Technical Paper

20164077

Adjoint Optimization for Vehicle External Aerodynamics

Georgios K. Karpouzas [1)] Evangelos M. Papoutsis-Kiachagias [2)] Thomas Schumacher [3)]
Eugene de Villiers [4)] Kyriakos C. Giannakoglou [5)] Carsten Othmer [6)]

1) ・3) ・4) Engys Ltd., Studio 20, RVPB, John Archer Way, London, SW18 3SX, U.K.
(E-mail: g.karpouzas@engys.com)
2) ・5) National Technical University of Athens, School of Mechanical Engineering, Parallel CFD & Optimization Unit, Heroon Polytechniou 9, Athens 15780, GREECE
6) Volkswagen AG, Group Research, Letter Box 1777, K-EFFG/V, D-38436 Wolfsburg

Received on June 15, 2015

Presented at the JSAE Annual Congress on May 20, 2015

ABSTRACT: Adjoint optimization is an exciting and fast growing field that has many applications in the automotive industry. We focus on the usage of adjoint for the optimization of aerodynamic performance. While adjoint methods have come to the attention of mainstream CFD through inclusion in prominent commercial codes, most of the available tools are severely limited, precluding productive use in this field. We detail a methodology that is based on the continuous adjoint method and is implemented in an Open Source framework. While more mathematically demanding in terms of its derivation, the continuous adjoint method requires orders of magnitude fewer resources without sacrificing significant accuracy. In this paper continuous adjoint methods are used for calculating gradients of aerodynamic objective functions (drag, lift, moments etc.) in applications with a huge number of design variables. Methodologies that accept either steady-state RANS[(1)(2)(3)] or time averaged DES[(4)(6)] as primal flow input are outlined and extensions to improve the accuracy of previously published methods are detailed. Finally, a novel methodology, based on volumetric B-splines to translate the surface sensitivities produced by the adjoint into optimized shapes, is introduced and showcased.

KEY WORDS: **Heat・Fluid, Aerodynamic Performance,** Optimization, Detached Eddy Simulation, Adjoint [D1]

1. Introduction

Computational Fluid Dynamics (CFD) is central to large portions of automotive design. Not only is it used prominently for the prediction of external aerodynamic performance but it also sees extensive use in a diverse set of applications including (but not limited too) HVAC (heating ventilation and cooling), combustion, exhausts, after treatment, air intakes, thermal management, head-lights, air bags, aero-acoustics, fuel supply, and the design of pumps/fans/blowers.

Much of the design process attached to these components and systems is still dependent on manual iterations between designer and computational analysts. Even in instances where automated stochastic optimization loops are in place, these are severely constricted by the fact that the costs are proportional to the number of parameters to be optimized. This places a limit on the degrees of freedom available to the optimization process and constrains the degree of optimization that can be achieved within the allowable parameter space. Further, the conceptually simple step of parameterizing the input CAD necessary for such a process can be a stumbling block in and of itself, especially considering the complexity of many of the components.

One of the fastest growing fields in CFD research is adjoint optimization. It is based on the Adjoint method, which allows for very efficient calculation of sensitivities of the objective function with respect to design variables. In fact, the computational effort in typical adjoint systems is dependent only on the number of objectives (drag, lift, etc.) rather than the number of design variables (surface node displacement). This makes adjoint methods ideal for optimizing problems with large numbers of variables and/or large design spaces.

In the context of the adjoint method, there are two main variants: discrete adjoint and continuous adjoint. In the discrete approach, the discretization of the partial differential equations happens before the adjoint differentiation. This can be done manually or via algorithmic differentiation tools such a TAPENADE[(7)]. In the continuous approach, the adjoint equations are formulated by differentiating the partial differential equations representing the flow (primal problem) directly. These “adjoint equations” are then discretized and solved. The two approaches perform similarly when applied to simple canonical test cases: the calculation of the sensitivity derivatives is fast, the procedure converges rapidly and the whole process is relatively straight-forward. However, in industrial scale applications the solution of the adjoint problem becomes significantly more challenging.

The discrete adjoint can theoretically provide the “exact” gradient w.r.t. the design variables for any objective specified, but it requires large amount of memory and computational cost. More precisely, all the intermediate states during the discrete calculation have to be stored in the memory. In discrete terminology, these intermediate states are referred to as the “tape”. Techniques such as revolve (check-pointing) have been developed to reduce the memory requirements but this significantly increases the computational cost of the method. The net effect is that the cost of discrete adjoint will typically be an order of magnitude or more in

excess of the primal solution, and can often place unmanageable burdens on the computational infrastructure in terms of memory requirements.

