisation and shape optimisation for ducted flows are introduced in the next section, along with an application of these two complementary methodologies to engine port flows. Section 3 focuses on external aerodynamics, comprising a full-vehicle validation study, an example of a component optimisation and a comparison of steady-state RANS sensitivities with those based on a time-averaged transient primal flow. The article finally closes with a comparatively young adjoint application area - car drag reduction by flow control.

2 Ducted flows

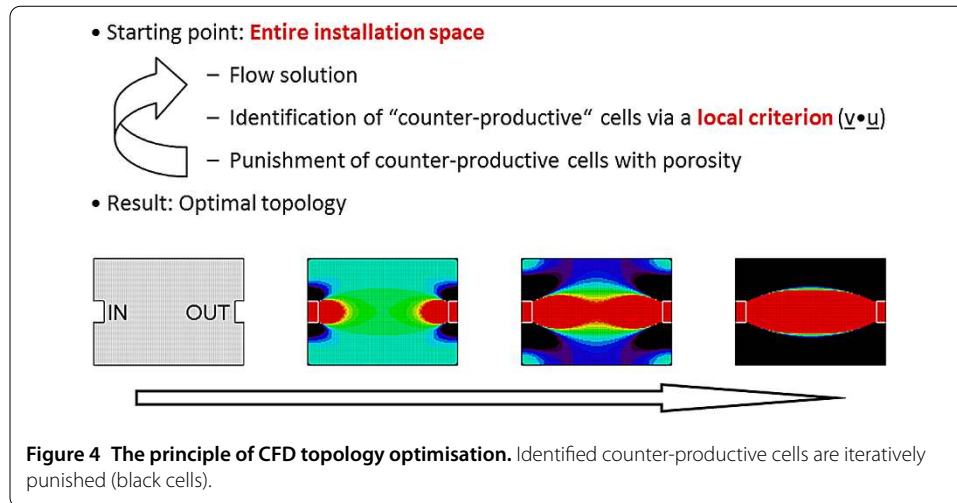
Geometric optimisation of ducted flows in cars, like airducts for cabin climatisation, engine air intakes or exhaust systems, are commonly subject to severe packaging constraints. This gave rise to the development and adoption of topology optimisation methods for automotive applications. After a concise retrospective of these development efforts, it will be shown that especially for ducted flows, topology optimisation is a perfect complement to the classical shape optimisation.

2.1 Topology optimisation

Topology optimisation is a well-established tool in computational structure mechanics [22] with widespread industrial use. In its simplest realisation, a topological optimisation starts from the available design domain filled up entirely with solid material of a certain density. In an iterative fashion, the given loads are applied, the stresses are computed all over the domain and the areas with low stresses are weakened by assigning a lower density to them. After several iterations, this method retains high-density material only in regions that are critical to fulfill the structural task, and in this way generates optimal lightweight structural designs. It delivers an un-biased design from scratch that automatically fulfills the installation space constraints.

It almost comes as a surprise that it took until 2003 until this elegant concept was transferred to fluid dynamics: Independently of each other Borrvall and Petersson [23] and Klimetzek [24] presented the first topological optimisation methods for ducted flows. Analogously to its structural mechanic archetype, their fluid dynamic topology optimisation starts from a completely flooded computational domain (Figure 4), uses some local criterion to identify those areas that are counterproductive w.r.t. the chosen objective function and iteratively 'punishes' or removes them from the fluid domain. As a result, the remaining areas constitute the optimal duct between inlet(s) and outlet(s) of the given installation space.

Despite being based on the same concept of flow optimisation, these two approaches exhibit significant differences. Klimetzek's method is a rather pragmatic solution: The local criterion he employs to identify counterproductive cells is the deviation between the actual velocity direction and a desired direction computed as polygonal traces between inlet and outlet. If this deviation exceeds a certain angle, the corresponding cells are blocked for the fluid. His criterion allows to efficiently remove unwanted recirculations from the flow and - albeit being used for other cost function as well - is therefore ideally suited to reduce the pressure drop. From the date of its original publication, Klimetzek's method has already been applicable to three-dimensional turbulent Navier-Stokes flows and is seeing a continuously increasing number of successful applications in industrial flow optimisation.



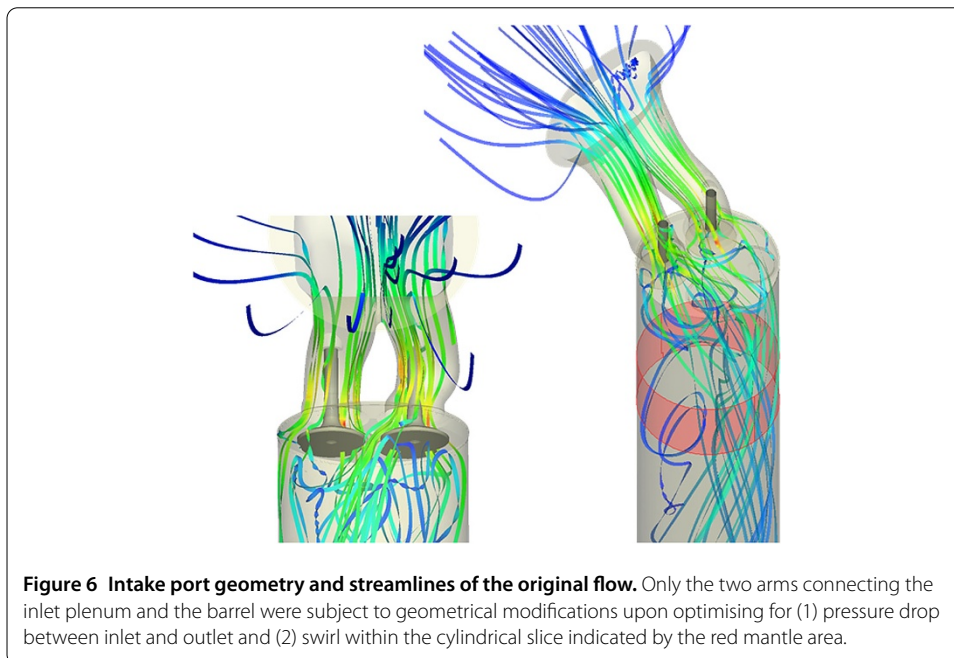
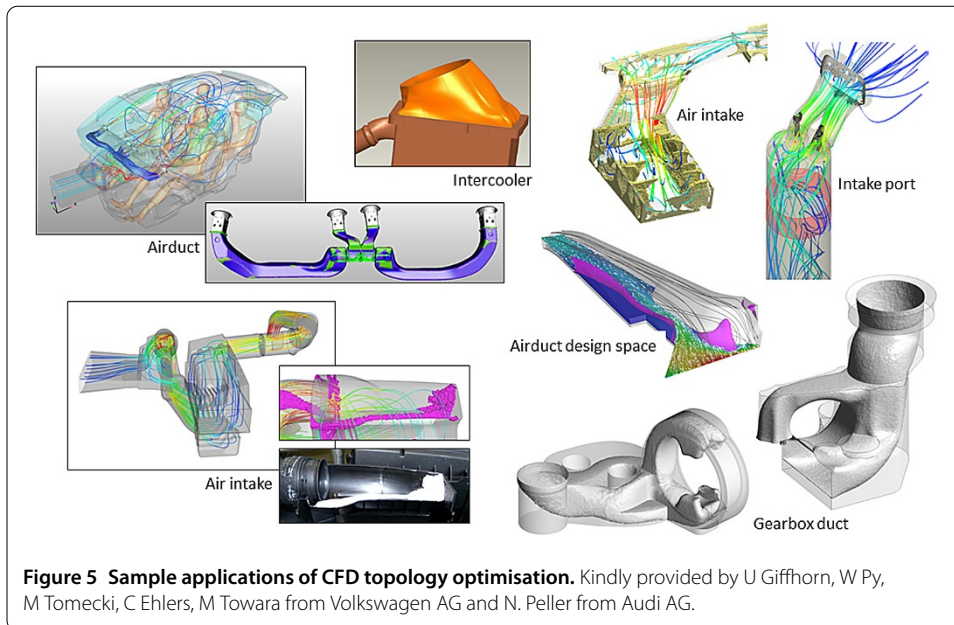
In contrast, the method proposed by Borrvall and Petersson was restricted to two-dimensional Stokes flows between parallel plates, and the punishment of counterproductive areas is performed by locally decreasing the distance between the plates - until zero upon convergence. What inspired a series of subsequent research works is, however, the versatility of their approach that lies in the way how they identify counterproductive fluid cells: by computing actual sensitivities via the adjoint method.

