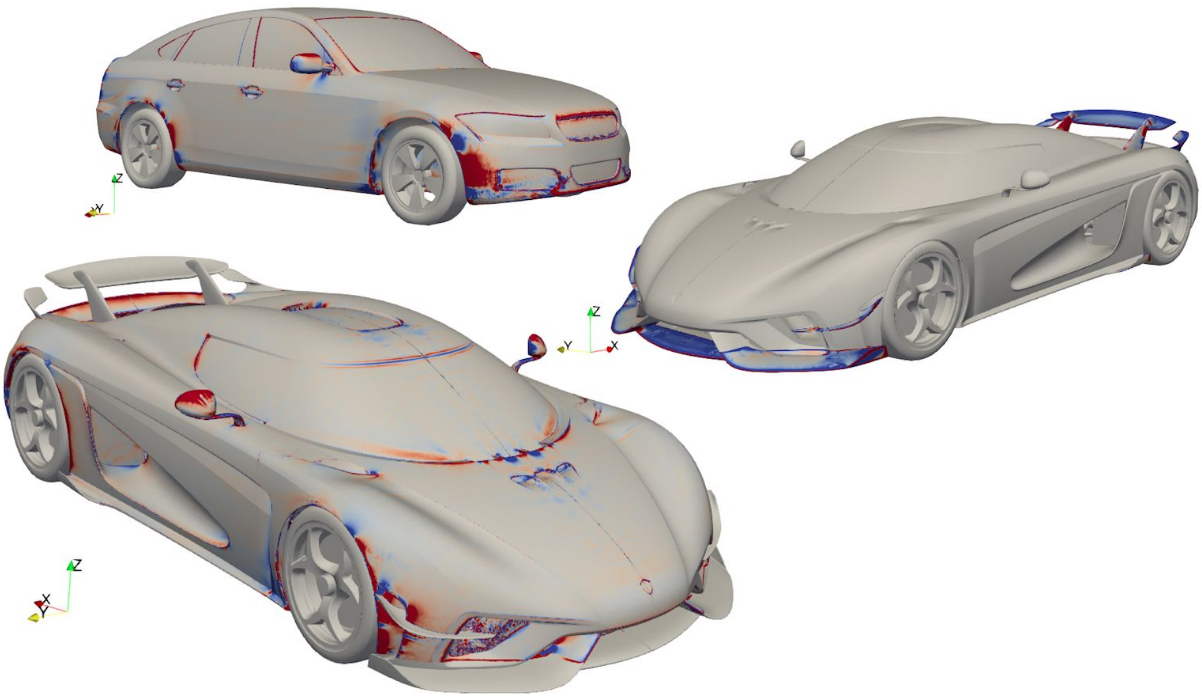

External aerodynamic optimization of ground vehicles using the adjoint method

Master's thesis in Automotive Engineering

GUSTAU LUNDEGAARD LANGE

DEPARTMENT OF MECHANICS AND
MARITIME SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2021
www.chalmers.se

MASTER'S THESIS IN AUTOMOTIVE ENGINEERING

External aerodynamic optimization of ground vehicles using the adjoint method

GUSTAU LUNDEGAARD LANGE

Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2021

External aerodynamic optimization of ground vehicles using the adjoint method
GUSTAU LUNDEGAARD LANGE

Master's Thesis 2021:01
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: + 46 (0)31-772 1000

Cover:
Sensitivity maps of a pre-production model of the Koenigsegg Regera with Ghost Package, and the DrivAer fastback model with smooth underbody, mirrors, and wheels.

Department of Mechanics and Maritime Sciences
Göteborg, Sweden 2021

External aerodynamic optimization of ground vehicles using the adjoint method
Master’s thesis in Automotive Engineering
GUSTAU LUNDEGAARD LANGE
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology

Abstract

The adjoint optimisation method within the field of ground vehicle aerodynamics has been a topic of continuous discussion in regards to advancing vehicle performance and decreasing development time. Conversely, the adjoint method has been limited by assumptions, stability issues, and increased computational cost. This thesis aimed to objectively demonstrate the capabilities of the newly implemented adjoint solver in OpenFOAM v2006, and its applicability to current development methods of Koenigsegg Automotive AB. First, a primal solver setup was validated on the DrivAer model by comparison of well-studied experimental force coefficients, pressure coefficients, as well as velocity and wall shear stress fields. The adjoint solver was investigated further on the DrivAer model by a parametric study of various mesh and solver settings, along with their impact on the sensitivity maps, which inform the engineer of potential areas of optimisation. Significant importance of the mesh quality and rotational boundary conditions have been observed on the sensitivity maps. Stability issues were encountered but mitigated, while retaining certain accuracy. Use was extended to a development CAD model of the Koenigsegg Regera to investigate stability and capability of the simulation at Re>20E6, and when adding porous zones. Steady optimisation with the Spalart-Allmaras turbulence model was validated by running three independent simulations with different two-equation turbulence models and notable drag minimisation was achieved within four cycles for the DrivAer model. Unfortunately, the steady optimisation proved to be unstable when morphing the finer mesh of the Regera and in areas of separated and highly rotational flow. Implementation of the adjoint method in external aerodynamic development of ground vehicles at Koenigsegg is possible under supervision of engineers. However, the adjoint method remains partially hindered by choice of adjoint turbulence modelling, mesh sensitivity, modelling accuracy, and compromised in regions of separated and highly rotational flow.

Key words: Adjoint method, aerodynamics, ground vehicles, DrivAer, Koenigsegg, optimisation, sensitivity maps

Preface

This Master's thesis work has been performed in collaboration with Koenigsegg Automotive AB with support and guidance from the Technical University of Denmark and Chalmers University of Technology. The Master's thesis concludes a two-year education for a degree in Master of Science in Mechanical Engineering at the Technical University of Denmark. The Master's thesis is equivalent to 30 ETCS point and has taken place from late August 2020 to early February 2021.
This Master's thesis is titled *External Aerodynamic Optimization of Ground Vehicles Using the Adjoint Method*. The Master's thesis is written and submitted by Gustau Lundegaard Lange. The Master's thesis is supervised by Marcus Lång, CFD / Aerodynamics and Mechanical Engineer at Koenigsegg Automotive AB, Professor Jens Honore Walther from the Technical University of Denmark, and Professor Simone Sebben from Chalmers University of Technology.

Gustau Lundegaard Lange, Gothenburg, February 2021

Acknowledgements

First of all, I would like to acknowledge the department of CFD and Aerodynamics at Koenigsegg Automotive AB for giving me the opportunity to do my Master's thesis in collaboration with such an exciting company. I would like to give special thanks to my industrial supervisor at Koenigsegg Automotive, Marcus Lång, who has made this collaboration possible and pleasant. Thank you for always being understanding, giving great guidance, and responding almost immediately to my many emails. Thank you, Harry Kounenis for your comments to the report and for stepping in during the final weeks of the Master's thesis. To professor Simone Sebben, I would like to express my gratitude for taking part in this thesis. With your preliminary guidance, we narrowed down an excellent Master's thesis topic. Your expertise and useful guidance within the field of Road Vehicle Aerodynamics have made this thesis capable of becoming even better. I would also like to express my sheer appreciation to Professor Jens Honore Walther. You have always been ready to support me or point me in the right direction in all matters regarding the thesis, including technical issues with the CFD and the cluster. I also want to thank DTU postdoc Jiun Cai Ong for always being ready to assist me in the technical issues of OpenFOAM and the DTU cluster. Moreover, thanks to Jakob Roar Bentzon for assisting me during several weeks of installing OpenFOAM v2006. Finally, I want to express gratitude to Louise Rugaard Lund, my girlfriend, who has been supportive and given me insightful advice on the process of writing a Master's thesis.

Gustau Lundegaard Lange, Gothenburg, February 2021

Nomenclature

Abbreviations and Acronyms

ATC	Adjoint Transpose Convection
BC	Boundary Condition
CFD	Computation Fluid Dynamics
Count	1 count is 0.001x a force coefficient
CP	Control Point
ESI	Enhanced Surface Integral formulation
FD	Finite Difference
FI	Field Integral
FSMW	Fastback, Smooth underbody, with Mirrors and with Wheels
GBM	Gradient Based Method
GS	Ground Simulation
IC	Initial Condition
kE	k-ϵ turbulence model
Ma	Mach number
MRF	Moving Reference Frame
PDE	Partial Differential Equations
PID	Property Identification
PVT	Volvo Cars wind tunnel
RANS	Reynolds Averaged Navier-Stokes
RKE	Realizable k-ϵ turbulence model
RWBC	Rotating Wall Boundary Condition
SA	Spalart-Allmaras turbulence model
SD	Sensitivity Derivatives
SI	Surface Integral formulation
SST	k-ω Shear-Stress Transport turbulence model
STL	Standard Tesselation Language
TI	Turbulence Intensity
URANS	Time-dependent Reynolds Averaged Navier-Stokes

Roman Symbols

A	Projected frontal area	$[\mathrm{m}^2]$
a	Speed of sound	$[\mathrm{m\,s^{-1}}]$
C_d	Coefficient of drag	$[-]$
C_f	Skin friction coefficient	$[-]$
C_{lf}	Coefficient of lift at front axle	$[-]$
C_{lr}	Coefficient of lift at rear axle	$[-]$
k	Turbulent kinetic energy	$[\mathrm{m^2\,s^{-2}}]$
L	Length of vehicle	$[\mathrm{m}]$
Ma	Mach number	$[-]$

p	Primal static pressure	[kg m^{-1} s^{-2}]
p	Free stream static pressure	[kg m^{-1} s^{-2}]
p_k	Kinematic pressure	[m^2 s^{-2}]
Q	Second invariant of the velocity gradient tensor	[s^{-2}]
q	Adjoint static pressure	[m^2 s^{-2}]
Re	Reynolds number	[−]
T	Temperature	[^{o}C]
t	Time	[s]
u^+	Dimensionless velocity	[−]
U_∞	Free stream velocity	[m s^{-1}]
u_i	Adjoint velocity	[m s^{-1}]
U_τ	Friction velocity	[m s^{-1}]
v_i	Primal mean velocity	[m s^{-1}]
v_i'	Primal fluctuating velocity	[m s^{-1}]
$\bar{v}_i$	Primal velocity	[m s^{-1}]
y	Wall distance	[m]
y^+	Dimensionless wall distance	[−]
Greek Symbols		
δ	Boundary layer thickness	[m]
ϵ	Turbulent kinetic energy dissipation rate	[m^2 s^{-3}]
ν	Kinematic viscosity	[m^2 s^{-1}]
ν_t	Kinematic turbulent, or eddy, viscosity	[m^2 s^{-1}]
$\tilde{\nu}$	Working variable of the Spalart-Allmaras turbulence	[m^2 s^{-1}]
ω	Specific dissipation of turbulent kinetic energy	[s^{-1}]
Ω	Vorticity	[s^{-1}]
τ_w	Wall shear stress	[m^2 s^{-2}]
Subscripts and Superscripts		
∞	Free stream value	
a	Adjoint field variable	
x	Along x direction	

List of Figures

Contents

1

Introduction

Chapter one presents a brief introduction to the Master's thesis background, purpose, problem statement, objectives, and limitations.

1.1 Background

This Master's thesis will investigate the new adjoint solver released in OpenFOAM v2006, and its applicability in external aerodynamic vehicle development at Koenigsegg Automotive AB, compared to conventional aerodynamic development. The adjoint method is to be validated on a reference model and tested on a Koenigsegg vehicle by conducting the adjoint simulations at conventional conditions of high Reynolds flow, porous zones, and ground simulation.

The adjoint method is a computational optimisation tool that can improve aerodynamic performance. The adjoint method can minimise objective functions such as lift, drag, or total pressure loss, independent of the number of design variables. Thus, the adjoint method is superior to other optimisation methods due to its reduced computational cost. The adjoint method can be used for generating sensitivity maps, i.e. the gradient of the objective function with respect to normal displacement of boundary nodes. This can aid the development team in observing areas of potential optimisation during early stage development or late stage optimisation. Alternatively, the adjoint method can be applied directly to perform iterative shape optimisation conducted in a semi-automated manner. The adjoint method has gained increasing interest within the field of road vehicle aerodynamics over the past 15 years and is currently applied within large automotive corporations, although performed at lower Reynolds numbers [3, 12, 15, 20, 21, 22, 30].

1.2 Problem Statement

During external aerodynamic development, different designs can be simulated using Computational Fluid Dynamics (CFD) and subsequently post processing by engineers. This development process is time-consuming and may be aided or automated to a certain extent by the adjoint method. However, the limitations of the adjoint method is still under research and development. The adjoint method may still be in need of instability mitigation and monitoring of flow field sensitivity. The oversight by an experienced engineer may be needed. Additionally, it shall be studied whether the adjoint method is worth the additional setup -and computational time, compared to the conventional CFD development method.

1.3 Purpose

This thesis aims at establishing knowledge of the advantages and limitations of the adjoint method along with an investigation of sensitivity, validity, and robustness of the open-source adjoint solver, released in OpenFOAM v2006. The investigation of the adjoint method may pave the way for implementation of the adjoint method in aerodynamic vehicle development of Koenigsegg vehicles.

1.4 Objectives

Careful analysis of mesh simplification, sensitivity of flow field properties, and the limitations of the simulation setup will be investigated. Koenigsegg Automotive AB will supply CAD models of interest. The objectives of this Master's thesis include:

- Establish the knowledge and investigate the potential of the adjoint solver included in OpenFOAM v2006 and its applicability to external aerodynamic flow simulation of a ground vehicle.
- Set up a fully functional, validated, primal - and adjoint solver.
- Optimize the solver setup in means of sensitivity, accuracy, time, and robustness.
- Discuss the advantages and limitations of adjoint optimisation vs. pressure coefficient reading of conventional CFD.
- Propose a method and write guidelines of mesh setting setup, primal solver setup, and interpretation of adjoint fields, testing different settings of mesh settings as well as for both the primal solver and the adjoint solver.
- The adjoint method shall be automated by introducing mesh morphing and optimized workflow for automated shape optimisation loops.

1.5 Limitations

- The time frame of this thesis work is limited to the autumn semester of 2020 (beginning late August 2020 and concluding early February 2021).
- Computational resources are limited to those available at the cluster of DTU and Chalmers.
- Post-processing is carried out with Paraview through client-server protocol with limited memory of 32GB or on a computer with no more than 16GB RAM. Thus, limiting the max memory and therefore max cell count of $\sim 40M$ cells.
- Two geometries are studied. (1) A virtual research model, namely, the DrivAer and (2) The Koenigsegg Regera with Ghost Package. More particularly, the DrivAer model used in this thesis work is denoted FSMW (Fastback, smooth underbody, with mirrors, and with Wheels) since found in several comparable studies.
- Simulations are conducted solely using symmetry models of both geometries and primarily limited to steady state, i.e. asymmetrical wake cross flows are not captured.
- Function objects are limited to drag and lift/drag during the generation of sensitivity maps and automatic morphing optimisation.
- Available software during this thesis work follows as SolidWorks, CATIA V5, ANSA, OpenFOAM, and ParaView.
- Only experimental data of the FSMW DrivAer model is available.

2

Theory

The sections in this chapter briefly present the theory, mathematical expressions, and physics of fluid dynamics and the adjoint optimisation utilized during the Master thesis.

2.1 Fluid Dynamics

Generally, the external flow around a ground vehicle experiences similar flow phenomena as that of blunt bodies. As blunt bodies, vehicle aerodynamics is governed by larger volumes of flow separation. As a ground vehicle travels forward, it is subjected to viscous forces and inertial forces. The relationship between viscous and inertial forces is given by the Reynolds number, $Re = U_\infty L/\nu$. The thickness of the boundary layer, δ, observed on Figure 2.1 depends on Reynolds number, i.e. the ratio of diffusing momentum. The boundary layer is formed due to the no-slip condition at the wall and high momentum free-stream interaction. Within the boundary layer, velocity components decrease from the value of the quasi-inviscid external flow at the outer edge of the boundary layer to zero at the wall, where the fluid fulfills a no-slip condition. Beyond this boundary layer, the flow can often be regarded as inviscid, given that the viscous forces are small or negligible compared to inertial forces. When the flow separates, as illustrated in Figure 2.1, the boundary layer is *diffused*.

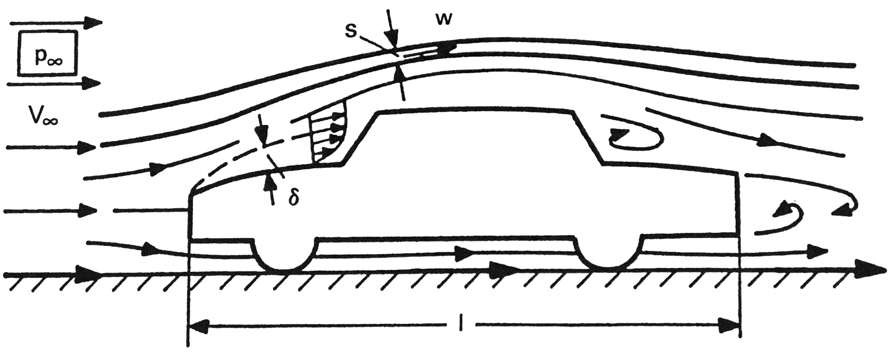

Figure 2.1: Schematic flow around a passenger vehicle [14].

Aerodynamic Drag

Aerodynamics is a field within Fluid Mechanics that studies the properties of moving air and the interaction between the air and the solid bodies moving through it [16]. Thus, aerodynamic analysis is an essential part of the development of a road vehicle, in terms of high-speed stability, handling, efficiency, etc. As a vehicle moves forward, the drag force is the largest aerodynamic force acting on the vehicle surface. The vehicle drag force is exerted by form (pressure) drag and friction drag. For external vehicle aerodynamics, the form drag, caused by the pressure distribution of the moving air, is by far the largest component of drag. The friction drag, due to molecular friction, causes the shear stress of the moving air, acting on the surface of the body. The drag force is generally formulated by the integration of the corresponding force components acting in the free stream direction according to:

$$D = \int_S p\, sin\varphi \; dS + \int_S \tau_w \, cos\varphi \; dS \tag{2.1}$$

where τ_w is wall shear stress and p is the static pressure. To make the drag force dimensionless, the coefficient of drag of ground vehicles can be formulated as:

$$C_d = \frac{D}{p_{dyn}\, A} \tag{2.2}$$

where $p_{dyn} = 0.5\,\rho\, U_\infty^2$ is the free stream dynamic pressure and A is the projected frontal area of the vehicle. In general, ground vehicles have a high-pressure region in the front and a low-pressure region in the rear. The pressure difference creates a net force opposite of the ground vehicle heading. According to Equation 2.2 the engineer may decrease the pressure difference, the shear force, and the frontal area to optimise the drag coefficient. Coefficients of drag can be challenging to compare in-between models with different frontal areas. The appropriate expression $C_d A$ allows for a comparison of the drag efficiency of different vehicles.

A second optimisation scalar may be the lift force, either on the whole vehicle or over the individual axles (front or rear). The lift force, L, is acting perpendicular to the drag force and side force. The coefficient of lift is expressed by:

$$C_l = \frac{L}{p_{dyn}\, A} \tag{2.3}$$

Note that various additional equations of lift exist in aerodynamics when studying lift devices [14, 16]. Next, the side force is accompanied by three moments (with a reference point in the center of the wheelbase) consisting of pitching, yawing, and rolling, which should be taken into account when developing a vehicle [16]. Furthermore, dimensionless variables include the coefficient of pressure and coefficient of skin friction, respectively stated in Equation 2.4, which are important measures during the post-processing of aerodynamic simulations.

$$C_p = \frac{p - p_\infty}{p_{dyn}}, \qquad C_f = \frac{\tau_w}{p_{dyn}} \tag{2.4}$$

where p_∞ is the free stream pressure. The coefficient of pressure is useful for observing areas of stagnation, lift, and drag. The skin friction coefficient is an indicator of flow separation, acceleration/deceleration, and stagnation. Post-processing of additional flow vectors include vector fields of velocity and vorticity, as well as scalar fields such as the Q-criterion and total pressure. For the sake of conciseness, visual post-processing of the flow field are primarily limited to the velocity field as well as surface plots of the pressure coefficient and the skin friction coefficient.

2.2 Computational Fluid Dynamics

This Master thesis is concluded using Computational Fluid Dynamics (CFD). This section should provide a general description of the theory and mathematics behind CFD.

The Governing Equations

The fluid flow can be described by equations based on the Continuity Equation 2.5 (conservation of mass) and the Navier-Stokes Equation 2.6 for steady state incompressible fluid flows, excluding heat transfer and body force, stated by indicial notation.

$$R^p = -\frac{\partial v_i}{\partial x_i} = 0 \tag{2.5}$$

$$R_i^v = v_j \frac{\partial v_i}{\partial x_j} + \frac{1}{\rho}\frac{\partial p}{\partial x_i} - \frac{\partial}{\partial x_j}\left[\nu\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i}\right)\right] = 0 \tag{2.6}$$

where v_i (velocity components), p is (static pressure divided by constant density), ν is the constant of bulk kinematic viscosity. The first term on the r.h.s. of Equation 2.6 is the convective acceleration. The second and third term are, respectively, the surface forces (pressure gradient) and the viscous diffusion. These equations, coupled with appropriate turbulence models, provide the foundation for solving the turbulent flow field with a certain accuracy.