In the continuous adjoint method on the other hand, the bottle-neck lies in the analytical differentiation of the flow equations, which can be complicated and time-consuming. This amount of effort will depend on the complexity of the primal equation system, but the full derivation only has to be performed once for each unique equation set. The memory requirements and computational cost is however similar to the solution of the primal equations. This makes the continuous approach ideal for industrial scale applications with intermediate to large computational meshes. Further, the derivation process that produces the continuous adjoint equations and the familiarity of the resulting equations and algorithms promotes a deeper understanding of the solution system. This in turn can be of significant benefit in the development of efficient adjoint solvers, as will be shown in this publication.

2. Methodology

2.1. Continuous Adjoint

In this paper, we demonstrate the use of the continuous adjoint method for efficient shape optimization of vehicle external aerodynamic problems. While transient adjoint formulations are possible, they are much more complex and costly to implement and maintain. We thus initially restrict ourselves to steady state: the primal flow RANS (Reynolds-Averaged-Navier-Stokes) equations are:

$$R^p = -\frac{\partial v_i}{\partial x_i} = 0 \quad (1)$$

$$R_i^v = v_j \frac{\partial v_i}{\partial x_j} + \frac{\partial p}{\partial x_j} - \frac{\partial}{\partial x_j}\left[(\nu + \nu_t)\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i}\right)\right] = 0 \quad (2)$$

$$R_i^z = \text{Convection} + \text{Diffusion} + \text{Production} + \text{Disipation} = 0 \quad (3)$$

where v_i is the primal velocity, p is the primal pressure, ν and ν_t are the kinematic and turbulent kinematic viscosity, respectively. R_i^z is considered to be any turbulence model with z_i representing the multicomponent turbulence vector.

Now, in order to derive the adjoint equations, let us start by defining an objective function F. The objective can be defined as a combination of surface and volume integrals.

$$F = \int_S F_S \, dS + \int_\Omega F_\Omega \, d\Omega \quad (4)$$

F is then augmented (extended) by the state equations, R^p and R_i^v.

$$F_{aug} = F + \int_\Omega q R^p \, d\Omega + \int_\Omega u_i R_i^v d\Omega \quad (5)$$

Here, q and u_i are the adjoint variables and due to the way they enter the solution algorithm can be interpreted as "adjoint pressure" and "adjoint velocity", respectively. In this example, derivation, the turbulent kinematic viscosity (ν_t) is assumed constant with respect to changes in the design variables (frozen turbulence assumption). This assumption can have an effect on sensitivity accuracy but in many instances the reduced complexity that results makes the assumption justifiable. The morphing cases in section 3, however, include the adjoint to the turbulence model equations.

Differentiating the augmented cost function, F_{aug}, (see[1] for details) produces the adjoint equations for incompressible flow with frozen turbulence:

$$R_\Omega^q = \frac{\partial u_i}{\partial x_i} - \frac{\partial F}{\partial p} = 0 \quad (6)$$

$$R_{i\Omega}^u = -v_j \frac{\partial u_i}{\partial x_j} + u_j \frac{\partial v_j}{\partial x_i} + \frac{\partial q}{\partial x_j} - \frac{\partial}{\partial x_j}\left[(\nu + \nu_t)\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)\right] + \frac{\partial F}{\partial v_i} = 0 \quad (7)$$

The adjoint equations are very similar to the primal equations (eqs. (1-3)). They both have convection, diffusion, a gradient of pressure and (possibly) a source term. The main difference is the additional term ($u_j \partial v_j / \partial x_i$), which appears in the adjoint momentum equations. This term is referred to as the Adjoint Transpose Convection (ATC).

After solving the adjoint equations, the surface sensitivities of the objective function with respect to the surface normal motion of the surface nodes (design variables) can be calculated with the following expression:

$$\frac{\delta F_{aug}}{\delta b} = -\int_{S_w} \left[(\nu + \nu_t)\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) - q n_i\right] \frac{\partial v_j}{\partial x_k} \frac{\partial x_k}{\partial b_m} dS \quad (8)$$

In shape optimization problems, the surface sensitivities are used as an input to a deformation tool, which morphs the computational domain. This process is typically iterative and there are many different methods that can be used to achieve the shape modification. These range from two-way coupling of the adjoint problem to parameterized CAD to direct deformation of the surface nodes.

2.2. Verification

In order to verify the basic methodology, we apply a naïve morphing operator to a previously published model vehicle – the DrivAer Estate[5]. In this case, the objective is drag minimization

and the primal is a steady state RANS using the Spalart-Allmaras turbulence model.