In order to extend the Borrvall and Petersson method towards three-dimensional flows, the punishment via plate distance was replaced by introducing a Darcy porosity term $-\alpha \mathbf{v}$ into the governing Eqn. (1) (see [25] and [26] for a finite-element and a finite-volume implementation, respectively). The individual cell porosities allow a continuous transition between fluid ($\alpha = 0$) and solid ($\alpha = \alpha_{\max}$) and act as the actual design variables of the optimisation problem. With this setting, topological sensitivity maps like those shown in Figure 3 are nothing but the derivative of the cost function w.r.t. an increase in the cell porosity α .

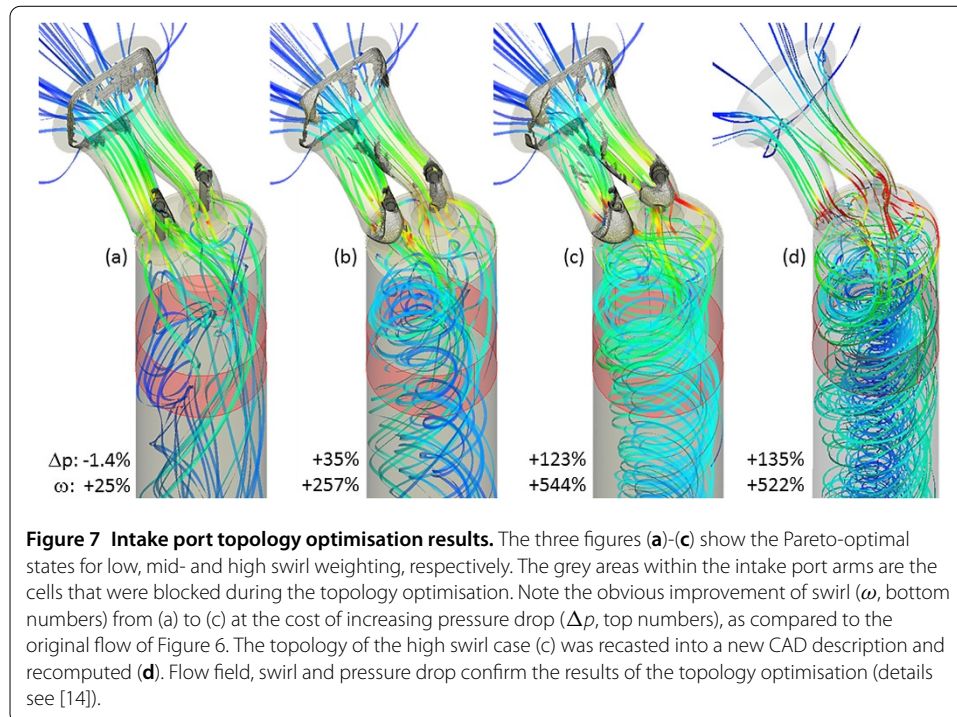
Another requirement for the industrial applicability of adjoint-based topology optimisation was its generalisation from pure Stokes flows to turbulent Navier-Stokes flows. After the principal feasibility of laminar Navier-Stokes topology optimisation was first demonstrated via Automatic Differentiation of an academic CFD code [27], the concept was implemented via a continuous adjoint also for turbulent flows [8] - albeit under the assumption of ‘frozen turbulence’, i.e. fully neglecting the variation of the turbulent flow quantities. Finally, the development of a continuous adjoint turbulence model [9] now allows to run topology optimisation for turbulent Navier-Stokes flows under full consideration of variations in the turbulence fields [13].

As topology optimisation is not only very fast but also mimics the classical way of part design when driven by installation space constraints, the uptake of this method in the regular automotive design process was quite straightforward. A number of ducted flow applications are today routinely optimised via adjoint-based topology optimisation (see Figure 5 for examples from the Volkswagen Group and [19, 20] for others). Exemplarily, one of them - the optimisation of an engine intake port - will be presented in more detail in the following.

The main components of the cylinder intake port static test configuration geometry that was the optimisation target for this study can be seen in Figure 6. Only the volume



enclosed by the two intake port arms was subjected to modification. The barrel, valves, valve seats and upstream inlet geometry (not shown) had to be kept clear of blockage via geometric and proximity based constraints. The working fluid for all the simulations was air modeled as a compressible perfect gas. Pressure at the outlet of the extended barrel was fixed, and at the inlet a fixed mass flow at ambient temperature was specified. For the adjoint system, density and turbulence was assumed ‘frozen’, thus no governing equations were solved for the adjoint systems beyond the momentum and pressure analogues. For the port flow the objectives were (1) to minimise the pressure loss between the inlet and



the outlet and (2) to maximise the swirl, i.e. the angular momentum $\mathbf{r} \times \rho \mathbf{v}$ around the barrel axis in the cylindrical sub-zone indicated in Figure 6.

One of the main issues with multi-objective adjoint optimisation is the choice of a sensible relative weighting when combining the sensitivity fields from the two objectives for input into the topology update. As a first step, we therefore performed an exploration by varying the weighting factors dynamically within a wide range. After choosing three design points with different weightings (low, mid- and high swirl) that represented a likely range for desirable fluid dynamic performances of the intake port, we actually ran the topology optimisation for these three trade-offs between swirl and pressure loss until convergence and thus obtained three Pareto-optimal states (Figure 7a-c, details see [14]).

For the low swirl case, it can be assumed that the geometry is only driven by the swirl sensitivities where the modifications have a very small detrimental effect on the pressure loss. As a result, we see a small reduction in pressure loss along with a modest increase in swirl. Predictably, the pressure loss starts to go up as the swirl weighting is increased and the fluid loses additional energy due to the highly swirling flow and the blockages required to produce it. What is encouraging is the extent to which the swirl can be reliably increased - by more than 500% - compared to the relatively modest relative increase in pressure loss (high swirl case of Figure 7c).