Turbulence Modelling

Turbulent flow and ground vehicle aerodynamics are nearly synonymous due to the nominal operating conditions in the high Reynolds regime of $Re > 10^6$. Ground vehicles are large objects travelling at relatively high speed in a gas with low viscosity. Common practice is to simulate fully turbulent flows, disregarding minor regions of laminar and transitional flow. To model turbulence during steady state simulations, the Navier-Stokes equation is decomposed. From such decomposition, the instantaneous pressure and velocity components are decomposed into a mean quantity (denoted by a superscript bar) and a fluctuation quantity (denoted by an apostrophe) given by indicial notation:

$$v_i = \bar{v}_i + v_i' \qquad , \qquad p = \bar{p} + p' \tag{2.7}$$

Decomposition of the Navier-Stokes Equation 2.6 mean flow yields the Reynolds Averaged Navier-Stokes (RANS). The Continuity Equation 2.5 and the steady state, incompressible RANS equations are formulated in indicial notation:

$$R^p = -\frac{\partial \bar{v}_i}{dx_i} = 0 \tag{2.8}$$

$$R_i^v = \bar{v}_j \frac{\partial \bar{v}_i}{\partial x_j} + \frac{1}{\rho}\frac{\partial \bar{p}}{\partial x_i} - \frac{\partial}{\partial x_j}\left[\nu\left(\frac{\partial \bar{v}_i}{\partial x_j} + \frac{\partial \bar{v}_j}{\partial x_i}\right)\right] - \frac{1}{\rho}\frac{\partial}{\partial x_j}\left[-\rho\,\overline{v_i'\,v_j'}\right] = 0 \tag{2.9}$$

Note here that Equation 2.6 is similar to Equation 2.9 apart from averaging and addition of the rightmost term in Equation 2.9, namely, the Reynolds stress, $-\rho\,\overline{v_i'\,v_j'}$. Hence, the author of this thesis have kept the density in the expression, by means of justifying the name *stress*. The kinematic viscosity, ν, can be decomposed into the constant bulk kinematic viscosity and the turbulent viscosity, ν_t. The turbulent viscosity, ν_t, is computed

simultaneously with the state equations by solving turbulence partial differential equations (PDEs). Equation 2.8 and Equation 2.9 comprise of four equations for the mean flow, namely one continuity equation and three RANS equations. The Reynolds stress tensor $\rho\,\overline{v_i'\,v_j'}$ results in the addition of six independent Reynolds stresses, three diagonal, and three off-diagonal elements [26]. Thus, the system is not closed, yet, with ten unknowns and just four equations. This problem is famously known as *the closure problem of turbulence*. Consequently, additional turbulence modelling is required.
Various kinds of turbulence models have been developed over the years. This thesis limits the scope of RANS-based turbulence-energy equations. If used correctly, applications of RANS-based turbulence-energy models are able to predict the exerted aerodynamic forces and fluid flow to a relatively high degree of accuracy. These models express Reynolds stresses as a product of a turbulent viscosity (also denoted eddy viscosity) and the mean strain rate. The turbulence viscosity is computed in terms of turbulence kinetic energy, $k = 0.5\,(\overline{v_i'\,v_i'})$. The turbulence-energy equations come in variants such as one-equation models, namely the Spalart-Allmaras Turbulence Model [24], and two-equation models, such as the k-ϵ model or the k-ω model. Various improvements to the original turbulence models, e.g. updated model constants and better wall treatment, have been performed over the years. The Realizable k-ϵ model is a popular choice in CFD turbulence modelling within the automotive industry. The k-ω SST model has proven to be the best choice when modelling flows with strong adverse pressure gradients [26]. This thesis work has shown that the Spalart-Allmaras is the most robust turbulence model during application of the adjoint method.

Near Wall Flow

Boundary layers form as moving fluid particles make contact with a stationary (no-slip) surface, as long as the flow is attached. Flows within the boundary layer are governed by gradients and viscous effects. The near wall flow can be defined by four layers with different flow properties stated according to a dimensionless height from the wall: the viscous sublayer, the buffer layer, the logarithmic layer, and the outer layer. The dimensionless height from the wall is denoted y^+.

$$y^+ = \frac{U_\tau\, y}{\nu} \qquad , \qquad U_\tau = \sqrt{\frac{\tau_w}{\rho}} \tag{2.10}$$

where U_τ is defined as the friction velocity. The viscous sublayer is governed by molecular viscosity within $y^+ < 5$. The buffer layer, a mix of the viscous sublayer and the logarithmic layer, is in the interval $5 < y^+ < 30$. Next, is the logarithmic layer is in the region of $30 < y^+ < 500$, also denoted the law of the wall. Due to computational limits, it is not possible to model the viscous sub-layer. Thus, wall modelling is required. OpenFOAM utilises wall functions specified for each model equation [19]. As a result, the target of modelling down to the logarithmic layer is sufficient. In the case of y^+ values approaching the viscous sublayer, OpenFOAM utilises blending functions to model wall functions accordingly [8]. The need to obtain accurate sensitivity derivatives for high-Re turbulence models in industrial applications dictates the necessity for differentiating high-Re turbulence models and the law of the wall [22].

Frozen Turbulence Assumption

The frozen turbulence assumption has been applied in various commercial adjoint solvers, e.g. Fluent 2019 R1, to simplify computations. The frozen turbulence assumption is mentioned in this thesis as it is applied for RANS-based turbulence-energy two-equations in the adjoint solver of OpenFOAM v2006. This thesis aims at excluding the application of the frozen turbulence assumption, if possible, due to limitations of the assumption. The frozen turbulence assumes that a coherent structure remains unchanged as it propagates downstream.

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \bar{u}\,\frac{\partial}{\partial x} = 0 \quad \rightarrow \quad \frac{\partial}{\partial t} = -\bar{u}\,\frac{\partial}{\partial x} \tag{2.11}$$

Frozen turbulence is applicable in steady flows with homogeneous turbulence and free of mean external pressure gradients. For the frozen turbulence assumption to be valid, the turbulence intensity needs to be small, $TI = \frac{\sqrt{\overline{u'^2}}}{\bar{u}} \ll 1$. Studies have proved erroneous results when applying the frozen turbulence assumption, meaning that the frozen turbulence assumption may lead to the computation of sensitivities with deviance in magnitude and in rare cases computation of wrong sign [22]. Furthermore, turbulent adjoint steady optimisation with frozen turbulence assumed has proven the need for additional optimisation cycles [22].

2.3 Theory of the Adjoint Method

The adjoint method is a gradient-based optimisation tool which can be used both for internal and external flow optimisation. The optimisation problem aims at the minimisation of a function object denoted F in the design space. A function object may be drag, lift, or total pressure loss depending on the application. The function object is minimised with respect to the vector of design variables $b_n = (b_1, b_2, .., b_N)$. Design variables may be the boundary nodes of a geometry. Generally, the function object is a function of flow variables, $\mathbf{U}$, and design variables, b_n. The advantage of the adjoint method, as opposed to alternative optimisation methods of Finite Difference and Direct Differentiation, is its ability to optimise a given function object, F, independently of design variables, b_n [22]. This allows for significantly decreased computational cost, and hence the option of performing shape optimisation on large scale industrial applications, with reference to external aerodynamic optimisation of ground vehicles with millions of design variables. The adjoint method in OpenFOAM is based on the assumption of incompressible steady state flow. Additionally, OpenFOAM v2006 is applying the continuous adjoint method. In the continuous adjoint approach, the primal PDEs are kept in continuous form and inserted into an augmented objective function. As a result the resulting adjoint equations are in the form of PDEs (together with their boundary conditions), subsequently to be discretized and numerically solved in order to compute the adjoint variables. The advantages and disadvantages of the continuous and discrete adjoint method have been summarised in section A.1.

Augmented Function

Formulation of the continuous adjoint method starts with the simplified augmented function, F_{aug}, [22].

$$F_{aug} = F + \int_{\Omega} u_i \, R_i^v \, d\Omega + \int_{\Omega} q \, R^p \, d\Omega \tag{2.12}$$

where Ω is the computational domain. F is the objective function to be minimised/optimised. R_p and R_i^v are the *primal* state equations, respectively, in Equation 2.8 and Equation 2.9, and yet married to a turbulence model. u_i stands for the adjoint velocity components and q stands for the adjoint pressure. Note that the simplified augmented function increases in number of terms of surface integrals when the adjoint boundary conditions are applied. Note that Equation 2.12 makes use of the frozen turbulence assumption. Thus, additional integrals of the turbulence state equations are excluded in Equation 2.12 for the sake of simplicity. The adjoint solver in OpenFOAM v2006 has the option of anisotropic turbulence modelling by the Spalart-Allmaras model. The reader is advised to study the augmented function formulated for the application of High-Re adjoint Spalart-Allmaras turbulence modelling in [22].

Differentiation of the augmented function is formulated by employing the Leibniz theorem for the differentiation of volume integrals with variable boundaries [22].

$$\frac{\delta F_{aug}}{\delta b_n} = \frac{\delta F}{\delta b_n} + \int_{\Omega} u_i \, \frac{\delta R_i^v}{\delta b_n} \, d\Omega + \int_{\Omega} q \, \frac{\delta R^p}{\delta b_n} \, d\Omega + \int_{S_{W_p}} (u_i \, R_i^v + q \, R^p) \, n_k \, \frac{\delta x_k}{\delta b_n} \, dS \tag{2.13}$$

where the boundary S of Ω are decomposed as $S = S_I \cup S_O \cup S_W \cup S_{W_p}$, whose constituents respectively correspond to inlet, outlet, fixed and parameterized wall boundaries of Ω. n_k denotes the components of the unit vectors normal to the surface [22]. As the residuals of the primal equations must be zero, $F_{aug} \equiv F$, hence $\frac{\delta F_{aug}}{\delta b_n} \equiv \frac{\delta F}{\delta b_n}$.

Adjoint Governing Equations

Further definition and derivations of the surface and volume integrals in Equation 2.13 are concluded in [22] to formulate the field adjoint equations of continuity, Equation 2.14, and momentum, Equation 2.15. The reader is referred to [22] for exact derivations.

$$R^q = -\frac{\partial u_j}{\partial x_j} = 0 \tag{2.14}$$

$$R_i^u = u_j \, \frac{\partial v_j}{\partial x_i} - \frac{\partial (v_j \, u_i)}{\partial x_j} + \frac{1}{\rho} \, \frac{\partial q}{\partial x_i} - \frac{\partial}{\partial x_j} \left[\nu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \right] = 0 \tag{2.15}$$

The adjoint continuity equation is similar to the primal Equation 2.5, being that the adjoint velocity is divergence-free. The adjoint momentum equation is similar to Equation 2.6 since it contains adjoint convection (although backward), adjoint pressure gradient, and adjoint viscous diffusion. Furthermore, the adjoint momentum equation is linear, since the adjoint velocity is convected by the primal velocity which should enhance convergence. Conversely, due to the first term, $u_j \, \frac{\partial v_j}{\partial x_i}$, denoted *adjoint transpose convection* (ATC), and adjoint boundary conditions, convergence and stability are compromised [22].

Sensitivity Derivatives

The final expression for the sensitivity derivatives, which are formulated by (1) satisfying the adjoint field equations Equation 2.14 and Equation 2.15 and their boundary equations, and (2) by adding terms derived from the adjoint boundaries. Full expressions of sensitivity derivatives are formulated in section 3.6, 4.4, and 5.2 in [22] depending on the Reynolds number and turbulence modelling.
The ability of the adjoint optimisation method to calculate sensitivity derivatives independently of number of design variables is due to introducing the Eikonal Equation for solving the distance variation field. The distance variation field with respect to design variables is found in the volume integrals of the sensitivity derivative formulations [22]. The Eikonal Equation is transformed into the Hamilton-Jacobi Equation for application in standard finite-volume schemes, avoiding numerical difficulties. The implementation of the Hamilton-Jacobi Equation was proved to to have a low error at a distance close to solid walls [22]. The Hamilton-Jacobi expression is inserted into the formulation F_{aug} for completion. Replacing the partial variation of the distance equation to make the sensitivity derivatives independent of the computation of design variables [22]. Since the distance variation may contribute to a significant amount of the sensitivity derivatives value, it is proposed in [22] to impose a zero Neumann condition on the adjoint distance field along the (parameterized) wall boundaries. Moreover, a zero Dirichlet condition should be applied on the inlet and outlet to satisfy the adjoint distance equation [22].
The adjointOptimisationFoam of OpenFOAM v2006 is explained and derived to a great detail in the papers of [17, 22, 23].

2.4 The Adjoint Method in OpenFOAM

The author is using OpenFOAM v2006. To the knowledge of the author, the adjoint optimisation algorithms in OpenFOAM provide potential for greater accuracy compared to other commercial codes. OpenFOAM v2006 offer the option of modelling adjoint turbulence (opposed to the frozen turbulence assumption), unfortunately limited to a single option of the adjoint Spalart-Allmaras turbulence model. One might expect an increased degree of accuracy in regions where turbulent effects dominate the flow, i.e. when the flow is separated. The software is still limited to a steady state adjoint solver [11]. When running steady state simulations, one should be aware that the external aerodynamics of passenger vehicles more often than not exhibit oscillatory flow. Therefore, it is impeccable that the primal solution does not oscillate such that the adjoint solver produces different results depending on when the simulation is ended. Transient regions are dealt with by appropriate window averaging. In addition, an oscillating primal solver may cause the adjoint solver to diverge. The OpenFOAM v2006 does also include the option of modelling high-Re flow with wall treatment of the viscous sublayer and the logarithmic layer of the boundary layer.

Flow Chart of the Adjoint Method

The computational steps of the adjoint method in OpenFOAM are visualized in Figure 2.2. Note that the flow chart describe two optimisation strategies, one being the calculation of sensitivity derivatives with respect to normal displacement boundary nodes, step #1 to #13. Subsequently, engineers can use the point sensitivity maps to morph the geometry manually, namely by the morphing tool found in ANSA. The second strategy is shape optimisation where step #1 to #8 is concluded, followed by shape optimisation by automatic morphing. Note that the shape optimisation cycle can be performed N times. $N > 1$ cycles will repeat step #3 to #8 followed by step #14 through #16. Note that both the primal and adjoint PDEs can be averaged. The adjoint residuals have been proven to decrease steadily for a converging solution, and should not be averaged. The primal solution does converge at a point in which averaging of the flow variables of pressure, velocity, and turbulence equations are averaged and subsequently fed to the adjoint solver. Additionally, the Eikonal equation is not solved during cases with Frozen Turbulence assumptions. Finally, it should be mentioned that the shape optimisation strategy requires the feeding of control points.

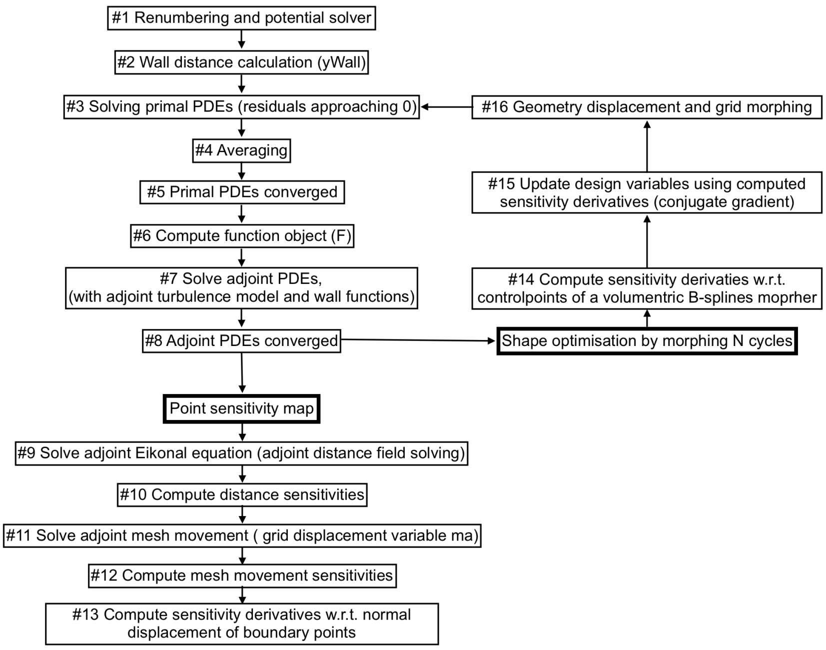

Figure 2.2: The steps of adjoint based optimisation algorithm is schemed for flow cases of Spalart-Allmaras turbulence modelling.

Note that for standard application of the adjointOptimisationFoam, the primal solver is based on the *simple* algorithm, for both the primal and adjoint PDEs. It is possible to conduct unsteady simulations for the primal flow, i.e. based on the piso algorithm. However, this requires appropriate temporal averaging when solving the primal PDEs. Furthermore, the use of unsteady algorithms requires manual input from the user, as the primal PDE solver will be set to false in the adjointOptimisationFoam, and hence an averaged unsteady solution will be fed manually into the adjoint solver. This process has been carried out by this author and has been deemed to work well. For semi-automated shape-optimisation cycles, this would be an inefficient choice, as an averaged unsteady primal field has to be solved manually for each cycle. Thus, modelling an unsteady primal field is recommended only for increasing accuracy of the sensitivity derivatives with respect to the displacement of boundary points. When solving the primal field manually, the options of using Detached Eddy Simulation (DES) are possible, again with attention to appropriate averaging and only with the RANS-based turbulence model of the Spalart-Allmaras in OpenFOAM v2006.

Literature on the Adjoint Method

The reputation of the adjoint method is partially synonymous with challenges of stability. However, the stability of the adjoint solver in OpenFOAM v2006, the *adjointOptimisationFoam*, shows potential for applying the adjoint method in various papers regarding external aerodynamic optimisation of vehicles [3, 11, 15, 17, 20, 22, 23]. The papers of [3, 11, 15, 17, 20, 22, 23] apply adjoint steady state and incompressible algorithms at $Re < 10E6$. This author expects to reveal the robustness of the adjointOptimisationFoam at $Re \geq 20E6$. Next, investigation of the adjoint method with respect to the DrivAer, briefly presented in [11], is performed. In fact, a parametric study of different mesh and simulation setup with respect to the generated sensitivity maps of the DrivAer will potentially be presented for the first time. The modelling of porous zone during the use of the adjointOptimisationFoam will possibly be concluded for the first time. A morphing study of the neW adjointOptmisationFoam solver will be conducted to investigate its' applicability to aerodynamic development of ground vehicles. For understanding the inner workings, including the derivation of adjoint formulations as well as assumptions and application of the adjointOptimisationFoam, see [11, 17, 22, 23]. These papers are a must-read, if trust worthy use of the adjointOptimisationFoam is expected.

2.5 Terminology

Figure 2.3 and Figure 2.4 are created to aid the reader in identifying regions during the written analysis and discussion in the chapter three and four. Automotive terminology may differ depending on British or American English. Stated terminology is used for this thesis only since different terms might be used academically or in the industry. A copy of the figure will be placed in the last appendix section A.6, should the reader wish a copy while reading. The *green house* is defining all windows. The *base* or *rear base* refers to the rear end with fully separated flow. A *count* should be defined as 0.001 times a force coefficient. For instance one drag count is $C_d = 0.001$.

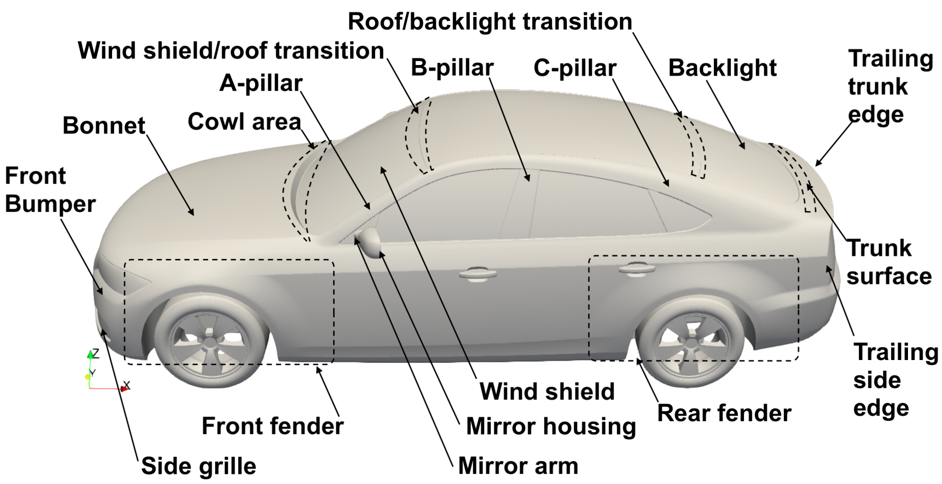

Figure 2.3: Terminology used for the describing areas of the FSMW.

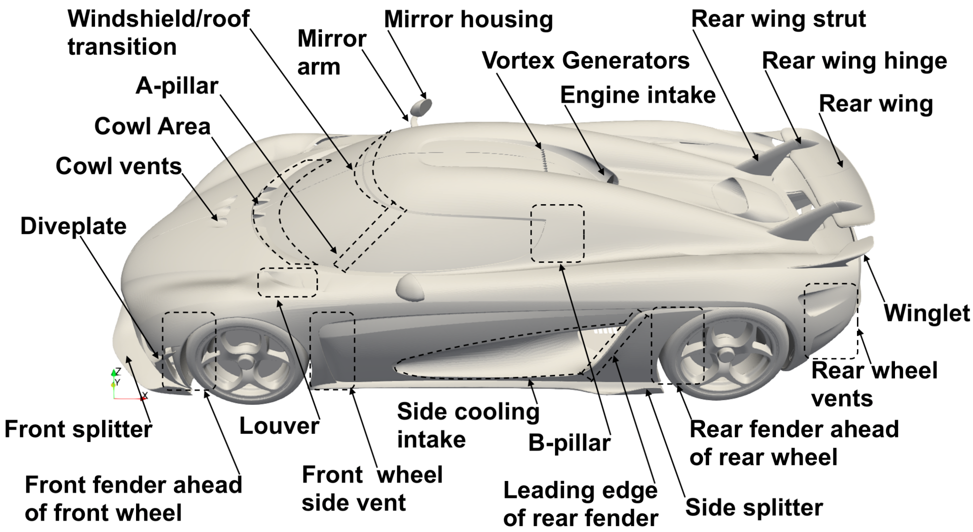

Figure 2.4: Terminology used for describing areas of the Koenigsegg Regera.

3
Method

Chapter three provides insight to the pre-processing, setup of the solvers, as well as a brief description of post-processing when performing adjoint simulation. The adjoint simulations aimed to approach similar model fidelity to current CFD development models for external aerodynamics at Koenigsegg. This ensured that the results and investigations were applicable to their current procedures.

3.1 Pre-processing

This section contains important information regarding the mesh generation of (1) the DrivAer reference fastback geometry, denoted FSMW, from the Technical University of Munich [6], and (2) a pre-production model of the Koenigsegg Regera with Ghost package.

FSMW Geometry

No pre-existing validated OpenFOAM simulation setups were provided for the author. Thus, a model with well-documented and realistic flow features was needed for the validation and bench marking of this thesis' simulation setup. The DrivAer model is selected for its ability to resemble realistic flow features of modern vehicle aerodynamics [7, 13, 25]. The DrivAer model is advantageous in the vicinity of the A- and C-pillars, the curved rear end, mirrors, wheels, and in the region of fenders, where the academic alternatives of the Ahmed Body and SAE model are impaired [13]. The DrivAer geometry is comprised of a combination based on the geometries of the Audi A4 and the BMW 3 Series. Several geometries of the DrivAer model are available. This thesis will apply a symmetry model of a DrivAer fastback version, with smooth underbody, with mirrors, and with wheels, here denoted $FSMW$, and first referenced in [13]. The FSMW model is chosen for the wider application in several papers [7, 10, 13, 25, 27, 28, 29]. The smooth underbody is chosen for simplicity and to decrease the amount of cells for the anticipated heavy adjoint simulation. The mirrors are chosen as they may be a region to study during adjoint optimisation. The geometry is found at the official site of the Chair of Aerodynamics and Fluid Mechanics at the Technical University of Munich [6]. Small simplifications have been carried out to increase mesh quality, such as removing small steep concave and convex surfaces, removal of gaps and holes, as well as translation of the body with respect to the origin. Additionally, the surfaces have been knitted and so-called watertight geometries have been created. STL files has been generated with the highest setting of tessellation in CATIA V5 and subsequently imported for meshing in OpenFOAM's snappyHexMesh. Unfortunately, snappyHexMesh proved itself insufficient since it snaps at the triangles of the STL files, which create artificial pressure changes, and hence wrongful sensitivity derivatives. Therefore, a shift to a commercial mesher, ANSA, was made to respect analytical model fidelity. The applied FSMW geometry and its overall dimensions can be observed in Figure 3.1. Note that the origin is displaced slightly above the lowest point of the wheels by 25 [mm] to form a contact patch that avoids exceeding the mesh quality of cell skewness, non-orthogonality, etc.