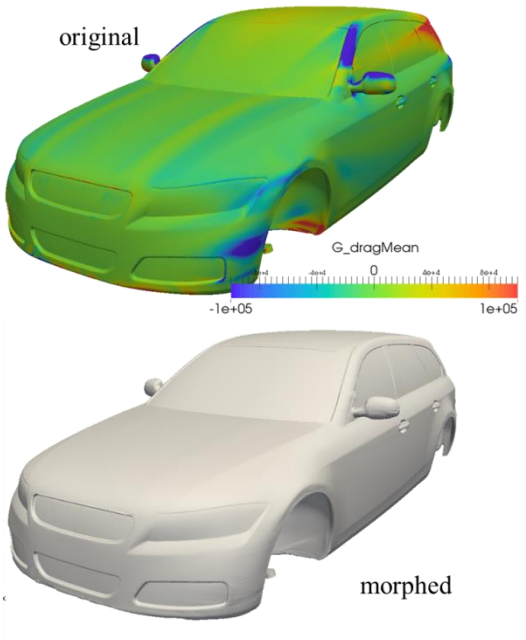

Fig. 1: DrivAer Estate showing the original model with adjoint surface sensitivities (blue = push in to reduce drag and red = pull out) and the morphed geometry. Wheels are omitted for clarity.

We apply a simple hyperbolic tangent filter to the adjoint surface sensitivities that limits the maximum single-step displacement to 2mm. Figure(1) displays the surface sensitivities generated for the case along with the result of the morphing operation. The morphed geometry was re-meshed and drag coefficient was found to reduce from C_D = 0.356 to 0.328 (~8%).

The difference between the two geometries is minimal and difficult to identify from the image. In this context, the large improvement achieved in the objective with such a small modification is all the more surprising. Unfortunately, morphing without constraints will not in general result in a manufacturable design, which means this approach has limited utility as an actual design tool. It does, however, showcase the effectiveness of the adjoint method in the context of a realistic vehicle geometry. Further and more rigorous validation of the method can be found in [1] [2].

2.3. Linearized ATC

The ATC term identified in sec. 2.1 is the main source of difficulty when attempting to solve the continuous adjoint equations with traditional segregated algorithms such as SIMPLE (Semi-Implicit Method for Pressure Linked Equations). The ATC cross couples the components of the adjoint velocity, which means it cannot be implemented implicitly in a segregated algorithm. The contribution thus has to be lagged, making it fully explicit. This results in weak coupling between the adjoint pressure and velocity, which in combination with imperfect cell quality or high surface normal gradients can lead to gradual self-reinforcing divergence of the solution.

In aerodynamic calculation, these conditions typically occur near the wall. The general practice in such cases is to limit the contribution of the ATC in the near-wall cells and to employ first order convection for the implicit part of the adjoint convection. The additional numerical dissipation introduced by the use of the upwind scheme and the excision of the strong source regions is typically enough to stabilize the calculation. The obvious drawback, however, is that the upwind scheme is less accurate than the second order schemes typically used for RANS primal calculations and, secondly, that by masking the near-wall contributions to the ATC, we are removing some of the largest source terms from the equation. A method that preserves second order accuracy while maintaining stability is thus highly desirable.

The obvious long term solution is to implement the primal and adjoint equations in a block coupled framework that allows implicit implementation of the full equation set. In the current context, an approximate approach based on the selective masking of instability inducing flow regions has been developed. A linear surrogate model is constructed for the ATC contribution to the matrix. The main purpose of this model is to identify the cells in which the ATC would give negative contributions in the diagonal of the matrix (left hand side) if it was implicit.

$$ATC_i = u_j \frac{\partial v_j}{\partial x_i} \quad (9)$$

$$A = ATC_i \frac{u_i}{u^2} \quad (10)$$

A is the projected diagonal contribution for an implicit ATC term. If A is negative, the ATC will reduce the matrix' diagonal dominance and make the solution more unstable. As a result, the system will become ill-posed and potentially diverge. Using this indicator, the ATC contribution can be damped on a cell-by-cell basis so that diagonal dominance of the matrix, and by inference stability, is guaranteed. This approach results in a minimal reduction in the ATC source term and allows the use of second order schemes for the adjoint convection term. Without this targeted damping a strong much more global damping of the ATC is required to ensure stable second order operation.

Figure(2) shows a comparison of drag minimization surface sensitivities for second order convection with linearized ATC and first order convection with global near-wall damping for the DrivAer Sedan model[5]. There are clear qualitative and quantitative differences. The second order method produces smoother fields with much more detailed small features. The overall trends are similar (as expected), but there are clear regional differences in the rear spoiler and front bonnet regions, which implies the use of near-wall ATC damping plus first order upwind for the adjoint convection can lead to suboptimal designs.

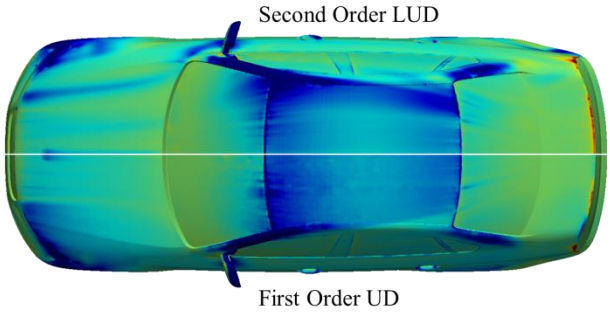

Fig. 2: Comparison of surface drag sensitivities for second order LUD plus adaptive linearised ATC (top) and first order UD with fixed near-wall blending (bottom). DrivAer sedan model [5].