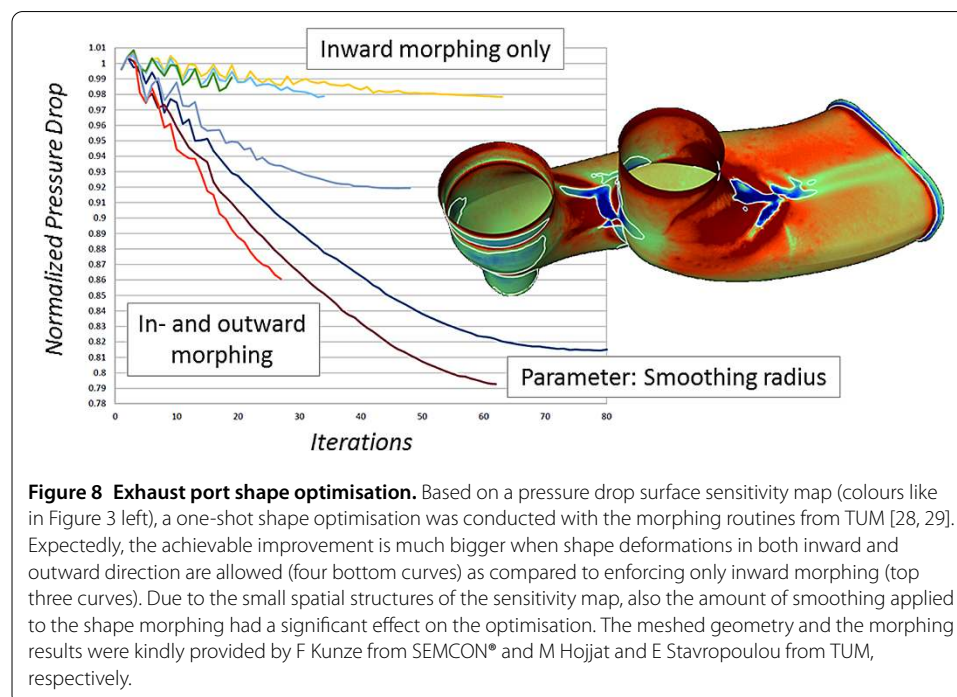
For topology optimisation, it is critical that the design extracted from the optimisation process be verified with a high accuracy simulation. To this end, the blocked element description from the high swirl case was extracted and incorporated into a new CAD description of the port flow geometry (Figure 7d). It is immediately apparent that the modified design is rather complex in its detail and not something a designer would likely create through intuition or experience. Neither could these geometries easily be arrived at via a process of parameter-based stochastic optimisation. The most important point to note is that the accurately rebuilt geometry very nearly reproduces the results from the topol-

ogy optimisation run. While we do not contend that this will consistently be the case, it certainly proves the efficacy of the process in this instance.

2.2 Shape optimisation

The topology optimisation method as introduced above, acts on the fluid domain like a sculptor on a piece of rock: It removes the unwanted portions and reduces the original geometry to its optimal shape. Since the original volume mesh is not changed during the optimisation process, the final geometry consequently consists of a subset of those cells that made up the original shape. Thus, the inherent geometrical resolution of topology optimisation corresponds to the cell size, and the final shape will inevitably have a ragged surface. Topology optimisation is therefore to be regarded as a *drafting* method. For *fine-tuning* geometries that are already close to their optimal shape, it is ideally complemented by so-called shape optimisation. Rather than on topological or volume sensitivities, shape optimisation is based on the detailed information contained in surface sensitivity maps in order to morph the shape towards a further improvement of the cost function.

An example of a surface sensitivity map of an exhaust port is shown in Figure 8. Red areas are those that have to be moved outwards (away from the air) in order to reduce the pressure drop while blue regions need to be pushed inwards. It seems straightforward to translate this information into a shape modification that improves the performance of the exhaust port: Simply move every surface node in a steepest descent manner inwards or outwards by a distance corresponding to its sensitivity magnitude. Due to the inherent 'roughness' of the sensitivity distribution this would, however, quickly result in a collapse of the surface mesh. Therefore, a clever smoothing process has to be applied to the surface sensitivities before they are actually translated into a shape update. This smoothing process does not only have to generate a good compromise between ironing out the sharpest discontinuities of the sensitivity map and at the same time conserving the essen-



tial sensitivity information, but also has to be capable of keeping certain feature lines of the geometry.

A surface regularisation method that comes up to these requirements has been developed by Prof. Bletzinger and his team at the Technical University of Munich (TUM) [28, 29]. The results of applying their method to the exhaust port above are shown in Figure 8. Depending on the size of the smoothing radius and on whether morphing is allowed only in the inward direction or both inwards and outwards, different pressure drop reductions between 2% and 20% are achieved. Since this shape optimisation process is set up such that primal and adjoint calculation are closely coupled to the shape update itself ('one-shot optimisation'), the overall cost of the optimisation amounts to not more than four primal computations.

As demonstrated above, topology optimisation and shape optimisation are complementary methods: the former for drafting, the latter for fine-tuning. Especially for the optimisation of car components that are subject to design space constraints - and in the increasingly tighter getting installation spaces within the engine compartment and the passenger cabin this is the case for the majority of components - the combination of both methods suits the requirements of automotive component design very well. Starting from the feasible installation space, they are capable of delivering a fine-tuned optimal geometry, thereby each of them making use of the full potential of the available design space.

3 External aerodynamics

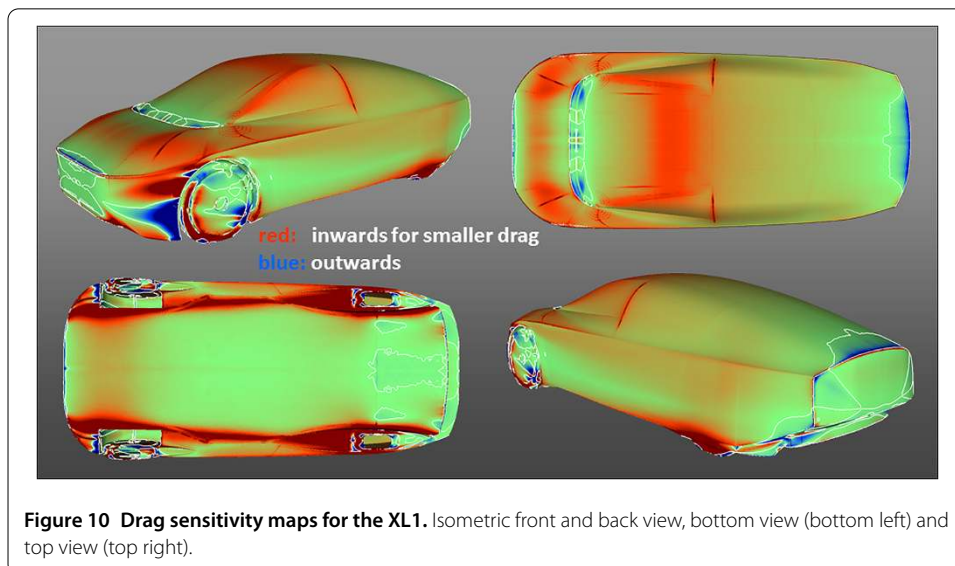

The adjoint-based topology and shape optimisation methods described above are obviously not restricted to ducted flows but can equally well be applied to external flows. The external aerodynamics of entire vehicles is, however, a peculiar application for optimisation methods: Except for dedicated low emission cars, vehicle shapes are to a high degree driven by aesthetic considerations rather than aerodynamic performance. This has two implications: (1) Given the rather limited design freedom, external aerodynamic optimisation is more about fine-tuning. Topology optimisation is therefore not an adequate tool. (2) Since aesthetic requirements can impossibly be casted into mathematical constraints, the *automatic* optimisation of vehicle shapes forbids itself - unless restricted to small portions like spoilers, mirrors or, of course, to the underbody.