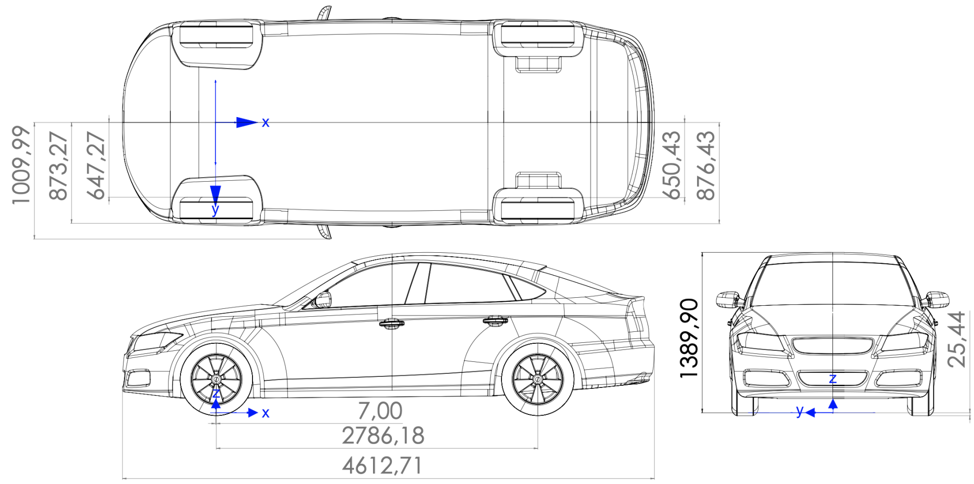

Figure 3.1: Geometry of full car *FSMW*. Length units in [mm].

FSMW Mesh

The FSMW mesh is constructed with consideration of base size refinement, volume refinement, and layer addition to secure a moderate simulation time, stability, and accuracy. The mesh is a hybrid mesh of structured and unstructured cells, respectively, to resolve the far-field domain and the complex geometry of FSMW. The mesh is primarily dominated by structured hexahedra cells, allowing for less computation effort and higher cell quality in the far-field domain.

Volume refinements of different levels are placed in regions where important flow phenomena, such as flow separation, are anticipated to be captured. Smaller cells are found closer to the surface, ensuring a cell growth ratio of 1.2. Prisms were applied to the surfaces of the body, wheels, and moving ground in order to capture the boundary layer formation and development. It is expected that strong vortices emanate from the A-pillar and the C-pillar, while the fluid domain between ground and underbody may be confined. Additionally, the mirror and wake region will create a completely detached flow. Consequently, refinements have been applied to increase cell density in the wake region, mirror region, underbody, A-pillar, and C-pillar.

Mesh settings are defined to fulfill the requirements of wall functions within each turbulence model. During the inspection of wall functions in the OpenFOAM API Guide, it was found that each turbulence model inherit individual blending functions which makes variable y^+ resolutions possible [8]. Therefore, a mean y^+ resolution target was set to resolve the logarithmic layer, $y^+ > 30$. As a result, several meshes were created, both with snappyHexMesh and ANSA. Figure 3.2 shows the final mesh of concern made in ANSA. This base mesh was varied during parametric studies of sensitivity maps found in section A.3.

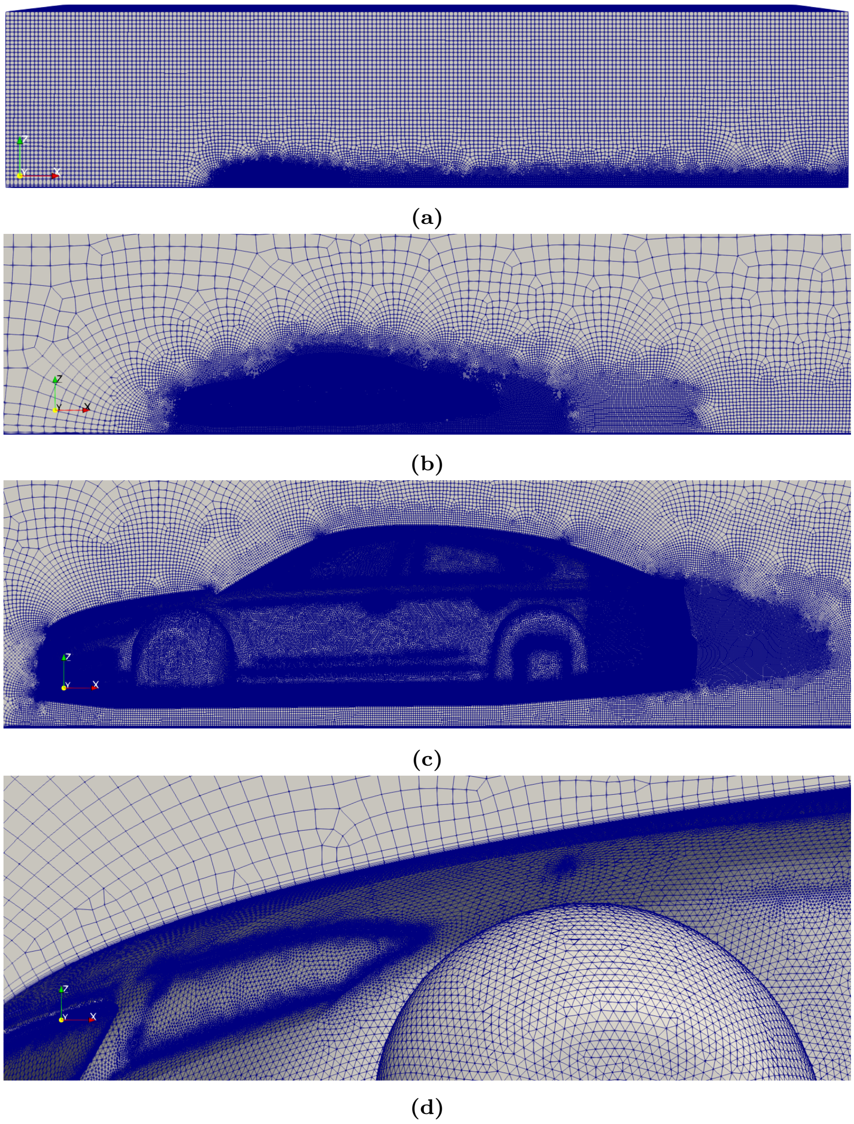

Figure 3.2: Figure (a) shows the full length and height of the flow domain at symmetry plane. Figure (b), (c) and (d) shows zoomed in snap-shots of the domain. Figure (d) shows resolved logarithmic layer on the bonnet.

Domain Boundaries

The mesh seen in Figure 3.2 is a large rectangular box of sufficient dimension to avoid interference effects from the walls. The far-field patches are modelled by slip and zero gradient boundary condition types. When the simulations are run with moving ground simulation, the ground is treated with a moving wall boundary condition with equal velocity to that of the inlet. The symmetry plane is treated with a symmetry condition.

Setup for Wheels

Two different rim geometries were studied, observed in Figure 3.3. The contact patches between ground and wheels have been facilitated by fusing the tire into the ground surface to avoid bad cell quality. In the case of moving ground simulation, the wheels have an angular wall velocity which, at the contact with the ground, gives a peripheral velocity with a magnitude equal to that of the vehicle speed. In the case of smooth rims, in

Figure 3.3a, the use of moving reference frames is not required. The moving reference frame zones are applied to the rims observed in Figure 3.3b. Dynamic mesh was not selected for simplicity and also for the ability to run steady state simulations. Both geometries were used for comparison with reference papers and during parametric studies in section A.3 to observe the effect of the rotating wheels and wheel geometry. The contact patch may create artificial flow phenomena in the vicinity and wake of the wheel. The contact patch changes the projected area marginally.

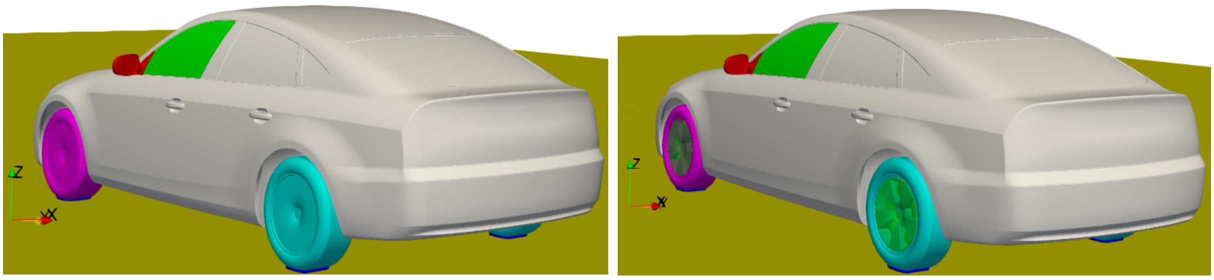

(a) FSMW with smooth rims.

(b) FSMW rim with spokes and MRF.

Figure 3.3: Applied wheel geometry variants.

Koenigsegg Regera Geometry

Koenigsegg provided detailed CATIA and STL files of a pre-production geometry of the Koenigsegg Regera with Ghost Package including front/rear frame, front splitter, rear wing, brakes, engine bay area, exhaust system, wheels, and more. Divergence of the adjoint solver, due to the meshes created in snappyHexMesh, deemed the STL files unfit for use. CATIA files were imported into ANSA where details were simplified, for instance tire threads, brake calipers, etc. Moreover, geometry import errors occurred and were corrected in ANSA to create a high-quality mesh. This was a temporal costly process. Measurement checks were performed to verify the exact location of the imported CATIA files with identical components in the STL files. This allowed for referencing this thesis' simulation with simulation results run by Koenigsegg. The final geometry in ANSA is observed in Figure 3.4 by property identification (PID). Notice the transparent volumes (internal volumes) in the wheels. These are the moving reference frames, ensuring rotational flow in the air pockets in between the spokes. Ideally, a dynamic mesh should be approached. Yet, steady state adjoint simulation is the only available option in OpenFOAM v2006. The open cooling was requested by Koenigsegg to include porous zones, and due to the fact that the external aerodynamics is explicitly influenced by cooling -and pressure ducts. Internal components on Figure 3.4 include front and rear frame, engine, differential, exhaust. External aerodynamic surfaces includes front splitter, front diveplates, sidesplitter, rear winglets, rear wing, vortex generators, etc. These components were all resolved by regions dense of prism cells. Hardware limitations forced this author to exclude the details of intake grills and fans to save cells. Modelling of the engine intake and exhaust outlet were excluded to simplify the adjoint simulations. Future simulations could include the engine intake and exhaust outlet, drag reducing devices, and additional components from the Koenigsegg production vehicles. Only one mesh version of the Regera was simulated due to time constraints. A left hand side symmetry model of the Koenigsegg Regera was simulated to save cells.

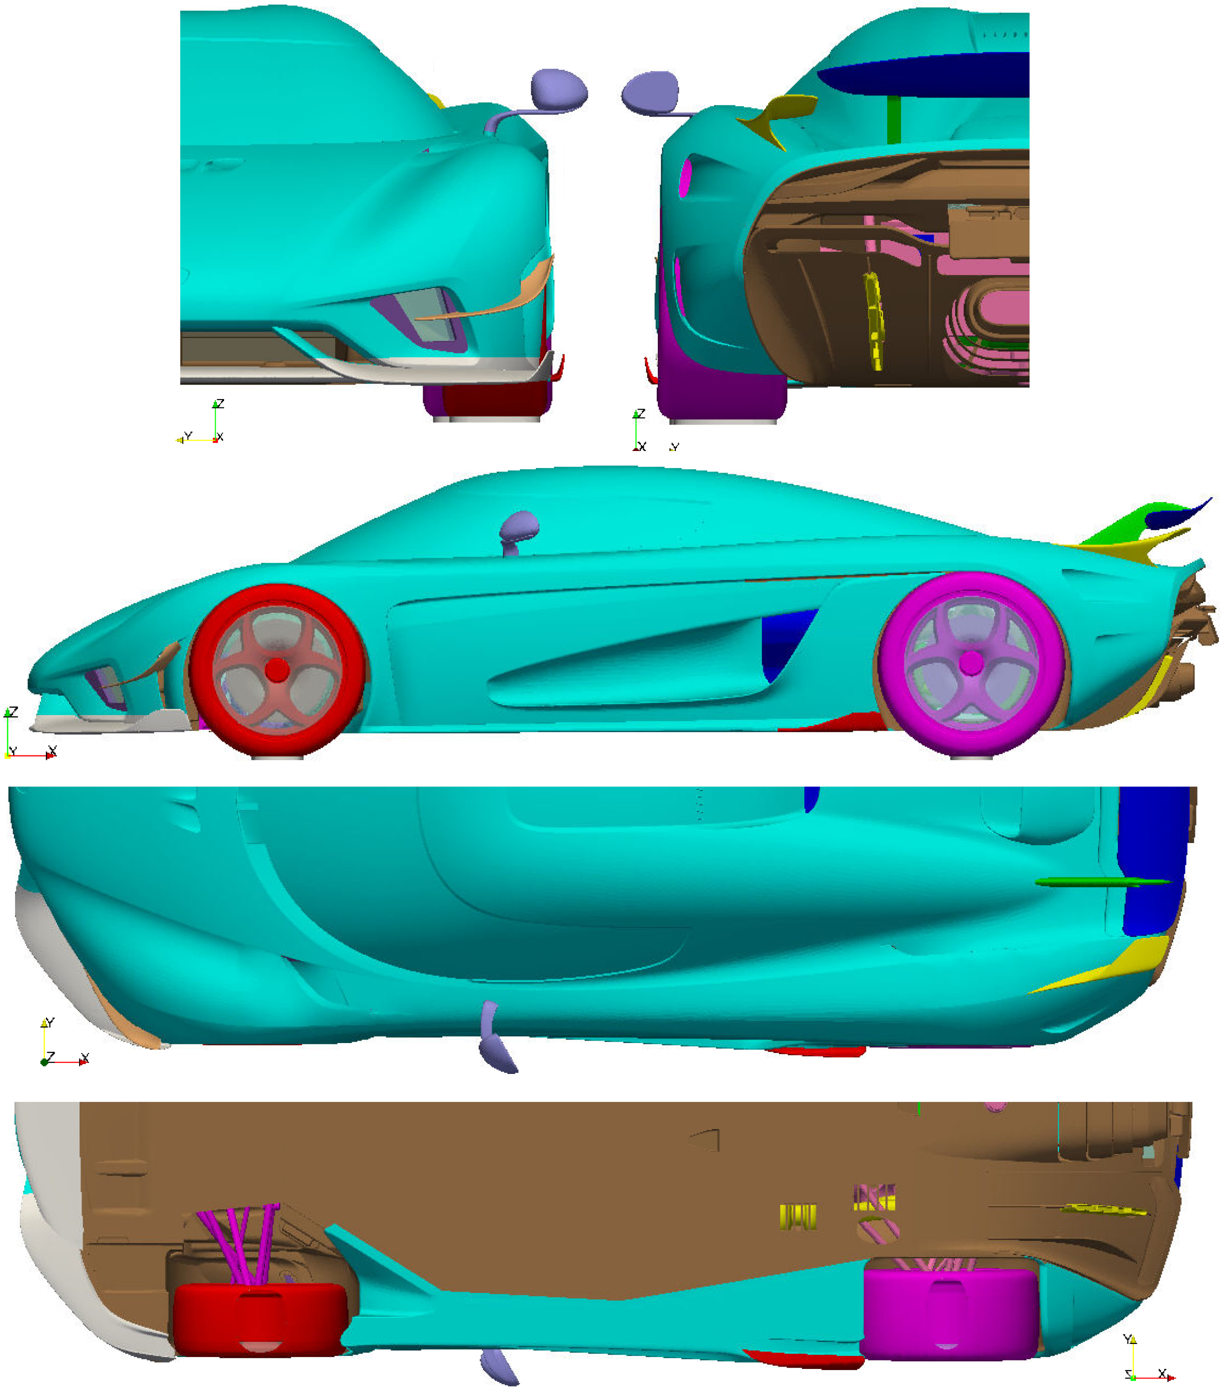

Figure 3.4: Pre-production Koenigsegg Regera with Ghost package by PID.

Regera Mesh

The Regera mesh too was composed of structured and unstructured cells. Hexa cells were applied in the far field domain. The tria surface mesh was chosen for building prism cells, while minimizing skewed and max angle cells on the geometry. The fluid domain is based on the dimension of the vehicle and can be observed in Figure 3.5.

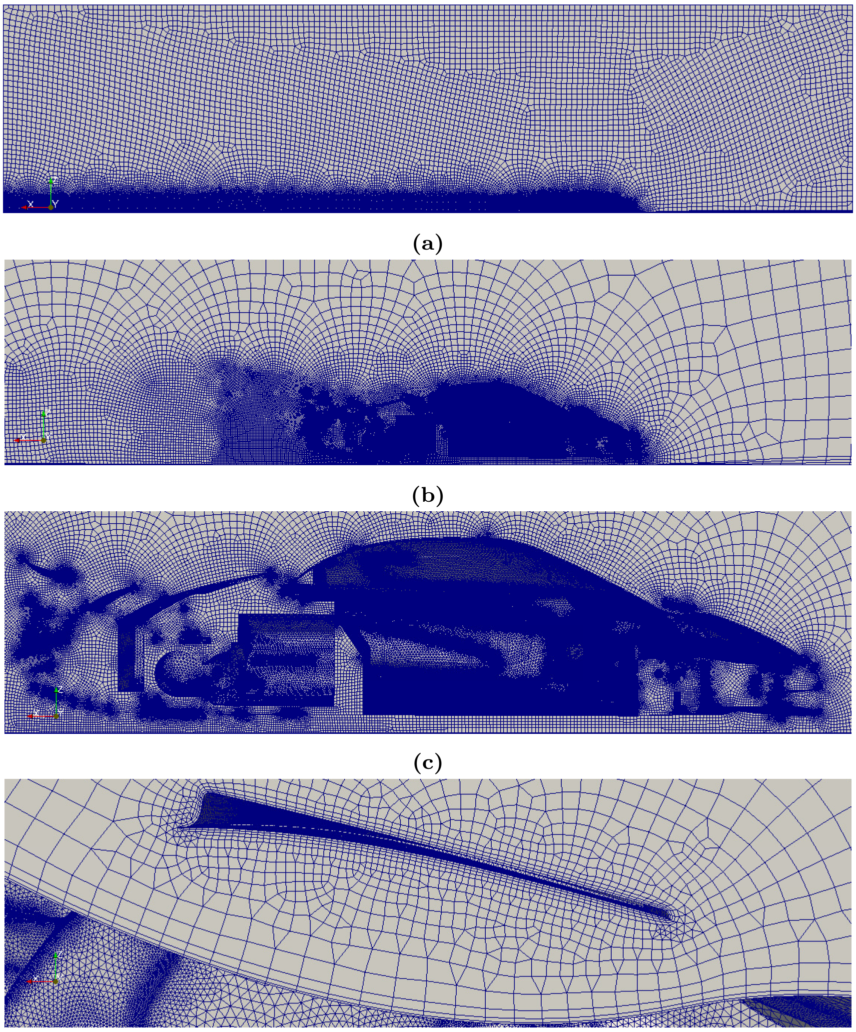
(a)
(b)
(c)
(d)

Figure 3.5: Figure (a) shows the full length and height of the flow domain at symmetry plane. Figure (b), (c), and (d) shows zoomed in snap-shots of the domain. notice Figure (c) shows convex and concave refinement. Figure (d) shows prism layers.

The proportion of the domain is much larger downstream of the vehicle to make sure that the wake flow properties are captured. The domain was revisited after computation to check that the pressure and velocity fields were not altered by wall interference. Figure 3.5b shows dense refinement in the area of curvature (convex/concave) as well as in the regions of free shear flows and high pressure gradient. This refinement was necessary to minimise non-orthogonality, skewness, max angles, and decreased chances of divergence. The refinement increased cell count significantly. Alternatively, the adjoint simulation would diverge due to lower mesh quality. Figure 3.5d shows four prism layers. The first prism layer height is 0.5 mm for most surfaces to reach a target value within $30 > y^+ > 60$. Variable prism layer height is applied to the different PIDs, i.e. mirrors, to capture an average y^+ of no more than 60. The height of the three outer prism layers are based on the aspect ratio between outer cells, with the fourth prism layer being 0.7 times the height of the neighbouring hexa cell. Ideally, more prism layers should have been applied. Yet, due to hardware limitations, only four prism cells were generated. The largest far flow cells had a cell length of 400 mm. This cell length could be decreased to resolve far-flow better, but it is of less importance than near surface flow when deriving sensitivity maps. Several additional small refinement volumes were created. Volume refinements were created in the shape of the wake with ANSA size boxes. The wake of the vehicle is highly unsteady, three-dimensional, irregular, and promotes a greater amount of diffusivity and mixing. Thus, a finer mesh is required in the wake region. Additional refinement was added where free shear flows were expected, in other words when high velocity external flow entered the engine bay region of low momentum flow. Attention to aerodynamic lift devices such as the front splitter, front diveplates, sidesplitter, vortex generators, rear winglet, and rear wing had the additional cell refined as well. Refinement was applied in the interface region between the wheels and the ground. Final refinement was done in areas where the adjoint simulation diverged. This was an iterative process. All refinement boxes ought to be finer if computational resources allow. The use of refinement volumes along with prism layers should ensure a high quality mesh keeping aspect ratio, skewness, non-orthogonality, and determinants to a minimum.

Domain Boundaries

Furthermore, the boundary conditions of far field walls were set to zero gradient and slip. The symmetry plane has a symmetry boundary condition. Inlet and outlet conditions were set as for the FSMW, although with a higher velocity and recalculated IC's for turbulence state variables. Only ground simulations were performed with the Regera mesh, namely moving ground, rotating wall on tires, and MRF on rims. Porous zones were applied over heat exhangers, fans, and intercooler.

3.2 Solver Setup

This section describes how the solver simulations were selected during this thesis work for both the FSMW and Regera mesh. This includes the type of solvers, discretization schemes, solver stability, and time steps. OpenFOAM v2006 were used for both solving primal and adjoint equations.