2.4. Time Averaged DES Adjoint

The steady state RANS adjoint approach has shown very positive results over a large range of applications[1][4][6]. The method however suffers from a shortcoming that is inherent to the underlying primal RANS equations in that many problems are not accurately described by this approach. Thus, any method that relies on RANS to generate sensitivities (adjoint or stochastic) will suffer from the same inaccuracies.

The DES (Detached Eddy Simulation) method is an approach that directly resolves most of the important scales of motion rather than approximating these scales with a turbulence model. Due to the reduced modelling dependence, it can be much more accurate and consistently accurate than any RANS based method[8]. The method is however inherently transient and as mentioned previously, transient adjoint is much more costly and cumbersome to evaluate than the steady equivalent. If we are however, only interested in the time averaged aerodynamic properties of the vehicle, i.e. the mean drag, then an interesting possibility presents itself: to formulate a steady adjoint for the time averaged transient primal flow.

To derive such an adjoint let us start from the basic LES/DES equations for an incompressible flow. In what follows, we neglect the continuity equation, since it's time-averaged form is identical to its instantaneous form. The turbulence quantities are assumed frozen to simplify the derivation (this might not be the case for actual simulations). The focus is then on the standard LES/DES momentum equation:

$$\frac{\partial \overline{v_i}}{\partial t} + \frac{\partial \overline{v_i}\,\overline{v_j}}{\partial x_j} + \frac{\partial \overline{p}}{\partial x_j} \qquad (11)$$

$$-\frac{\partial}{\partial x_j}\left[(\nu + \nu_{SGS})\left(\frac{\partial \overline{v_i}}{\partial x_j} + \frac{\partial \overline{v_j}}{\partial x_i}\right)\right] = 0$$

The over-bar denotes a spatial filtering operation and ν_{SGS} represents the sub-grid scale turbulence viscosity, which is typically much smaller than a RANS based turbulence viscosity. In order to recover a steady flow field, we apply a time-averaging operator to eq. (15) denoted by the caret (^). The time-averaged DES equation results:

$$\frac{\partial \widehat{\overline{v_i}}}{\partial t} + \frac{\partial \widehat{\overline{v_i}\,\overline{v_j}}}{\partial x_j} + \frac{\partial \widehat{\overline{p}}}{\partial x_j} \qquad (12)$$

$$-\frac{\partial}{\partial x_j}\left[\widehat{(\nu + \nu_{SGS})}\left(\frac{\partial \overline{v_i}}{\partial x_j} + \frac{\partial \overline{v_j}}{\partial x_i}\right)\right] = 0$$

Assuming the difference operators are constant in time, the flow is statistically steady and the averaging is long enough, several simplifications can be made:

$$\frac{\partial \widehat{\overline{v_i}}}{\partial t} \to 0, \frac{\partial \widehat{\overline{p}}}{\partial x_j} = \frac{\partial \widehat{\overline{p}}}{\partial x_j} \qquad (13)$$

$$\frac{\partial \widehat{\overline{v_i}}\,\widehat{\overline{v_j}}}{\partial x_j} = \frac{\partial \widehat{\overline{v_i}}\,\widehat{\overline{v_j}}}{\partial x_j} - \frac{\partial T_{ij}}{\partial x_j} \qquad (14)$$

$$T_{ij} = \widehat{\overline{v_i v_j}} - \widehat{\overline{v_i}}\,\widehat{\overline{v_j}} \qquad (15)$$

For the viscous term we assume that:

$$\nu_{SGS}\left(\widehat{\frac{\partial \overline{v_i}}{\partial x_j}} + \frac{\partial \overline{v_j}}{\partial x_i}\right) \approx \widehat{\nu_{SGS}}\left(\frac{\partial \widehat{\overline{v_i}}}{\partial x_j} + \frac{\partial \widehat{\overline{v_j}}}{\partial x_i}\right) \qquad (16)$$

and collect the different into the resolved Reynolds stress tensor T_{ij}. Substituting eqs. (13-16) back into (12), the time-averaged DES momentum equation becomes:

$$\widehat{R_i^{\bar{v}}} = \frac{\partial \widehat{\overline{v_i}}\,\widehat{\overline{v_j}}}{\partial x_j} - \frac{\partial \widehat{\overline{p}}}{\partial x_j} - \qquad (17)$$

$$\frac{\partial}{\partial x_j}\left[(\nu + \widehat{\nu_{SGS}})\left(\frac{\partial \widehat{\overline{v_i}}}{\partial x_j} + \frac{\partial \widehat{\overline{v_j}}}{\partial x_i}\right) + T_{ij}\right] = 0$$