Under these circumstances, the mere computation of surface sensitivity maps - with neither a topological nor a shape update - turned out to be a tool that fits very nicely into the specific requirements of the development of external vehicle shapes: The information contained in these maps provides the designer with concrete suggestions for aerodynamic improvements, and it is in his/her hand to translate this information into an aerodynamically improved shape without compromising the aesthetics.

3.1 Validation of RANS shape sensitivities for external aerodynamics

In order to validate sensitivity maps and to explore their potential for external aerodynamics, the adjoint method was first applied to the Volkswagen XL1 - a dedicated low emission vehicle developed by Volkswagen in 2011 (Figure 9). Based on a steady-state RANS solution on a low-Reynolds mesh of a half-model of the XL1, our adjoint OpenFOAM code with adjoint turbulence model according to Zymaris et al. [9] was employed to compute a sensitivity map for drag (Figure 10).

Various areas are identified as being very sensitive (i.e. dark blue or dark red). In relation to the experience of a well-versed aerodynamicist, these areas are of different qualities:



First of all, the ‘usual suspects’ of external car aerodynamics (the roof, the high-curvature regions of the bonnet, the areas around the wheels and the rear sillboard) are confirmed. For other sensitive areas like the spanwise blue stripe at the transition between the front-end and the bonnet (region 2 in Figure 11), it can be seen in the primal solution already that the flow detaches and requires the curvature to be reduced *locally* - an adjoint is not needed to get this piece of information. The added-value of the adjoint, however, becomes obvious for regions 1, 3 and 4: the rear end of the car where the adjoint ‘suggests’ a spoiler, the blue area of the front wing which should be further exposed, and the wheel spoiler area in front of the front tyres. This information cannot be deduced by looking only at the primal. The effect of changing the car shape here is *non-local* and happens further downstream of the sensitive regions themselves. It can therefore only be identified by the adjoint.

For the validation study, the four regions already mentioned and shown in Figure 11 were parameterised with free-form deformation boxes (Figure 12), and the car shape was




Figure 11 Selected areas for the validation of RANS surface sensitivity maps. Rear spoiler (1), front end of bonnet (2), front wing (3) and wheel spoiler area (4).


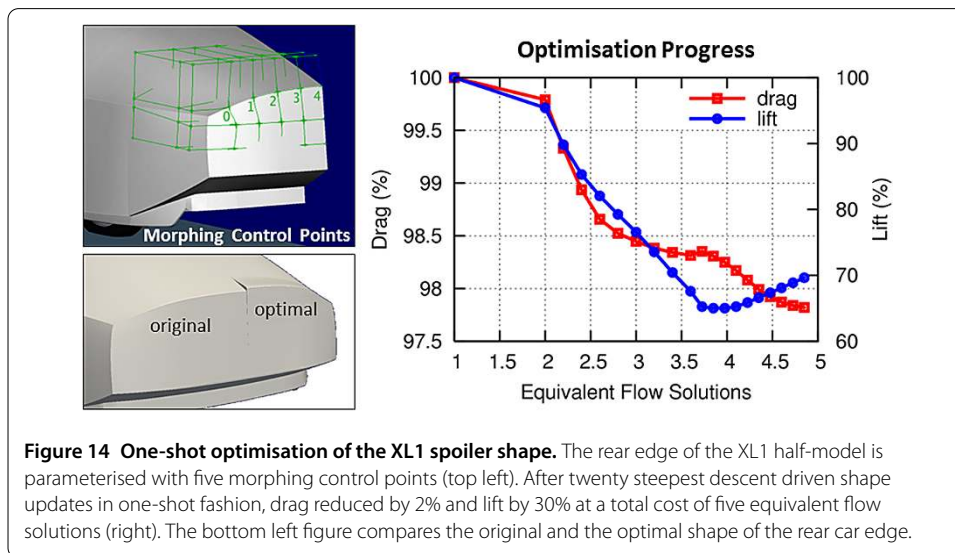
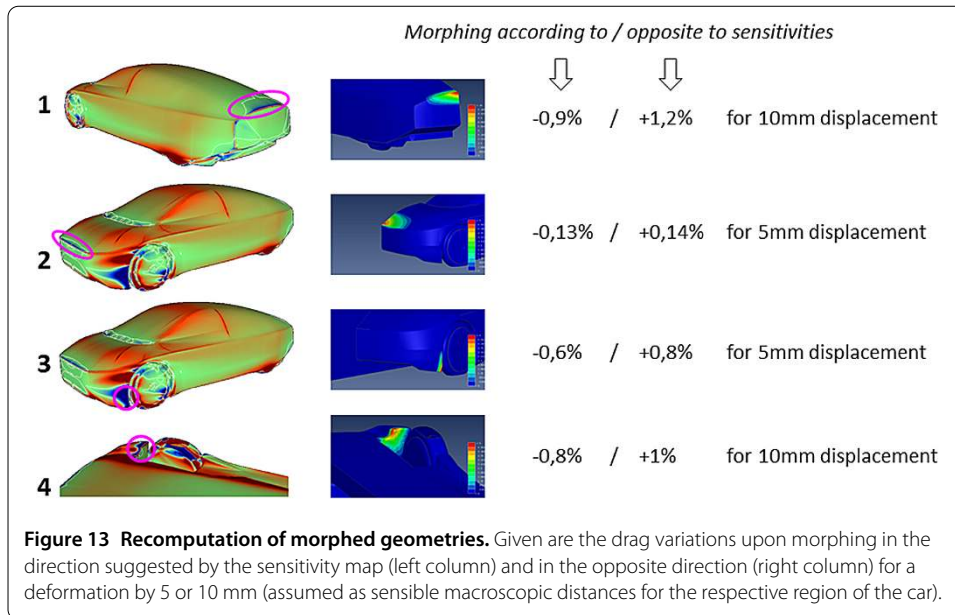


Figure 12 Definition of morphing boxes. The numbers refer to the region numbering of Figure 11 (kindly set up by BETA CAE Systems within their pre-processor ANSA®).

morphed by ‘macroscopic’ distances (5 or 10 mm, respectively) both according to and opposite to the computed sensitivity map. Recomputations of the morphed car shapes confirmed the correctness of the sensitivity signs in all cases: Morphing in the direction of the sensitivities resulted in a decrease of drag, while drag increased upon morphing opposite to the local sensitivities (Figure 13). Additionally, the drag reduction resulting from a rear spoiler (region 1) and from a wheel spoiler (region 4) were confirmed in wind tunnel experiments.

The automation of vehicle shape morphing - at least for limited portions of the entire vehicle shape - can be accomplished by coupling the adjoint OpenFOAM solver with the free-form deformation tool of ANSA® [30]: Via the chain rule, the node sensitivities as computed by the adjoint solver are mapped onto the morphing control points of ANSA (see [10]). Their positions can then be optimised by any gradient-based algorithm. Exemplarily, this method was applied to the rear spoiler area of the XL1 (region 1). The rear edge of the half-model was parameterised by 5 morphing control points whose move-



ments were restricted to the z -direction (i.e. up and down, Figure 14). Within a one-shot optimisation driven by a simple steepest descent algorithm the z -coordinates of the morphing control points were moved towards their optimal position after roughly 20 iterations (Figure 14). The overall cost of the optimisation corresponded to less than five primal flow solutions and resulted in a drag decrease of 2%. Given the fact that the XL1 with a drag coefficient of less than 0.2 was aerodynamically nearly perfect already, this was a significant reduction. As a beneficial side-effect, lift decreased by 30%.