Boundary Conditions

A fixed-value inlet velocity normal to the inlet surface of 140 [kph] were applied for the FSMW, while 250 [kph] was applied for the Regera. The inlet velocity of the FSMW was chosen to enable comparison with other studies, while the simulation velocity of 250 [kph] was requested by Koenigsegg, since the Regera is designed to travel at an average of far higher velocities. Turbulence intensity, TI, was set at 1% and the eddy viscosity ratio, $\frac{\nu_t}{\nu}$, was set at 1.5, unless otherwise mentioned. These values were set as initial conditions, but developed independently towards a constant value for each cell. The outlet was set to a pressure outlet with a gauge pressure at 0 [Pa]. All surfaces of the vehicles were set to the no-slip wall condition. The wall under the vehicles was set to a moving no-slip wall with the same velocity components as given at the inlet. Additional boundary conditions were applied to imitate the rotating spokes, namely, the moving reference frames, MRF. These were volumes of fluid encapsuling the cells in between each spoke. These volumes were given rotational velocity around each hub center of, respectively, the front and rear axle. Additionally, the walls of tires and rims had rotational walls applied with the correct rotational frequency compared to the wheels circumference and the velocity condition given by the moving ground. In simulations of the FSMW with no MRFs, and smooth rims, the tire and rim boundary condition were set by angular wall velocity. During the exclusion of ground simulation, the tires and rims were set to stationary walls, and the ground wall was set as no-slip and stationary. During the exclusion of ground simulation, the tires and rims were set to stationary walls, and the ground wall was set as no-slip and stationary.
Naturally, the fluid was set to air with a kinematic viscosity, ν, of $1.5E\text{-}5\,[\mathrm{m^2s^{-2}}]$, with a reference pressure of 0 [Pa] for the FSMW. Further simulation constants of the Regera are kept confidential.

Primal Solver Setup

Turbulence-energy equation models were run for the FSMW mesh, including Spalart-Allmaras, k-ϵ, Realizable k-ϵ, and k-ω SST, each with a different purpose. The Spalart-Allmaras model was chosen for proving the stability and robustness during the operation of the adjointOptimisationFoam. Two-equation models were run to test the impact of the frozen turbulence assumption when solving the adjoint equation. The frozen turbulence assumption is the only available option when modelling adjoint turbulence by two-equation turbulence models in OpenFOAM v2006. Experience proved instability of the two-equation turbulence model when solving the adjoint PDEs. The primal and adjoint equations were solved as incompressible steady state equations using the simple algorithm. However, one attempt of incompressible unsteady state equations with the piso algorithm was conducted to investigate the impact of unsteady flow on the sensitivity maps and to achieve the monitoring of oscillating force coefficients (unsteady), as opposed

to discontinuous fluctuating (steady). The unsteady simulation ran subsequently after a fully converged steady state simulation with a time step of $4E$-4 [s] for 10,000 iterations, where the convergence of force coefficients and residuals were achieved. The unsteady simulations were below a Courant number of 5 at all times.
All primal state variables presented in this report were run by second-order schemes, when using the Spalart-Allmaras turbulence model. First order schemes are anticipated to be numerically diffusive, which means they will underestimate the forces and smear the gradients. Nonetheless, first-order schemes are robust.
Wall distance calculation was performed with the advection-diffusion method. Averaging of the primal flow was performed. All primal steady state simulations approach a state of convergence in the order of 7,000-8,000 iterations. Averaging of the steady state primal flow began at iteration 7,000 and ran to iteration 10,000 for the primal simulation, unless otherwise mentioned. Averaging of the field equations of the unsteady primal simulation was performed over the final 2 [s] of the 4 [s] flow simulation. All simulations were run with OpenFOAM's mesh renumbering and potentialFoam solver, prior to solving the primal flow equations.

Simplification

One must be aware of compressibility effects, which occur at Ma $>$ 0.3 [18]. The inlet velocity, of the medium air, will be at a fixed value to Ma = 0.215 which should cause incompressible conditions in the far-field domain. The compressibility effect in the vicinity of the Regera will be monitored continuously. The fluid may be considered an incompressible Newtonian fluid. The flow field is assumed to be steady. Still, this author would like to stress once again that the flow around a ground vehicle is unsteady, due to the flow separation in the wake, external mirrors, rotating wheels, and vortex shedding of aerodynamic devices. Also, it was assumed that the flow is symmetric along the longitudinal axis to reduce the computational domain, although the cross flow due to trailing vortices makes for a highly three-dimensional and asymmetric flow. Finally, for the sake of reducing the number of simulations, cross flow analysis and the like shall be postponed to another study. Thus, this study will have an inlet velocity with zero angle of cross flow.

Primal Convergence Assessment

A high-quality mesh is of utmost importance and may be the deciding factor of whether an adjoint simulation can run or not. Additionally, a fully converged primal and adjoint equation is a crucial criteria to provide reliable sensitivity maps or shape optimisation. Theory also proved that residuals should be zero for the derived formulations of the adjoint equations to be true [22]. Convergence was based on two conditions: (1) Both primal residuals of velocity, continuity, pressure, and turbulence variables have reached a quasi-constant value below $1E$-4. Primal pressure residuals were larger due to the actual flow being unsteady and were assumed converged when reaching a quasi-constant value around $1E$-3. (2) Force coefficients, C_d, C_{lf}, and C_{lr} should be leveled out, only leaving a minimum fluctuation from a converged mean value. During the work of this thesis, the simulations were primarily run 10,000 iterations, unless otherwise mentioned, to facilitate converged residuals and force coefficients. It was also found that convergence of the primal solver impacted directly on the stability of the adjoint solver. In other words, if less than 5,000 primal iterations were run, the adjoint solver could diverge.

Adjoint Solver Setup

The adjoint ATC-term, described in Equation 2.3, makes solving the adjoint solver diverge for for complex applications. OpenFOAM v2006 has several options available to deal with the unstable ATC term. First, the ATC term can be modelled by three options. The most robust modelling is done by excluding the ATC term from the adjoint momentum equation during the solution of the the adjoint PDEs. This is at the cost of accuracy. A second option is the standard model which computes the ATC term by $u_j \frac{\partial v_j}{\partial x_i}$, where v_j and u_j respectively are the primal and adjoint velocity vectors [11]. This option has been used for all adjoint simulations during this thesis work, as it offers the best trade-off between stability and accuracy. The most accurate ATC model is denoted $UaGradU$ in OpenFOAM, and was not used during this thesis work, since it is inherently unstable. The reader is referred to [11] for further details.
The standard ATC model proved itself unstable. Further functions related to the ATC term could be activated in order to facilitate convergence. The first was *extraConvection* which adds numerical dissipation. This filter was not used during this thesis work. Next, zeroing of patch types or zones was possible. Zeroing of patch types of walls and patches were performed since significant divergence was observed at these patch types. Next, a Laplacian-like filter were applied to the walls and patches by smoothing three times[11].
Attempts to solve the FSMW mesh by two-equation adjoint turbulence models were made. Though, this requires additional numerical filtering of the ATC-term in the adjoint momentum equation and first order schemes of the turbulence equations, even in the primal flow, due to the frozen adjoint turbulence assumption. Consequently, two-equation adjoint simulations were considered less accurate. Instead, two-equation models were applied in validation post-morphing to study whether morphing with the Spalart-Allmaras turbulence model was accurate.
All adjoint state variables presented in this report were run on second-order schemes, apart from the adjoint turbulence variable of the Spalart-Allmaras model, $\tilde{\nu}_a$, to ensure convergence [11]. The reader is advised to read the User Manual of adjointOptimisationFoam in case of recommendation of solver settings [11]. Specific findings of crucial adjoint settings will be mentioned in the following chapters of Results and Discussion.

Adjoint Convergence Assessment

Convergence of the adjoint solver was monitored by adjoint residuals, as well as checking the adjoint state variables, the adjoint distance field, and adjoint mesh movement field. It was found that low adjoint residuals were not equal to convergence as the sensitivity maps were sensitive to further iterations, even if the residuals were below $1E$-5, see subsection A.3.3.
Another measure of convergence would be to visually inspect the adjoint state variables to see if they were fully developed. During visualisation of adjoint fields, a so-called adjoint wake was present in front of the vehicle, as observed in Figure 3.6. This was due to the backward calculation adjoint convective term. The adjoint wake is also expected to cause a region of several cells with higher residuals, due to coarser refinement in the upstream direction. This was less of a concern as the adjoint wake developed far from the surface of interest. During the testing of the meshes, it was found that the adjoint solver reacted delicately to imperfections in the mesh. This was in the region of flow separation, i.e. in regions transitioning from wall-bounded flows to free shear flows. This observation was also made by [20].

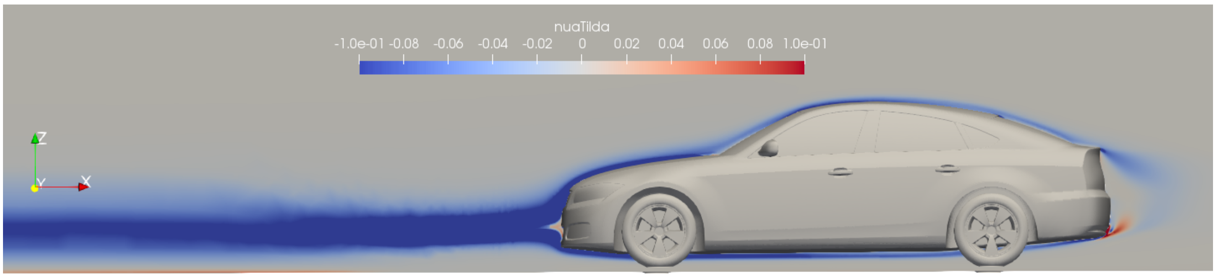

Figure 3.6: Adjoint Spalart-Allmaras turbulence state variable, $\tilde{\nu}_a$, field of the FSMW.

3.3 Post-Processing

The primal flow fields of the pressure coefficient, skin friction coefficient, and velocity, along with the adjoint flow fields, were analysed. Max values were monitored to determine whether the adjoint results seemed physically reasonable. Aforementioned precautions of measuring convergence were taken.

Sensitivity Maps

A sensitivity map visualises the variation of the objective function with respect to the wall-normal displacement of the boundary nodes of a given geometry. In this thesis, sensitivity maps are distributed over the surface of the vehicles. Thus, the sensitivity maps may highlight regions of aerodynamic optimisation potential. The sensitivity maps of the FSMW model seemed partially unreliable at first sight. However, through a parametric study, in section A.3, flow features and patterns were classified according to the mesh and solver setup.

The colour bars of the sensitivity maps in this Master's thesis have red contours identifying wall-normal outwards (positive) displacement and blue contours indicating wall-normal inwards (negative) displacement to minimise a given function object. All colour bars of sensitivity maps will be denoted by *pointSensNormalas1ESI*. *point* stands for derivation of sensitivity derivatives with respect to points of the boundary nodes, and not the faces, to respect the analytical geometry to the highest possible degree. *SensNormal* describes the sensitivity derivatives are visualised wall-normal to the boundary nodes. *as1* is the adjoint solver name. In this case there is only one adjoint solver. *ESI* stands for the use of the Enhance Surface Integral formulation. All sensitivity derivatives were computed with the Enhanced Surface Integral formulation. The Enhanced Surface Integral formulation keeps the speed of the Surface Integral formulation and approaches the accuracy of the Field Integral formulation, while still being independent of the number of design variables [17].

3.4 Morphing

As an alternative to the sensitivity map, which guides the engineer to morph the geometry manually, a semi-automated shape optimisation approach can be utilised to morph the mesh. In this case, the shape optimisation runs in cycles, visualised in Figure 2.2. During this approach, the variation of the objective function is calculated with respect to the normal displacement of the control points parameterisation of a selected region with

optimisation potential. The control points may be defined on the basis of a sensitivity map or by analysis of the primal flow by an experienced aerodynamicist.
Shape optimisation performed during this thesis work the is based on a volumetric B-splines morpher using Field Integrals. The volumetric B-spline morpher moves the mesh using an analytic/algebraic formula [11]. Using the Field Integral formulation yield higher accuracy sensitivity derivatives, and in some case avoid error of wrong sign and magnitude [17]. Consideration of the number of control points defined in the three spatial directions must be done, as the computation of too many control points with the Field Integral formulation are not temporally feasible.
Inconsistencies between the sensitivities calculated with respect to normal displacement of (1) the wall boundary nodes or (2) the control points were observed. This is due to the fact that the sensitivity maps are computed by the Enhanced Surface Integral formulation, assuming a grid displacement model of a set of Laplace-based PDEs. The max displacement of the volumetric B-spline morpher is defined to limit a shape deformation parameter, η, to control the morpher in terms of robustness and speed.
The *Conjugate Gradient* method is chosen when updating the design variables. The conjugate gradient method is a gradient based iterative algorithm. It is significantly faster than the *Steepest Descent*, but it can still tolerate discrepancies of the sensitivity derivatives [11]. The *Line Search* method is used for a so-called inner iterative optimisation that calculates the most optimal step, η, in each iteration of updating the design variables. This is useful to speed up the convergence towards a local minimum of the given function object.

4

Results and Discussion

This chapter provides validation, analysis, and discussion of the FSMW model in terms of the primal field and sensitivity maps. Subsequently, the Regera simulation results of the primal field and sensitivity maps are analysed and discussed. Morphing results of the FSMW and Regera are presented and discussed. Finally, a discussion of the adjointOptimisationFoam application is discussed and compared to conventional aerodynamic development with CFD. Results are based on ANSA meshes, unless otherwise mentioned.

4.1 FSMW Primal Analysis

Validation of the primal solution established a foundation for the adjoint optimisation method. Post-processing of the primal field was crucial for the credibility of the adjoint sensitivity maps. The validation was computed as steady state at 140 kph, 20[^{o}C], with moving road and wall-tangential velocity boundary conditions at wheels.

Experimental vs Numerical Setup of the FSMW

References of numerical and experimental studies have been investigated. The reader is referred to read each paper which explains uncertainty and individual setups to a greater detail of [7, 10, 13, 25, 27, 29]. Experimental uncertainties in the referred papers include:

- Blockage effect of the wind tunnels
- Efficiency of the boundary layer scoops
- Disturbances from models support systems (struts on body and wheels)
- Different rim configurations
- Interaction of rolling road system
- General measurement error
- Stationary or rotating wheels
- Wind tunnels capable of different max Reynolds numbers

The numerical studies also include differences in setups such as:

- Application of rotating walls and moving reference frames
- Different cell counts (herein variable surface quality of the driver model)
- Boundary layer addition of different quality
- Full car versus symmetry car
- Different turbulence models
- Unsteady versus steady
- Modelling of wind tunnel domain or not
- Different Reynolds number

The largest sources of deviations between experimental and numerical solutions were found to be artificial pressure recovery over the roof and sides of the FSMW respectively due to central supporting struts and wheel supporting struts found in [7, 13]. Additionally, the type of moving ground simulation as well as rotating wheels, or not, caused deviation of force coefficients. This author has not expect to match exact qualitative correlation of force coefficients down to third decimal, nor has perfect quantitative correlation of C_p distribution been expected. Rather, this author has aimed for a solid simulation model capable of predicting physical flow properties and force coefficient, similar to those found in well documented studies.

FSMW Mesh Convergence

A preliminary mesh convergence study was conducted with steady state, incompressible, *Realizable* k-ϵ turbulence model, rotational BC's at tires, open rims with MRF, and moving ground simulation at 140 kph and $20\,^oC$. Force coefficients include C_d, C_{lf}, and C_{lr}, in Figure 4.1. The subscripts f and r stands for front and rear, respectively. C_d, C_{lf}, and C_{lr} values are averaged over the last 3,000 iterations out of 10,000 iterations, denoted by the solid lines in Figure 4.1. The bars serve to indicate maximum and minimum values recorded during this averaging interval.

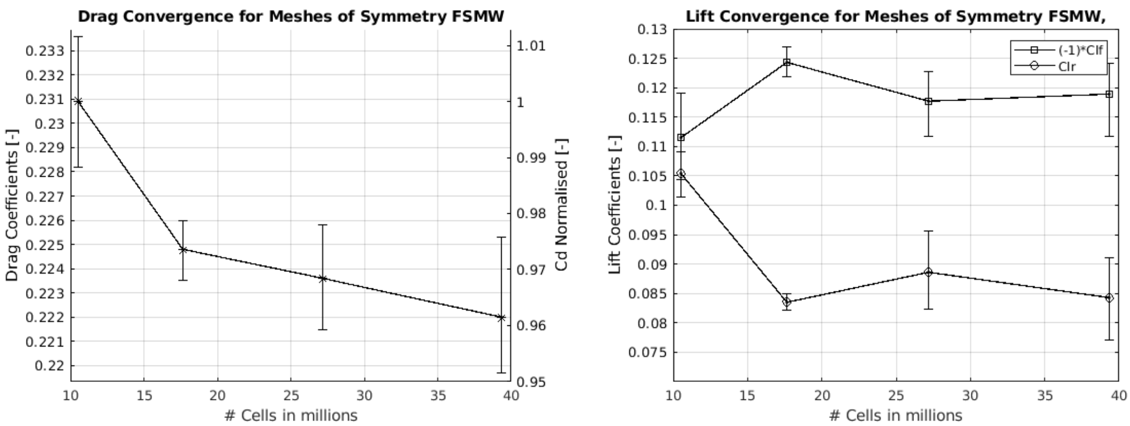

(a) Mesh convergence study of drag coefficients.

(b) Mesh convergence study of lift coefficients.

Figure 4.1: Mesh convergence study, symmetry FSMW model.

The mesh convergence study represents meshes with 10M-40M cells. Meshes exceeding 40M cells are not included due to stated thesis limitations. Additionally, it has been found that more than 50M cells of the FSMW model affect the C_d value marginally [29]. The mesh convergence was carried out for the right hand side of the FSMW found in Figure 3.2 and in steady state conditions. Thus, asymmetric temporal flow properties, i.e. in the wake region, are excluded and hence accuracy is decreased. The C_d, C_{lf}, and C_{lr} values in Figure 4.1 inherit uncertainty due to certain meshes slightly exceeding the internal cell skewness value of four. Finally, the meshes inherit different resolution of boundary layers due to variable cell base size, even though the first layer height was kept constant. It is observed that meshes above 20M cells approach a final value closely. During the thesis work, this author shifted from OpenFOAM's snappyHexMesh to the commercial mesher, ANSA. No mesh convergence was carried out for ANSA meshes due to time-constraints. Still, experience with meshes created in ANSA of the FSMW model proves close correlation to Figure 4.1.

FSMW Reynolds Sweep

A parametric study of various Reynolds numbers was performed in order to compare the simulation model with reference values and test the solver at high Reynolds flow. The Reynolds sweep of the symmetry FSMW was performed with a snappyHexMesh mesh of 22M cells with incompressible steady state *Realizable* k-ϵ turbulence modelling, again with rotational BC's at tires, open rims with MRF, and moving ground simulation. The same mesh was used for all Reynolds sweep computations. C_d, C_{lf}, and C_{lr} values are averaged over the last 3,000 iterations out of 10,000 iterations and is denoted by the solid lines in Figure 4.2. The bars indicate maximum and minimum values recorded within the averaging interval.

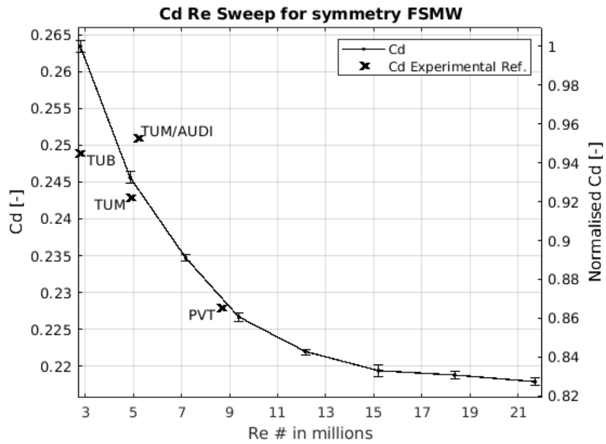

Figure 4.2: Reynolds sweep of FSMW model.

At the highest Reynolds number simulation, y^+_{max}~140 are observed on surfaces of the FSMW, and y^+_{max}~250 are observed on the wheels. It is questionable whether the boundary layer is sufficiently resolved for higher Reynolds numbers. The Reynolds sweep in Figure 4.2 serve as an indicator of Reynolds sensitivity for the symmetry FSMW. It is observed in Figure 4.2 that the experimental values are showing a similarly decreasing trend of C_d values as the Reynolds number increase. Good correlation has been found with experimental values of the Technical University of Munich (TUM) and TUM/AUDI respectively [7, 13] and experimental PVT value from Volvo Cars Wind Tunnel [5, 29]. Note that the Reynolds number from the PVT source is assumed to be around the order of Re 8.7M-12.2M, as simulated in [29]. Additionally, Volvo Cars Wind Tunnel tests are usually carried out at free stream speeds of 140 kph, for full scale cars. The papers of TU Berlin (TUB) and TUM/AUDI [7, 25] show similar Reynolds sweeps which converge at lower Reynolds numbers of $2.5M$-$5.5M$. This particular numerical setup may over-predict C_d values slightly. Different flow structures are expected for different Reynolds numbers. Although, at increasing Reynolds numbers the flow structures may approach self-similarity, as observed in Figure 4.2.

Yplus

The following figures show the primal solution of a symmetry FSMW model simulated with steady state, incompressible, Spalart-Allmaras turbulence model, moving ground simulation and rational wall boundary condtions, at 140 kph and 20$[^oC]$. The logarithmic layer, and potentially the viscous sub-layer, have been modelled by turbulence wall functions and resolved by blending functions specific to each turbulence model [8]. The reader is advised to read the OpenFOAM API Guide of wall functions for further details of blending and wall functions. Average y^+ values are observed on the surface of the FSMW in Figure 4.3. y^+ values are low in stagnated areas and where flow has separated. y^+ values are in the range of 30-70 on the body of FSMW. Increase in y^+ occurs in regions with curvature, due to velocity gradient.

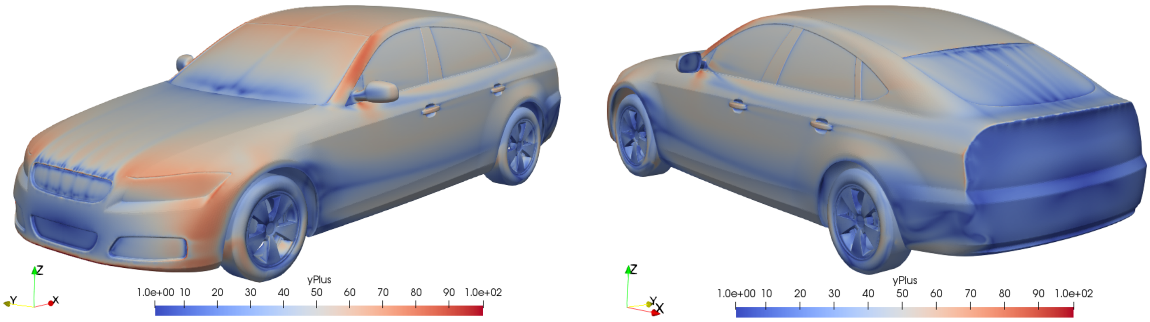

(a) **(b)**

Figure 4.3: y^+ distribution of symmetry FSMW model.