All the time-averaged quantities, $\widehat{\overline{v_i}}$, $\widehat{\overline{p}}$, $\widehat{\nu_{SGS}}$ and T_{ij} can be calculated during the primal run. The averaged resolved Reynolds stress T_{ij} is not however differentiable, so in order to include its effect in the adjoint, we need to define a more implicit approximation. Taking a cue from RANS and as a first approximation, we formulate the stress as the product of the strain with some "turbulent" viscosity, ν_t, plus a residual stress.

$$\nu_{SGS}\left(\widehat{\frac{\partial \overline{v_i}}{\partial x_j}} + \frac{\partial \overline{v_j}}{\partial x_i}\right) + T_{ij} = \nu_t\left(\frac{\partial \widehat{\overline{v_i}}}{\partial x_j} + \frac{\partial \widehat{\overline{v_j}}}{\partial x_i}\right) + Q_{ij} \qquad (18)$$

where Q_{ij}, is the uncorrelated residual stress assumed to be independent of the mean flow and design variables (frozen assumption). There is a degree of freedom in choosing ν_t as any residual will simply be absorbed into Q_{ij}. However, choosing a model that minimizes Q_{ij} in the context of the formalism will lead to a more implicit and, therefore, accurate adjoint.

The easiest approach is to find ν_t by simply solving a standard RANS model with the time-averaged primal velocity $\widehat{\overline{v_i}}$, as input. The adjoint equations for the time-averaged DES primal are identical to those for the steady RANS primal, other than that time-averaged inputs are used.

Figure(3) shows a comparison between the surface sensitivities on an Audi A7 of a RANS steady state adjoint and a time-averaged DES adjoint using the Spalart-Allmaras model to calculate the turbulent viscosity. Note the similarity between RANS- and DES-based drag sensitivities for the largest part of the car roof and the dramatic difference at the rear part. While the RANS computation misses the favorable effect of a rear spoiler, it is clearly present in the DES results. Tests with a variable spoiler on the actual vehicle have corroborated the DES sensitivity result.

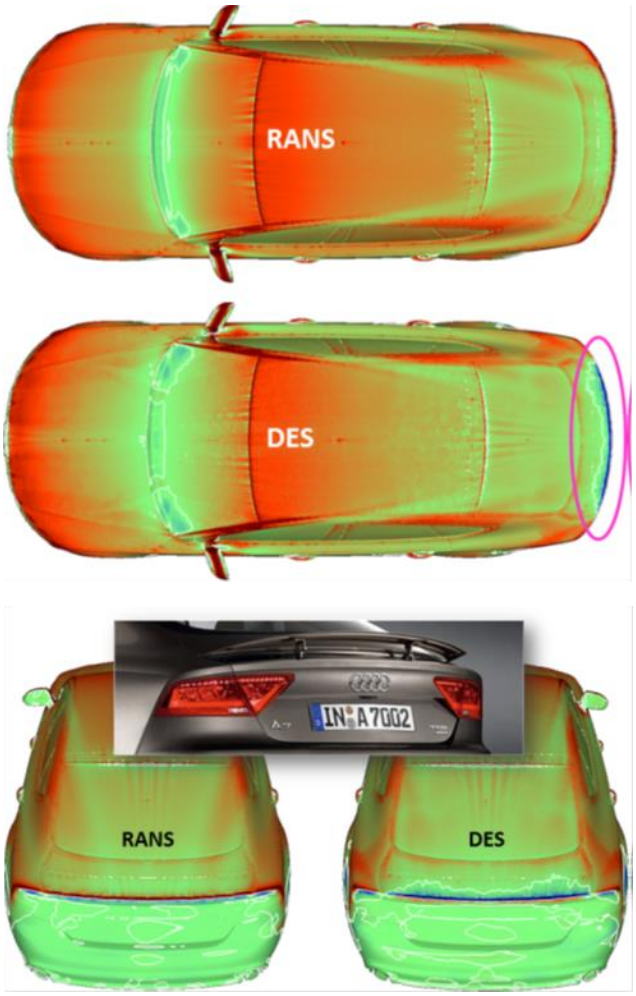

Fig. 3: RANS vs. approximate DES sensitivities for Audi A7. The actual shape of the active spoiler in the current Audi A7 is shown in the bottom of the image.