Rather than with free-form deformation boxes, which allow to modify shapes in a very controlled manner but are quite tedious to set up, another option of translating surface sensitivities into an improved shape are of course the node-based morphing algorithms from TUM as introduced in the previous section. Since they are capable of moving each




Figure 15 Shape optimisation of an external mirror. Displayed are RANS-based drag sensitivities in the mirror region of the previous Volkswagen Passat. The colouring is as usual - red regions have to be pushed inwards and blue to be pulled out in order to reduce drag.

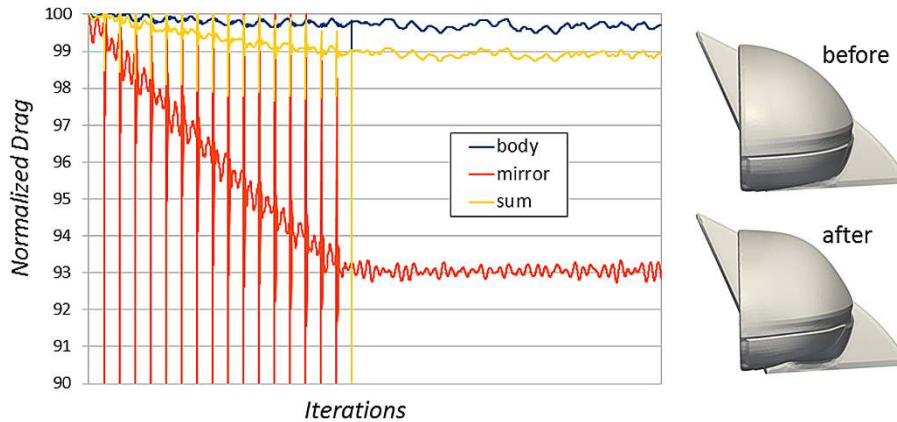


Figure 16 Shape optimisation of an external mirror. Based on the drag sensitivities of Figure 15, the TUM morphing routines were employed to conduct a one-shot shape optimisation. Only the mirror and the triangle mounted to the car were subject to deformation. The evolution of drag on the mirrors themselves (orange curve), the remaining car body (blue) and their sum effect on the entire vehicle (yellow) are shown. Due to a missing constraint on the mirror glass area, the optimisation would ultimately generate inapplicably small mirror glasses and was therefore aborted before that happened. A constrained version of this optimisation will be published in [29]. Note how the morphing affected the overall shape of the mirror (e.g. the bulgier upstream shape and the more horizontal upper edge), but was capable of keeping the essential feature lines of the original design. An illustrative movie of the shape evolution during optimisation is provided as *Additional file 1*.

and every node into its optimal position, they are ideally suited for the optimisation of vehicle details - like external mirrors.

Figure 15 depicts an extract of the drag sensitivity map of the previous Volkswagen Passat. It was used in combination with the TUM morphing routines in order to optimise the shape of the external mirrors. The evolution of drag against the iterative progress of the one-shot optimisation is shown in Figure 16. A remarkable 1% drag reduction of the entire car was achieved by only optimising the mirror shapes and leaving the rest of the car unchanged. The bulgier upstream part and the more horizontal top edge reduce the size of the mirror wake and thus the drag (Figure 16).





Figure 17 The current Audi A7. As simplifications for the primal flow model, the engine bay was completely sealed (i.e. no underhood flow) and the floor was assumed stationary.

3.2 RANS sensitivities vs. approximate time-averaged DES sensitivities

The preceding section summarised the achievements in the development of adjoint-based shape optimisation methods for external aerodynamics of vehicles. These methods rest upon steady-state RANS simulations. Meanwhile, however, the common practice of external aerodynamics computations has moved - at least in the largest part of the Volkswagen Group - from steady-state RANS to unsteady Detached-Eddy Simulations (DES, [31]). To create an adjoint to a time-varying primal like DES entails particular difficulties: Since the adjoint runs backwards in time, all primal states for the time period of interest have either to be stored or recomputed when running the adjoint. The huge amount of storage space and computational time can, in principle, be reduced to a manageable amount by clever checkpointing algorithms [32, 33], but for the typical model sizes of car external aerodynamics an adjoint to a DES run is still not feasible yet.

Our provisional solution to compute sensitivity maps for primal DES computations therefore looks as follows: We plug the time-averaged velocity \mathbf{v} and pressure p from the primal DES computation into a RANS turbulence model in order to solve for a turbulent viscosity ν_t . With these quantities - time-averaged \mathbf{v} and p as well as RANS- ν_t - we run the existing RANS adjoint code to obtain the sensitivities. A validation study similar to the one above for the XLI revealed that we can only expect qualitative accuracy from this procedure. Quantitatively, the sensitivities are not correct. Still, the following comparison will demonstrate that these approximate DES sensitivities are useful and actually have an added value against pure RANS sensitivities.

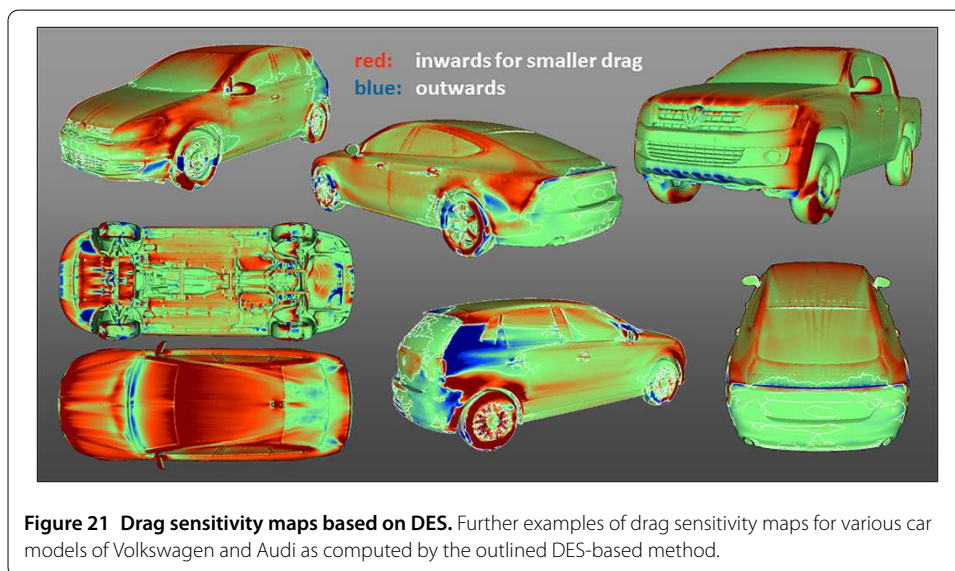

Drag sensitivities for the current Audi A7 (Figure 17) as computed by two different procedures are compared in Figures 18 to 20: RANS adjoint based on a RANS primal on the one hand, and the same RANS adjoint but run on a DES primal as described above on the other hand. For both RANS and DES primal the Spalart-Allmaras turbulence model was employed, while the adjoint made use of the frozen turbulence assumption. On the largest part of the car body, the sensitivities can be seen to be very similar. This does not come as a surprise, because here the flow is attached to the surface and RANS and DES flows do not differ dramatically for attached flows. Significant differences occur in the rear part of the car where the flow is close to detachment: While the RANS-based adjoint fails to predict the experimentally well-confirmed beneficial effect of a rear spoiler, it is clearly indicated by the DES-based adjoint. Moreover, the productive effect of a further 'boat-tailing', i.e. tapering of the car back, is almost missed by the RANS-based adjoint while



clearly present in the DES counterpart. It is this kind of results that - despite the lack of their quantitative accuracy - gave a lot of credibility to our provisional solution of computing sensitivities based on DES primals and supported their integration into the regular computational processes for external aerodynamics (see Figure 21).