Skin Friction Coefficient

The y^+ distribution correlates with the wall shear stress, and hence the skin friction coefficient. The skin friction coefficient, plotted in Figure 4.4, gives an indication of separation or stagnation regions.

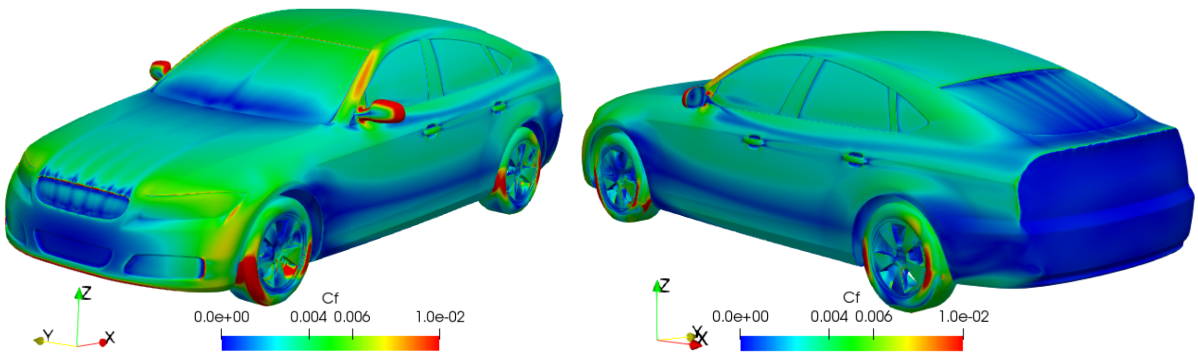

(a) **(b)**

Figure 4.4: Skin friction distribution of FSMW model.

Low or zero skin friction coefficient is observed in Figure 4.4a at the front bumper and windshield base area due to stagnation. An increase in the skin friction coefficient occurs at the mirror housing, A-pillar, front fenders, and tires due to curvature increasing the velocity gradient. Vortex shedding from the A-pillar and mirror induce changes to the skin friction coefficient in Figure 4.4b. The A-pillar vortices exist across the side window, and roof, followed by merging with C-pillar vortices. The light blue longitudinal lines over the backlight are due to vortices, created over the sharp trailing edge of the roof. The vortices rotate around the transverse axis, re-energizing the flow by mixing momentum into the boundary layer. The A-pillar, C-pillar, and trailing roof vortices may delay separation over the backlight and trunk surface during an adverse pressure gradient flow.

The flow over the backlight and upper surface of the trunk might be separated partially in Figure 4.4b. Next, the flow detaches in the cowl area, and aft of the front and rear wheels. The mirror wake and door handles leave streaks of low momentum flow. The flow may recover partially aft of the front wheels, mirror and door handles, as the skin coefficient increase downstream. Strong separation occur aft of the rear wheel. Thus, the FSMW may have too much boat-tailing (inward slope of surface). Finally, the rear base is completely separated, as expected. The trunk trailing edge aids in a clean separation zone where skin friction increases. Therefore, the trunk edge may be a source of potential drag optimisation. Finally, drag optimisation may be achieved by altering surfaces of the mirror arm and mirror housing, windshield/cowl area, wheel arches and fenders, rear window slope, boat-tailing and A-pillar shape. These areas are expected to be highlighted in the sensitivity maps.

Static Pressure Coefficient

Figure 4.5 shows static pressure distribution over the FSMW surface and wheels and the velocity distribution of a plane normal displaced 5 [cm] from the symmetry plane. Note that Figure 4.5a and Figure 4.5b have different scaled C_p bars, to highlight relevant flow features of the FSMW.

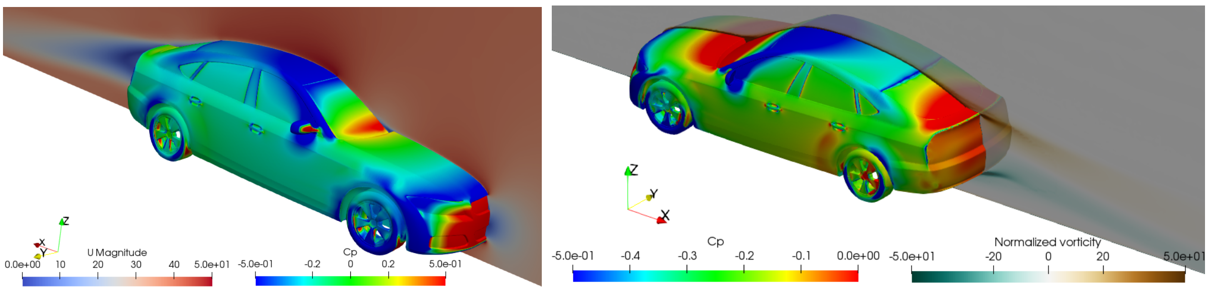

(a) (b)

Figure 4.5: Static pressure distribution with symmetry-plane velocity in Figure 4.5a and symmetry plane normalised vorticity in Figure 4.5b.

Areas of stagnation are observed on mirror housing and arm, windshield base, front bumper and grill area, and tires. Separation is observed at the cowl area of Figure 4.5a and over the lower backlight in Figure 4.5b. Full separation occur at the base of the FSMW. The downwash of the wake flow is observed. Balancing the wake could lead to drag improvements and decrease lift over the rear axle. In Figure 4.5b, the flow over the backlight remains partially attached and a shear layer is generated in the boundary layer seen by the dark orange area over the rear window in Figure 4.5b. The shear layer convects downstream, thickens, followed by a separation at the trailing edge of the trunk. Two free shear layers emerge, visualised by the brown and green regions, which roll up and form vortices in Figure 4.5b. The upper and lower shear layer merge together and form a common wake profile downstream at this point. The upper shear layer may be stronger than the lower one, confirming the unbalanced wake profile. Additionally, vortex formation is observed in low pressure zones of the A-pillar and over the C-pillar. Low pressure over the A-pillar may decrease drag, since the pressure is partially normal to the heading of the FSMW. Low pressure zones are also observed on the leading zones of the front fenders, hood, and roof where lift is induced. In Figure 4.5a front and rear fenders,

aft of the wheels has a low pressure distribution due to vortex shedding of wheels. Finally, the leading zone of the rear fender has a higher pressure zone than the remaining side of the FSMW and may be a source of drag, if not enhancing flow over the rear wheels. The base pressure is nearly uniform.
Validation of the pressure distribution of the FSMW simulation can be seen in section A.2.

Velocity Field

Figure 4.6 visualises wake flow and velocity distribution around the FSMW. Wakes exist downstream of the mirror, front and rear wheels, and the base. The mirror wake is not recovered and convect downstream to merge with the rear wake. The sensitivity map is expected to decrease the size and shape of the mirror housing. The wake of the front wheel is not diffused before reaching the rear wheel. Therefore, wake flow along the lower vicinity of the side of the FSMW inherits partially separated flow and should be visualised in the sensitivity map. The rear wheels are forming a stronger wake profile than the front wheels. The rear wheel wake merges with the base wake to form a region of separated and highly three-dimensional flow. The aforementioned partially separated flow over the backlight and trunk is confirmed. Hence, the sensitivity map is expected to highlight this area for drag optimisation.

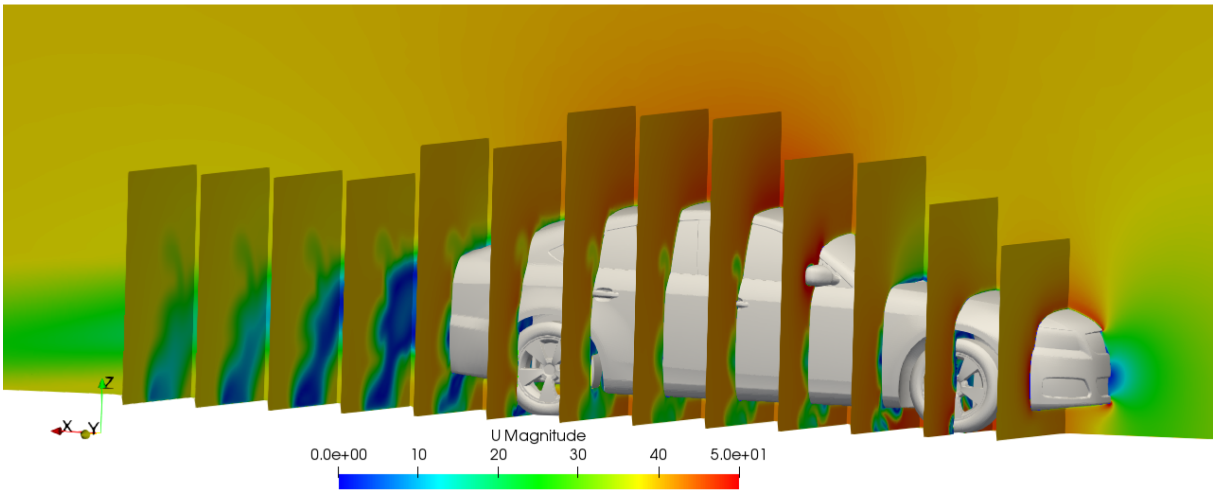

Figure 4.6: Velocity distribution of FSMW model.

Forces and Pressure Coefficients

The pressure distribution of the FSMW was monitored for all turbulence models and compared to experimental C_p plots of [13]. The pressure distribution was averaged over the last 3,000 iterations out of 10,000 iterations during primal solving. The experimental setup causes deviation from the numerical simulations. Most notably, the experimental pressure over the roof in Figure 4.7a is higher due to stagnated flow ahead of the central supportive strut, fixing the scaled model in the wind tunnel in [13]. This is confirmed by sources of numerical simulations [1, 2, 5, 9]. Measurement of experimental pressure distribution in the wheel housing region is not performed. Interference from struts supporting the wheels is expected to create artificial experimental pressure conditions. Next, sharp numerical pressure drops have been observe on all figures of Figure 4.7. This is due to discontinuities in the mesh caused by creases in the surface, i.e. transition from windshield trailing edge to roof leading edge. These discontinuous pressure drops should be disregarded when comparing to experimental data. The turbulence models predict the trends of the pressure distribution correctly, especially in continuous regions. Yet, aft of $x = 2.5\,[\text{m}]$, the turbulence models differ.

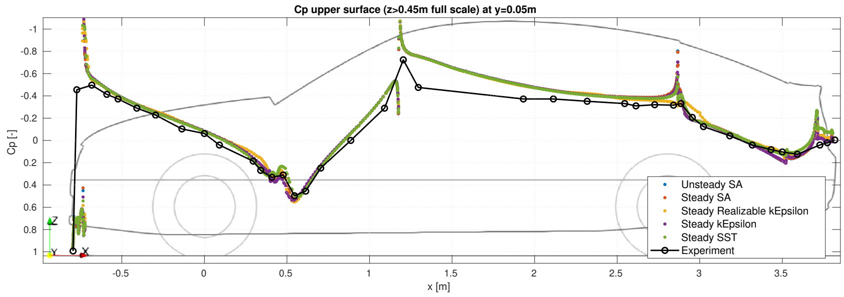

(a) $y = 5\,[\text{cm}]$ and $z \geq 45\,[\text{cm}]$.

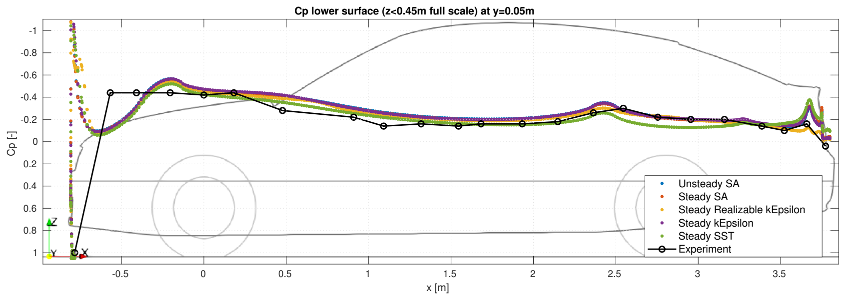

(b) $y = 5\,[\text{cm}]$ and $z \leq 45\,[\text{cm}]$.

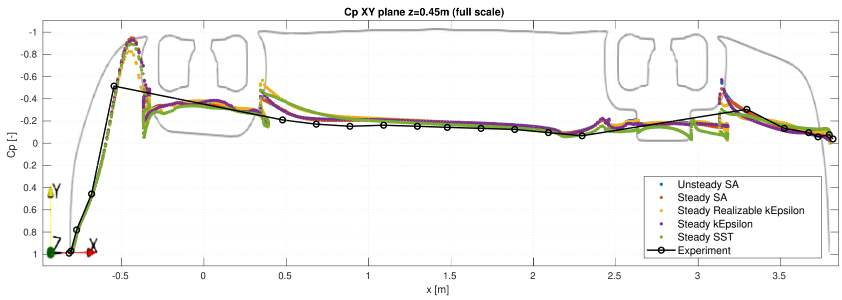

(c) $z = 45\,[\text{cm}]$.

Figure 4.7: Numerical pressure distribution of different turbulence models compared to experimental data from [13].

An unsteady Spalart-Allmaras simulation was run and can be seen in Figure 4.7, Figure 4.8, and Figure 4.9. Studies revealed an increase of accuracy of the sensitivity maps when using an unsteady averaged primal flow [3, 20]. The unsteady pressure distribution does not differ significantly from the steady. The quantitative difference between a steady primal flow and the unsteady averaged primal flow is studied in a cumulative force coefficient plot, as observed in Figure 4.8. The cumulative drag coefficient predict, sources of drag in regions of the rear wheels, the backlight, and the trailing edge of the FSMW. More prominent qualitative differences in lift coefficients are observed between the unsteady and steady simulation. The unsteady solution reach lift coefficient closer experimental values aft of the front wheels. Higher accuracy of the unsteady simulation is expected in regions of adverse pressure gradient. Trends are similar with downforce over the front axle and lift over the rear axle.

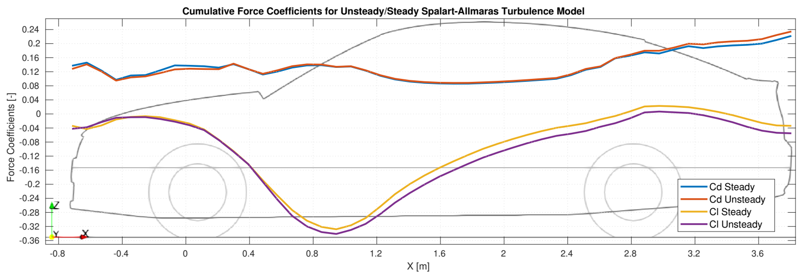

Figure 4.8: Cumulative force coefficients of lift and drag of unsteady/steady.

Figure 4.9 present numerical force coefficient solutions of the different turbulence models with a dot. This is useful when comparing the Spalart-Allmaras turbulence model with two-equation turbulence models. Turbulence models are denoted *U-SA*, *SA*, *RKE*, *KE*, and *SST*, which respectively are short for Unsteady *Spalart-Allmaras*, steady *Spalart-Allmaras*, steady *Realizable* k-ϵ, steady k-ϵ, and steady k-ω *SST*. Mean values are denoted by a dot. Minimum and maximum deviation within the averaging of last 3,000 iterations are plotted by an error bar.

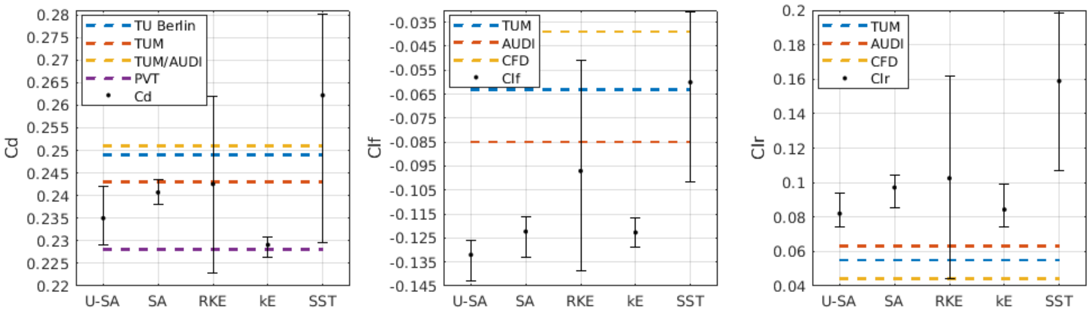

Figure 4.9: Force coefficients of primal flow averaged over last 3,000 iterations out of 10,000 iterations.

The k-ω *SST* model is the only turbulence model to exceed experimental C_d value bounds. The *Realizable* and *SST* turbulence models have a large deviation of minimum and maximum values recorded within the averaging interval. This may be due to less numerical damping of the *Realizable* and *SST* turbulence models compared to the *Spalart-Allmaras* and k-ϵ turbulence models. The large maximum and minimum deviations from the mean indicate that the actual flow is unsteady. The Spalart-Allmaras turbulence model corresponds well with all experimental C_d values. The Spalart-Allmaras turbulence model predicts lift coefficients with a less accuracy. Generally, all turbulence models predict lift coefficients with some deviation from the experimental values. This trend too is observed in [2, 9]. The deviation from the experimental values could be measurement related since the lift values are small in magnitude.

Understanding the primal flow solution is crucial for understanding the results of the adjoint flow solution and the sensitivity maps.

4.2 FSMW Adjoint Analysis

With the validation and understanding of the primal solution, the thesis work continued with the adjoint solver and generation of sensitivity maps. Sensitivity maps posed in this thesis are given as point sensitivity maps, not faces. Normal displacement of the boundaries are calculated with respect to points. This is a more accurate option, as points may respect the analytical geometry behind the mesh to a higher degree. All sensitivity maps presented in this chapter is based on primal and adjoint steady state, incompressible, Spalart-Allmaras turbulence modelling.
Note that red colours of the sensitivity maps in this thesis signify wall-normal outwards (positive) displacement. Vice versa, blue colours signify wall-normal inwards (negative) displacement. Sensitivity maps presented of the FSMW are a result of the objective function set to drag minimisation. The turbulence modelling for both the primal and adjoint field is the Spalart-Allmaras turbulence model. Furthermore, the Spalart-Allmaras turbulence model is the only turbulence model capable of modelling adjoint turbulence equations in OpenFOAM v2006. Two-equation models such as k-ϵ and k-ω SST are also available, although with frozen turbulence assumed. Therefore, this is not the only solution of a sensitivity map to the FSMW, as the reader will see in the following subsection 4.2.1.

Figure 4.10 shows a sensitivity map of the FSMW, with steady state, incompressible, Spalart-Allmaras turbulence modelling, with open rims, moving ground simulation, rotating wheels boundaries, and moving reference frames in each wheel set for 140 kph. The sensitivity map is from the same simulation as in the aforementioned analysis/validation of the primal flow.

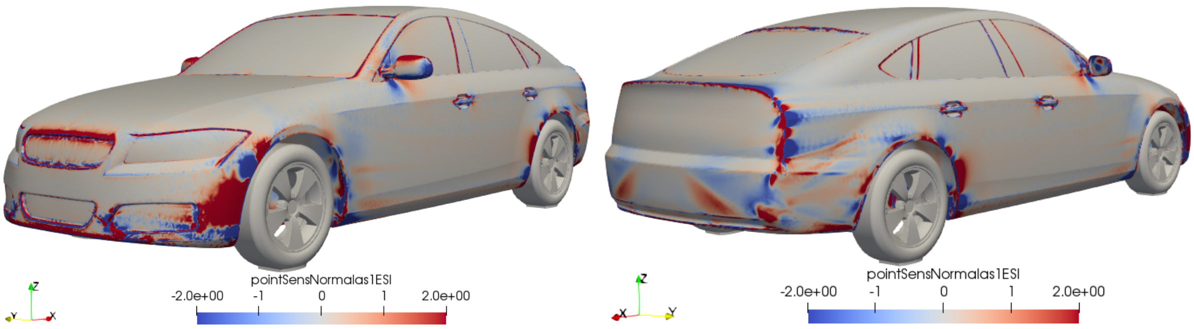

(a) (b)

Figure 4.10: Sensitivity map of the FSMW model.

Firstly, the reader might notice the red and blue highlights of the creases surrounding all edges of the greenhouse, as well as head-light/fender edge, and grill/bonnet edge. As expected, the sensitivity map is informing that these creases are a source of drag, and should be smoothed. This observation was also made in Figure 4.4, where the skin friction coefficient was approaching zero at the creases. More parallels to the skin friction plots can be drawn to the trailing edge/crease of the side grill. The sharp edge at the side grill causes increased skin friction and changes to the flow downstream. The sensitivity map shows that this region should be displaced wall normal outwards possibly to decrease separation which can be observed at the front fender, upstream of the front wheel, and downstream of the side grill, observed in Figure 4.6. Next, the sensitivity maps in the vicinity of the wheel arches seem chaotic. This is due to chaotic flow when simulating rotating wall boundaries and moving reference frames on the wheels, proved in subsection A.3.5 and subsection A.3.6. The leading edge of the rear fender shows larger region of wall-normal outwards (red) displacement. This may be related to separated flow over the lower part of the doors observed in the skin friction and velocity plots, respectively Figure 4.4 and Figure 4.6 of the primal solution. The cowl area and region aft of the rear wheel inherit very chaotic flow. As a result, it is difficult to justify reasonable changes. To get a higher quality sensitivity map of regions with separated flow, unsteady full car simulation with averaging should be conducted. The sensitivity map in Figure 4.10b recommends an extension of the trailing edge of the car along with a partial decrease in trunk slope/angle, to decrease drag. This is plausible as these changes would decrease the base area of the rear end. The mirror arm and mirror housing are subjects to change in both wall-normal outwards and inwards displacement. The wake of the mirror, observed in Figure 4.6, shows a greater loss of momentum. To minimise drag, the mirror sensitivity map was expected to show wall-normal inwards displacement, to decrease mirror frontal area. Conversely, monitoring forces acting on the mirror revealed upwash in the wake. This may explain why the sensitivity map distributes the colours differently.

Altogether, the sensitivity map highlights areas of the FSMW, where drag optimisation may be improved, namely creases, mirror, covering of wheels, and extension of trailing edge. Yet, the sensitivity map also proved counter intuitive patterns on the wheel arches, and in the regions aft of wheels. Skepticism in these areas are advised. Next, the sensitivity map is influenced by the mesh and simulation setup, as concluded in the following subsection 4.2.1.