2.4.1 Tensorial Viscosity

Using a scalar viscosity calculated via a RANS turbulence model might be straight forward but it completely neglects the availability of the time averaged resolved Reynolds stress (T_{ij}). In order to reduce the magnitude of the residual stress , Q_{ij},. we can instead derive a tensorial viscosity. Choosing a tensor coefficient provides more degrees of freedom when trying to approximate T_{ij}, since we do not have to assume isotropy nor alignment of the stress and strain rate. The proposed method uses the same relation as eq. (18) but instead of a scalar viscosity (ν_t) we calculate a tensorial viscosity (N) from the averaged strain ($\hat{S}_{ij}$), averaged SGS stresses ($\hat{B}_{ij}$) and averaged resolved Reynolds stresses (T_{ij}). The necessary averaged quantities are calculated with the following formulas:

$$\hat{B}_{ij}{}^{n+1} = a\hat{B}_{ij}{}^{n} + (1-a)B_{ij} \quad (19)$$

$$\hat{S}_{ij}{}^{n+1} = a\hat{S}_{ij}{}^{n} + (1-a)S_{ij} \quad (20)$$

$$T_{ij}{}^{n+1} = a\left(T_{ij}{}^{n} + \bar{v}_i{}^{n}\bar{v}_j{}^{n}\right) + (1-a)v_iv_j - \bar{v}_i{}^{n+1}\bar{v}_j{}^{n+1}$$
$$a = \frac{t-\Delta t}{t} \quad (21)$$

where t is the current time and Δt is the time step. The total mean stress, E_{ij}, can be found from:

$$E_{ij} = \hat{B}_{ij} + \mathrm{T}_{ij} = N_{ik}\hat{S}_{kj} + Q_{ij} \quad (22)$$

We still need to account for a residual stress Q_{ij}, since it entirely possible for a mean stress to exist in the absence of mean strain. We seek a value of N_{ij} such that Q_{ij}, is minimized while avoiding singularities. If the strain $\hat{S}$ is invertible (i.e. $\hat{S}_{ij}{}^{-1}$ exists), which is the case for the majority of the mesh elements, then N_{ij} is can be calculated directly.

$$N_{ij} = E_{ij}\hat{S}_{ij}{}^{-1} \quad (23)$$

The process becomes more problematic if $\hat{\mathrm{S}}$ is one or two dimensional or $\hat{\mathrm{S}}$ and E_{ij} are uncorrelated. In such elements, a tensor component analysis is applied in order to invert the $\hat{\mathrm{S}}$ tensor with minimum amount of modification. The transformation analysis is based on the fact that the mean strain is a symmetric tensor. Symmetric tensors are diagonalizable and have orthogonal eigenvectors.

$$E_{ij}{}^{*} = gE_{ij}g^{T} \quad (24)$$

$$\hat{S}_{ij}{}^{*} = \lambda\hat{S}_{ij}\lambda^{T} \quad (25)$$

After making $\hat{\mathrm{S}}$ diagonal, the elements with 1D and 2D strains will have zero or near-zero entries in one or more positions on the diagonal. In the locations that the diagonal of $\hat{\mathrm{S}}_{\mathrm{ij}}{}^{*}$ is not zero or very small, the diagonal of $\mathrm{N}_{\mathrm{ij}}{}^{*}$ is calculated.

$$N_d{}^{*} = E_d{}^{*}/\hat{S}_d{}^{*} \quad (26)$$

For degenerate entries heuristics are used to find an approximate value for the corresponding diagonal element to produce $\mathrm{N}_{\mathrm{ij}}{}^{**}$ and $\hat{\mathrm{S}}_{\mathrm{ij}}{}^{**}$ is made spherical using the mean of the finite diagonal entries. The resulting stress tensor is calculated.

$$E_{ij}{}^{**} = N_{ik}{}^{**}\hat{S}_{kj}{}^{**} \quad (27)$$

The $\mathrm{E}_{\mathrm{ij}}{}^{**}$ and the $\hat{\mathrm{S}}_{\mathrm{kj}}{}^{**}$ are transformed using the eigenvectors g and λ, respectively.

$$E_{ij}^{***} = g^{-1} E_{ij}^{**} g^{T^{-1}} \quad (28)$$

$$\hat{S}_{ij}^{***} = \lambda^{-1} \hat{S}_{ij}^{**} \lambda^{T^{-1}} \quad (29)$$

The new tensorial viscosity is calculated by inverting the $\hat{S}_{ik}^{***}$ tensor.

$$N_{ij}^{***} = E_{ij}^{***} \hat{S}_{ij}^{***^{-1}} \quad (30)$$

The resulting tensorial viscosity can be used directly in the adjoint equations. Figure(4) shows a comparison of the viscosity generated by the new scheme with that from the Spalart-Allmaras RANS equation. Pronounced differences can be observed. The effects on the adjoint sensitivities are significant but do not change the general character of the results. The validation of the new method will be detailed in future publications.