3.3 Adjoint methods for flow control


Geometrical modifications, be it via shape or via topology optimisation, are not the only option of improving the fluid dynamic performance of car components or the entire vehicle. Recently there is a growing interest in qualifying active flow control devices for vehicle aerodynamics: Brunn et al. [34] demonstrated how steady blowing and oscillating jets can be successfully applied to reduce drag of the well-known Ahmed body [35]. Bideaux et al. [36] used the same bluff body shape to investigate the effect of a line of pulsed jet actuators



on drag and lift, while Heinemann et al. [37] experimentally investigated continuous jets on a 1:4 model of a passenger production car.

Since flow control of entire vehicles is a rather young subject, experience for the optimal layout of jet configurations is still very limited. The questions to be answered when designing a flow control concept are: (1) What kind of jets should be applied - blowing or suction? (2) Where should the jets be positioned to have a maximum effect? (3) What is the achievable aerodynamic improvement?

To answer these questions, the adjoint method is again the tool of choice, as it allows to compute a sensitivity map of the whole car for the positioning of jets. While for shape sensitivity maps, the design variables were the normal displacements of all surface nodes, it is now the normal flow velocity v_n on each surface element that constitutes the design space. In other words, the car is regarded as having a perforated surface with initially zero normal

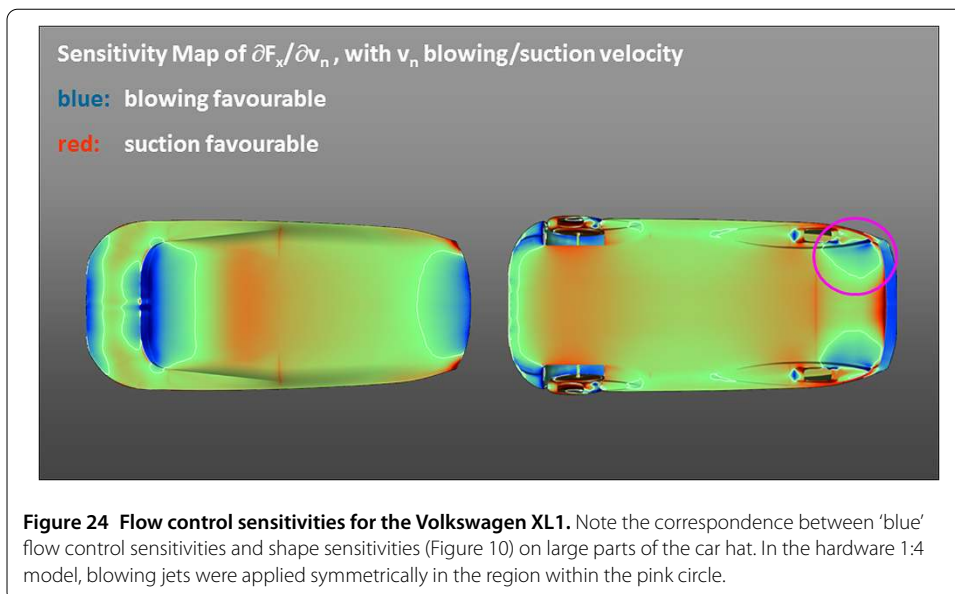
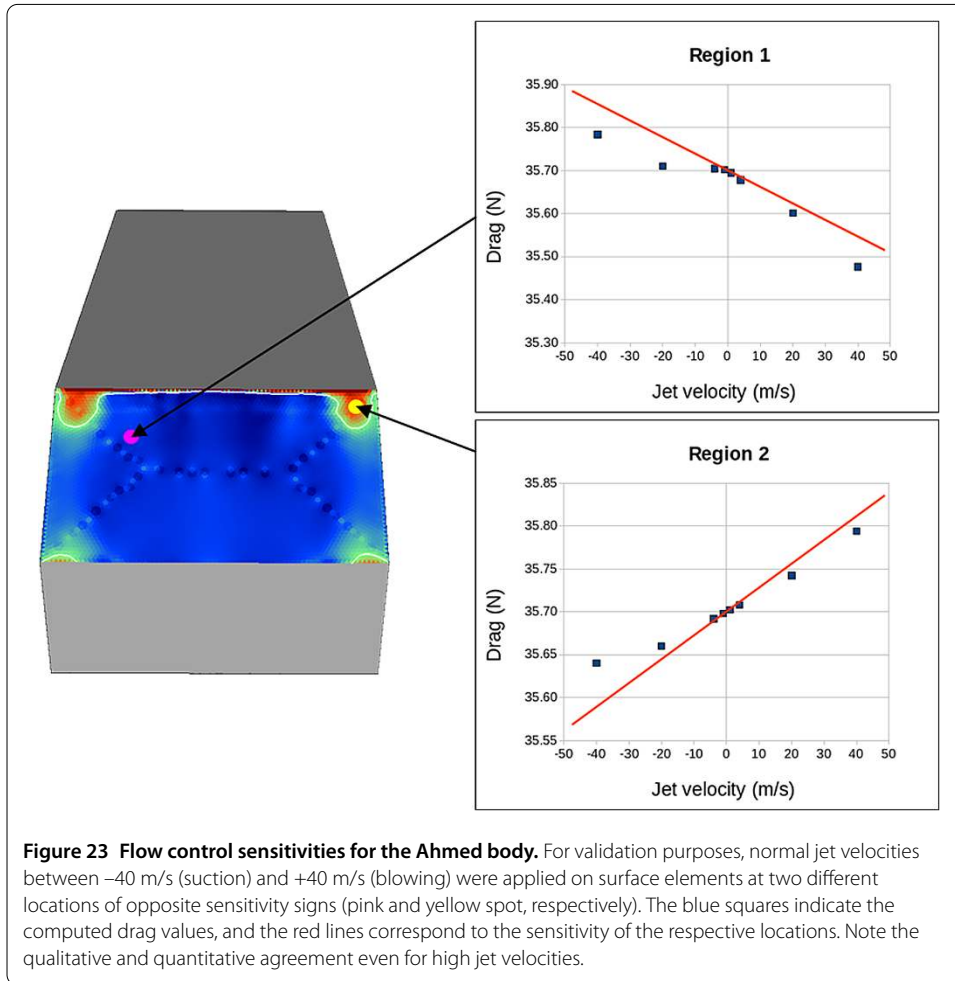