4.2.1 Understanding the Sensitivity Map

A parametric study of the sensitivity map was conducted in order to investigate the robustness and sensitivity.

Iterations and Averaging Window

The sensitivity maps are closely related with number of iterations performed along with averaging of the primal PDEs. Thus, it is crucial to obtain a self-similar flow and low residuals during the primal solver, if consistent and reliable sensitivity maps are expected. The parametric study, found in subsection A.3.3, shows that great changes of sensitivity derivative magnitude is occurring respectively within the 5,000 first iterations of the primal and adjoint solver. A second and third simulation was run for reference of respectively 10,000 and 15,000 iterations each when solving the primal and adjoint PDEs. Significant changes in magnitude and position of the sensitivity gradients were observed in the areas of front and rear fenders, creases, and trailing edge of FSMW.

Another Chalmers thesis reviewed sensitivity maps vs. iterations [28]. The thesis found that at least 5,000 iterations of the adjoint simulation is necessary before reaching a quasi-converged state for external aerodynamic application of vehicles. The source ran adjoint simulation with ELEMENTS. However, this is also true for when solving adjoint equations with OpenFOAM v2006.

Moving Ground simulation with Rotating wheels vs. Stationary

The largest deviation of sensitivity maps were observed during the parametric study in section A.3, when applying wall tangential velocity on the ground on the domain and adding rotational wall boundary conditionals on the wheel. Figure 4.11 shows a side view of sensitivity map D, which has stationary road and stationary wheels boundary conditions, and sensitivity map E, which has moving ground and wall-tangential velocity boundary conditions at wheels. Sensitivity map D and sensitivity map E have the same mesh and identical solver setup apart from the boundary conditions at the wheels and road. The rotational wall boundary conditions were the sources of a wavy sensitivity maps pattern, observed on the wheel arches in Figure 4.11E. Additionally, changes are observed downstream of the rear wheels in the region of the trailing edges. This is expected due to changes of the rear wheel wake. The change of the rear wheel wake induce changes to the base wake, and hence changes to the sensitivity maps at the side trailing edge, see Figure 4.11. The leading/upstream region of the car and mirrors were relatively unaffected by the change of moving ground simulation. Further analysis is presented in subsection A.3.5. Naturally, simulation setup (E) is closer to real driving conditions.

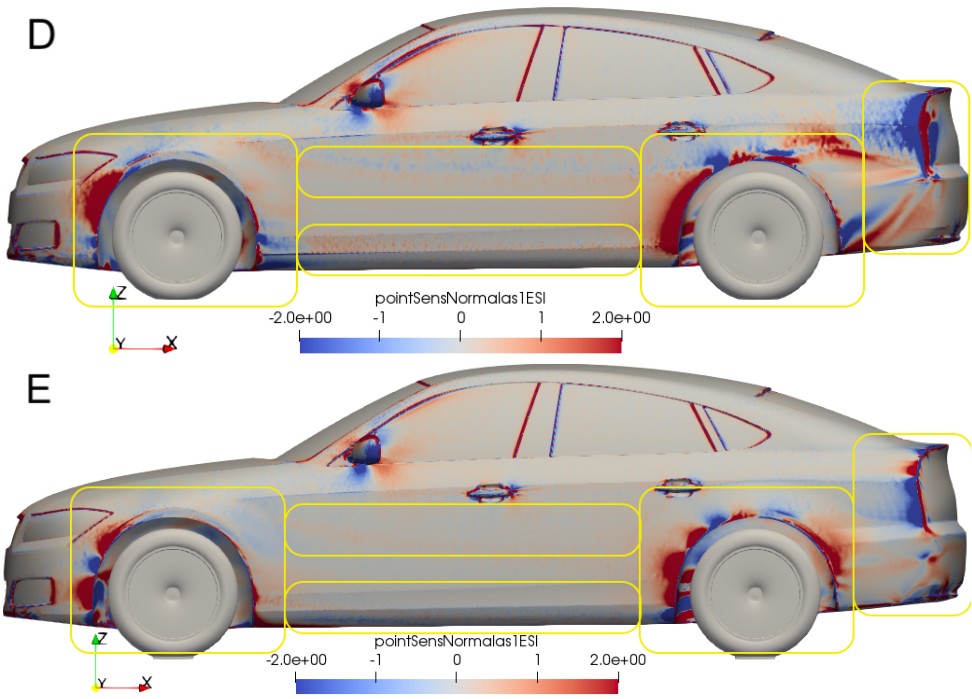

Figure 4.11: Sensitivity maps of sides of (D) with stationary road and wheels vs. (E) with moving ground and rotating wheels.

Open Rims vs. Smooth Rims

Analysis of changing rim geometry on sensitivity map was conducted. Sensitivity maps of the FSMW with (E) smooth rims vs. (F) five spoke rims, are observed in Figure 4.12. This study utilised the same mesh and identical solver setup apart from the boundary conditions and mesh of wheels. More specifically, the smooth rims had a rotating wall boundary condition applied to both rims and tires. The open rims had rotating wall boundary conditions applied to the tires and rims, while the pockets of fluid in between spokes were modelled by a moving reference frame. Both simulations had the ground simulation active. Further changes of the fender areas, dominated by wheel flow interaction, were observed when changing from smooth rims to open rims. Higher magnitudes of sensitivity derivatives were observed, adding to the chaos of the moving ground simulation and rotational wall wheel boundary conditions. Areas of opposite sign are observed on the wheel arches of Figure 4.12E and Figure 4.12F. The interaction of the rear wheel wake and base wake caused a trailing side edge waviness, see Figure 4.12F. Further analysis have been conducted in subsection A.3.6.

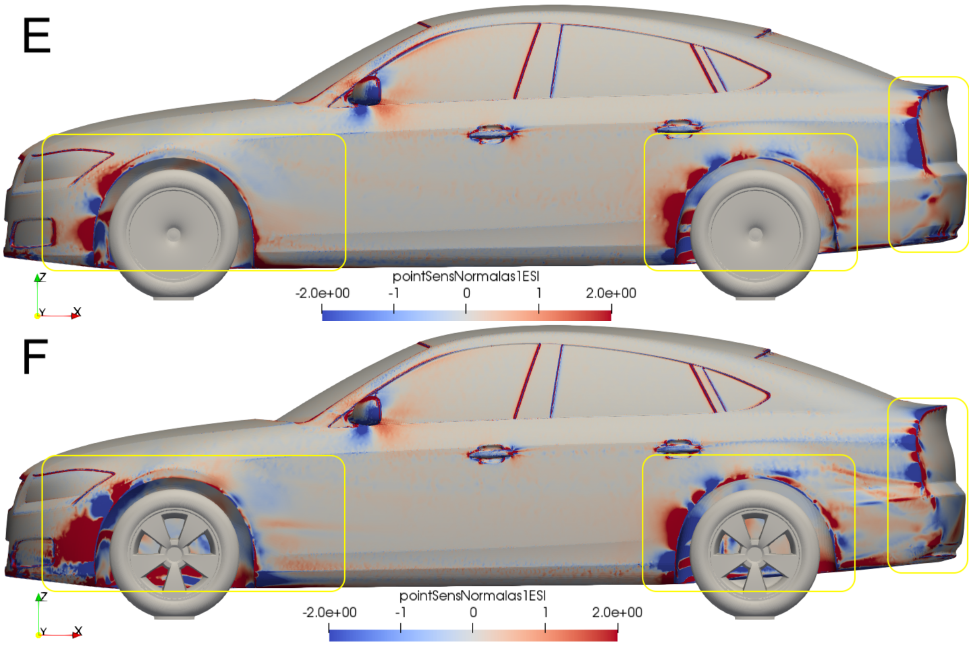

Figure 4.12: Sensitivity maps of sides (E) Smooth vs. (F) Open Rims.

Hexa vs. Tetra Interior Mesh

The importance of choosing interior cell types when generating a mesh for adjoint simulation was proven in subsection A.3.1. Areas were affected where change in pressure was observed, primarily in the vicinity of the wheel arches and trailing edge of FSMW.

Prism vs Hexa Boundary Layer

Near surface cell types of prisms and hexahedra were investigated in respect to sensitivity maps. Changes in sensitivity derivative magnitudes were observed with few areas of opposite signs. The overall trend of sensitivity maps of prisms and hexahedra cells were self-similar. Prisms are recommended to resolve the curvature of continuous surfaces by keeping quasi-equidistant cell center height in the boundary layer.

Boundary Layer First Cell Height

Negligible differences to the sensitivity maps were observed for different meshes with average y^+ values of respectively 19, 32, and 38 over the body of the FSMW. Different first cell height, constant growth ratio, and same number of prism layers amongst the different meshes caused minor differences in total height of prism cell layers. The sensitivity map did tolerate variation in y^+ values, as long as correct wall functions were applied.

Boundary Conditions and Sensitivity Maps

Boundary conditions should be applied similarly to that of a conventional CFD simulation. Adjoint boundary conditions are recommended in the adjointOptimisationFoam User Manual [11]. Symmetry boundary condition is recommended for cases of symmetry flow cases, as zero gradient/slip conditions will yield artificial sensitivity derivatives in the vicinity of the symmetry-plane.
During adjoint simulations, the adjoint state variables would develop gradually towards a constant value. This gradual development was partially hindered by bounding the turbulence equations. In turn, bounding was completely necessary along with application of under-relaxation factors to control divergence of the adjoint state variables. Lessons were learned during adjoint simulations with various turbulence initial conditions. For instance if $\tilde{\nu}$ was set too low and bounding was performed, the initial conditions changed the final result of the sensitivity map, if not an adequate amount of iterations is run of more than 5,000 primal iterations. An example of deviation by initial conditions are given in subsection A.3.4, where the eddy viscosity ratio is increased (simulation Q).

Turbulence Models and Frozen Turbulence Assumption

Impact on the sensitivity maps when changing to the frozen turbulent assumption and turbulence models are yet to be investigated further. Simulations of the aforementioned turbulence models including the k-ϵ, *Realizable* k-ϵ, and k-ω *SST* turbulence model have been performed, but diverged when solving the adjoint PDEs. Additional numerical filtering and relaxation of the adjoint state variables have not proved to solve the divergence issue completely. Further numerical diffusivity and filtering along with first order schemes of the two-equation turbulence equations have proved convergence. Thus the overall validity of the adjoint solution has been compromised and ought to be investigated further.

Unsteady vs. Steady

A parametric study of steady (F) vs. unsteady (S) averaged Spalart-Allmaras turbulence modelling was performed and studied in further detail in subsection A.3.7. This study utilises the same mesh and identical solver setup apart from running simpleFoam (F) and pisoFoam (S). Both the steady and unsteady simulation has been averaged over the last 5,000 iterations of the primal flow. This equals to two seconds of flow averaging for the unsteady flow. The unsteady simulation may have the advantage of being more accurate, as the unsteady primal simulation is run based on a fully converged steady state primal solution. Distinguisement of side sensitivity maps in the region aft of the rear wheel is observed in Figure 4.13. The sensitivity map in Figure 4.13S shows contours of unsteady flow properties, formed by an unsteady rear wheel wake. The region of the rear fender in front of the rear wheel is also affected where the upper rear wheel arch is displaced opposite, when comparing (F) with (S). This is also the case in the lower front wheel arch aft of the front wheel. Major changes to the sensitivity maps when comparing (F) and (S) are observed on the front fender upstream of the front wheel. Here, the unsteady simulation creates a wavy pattern compared to the steady. To cover the front wheel by a wavy front fender is viewed with some skepticism. Note also that the rotating wheel boundary condition was responsible for the wavy pattern. The unsteady simulation may amplify this wavy pattern further. Finally, the sensitivity maps of the mirror housing is kept constant. This proves that that the mirror is a component suited for automatic morphing, as it is relatively unaffected by changes of various simulation parameters.

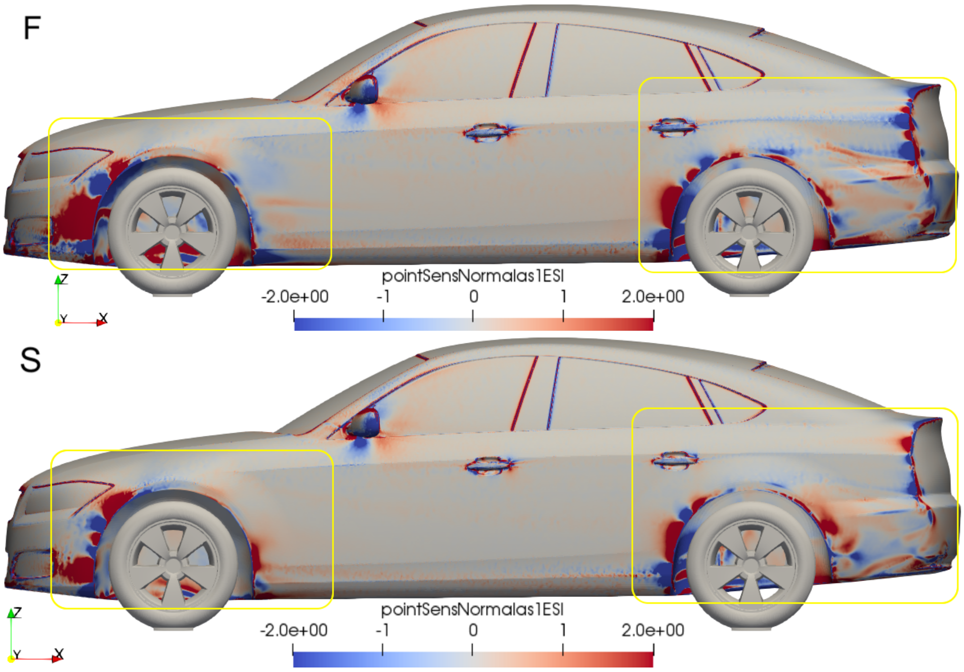

Figure 4.13: Sensitivity maps sides of (F) Steady vs (S) Unsteady.

Mesh Sensitivity of adjointOptimisationFoam

During the thesis work, it was found that OpenFOAMs adjointOptimisationFoam is sensitive to meshes which inherit non-orthogonal values above 60. Non-orthogonal values above 70 correlated with adjoint divergence. Experience with simpleFoam and pisoFoam algorithms have proved more stability with higher non-orthogonal cells. Efforts were made to decrease mesh-related divergence of adjointOptimisationFoam with non-orthogonal correctors and relaxation factors. Correctors did not improve divergence, whilst relaxation factors proved to slightly decrease and delay the divergence.

Skewness should be minimized. AdjointOptimisationFoam converged as long as the strict quality criterion of internal skewness 4 was met. Special care was taken in areas of model self-proximity and double-curvature where skewed cells were prone to develop when generating prism layers. In fact, cells in the prism layer caused most sources of instabilities during all simulations in this thesis work. Generation of prism layers with adequate growth ratio, correct first layer thickness, and aspect ratios close to hexahedra cells, whilst adhering strict mesh quality criteria and keeping the cell count to a minimum, turned out to be the largest obstacle during this thesis work.

Minimum cell determinants should be as large as possible. Experience with different meshes of the FSMW proved that minimum cell determinant, and should not reach below order of 1E-8. Although, the strict quality criteria of ANSA recommends a minimum cell determinant value of 1E-3. The strict quality criterion was set but not prioritised by ANSA. Prism layers with a high aspect ratio caused the volume to area fraction to be disproportional. Hence, excessively small determinant numbers were encountered during fine prism layer generation.

This author strongly recommends a high-quality mesh to begin with, especially when running adjointOptimisationFoam, to avoid divergence issues.

4.2.2 Adjoint Solver Divergence Due to SnappyHexMesh

Divergence of the adjoint solver occurs when generating meshes with snappyHexMesh. SnappyHexMesh create discontinuous surfaces by constructing a so-called Lego surface, Figure 4.14a, creating artificial oval contours on analytically continuous surfaces. Alternatively, ANSA creates a quasi-continuous surface mesh, Figure 4.14b. It has been found that snappyHexMesh does not respect the analytical geometry, but snaps on the STL triangles. As a side-note, this explains that tessellation has proven itself to be of the utmost importance when generating a mesh with snappyHexMesh. The tessellation quality has been set to max for the FSMW in Figure 4.14a. Thus, tessellation quality is not the final reason of the divergence with snappyHexMesh meshes. The snappyHexMesh surface mesh cause artificial pressure gradients and makes the time step continuity error residual diverge when solving the adjoint PDEs. Furthermore, the magnitude of the sensitivity derivatives may increase to the order of thousands. Confirmation of snappyHexMesh being root of the divergence issue, has been received from the team behind the adjoint solver introduced in OpenFOAM v1906-v2006.

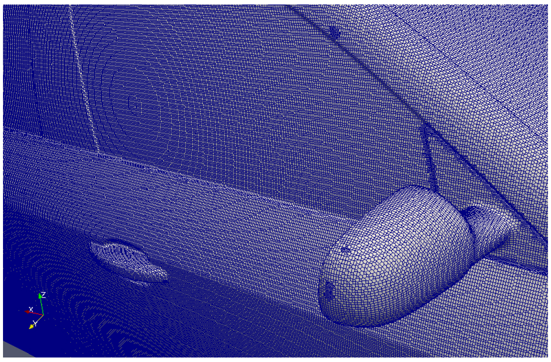

(a) SnappyHexMesh surface mesh causing divergence.

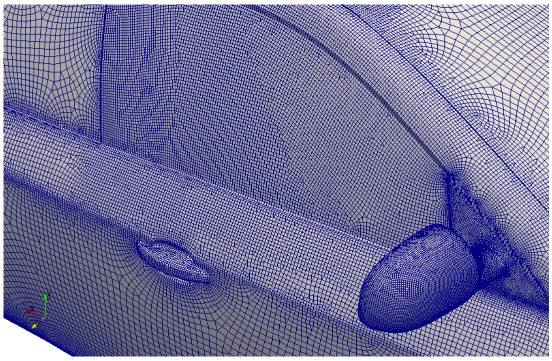

(b) ANSA surface mesh facilitating convergence.

Figure 4.14: Mesh impact on the adjoint solver.

To further prove that snappyHexMesh was the cause of divergence, the two meshes, shown in Figure 4.14a and Figure 4.14b, were run with identical solver settings and boundary conditions. This yielded the following residual plots in Figure 4.15. The ANSA mesh was even coarser in terms of fewer cells than the snappyHexMesh mesh. Therefore, the divergence was not due to the cell count. The continuity error residual ran continuously throughout the primal and adjoint simulations. Thus adjoint continuity error residuals can be observed in the interval of 10,000-20,000 iteration in Figure 4.15a and 4,000-9,000 iterations in Figure 4.15b. The simulation in Figure 4.15a and Figure 4.15c ran 10,000 primal iterations and 10,000 adjoint iterations. Whilst the residuals in Figure 4.15b and Figure 4.15d converged respectively within 4,000 primal iterations and 5,000 adjoint iterations. This proved that the snappyHexMesh mesh simulation would not converge further even when increasing or decreasing the number of iterations for both primal and adjoint solver. On the other hand, the residuals in the simulation with the ANSA mesh could have run even longer to decrease residuals, run field averaging, and increase accuracy of the sensitivity maps. The primal solution for snappyHexMesh converged. Yet, the adjoint residuals and continuity error diverged rapidly from approx. 150 iterations within the adjoint simulations as observed in Figure 4.15a and Figure 4.15c. The adjoint residuals in Figure 4.15c were smoothed, relaxed, and bounded to hinder further divergence.

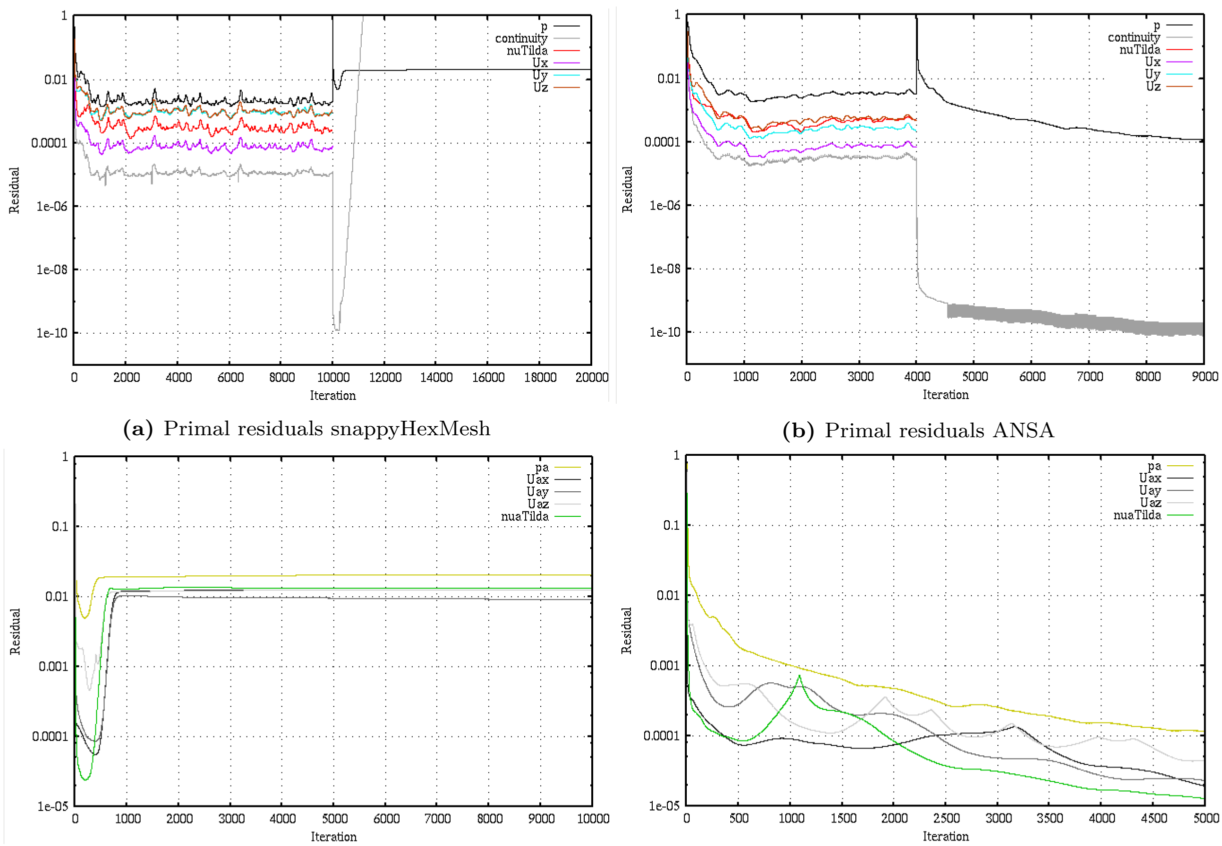

(a) Primal residuals snappyHexMesh

(b) Primal residuals ANSA

(c) Adjoint residuals snappyHexMesh

(d) Adjoint residuals ANSA

Figure 4.15: Residuals for snappyHexMesh and ANSA meshes.