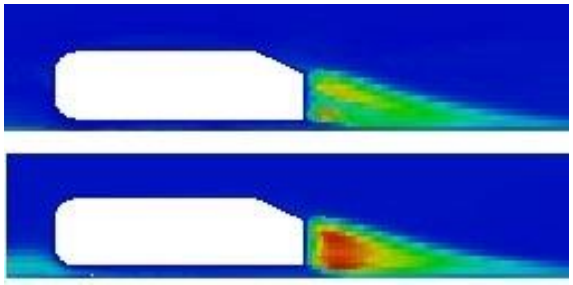

Fig. 4: Comparison of tensorial viscosity magnitude (top) with RANS viscosity for a time averaged DES flow field (bottom). Ahmed body 35° rear slant.

3. Morphing

This section presents the coupling of the continuous adjoint method with a morphing tool for shape optimization in car aerodynamics.

3.1. Morphing Methodology

All cases that follow use a mesh parameterization and displacement strategy based on volumetric B-splines, which can be seen as a Free Form Deformation (FFD) method, and its coupling with the continuous adjoint solver. The method uses a set of control points in 3D space, in the form of a structured control grid. CFD mesh points residing inside the boundaries of the control grid are displaced following the displacement of the control grid points. The use of this software in aerodynamic shape optimization is twofold: (a) for the parameterization of the surface of an aerodynamic body by defining arbitrary control points in 3D space, to be used as the design variables and (b) for the displacement of the surface and volume nodes of the CFD mesh within each optimization cycle. The method exhibits great potential since the cost of each mesh displacement is extremely small, the minimum degree of surface continuity can be defined a-priori and the setup of each case is not cumbersome.

The version of the continuous adjoint solver employed for the morphing does not neglect variations in turbulent viscosity: both the mean–flow and turbulence equations are differentiated. The differentiation has been performed for the most widely used turbulence models (Spalart–Allmaras[1], k-ϵ[2] and k-ω SST[3] models). While it is case dependent, the omission of the turbulence differentiation can result in incorrect or even incorrectly signed sensitivities. This will, in turn, result in an erroneous optimization outcome[4].

The steps of the shape optimization algorithm are listed below:

(a) Define the control grid to enclose the part of the geometry to be optimized. The designer may choose to increase the control points number or lower the basis functions degree, leading to more localized (but less smooth) geometry changes.
(b) Locate the CFD mesh points that reside within the boundaries of the control grid.
(c) Compute the parametric coordinates of the points found in the previous step. The computational cost of this step increases with the number of control points and the number of mesh points to be parameterized.
(d) Solve the flow equations.
(e) Compute the objective function value and apply the termination criterion.
(f) Solve the adjoint equations.
(g) Compute the objective function gradient w.r.t the boundary CFD mesh nodes to be displaced, $\delta F / \delta x_m$, where x_m are the coordinates of the surface nodal points.

Project the surface sensitivities to control points, in order to compute the control points sensitivities, by applying the chain rule

$$\frac{\delta F}{\delta b_l} = \frac{\delta F}{\delta x_m} \cdot \frac{\delta x_m}{\delta b_l} \quad (31)$$

where b_l are the control point coordinates.

(h) Update the control point positions.
(i) Compute the new surface and volume mesh points positions through the volumetric B-splines equations, using the already computed parametric coordinates associated with each one of them.
(j) Move to step (d).

Three test cases are presented.

3.2. Side Mirror Morphing

The first case is the optimization of a side-mirror of a passenger car, with mirror-induced drag minimization being the target objective. The Spalart-Allmaras turbulence model with wall functions and its adjoint solver[4], were used. The starting and final mirror geometries are shown in Figure(5). The mirror-induced drag is reduced by ~7%.

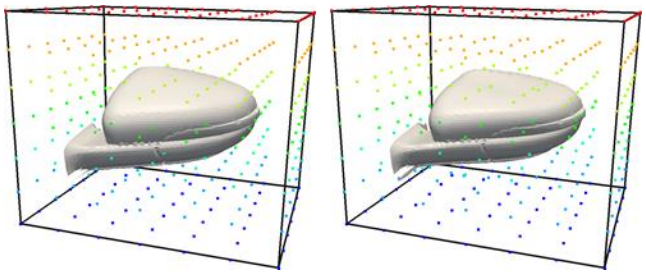

Fig. 5: Boundary of control box and control points coloured based on their z coordinate, for the initial (left) and optimized (right) geometries.

3.3. Vehicle Rear Spoiler Morphing

The second case deals with the drag minimization of the fast-back configuration of the DrivAer car geometry[5]. The morphing box was placed around the rear part of the car only, targeting the formation of a rear spoiler. Again, the Spalart-Allmaras turbulence model with wall functions is employed. Figure(6) shows the comparison between the baseline and the optimized rear part of the car. The spoiler reduces the drag by 0.2%.