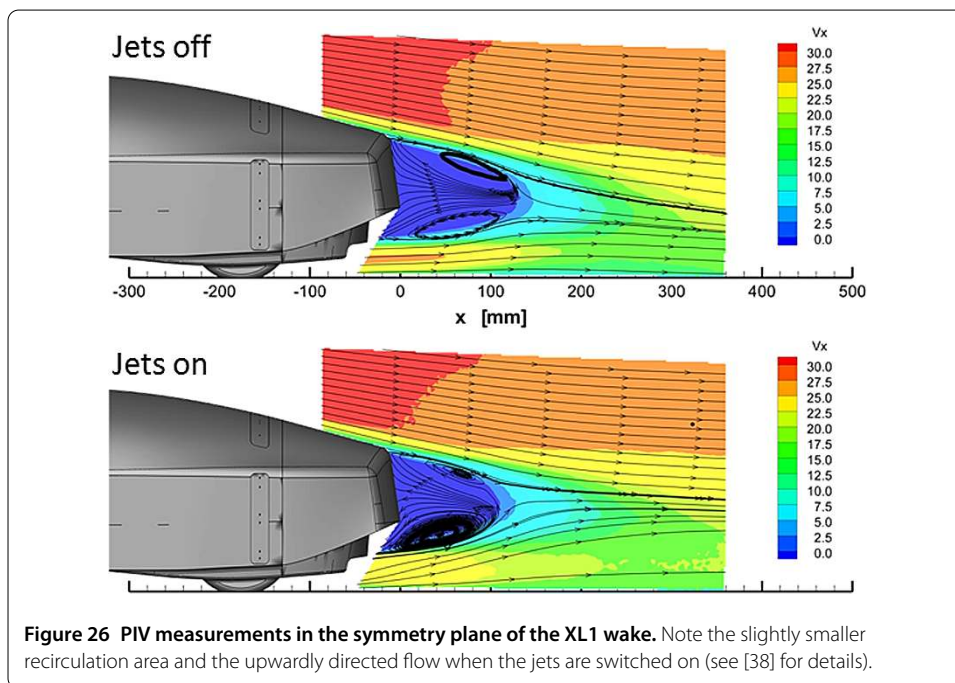

mass flux through it, and the adjoint method is then applied to compute the sensitivity of drag F_x or lift F_z w.r.t. changing the mass flux through each of the ‘holes’ [11].

Figure 22 shows such a sensitivity map for the Ahmed body, i.e. the distribution of $\partial F_x / \partial v_n$. Blue areas are those where blowing ($v_n > 0$) would result in drag reduction, while in red areas a suction jet would help to reduce drag. The colour intensity indicates the magnitude of the effect on the drag F_x . The computation was carried out on a low-Reynolds mesh ($y^+ \approx 1$) using a RANS primal with the Spalart-Allmaras turbulence model and its adjoint counterpart.

The validation of these sensitivities is shown in Figure 23. At two different locations, one in the blue and another in the red region, various normal velocities ranging from -40 m/s (suction) to $+40$ m/s (blowing) were applied in order to compute their effect on drag. As expected, for small jet velocities of up to 5 m/s, the sensitivity captures the gradient $\Delta F_x / \Delta v_n$ very well. But surprisingly, still for higher jet velocities the sensitivity constitutes a good approximation to the actual effect of the jet. In conclusion, flow control sensitivity maps do not only indicate whether a suction jet or a blowing jet is favourable and where to place it, but also give a useful estimate of the achievable effect.

For the Volkswagen XL1, apart from pure shape optimisation as shown above, also the possibility of further drag reduction by blowing jets was investigated. A flow control sensitivity map revealed several promising areas for the application of blowing jets (blue areas in Figure 24). On the car hat, they correspond to a large extent to the blue areas of Figure 10, i.e. where an outward perturbation of the surface was found beneficial. The actual application of flow control, however, concentrated on the underbody - more precisely in a region behind the rear axis (pink circle in Figure 24), where we installed a number of blowing jets. Force measurements, oil film visualisations as well as Particle-Image-Velocimetry (PIV) measurements on a 1:4 model (Figure 25, details see [38]) revealed the effect of these jets on the flow around the XL1: The recirculation area behind the car became smaller, and the flow in the symmetry plane was directed in the upward direction (Figure 26). This was found to be due to a weakening of the longitudinal vortices that emanate from the rear

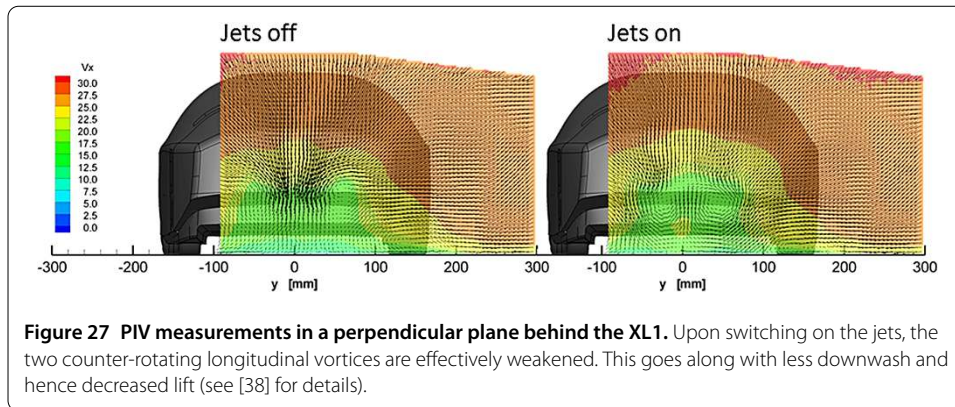




corners of the car (Figure 27). As a result, lift dropped by 20%, and drag - in accordance with the computed sensitivities - decreased by about 2%. Even though this additional drag reduction of 2% is significant for a car like the XL1, it is unfortunately not enough to incorporate a jet system into a car and run it productively.

4 Summary and outlook

Starting from the obstacles of introducing the adjoint method into automotive CFD, we have reported on the efforts undertaken with various partners to develop an industrially viable adjoint solver. Based on OpenFOAM, powerful optimisation methods - topology optimisation for ducted flows, and shape optimisation for ducted as well as external flows - have been developed. While topology optimisation is being used systematically already, the implementation of shape optimisation into the regular automotive development pro-



cess has just started with some initial promising results. Both methods will most surely have an impact on overall vehicle performance and consumption in the future.

The adjoint development efforts have so far concentrated on steady-state phenomena. For an extension towards inherently transient flow applications, like aeroacoustics, flow control with pulsating jets and transient aerodynamics with DES, the adjoint has to become transient as well. This is a subject that is under investigation by several research groups already (e.g. [39–41]) and is currently addressed in the Marie-Curie International Training Network ‘aboutFlow’ [42]. Its industrialisation can be regarded as the next challenge in developing adjoint methods for the automotive industry.

Additional material

Additional file 1: External mirror shape optimisation. How the shape of the external mirror of Figures 15 and 16 evolves from the original to the optimal shape and back again is illustrated by this movie. The colour represents the magnitude of node displacement as compared to the original (details see [28, 29]).

Competing interests

The author declares that they have no competing interests.

Acknowledgements

The author is truly grateful to E de Villiers and his colleagues at Engys, KC Giannakoglou and his team at NTUA (above all EM Papoutsis-Kiachagias and AS Zymaris), E Stavropoulou, M Hojjat and KU Bletzinger from TUM, JD Müller (Queen Mary University London) and the FlowHead team, as well as E Skaperdas and K Haliskos from BETA CAE Systems. The support from many CFD colleagues of the Volkswagen Group in the form of continuous interest in the adjoint method, provision of resources and of valuable feedback is also gratefully acknowledged.

Received: 30 March 2013 Accepted: 5 March 2014 Published: 3 June 2014

References

1. Lions JL: *Optimal Control of Systems Governed by Partial Differential Equations*. New York: Springer; 1971.
2. Pironneau O: **On optimum design in fluid mechanics**. *J Fluid Mech* 1974, **64**:97-110.
3. Jameson A: **Aerodynamic design via control theory**. *J Sci Comput* 1988, **3**:233-260.
4. Löhner R, Soto O, Yang C: **An adjoint-based design methodology for CFD optimization problems**. *AIAA-03-0299* 2003.
5. Nadarajah S, Jameson A: **A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization**. *AIAA-00-0667* 2000.
6. OpenFOAM® - The Open Source CFD Toolbox [http://www.openfoam.com]
7. Othmer C, de Villiers E, Weller HG: **Implementation of a continuous adjoint for topology optimization of ducted flows**. *AIAA-2007-3947* 2007.
8. Othmer C: **A continuous adjoint formulation for the computation of topological and surface sensitivities of ducted flows**. *Int J Numer Methods Fluids* 2008, **58**:861-877.
9. Zymaris AS, Papadimitriou DI, Giannakoglou KC, Othmer C: **Continuous adjoint approach to the Spalart-Allmaras turbulence model for incompressible flows**. *Comput Fluids* 2009, **38**(8):1528-1538.
10. Othmer C, Papoutsis-Kiachagias EM, Haliskos K: **CFD optimization via sensitivity-based shape morphing**. In *Proceedings of the 4th ANSA & μETA Internat. Conf., Thessaloniki, Greece, BETA CAE Systems S.A.* 2011.