Decreasing divergence of residuals, mainly the time-step continuity error as well as adjoint pressure and velocity residuals, may be accommodated by the following:

- Decreasing adjoint div-schemes to 1st. order
- Bounding adjoint div-schemes of 1st. order
- Reducing relaxation factors
- Changing ATC term modelling
- Adding numerical diffusivity
- Smoothing and zeroing the ATC term in the adjoint PDE on walls and patches

The compromise of these settings proved to create misleading sensitivity maps, suggesting to morph surfaces opposite of those found for the ANSA mesh and results published in [11]. Consequently, controlling adjoint residuals due to poor mesh quality compromises accuracy and may induce error.

This author attempted to increase mesh surface quality in snappyHexMesh by increasing snapping quality parameters of tolerance, nSmoothPatch, nSolveIter, and nFeatureSnap to maximum or very high values. In other words, the snapping quality parameter, resolveFeatureAngle, was decreased to resolve all angles above one degree. Still, oval contours, though changed by the setting, appeared. Additionally, by refining the surface to a high degree, the cell count increased significantly, exceeding limits of 40M cells, and subsequently increasing solver time. Thus, this author, along with approval from supervisors from DTU, Chalmers, and Koenigsegg, decided to switch to generation of meshes in ANSA.

4.3 FSMW Steady Optimisation

Two semi-automatic shape optimisations were conducted, based on the parametric study of the FSMW sensitivity maps in section A.3, and the analyses of the primal and adjoint fields. Drag minimisation were selected as the function object for the two shape optimisations, namely at the mirror and at the trunk. Optimisation of the trunk surface was performed, since the change of base pressure was expected to affect total vehicle drag greatly. Figure 4.16a shows the normal displacement of boundary nodes achieved after six iterations, in the color bar to the left. Normal displacement is stated in [mm] for Figure 4.16, Figure 4.17, and Figure 4.18b. Max displacement per cycle was set to 2 [mm]. The max cumulative displacement of 12 [mm] was achieved in the positive direction (wall-normal outwards). The normal displacement distribution to the left in Figure 4.16a showed a longitudinal extension of the trailing edge to reduce drag. Furthermore, the outwards displacement of the trailing trunk edge was observed in the vertical direction, thus decreasing lift over the rear axle and approaching a balanced wake in Figure 4.16a. The negative (wall-normal inwards) displacement just upstream of the trailing trunk edge was bound by confinement of control point boundaries. Evidently, tapering of the FSMW rear end may decrease drag. Figure 4.16a also included the sensitivity map, to the right, based on the same simulation model analysed in the primal flow analysis. The sensitivity map displayed similar distribution as the normal displacement distribution. The differences in magnitude and sign were attributed to the aforementioned calculation with different formulations of sensitivity derivatives.

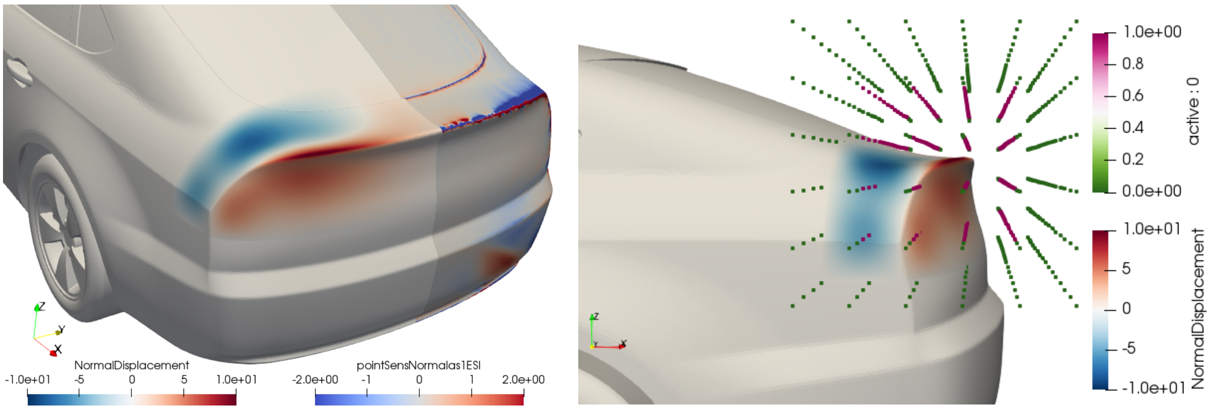

(a) **(b)**

Figure 4.16: Trailing trunk edge after six steady optimisation cycles.

Control points (CPs) were defined for the trunk surface in Figure 4.16b. The *active:0* color bar in Figure 4.16b revealed confinement in a binary manner. For instance the magenta CPs, numbered one, were free to move whilst the green CPs, numbered zero, were confined in a given spatial direction. The confinement of these CPs were applied to prevent the mesh externally of the CPs to be deformed, and mesh quality compromised. Confinement in spacial directions were possible in case of critical tolerances or self-intersecting geometries. In hindsight this author could have distributed more points in the z and x direction and less in the y direction, according to Figure 4.16b. Conversely, the volume B-spline morpher proved its' capability even in case of few CPs in certain spacial directions. Onwards, the self-similarity of the mirror sensitivity maps made the mirror an ideal region of shape optimisation, despite of the solver changes. The steady optimisation of the mirror yielded max displacement of 14 [mm], in Figure 4.17, with max displacement of 3 [mm] set for each cycle. Note that the frontal area of the mirror decreased. Radius of curvature on the top surface of the mirror was morphed. Tapering towards the trailing mirror edge, in Figure 4.17, yielded a minor pressure recovery, resulting in less drag. Teardrop formation of the mirror arm was initialised and was expected to fully develop if further unconstrained optimisation cycles were performed.

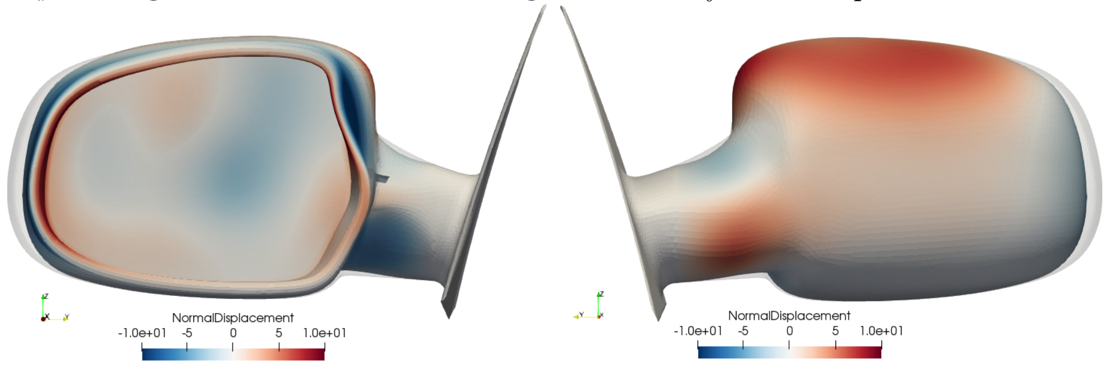

(a) (b)

Figure 4.17: Mirror after five steady optimisation cycles.

CPs of the mirror were defined as observed in Figure 4.18b, with the same convention for CP confinement as in Figure 4.16b. Figure 4.18a showed total drag reduction over steady optimisation cycles of the mirror and the trailing trunk surface. Optimisation cycle zero was equivalent to the reference model. Turbulence models of *Spalart-Allmaras*, k-ϵ, *Realizable* k-ϵ, and k-ω *SST* were run as a reference for the undeformed mesh and subsequently for the deformed mesh. This should validate the morphing with the Spalart-Allmaras turbulence model. A total drag reduction of three counts over five steady optimisation cycles (with the Spalart-Allmaras turbulence model) was achieved equivalent to a total 1.25% drag reduction of the FSMW. More optimisation cycles would yield further drag reduction, eventually converging towards a constant value. Unfortunately, cell non-orthogonality > 70 in the vicinity of the mirror made the other turbulence models diverge for cycle five, as observed in Figure 4.18a. A total drag decrease of less than two counts (of the Spalart-Allmaras turbulence model) was achieved by deforming the trunk surface over six optimisation cycles. All turbulence models but the k-ϵ shows decrease in total drag. The increase in drag for the k-ϵ turbulence model was less than one count. This author recommends further shape optimisation of the trunk surface to justify the numerical drag reduction.

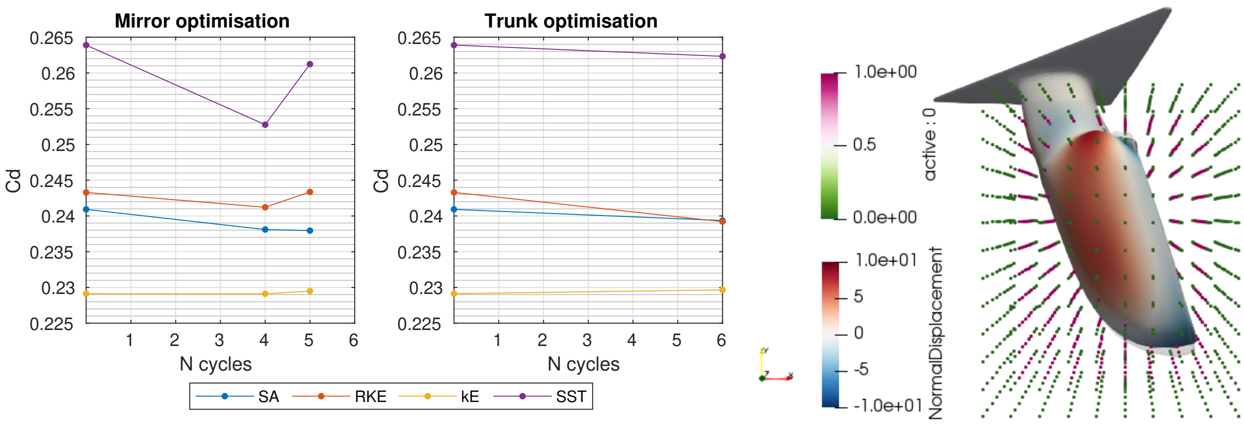

(a) (b)

Figure 4.18: C_d related turbulence modelling vs. optimisation cycles and mirror CPs.

4.4 Koenigsegg Regera Primal Analysis

Further investigation of the adjointOptimisationFoam's applicability to real cars is conducted. All figures of the Regera in the following sections are of a left hand side pre-production model of the Koenigsegg Regera with the Ghost Package. All simulations of the Regera are extended to 250 kph with open cooling, porous zones, and wall-tangential velocity at road and wheels active. The adjoint Regera simulations will be performed with the steady state, incompressible Spalart-Allmaras turbulence model, omitting possible instability issues experienced during the FSMW simulations with two-equation turbulence models.

Mesh Convergence, Correlation, and Reynolds Sweep

The creation of a high-quality mesh included full control over boundary layer refinement of the logarithmic layer, no cells exceeded non-orthogonality of 60, no cells exceeded internal skewness of four, whilst keeping the cell count to a minimum. Experience proved that when approaching advanced user-level in ANSA, this process was significantly reduced in turnaround time. Several meshes were created during the development of a converging adjoint simulation for the Regera. An explicit mesh convergence study was not performed due to time limitations. Furthermore, the minimum cell count of a converging mesh of the Regera nearly exceeded the available hardware limits. A mesh convergence study is recommended in future work. Note that post-processing of the simulation presented in this thesis correlated with simulations performed at Koenigsegg. Thus this simulation setup of the Regera was expected to provide credible sensitivity maps. During the Reynolds sweep, observed in Figure 4.2, simulations indicated that an increased Reynolds number decreases the drag coefficient. Correctly, this may be due to an increase in mixing of free stream turbulent kinetic energy with the low momentum flow in the boundary layers, delaying separation, and eventually decreasing drag. Therefore, an additional Reynolds sweep for the Regera may be unnecessary for the scope of this thesis, as the solver has proved its capabilities at different Reynolds numbers.

Residuals

The residuals and force coefficients have been monitored as a measure of convergence of the Regera simulations. The continuity error is plotted continuously for both the primal and adjoint simulation. Hence, the adjoint continuity error is observed in Figure 4.19b in the iteration interval of 10,000-20,000. The primal pressure residual, in Figure 4.19a, is observed to be relatively high. The high primal pressure residual is due to the flow being inherently unsteady. The primal residuals approach residual convergence after 3,000 iterations. However, the primal flow field variables and monitored force coefficients converge after 6,000-7,000 iterations, excluding additional iterations for averaging. Attempts were made to stop the primal solver at 5,000 iterations. This caused instability and divergence for the adjoint solver. A fully converged primal flow with a sufficiently large averaged window is highly recommended. The adjoint residuals in Figure 4.19b decrease continuously and do not converge fully within 10,000 adjoint iterations. Changes to the sensitivity maps beyond 10,000 adjoint iterations shall be minor, similarly observed in subsection A.3.3.

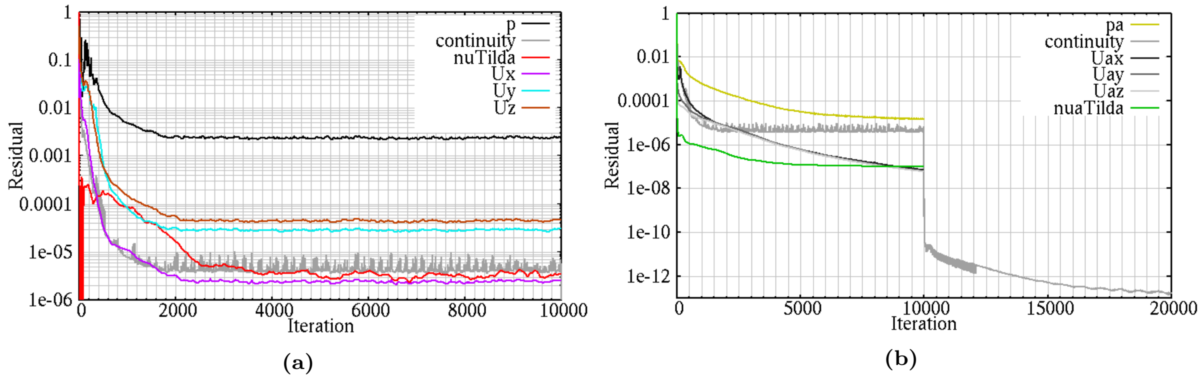

Figure 4.19: Primal (left) and adjoint (right) residuals of the Spalart-Allmaras turbulence model.

Porous Zones and MRF

MRF were added, as in the FSMW simulations, to model the pockets of fluid in-between the spokes of the rims. The MRF were successfully implemented for the adjoint simulations of the Regera, and coped without signs of instability at rotational frequency equal to the free stream velocity of 250 kph. Additionally, porous zones were added. The porous zones were meshed as individual internal volumes in ANSA. The porous zones proved some instability, as internal volume cells initially were modelled too large compared to surrounding cells. Furthermore, the intersection of surface geometry (walls) and the porous zones (internal volume) are suspected to add to the instability. This is to be investigated further. It is recommended that internal volumes are only in contact with the domain volume. Furthermore, local change of cell size between the porous zone and the domain should not exceed a growth ratio of 1.3.

Yplus

Figure 4.20 shows the y^+ distribution of the Koenigsegg Regera mesh when running 250 kph. Overall, the y^+ distribution is suited for wall treatment of the logarithmic layer. Note that variable prims layer height is applied. Narrow prism cells are applied on components where velocity exceed average values, e.g. front-splitter, front diveplate, mirror, wheels, winglet, and rear wing. As observed in Figure 4.20a, the front wheel shoulder region ought to be treated with more narrow prism cells. Yet, this is omitted since additional cells are required for such a refinement. Next, prism cells are not applied on rear end panels, observed in Figure 4.20b. The exclusion of boundary layer modelling in this region was due to avoiding bad cell quality which occurs when growing prism cells in regions of self-intersecting geometry with high convex/concave curvature. Additionally, the sensitivity map is not expected to generate reliable sensitivity maps in this region due to effects of completely separated and three dimensional flow. For this reason, this region is excluded in analysis of the sensitivity maps.

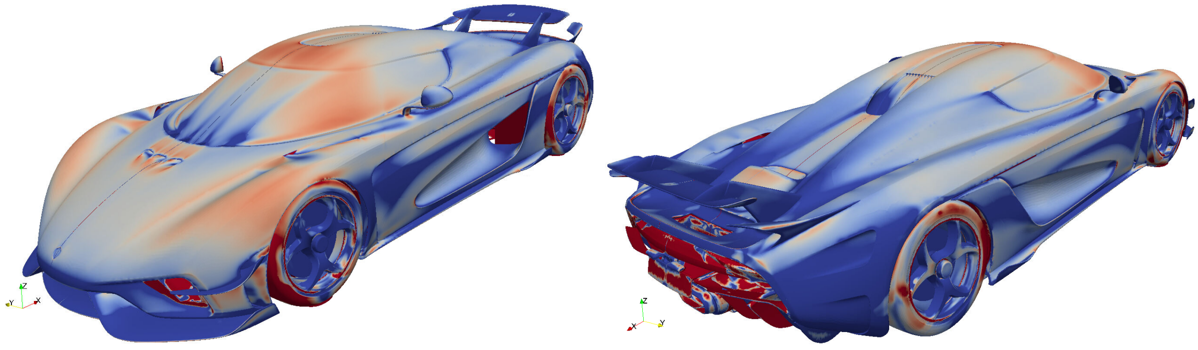

(a) (b)

Figure 4.20: y^+ distribution.

Skin Friction Coefficient

The skin friction coefficient distribution, seen in Figure 4.21, reveals similar local flow conditions as for the y^+ distribution, in Figure 4.20. Few areas of unexpected separation are observed on Figure 4.21. These are revealed by dark blue areas, namely on the front wheel side vent, and downstream of the front wheel and at rear wheel vents. Consequently, a wall-normal outwards displacement of the front wheel side vent is expected to increase skin friction and increase mass flow of the front wheel side vent. Also, the leading edge of the rear wheel fender shows signs of flow separation as it is designed to feed high mass flow into the side cooling intake. This may be at the expense of drag. Separation is also observed in the region aft of the rear wheels. This is partially mitigated by blowing air into this region by rear wheel vents. The base of the Regera is dominated by separation forming a rear wake. Additional separation is observed at the engine intake, just downstream of the vortex generators on the roof, in Figure 4.21b. The vortex generators are observed to increase skin friction and aid in keeping a portion of the flow attached during an otherwise strong adverse pressure gradient. Excluded suction from the engine intake would have aided in attaching the flow behind the roof vortex generators.

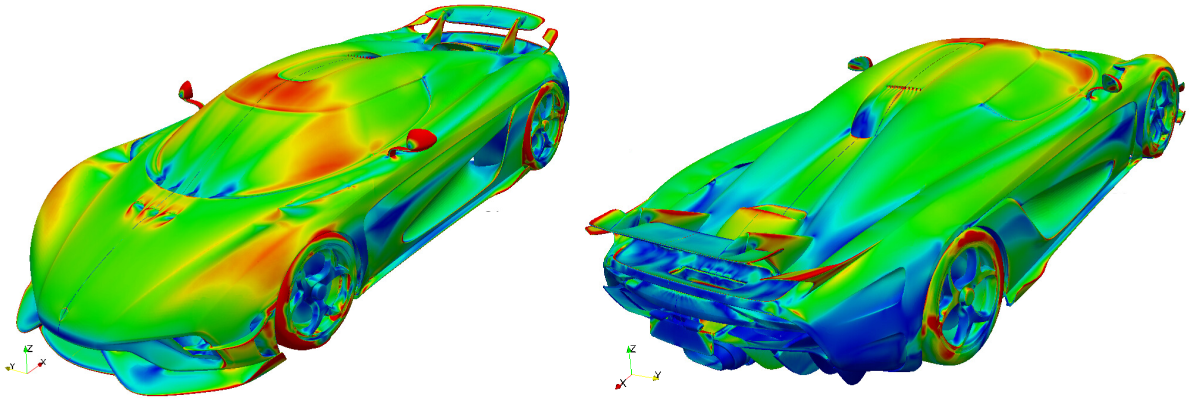

(a) (b)

Figure 4.21: Skin friction coefficient distribution.

A close-up on the mirrors in Figure 4.22 reveal expected local stagnation areas at leading and trailing edges. The rear view of the mirror, in Figure 4.22, show unclean separation during the primal simulation. Thus, changes to the mirror housing is expected to be proposed by the sensitivity maps.

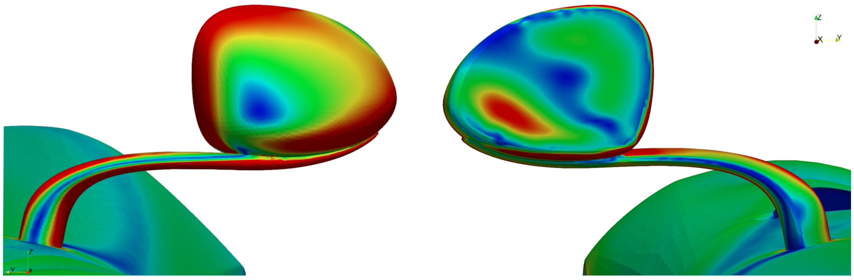

Figure 4.22: Skin friction coefficient distribution at front (left) and rear (right) view of the mirror.

Top and bottom view of the front diveplate, representing local skin friction are observed in Figure 4.23. Blue areas observed in Figure 4.23a are due to flow stagnation on adjacent surfaces. Green uniform area represents desired attached flow on the pressure side on the diveplate. The suction side, see Figure 4.23b, has a uniform red region, proving high skin friction. This is expected as a diveplate is designed to induce negative lift by vortex generation and vortices interacting with the leeward side of the diveplate [16]. The sensitivity map may propose changes to regions where local flow stagnation occur, or balance lift and drag, if optimised for 50/50 downforce/drag function objects. The latter may yield organic aerodynamic designs in the spirit of recent days of Formula One.

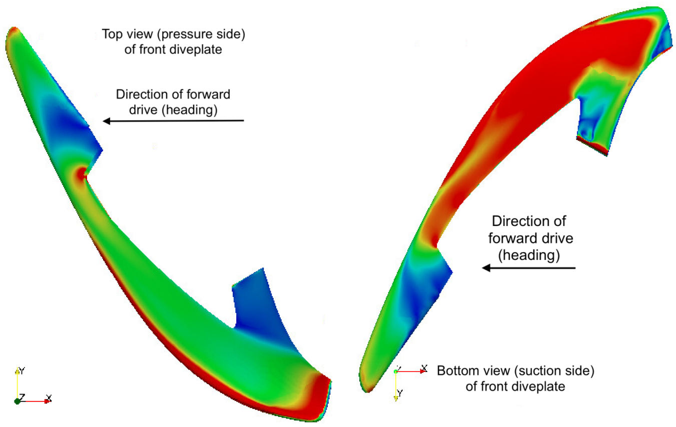

(a)

(b)

Figure 4.23: Skin friction coefficient distribution at the diveplate.