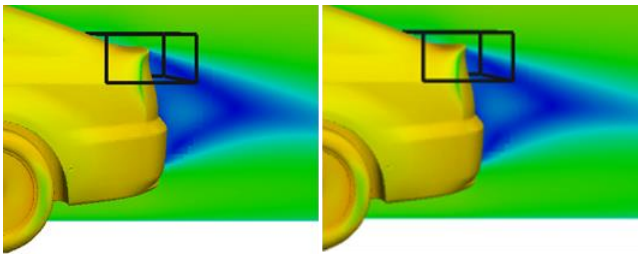

Fig. 6: Initial (left) and optimized (right) rear parts of the DrivAer geometry, targeting min. drag force. The boundaries of the control box are painted in black. Car surface colored by pressure and symmetry plane by velocity magnitude.

3.4. S-Bend Duct Morphing

The third case demonstrates the reduction of total pressure losses in an S-bend duct used in the HVAC system of a passenger car (an EU ITN AboutFlow test case). Here the flow is laminar and the optimization has led to a reduction in total pressure losses by more than 60%, (Figure(7)).

4. Conclusion

We have introduced a set of methods and tools for applying adjoint-based optimization methods to vehicle external aerodynamics. While not unique in concept, it is the first such method with sufficiently robust and efficient algorithms to allow it to be confidently applied within the context of existing vehicle simulation workflows. Specifically, the nature of the method allows it to be used in conjunction with geometries and grids of industrial scale and complexity. The introduction of the time-averaged DES adjoint approach allows the use of the type of high accuracy primal input that is required for reliable and dependable aerodynamic prediction and optimization. The adjoint solvers, when coupled with appropriate free form deformation methods, result in powerful automatic optimization methods able to radically improve vehicle aerodynamic characteristics at a fraction of the cost of equivalent stochastic methods.

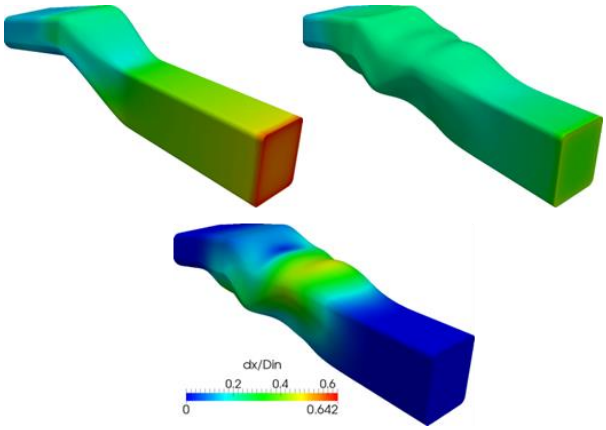

Fig. 7: Total pressure losses for the initial (top-left) and optimized (top-right) ducts. Flow from right to left. The large mesh displacements (up to 64% of the inlet diameter, bottom) can be efficiently handled by the morphing tool.

References

(1)A.S. Zymaris, D.I. Papadimitriou, K.C. Giannakoglou, C. Othmer: Continuous Adjoint Approach to the Spalart--Allmaras Turbulence Model for Incompressible Flows, Computers & Fluids, Vol. 38, Issue 8, p. 1528-1538, (2009).

(2)E.M. Papoutsis-Kiachagias, A.S. Zymaris, I.S. Kavvadias, D.I. Papadimitriou, K.C Giannakoglou: The Continuous Adjoint Approach to the k-ε Turbulence Model for Shape Optimization and Optimal Active Control of Turbulent Flows, Engineering Optimization, Vol. 47, Issue 3, p. 370-389, (2015).

(3)I.S.Kavvadias, E.M. Papoutsis-Kiachagias, G. Dimitrakopoulos, K.C. Giannakoglou: The Continuous Adjoint Approach to the k-ω Turbulence Model with applications in shape optimization, Engineering Optimization. Vol. 47, Issue 11, p. 1523-1542, (2015).

(4)E.M. Papoutsis-Kiachagias, K.C. Giannakoglou: Continuous Adjoint Methods for Turbulent Flows, Applied to Shape and Topology Optimization: Industrial Applications, Archives of Computational Methods in Engineering, p. 1-45, (2014)

(5)A. Heft, T. Indinger, N. Adams: Introduction of a New Realistic Generic Car Model for Aerodynamic Investigations, SAE Paper 2012-01-0168, (2012)

(6)C.Othmer. Adjoint Methods for car aerodynamics. Journal of Mathematics in Industry. Vol 4, No. 6, (2014)

(7)L. Hascoet, L. and V. Pascual. 2004. TAPENADE 2.1. INRIA. http://www-tapenade.inria.fr:8080/tapenade/index.jsp (accessed 2015.11.30)

(8)E. de Villiers, A. Lock, P. Geremia, Todd Johansen: ELEMENTS, A New Aerodynamics Analysis Software Aerodynamics, UHTM, HVAC and Cabin Comfort, Soiling and Water Management. JSAE International Journal of Automotive Engineering, Vol. 6, No. 2, pp.39-44 (2015)