11. Zymaris AS, Papadimitriou DI, Giannakoglou KC, Othmer C: **Optimal location of suction or blowing jets using the continuous adjoint approach**. In *Proceedings of the 5th European Conf. on Computational Fluid Dynamics, Lisbon, Portugal, ECCOMAS CFD 2010*.
12. Zymaris AS, Papadimitriou DI, Giannakoglou KC, Othmer C: **Adjoint wall functions: a new concept for use in aerodynamic shape optimization**. *J Comput Phys* 2010, **229**(13):5228-5245.
13. Kontoleonos EA, Papoutsis-Kiachagias EM, Zymaris AS, Papadimitriou DI, Giannakoglou KC: **Adjoint-based constrained topology optimization for viscous flows, including heat transfer**. *Eng Optim* 2013, **45**(8):941-961.
14. de Villiers E, Othmer C: **Multi-objective adjoint optimization of intake port geometry**. *SAE Technical Paper 2012-01-0905* 2012.
15. Helgason E, Krajnovic S: **Aerodynamic shape optimization of a pipe using the adjoint method**. In *Proceedings of the ASME 2012 International Mechanical Engineering Congress & Exposition, Houston, Texas, USA, IMECE 2012*.
16. Lincke A, Rung T: **Adjoint-based sensitivity analysis for Buoyancy-driven incompressible Navier-Stokes equations with heat transfer**. In *Proceedings of the Eighth Internat. Conf. on Engineering Computational Technology, Dubrovnik, Croatia 2012*.
17. Jakubek D, Wagner C: **Shape optimization of train head cars using adjoint-based computational fluid dynamics**. In *Proceedings of the First International Conference on Railway Technology, Las Palmas de Gran Canaria, Spain 2012*.
18. Towara M: **Numerical optimization of an oil intake duct with adjoint topological methods**. *Diploma thesis*. RWTH Aachen University, Dept. of Software and Tools for Computational Engineering; 2011.
19. Hinterberger C, Olesen M: **Automatic geometry optimization of exhaust systems based on sensitivities computed by a continuous adjoint CFD method in OpenFOAM**. *SAE Technical Paper 2010-01-1278* 2010.
20. Hinterberger C, Olesen M: **Industrial application of continuous adjoint flow solvers for the optimization of automotive exhaust systems**. In *Proceedings of the ECCOMAS Thematic Conf. - CFD & Optimization, Antalya, Turkey, ECCOMAS 2011*.
21. Giannakoglou KC: **Continuous adjoint methods in shape, topology, flow-control and robust optimization**. In *Open Source CFD International Conference, London, UK, ICON-CFD 2012*.
22. Bendsoe MP, Sigmund O: *Topology Optimization: Theory, Methods and Applications*. Berlin: Springer; 2004.
23. Borrvall T, Petersson J: **Topology optimization of fluids in Stokes flow**. *Int J Numer Methods Fluids* 2003, **41**:77-107.
24. Moos O, Klimetzek F, Rossmann R: **Bionic optimization of air-guiding systems**. *SAE Technical Paper 2004-01-1377* 2004.
25. Gersborg-Hansen A, Sigmund O, Haber R: **Topology optimization of channel flow problems**. *Struct Multidiscip Optim* 2005, **30**:181-192.
26. Othmer C, Grahs T: **Approaches to fluid dynamic optimization in the car development process**. In *Proceedings of the EUROGEN Conference, Munich, Germany 2005*.
27. Othmer C, Kaminski T, Giering R: **Computation of topological sensitivities in fluid dynamics: cost function versatility**. In *Proceedings of the ECCOMAS CFD Conference, Delft, The Netherlands 2006*.
28. Stavropoulou E, Hojjat M, Bletzinger KU: **In-plane mesh regularization for node-based shape optimization problems**. *Comput Methods Appl Mech Eng*, in press [<http://www.sciencedirect.com/science/article/pii/S004578251400070X>]
29. Hojjat M, Stavropoulou E, Bletzinger KU: **The vertex morphing method for node-based shape optimization**. *Comput Methods Appl Mech Eng* 2014, **268**:494-513.
30. ANSA® pre-processor – BETA CAE Systems S.A. [<http://www.beta-cae.gr/ansa.htm>]
31. Islam M, Decker F, de Villiers E, Jackson A, Gines J, Grahs T, Gitt-Gehrke A, Comas i Font J: **Application of detached-Eddy simulation for automotive aerodynamics development**. *SAE Technical Paper 2009-01-0333* 2009.
32. Griewank A: **Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation**. *Optim Methods Softw* 1992, **1**:33-54.
33. Wang Q, Moin P, Iaccarino G: **Minimal repetition dynamic checkpointing algorithm for unsteady adjoint calculation**. *SIAM J Sci Comput* 2009, **31**(4):2549-2567.
34. Brunn A, Wassen E, Sperber D, Nitsche W, Thiele F: **Active drag control for a generic car model**. *Notes Numer Fluid Mech Multidiscip Des* 2007, **95**:247-259.
35. Ahmed SR, Ramm R, Falting G: **Some salient features of the time-averaged ground vehicle wake**. *SAE Technical Paper 84-0300* 1984.
36. Bideaux E, Bobillier P, Fournier E, Gillieron P, El Hajem M, Champagne JY, Gilotte P, Kourta A: **Drag reduction by pulsed jets on strongly unstructured wake: towards the square back control**. *Int J Aerodyn* 2011, **1**(3/4):282-298.
37. Heinemann T, Springer M, Lienhart H, Kniesburgess S, Becker S: **Active flow control on a 1:4 car model**. In *Proceedings of the 16th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal 2012*.
38. Petzold R: **Experimentelle Untersuchungen zur Strömungsbeeinflussung an einem 1:4-Fahrzeugmodell**. *Diploma thesis*. TU Braunschweig, Inst. für Strömungsmechanik; 2011.
39. Torbert S, Löhner R: **Checkpointing schemes for adjoint methods and strongly unsteady flow**. *AIAA-2012-0572* 2012.
40. Economou TD, Palacios F, Alonso JJ: **A coupled-adjoint method for aerodynamic and aeroacoustic optimization**. *AIAA-2012-5598* 2012.
41. Carnarius A, Thiele F, Özkaya E, Nemili A, Gauger N: **Optimal control of unsteady flows using a discrete and a continuous adjoint approach**. In *Proceedings of the 25th Conference on System Modeling and Optimization, Berlin, Germany; 2011*:318-327.
42. aboutFlow - **Adjoint-based optimisation of industrial and unsteady flows** [<http://aboutflow.sems.qmul.ac.uk>]

doi:10.1186/2190-5983-4-6

Cite this article as: Othmer: Adjoint methods for car aerodynamics. *Journal of Mathematics in Industry* 2014 **4**:6.