Skin friction over the winglet is pre in Figure 4.24, and does not reveal local flow separation. Increased skin friction i ved at the top left and top right of the winglet due to high pressure flow from the re side which travels towards the low pressure at the suction side, shedding small dinal tip vortices.

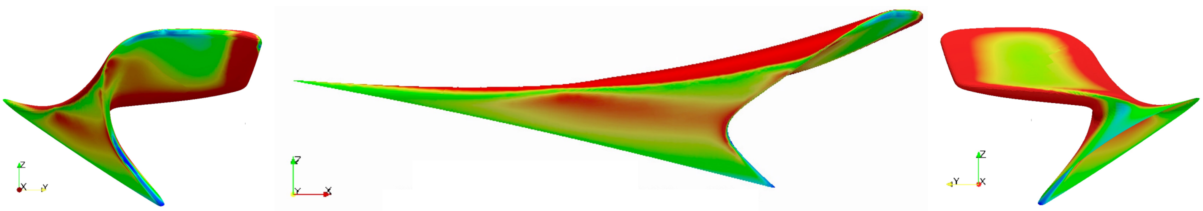

Figure 4.24: Skin friction coefficient distribution on the winglet.

Static Pressure Coefficient

The mean static pressure coefficient distribution in Figure 4.25 is analysed to discover potential areas of drag or lift optimisation. A smooth favourable pressure gradient is observed during the transition from the front bumper stagnation area to the bonnet in Figure 4.25a. Next, a larger low pressure region around the A-pillar is observed in Figure 4.25a despite the curvature of the wind shield and A-pillars. The leading region of the roof inherits a significantly low pressure zone. The low pressure region on the roof adds undesired lift due to pressure acting in the wall normal direction. The sensitivity map is expected to smooth the leading roof region, if a function object of lift minimisation is set. Local stagnation on the windshield is observed, and expected visualised in sensitivity maps of drag minimisation. High pressure flow on top of the engine bay is due to stagnation on the rear wing struts and cooling intake in Figure 4.25b. Note the scaling of the colour bar in Figure 4.25b.

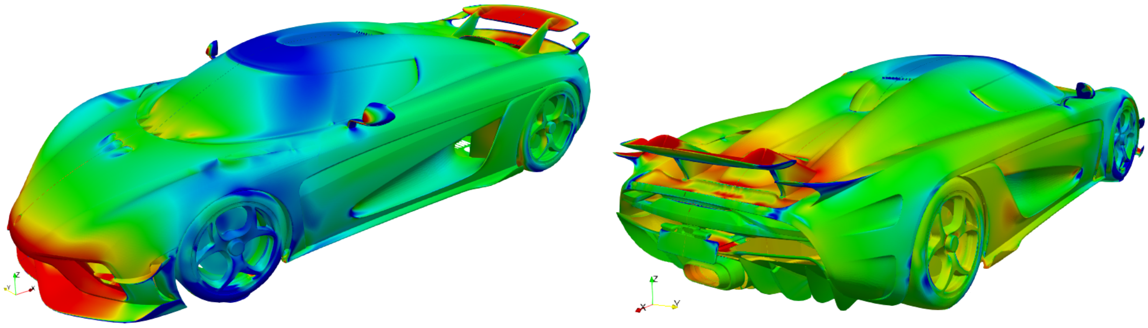

(a) (b)

Figure 4.25: Mean static pressure coefficient distribution.

The pressure distribution of the mirrors is presented in Figure 4.26. Note the different colour bars. The mirror back was expected to have uniform low pressure distribution. The non-uniform pressure distribution confirms unclean flow separation. The sensitivity map is expected to show drag optimisation potential in this area.

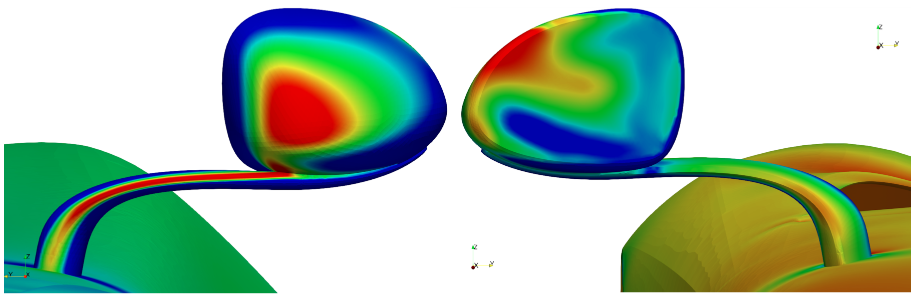

Figure 4.26: Mean static pressure coefficient distribution at the mirror.

The pressure distribution of the diveplate in Figure 4.27a shows a potential of increasing pressure, hence downforce, on the pressure side. A more uniform low pressure zone is observed on the suction side diveplate in Figure 4.27b. High pressure zones on the inner tip of the diveplate is due to stagnation on adjacent surfaces.

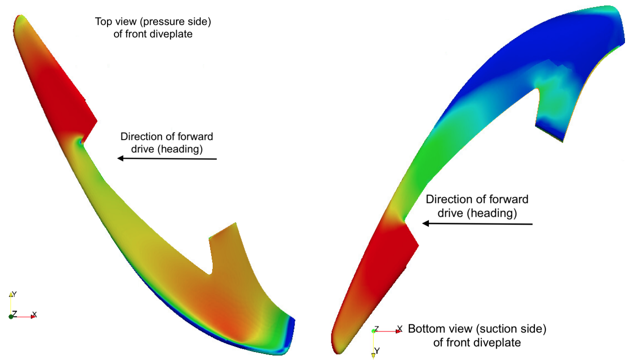

(a) (b)

Figure 4.27: Mean static pressure coefficient distribution on diveplate

The low pressure distribution of the winglet tips confirms tip vortex shedding, which induce downforce, in Figure 4.28. The magnitude of pressure on the pressure side of the winglet offer a potential downforce optimisation. Changing the pitch of the winglet may aid in decreasing pressure on the suction side.

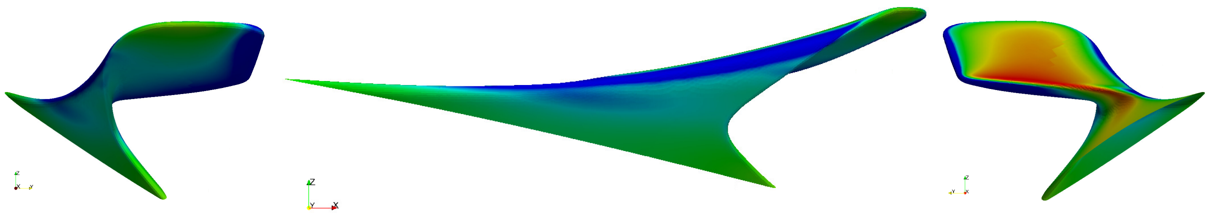

Figure 4.28: Mean static pressure coefficient distribution on the winglet

Velocity Field

Several cross-sectional planes of the velocity field yield an overview of the flow field of the Regera, seen in Figure 4.29. A region separated flow over the front wheel side vent is confirmed by low momentum flow downstream of the front wheel. The mirror wake is observed but recovers partially downstream. A rear base wake is observed in Figure 4.29.

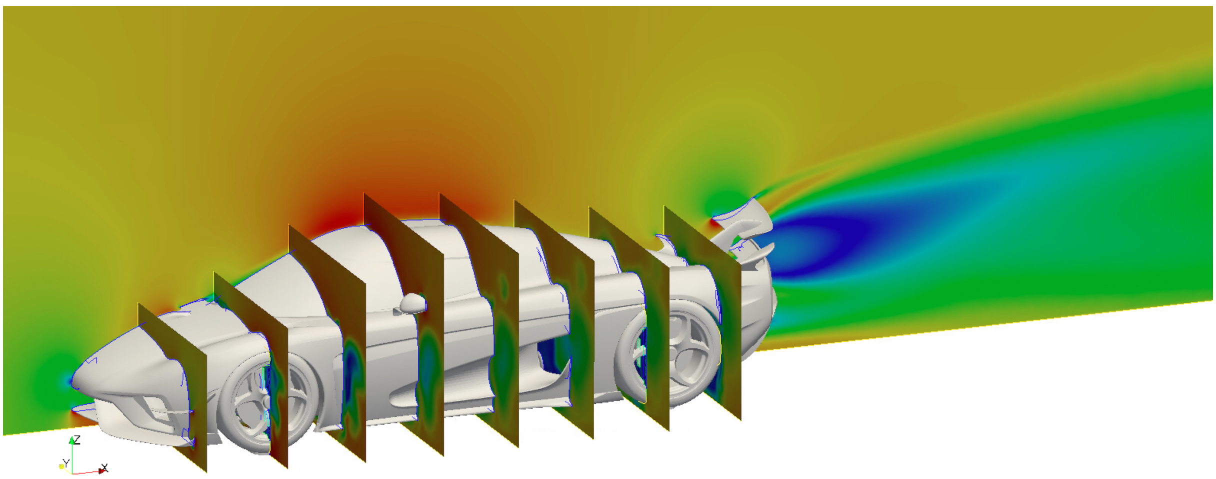

Figure 4.29: Normalised velocity distribution.

The local rear wake is shaped by the rear number plate and inverter cooling. Globally, the wake is large and indicate upwash generated by the rear wing and diffuser. The rear wing may aid in keeping the flow on the diffuser attached. The wake of the rear wing indicates that the angle of attack or camber might be too high for the flow to be attached on the suction side. The velocity of the roof is approaching high velocities, although not compressibility effects.

4.5 Koenigsegg Regera Adjoint Analysis

The following section presents sensitivity maps with function objects of drag minimisation, and downforce/drag optimisation. The sensitivity maps serve to investigate the robustness and value of the adjointOptimisationFoam, related to the primal flow field just post-processed.

Drag Sensitivity Maps

The drag sensitivity derivatives are observed in Figure 4.30 on all surfaces apart from those of the lift devices and rear end. To minimise drag, the surfaces in Figure 4.30 with red contours identify wall-normal outwards (positive), and blue contours indicate wall-normal inwards displacement.

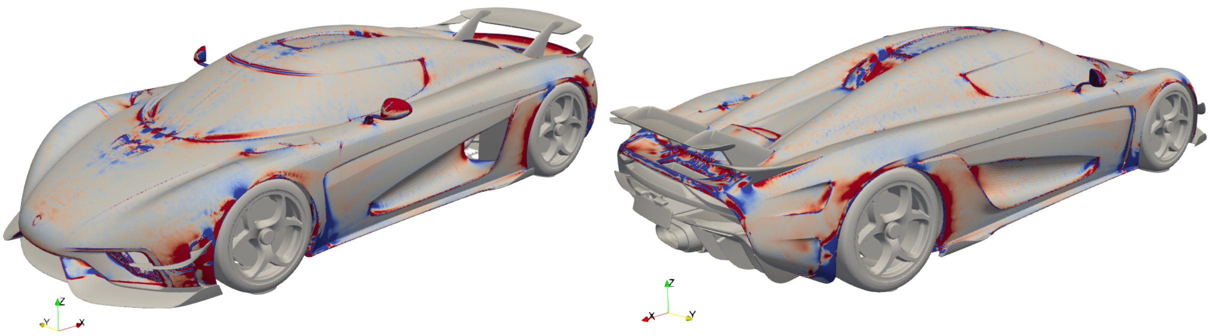

(a) (b)

Figure 4.30: Drag optimised sensitivity map.

Correctly, the sensitivity derivatives indicate removal of edges that perimeter the greenhouse, specifically at the windshield to roof transition and at the side window to B-pillar transition. The sensitivity map identifies the same compromised lower surfaces of the windshield as in Figure 4.21a and Figure 4.25a. Next, the vortex generators create a highly rotational flow, which makes the sensitivity derivatives inconclusive. As for the FSMW, contours on the front fender show wavy contours of wall normal displacement due to rotational boundary conditions of tyres and rims. The red contour on the low part of the front fender, recommend wall normal outwards displacement, covering the front wheel. The rear wheel arch display a more uniform distribution of sensitivity derivatives than expected. This indicates trustworthy wall-normal displacement, in other words outwards displacement of the rear fender upstream of the rear wheel to cover the rear wheel. Next, the front wheel side vent shows contours of wall normal inwards displacement in Figure 4.30a. Hence, the lowest surface of the external front wheel side vent should be displaced inwards, possibly evacuating the front wheel wake flow better. This region was observed in the velocity plot of Figure 4.29. The internal surfaces of front wheel side vent itself should increase mass flow by wall normal inwards displacement.

The leading edge of side cooling intake should be reshaped and rounded for drag optimisation, also observed in Figure 4.21 and Figure 4.25. The trailing edges of the internal side cooling surfaces show drag optimisation potential by uniform contours. However, this region is affected by the intercooler, which has only been represented by a porous zone, with the actual fan missing. Certain skepticism should be applied in this area. The rear wheel vents are compromised by partially separated flow and do not represent reasonable wall normal displacement.
A close-up of the mirror sensitivity derivatives, in Figure 4.31, shows counter-intuitive wall normal outwards displacement at the top surface of the mirror. The mirror projected area is expected to be minimised for drag minimisation. Monitoring of forces acting on the mirror reveal downforce on the mirror, due to the drooped leading edge of the mirror housing. This may induce upwash in the mirror wake. Scaling of the colour bar reveals the top surface to be shaped towards an adverse pressure gradient. The front view of the mirror identify wall-normal inwards displacement of the inner side of the mirror housing in Figure 4.31, minimising frontal area. The sensitivity derivatives reveal a more teardrop shape on the mirror arm. The rear view sensitivity derivatives of the mirror are of high order, and identify formation of a non-planar surface. This should be disregarded. However, nearly the complete circumference of the planar mirror surface proposes longitudinal extension to aid clean separation and perhaps minor pressure recovery. The highest order sensitivity derivatives were found on the surface of the mirror. Thus, the actual displacement magnitude of the mirror could be interpreted with a different colour bar than the one observed in Figure 4.31. If increasing interval of the colour bar, smaller sensitivity derivatives will fade.

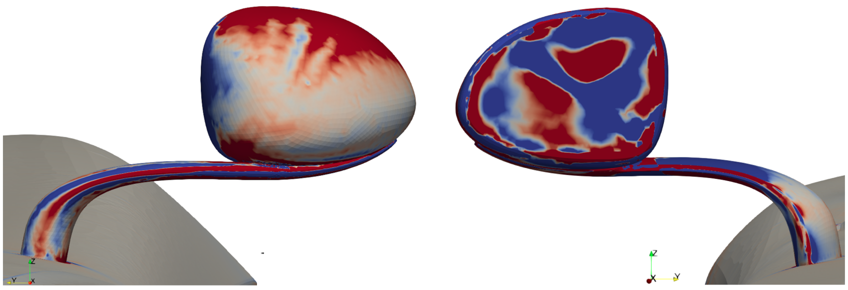

Figure 4.31: Drag optimised sensitivity map, mirror.

Downforce/Drag Sensitivity Maps

The downforce/drag function objects are performed on lift devices only. A weighting of 50/50 downforce/drag is applied, and can be changed arbitrary if desired. The sensitivity map colours in Figure 4.32 continuously signify that red contours shall be displaced wall-normal outwards (positive), and blue contours indicate wall-normal inwards displacement. The frontsplitter on Figure 4.32a and Figure 4.32c indicate displacement of the frontsplitter in the negative Z (vertical) direction. This is expected as ground proximity with the front splitter increase downforce non-linearly [16]. Drooping the frontsplitter towards the ground may alter the high pressure region of the frontsplitter pressure side (top surface). However, this will increase the frontal area, which may increase drag. Conversely, the high pressure region of the pressure side (top surface) of the frontsplitter is aiding in downforce, if pitched. The pressure side of the rear wing in Figure 4.32b indicates a reducing camber and angle of attack of the rear wing. This is expected due to the high angle of attack of 19.3^o and since the shape optimisation is set to 50/50 downforce/drag.

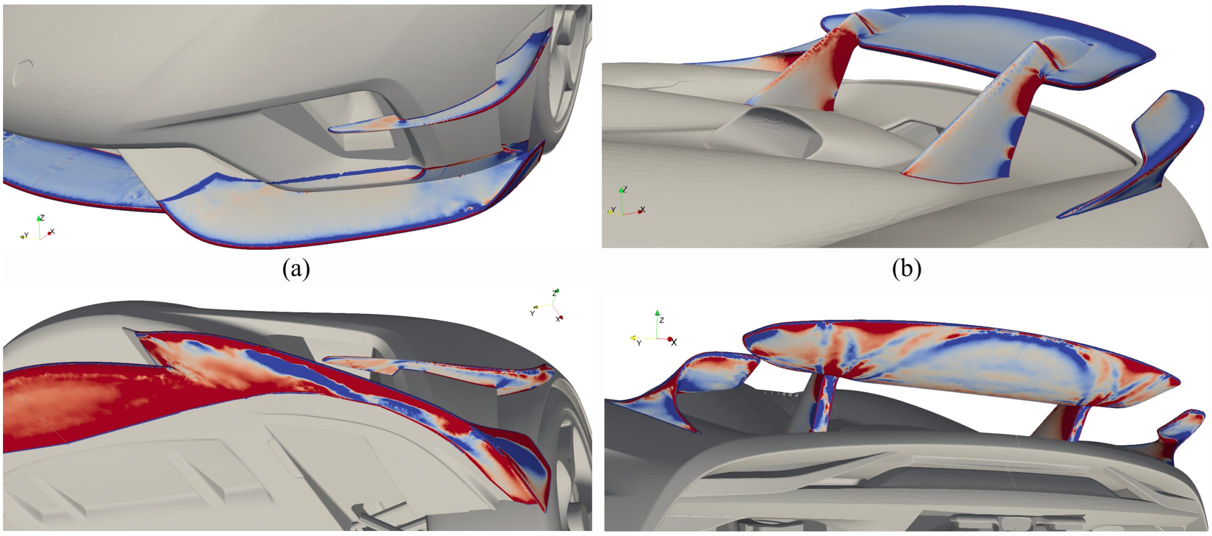
(a) (b) (c) (d)

Figure 4.32: Downforce/drag (50/50) optimised sensitivity maps.

The sensitivity map on the suction side of the rear wing in Figure 4.32d is more chaotic, due to separated flow aft of the rear wing struts. The sharp edges of the wing top mounted hinges, observed in Figure 4.32b, are recognised by the sensitivity map and should be smoothed. The rear-wing strut has a uniform area of wall-normal outwards displacement on the inner side. This is counter-intuitive as the struts in theory should be as thin as possible to decrease drag and should not aid in producing downforce. A wavy pattern on strut is observed on the left hand side trailing edge surface. This is due to separation of the flow. The sensitivity map in Figure 4.32b identifies optimisation by changing the angle of the winglet root/arm where transverse displacement is recommended. The pitch of the winglet is proposed to be reduced. A change of curvature in the region of transition from the top surface of the winglet to the side surface is also identified. The leading edge of the diveplate should be extended and thinned, according to Figure 4.32a, for optimisation in downforce and drag reduction. The trailing region of the diveplace pressure side should be decreased in curvature. Finally, vortex shedding from the pressure side to the suction side create irregular sensitivity derivatives on the suction side of Figure 4.32c.

4.6 Koenigsegg Regera Shape optimisation

Drag optimisation of the mirror is still under investigation along with downforce/drag optimisation of the winglet, diveplate, and rear wing. Unfortunately, this process of trial and error is temporally costly, for a fine mesh as the Regera. The Regera mesh is more unstable during morphing, due to the thin prism layers. Efforts includes search for optimum trade-off between max displacement of each optimisation cycle and stability. Placement, density, and confinement of control points is adjusted to aid in stability. At the time of the submission of this report, morphing of the winglet was running successfully, whilst optimisation of the mirror, diveplate, and rear-wing is still under investigation. This author will continue this investigation until the Master thesis presentation.

4.7 Discussion of the Adjoint Method

This section aims to sum up the key findings, limitations, and advantages of the adjoint method in OpenFOAM v2006. The value and potential of the adjoint method will be discussed and compared to the current development procedure.

Accuracy Limit

The solving of adjoint PDEs is limited to steady state incompressible flows which can be modelled by turbulence-energy equations. RANS-based turbulence models and steady-state solution algorithms may fail to capture time and length scales important to vehicle aerodynamics. Fundamentally, the adjointOptimisationFoam is limited in accuracy. Although, there are workarounds to these limitations. For instance, the primal PDEs can be solved manually by unsteady solution algorithms, and eventually averaged over an appropriate window for the insertion of field variables when solving the adjoint PDEs. This process has been tested during this thesis and proves an unsteady flow structure formation in the sensitivity map, increasing the sensitivity map credibility. The unsteady simulation of the primal PDEs can be extended to Detached Eddy Simulation, although only with the Spalart-Allmaras turbulence model for adjointOptimisationFoam. Therefore, the limit of the adjoint method in OpenFOAM may be due to the choice of Spalart-Allmaras turbulence modelling, excluding two-equation models due to the frozen turbulence assumption and instability. When solving the adjoint PDEs, the adjoint turbulent state variable, $\tilde{\nu}_a$, of the Spalart-Allmaras turbulence model is limited to first-order scheme as a necessity to converge. Furthermore, it was established that solving the adjoint transpose convection term (ATC) in the adjoint momentum equation, observed in Equation 2.3, may cause divergence and instability. Therefore, the ATC term must be treated with filtering and zeroing of patches. This may lead to further deviation from the exact solution. The convergence sensitivity of the primal field has proven to be an important factor for computing sensitivity maps accurately. Yet, the pressure residuals for both the FSMW and Regera were above 0.001. This may be due to the simulation of an inherently unsteady flow, modelled as steady. The residuals should be approaching zero for the formulations behind the adjoint solver to be valid. On the contrary, the mean pressure coefficient was validated against the averaged experimental pressure coefficient of the FSMW. As a result, it remains unclear whether the primal pressure residual decreases the accuracy significantly. Finally, the stability and accuracy of the adjoint method is bound by the mesh quality. A high-quality mesh of a complex geometry, such as the Koenigsegg Regera, may take additional time to produce compared to the mesh quality needed for convention CFD. Additionally, such a high-quality mesh of a complex geometry may become computationally heavy. Production of a mesh with even more refinement in the necessary regions to ease convergence may take less time to generate but more time to compute.

Key Findings

A parametric study of sensitivity maps, based on a validated primal solution of the FSMW, proved significant dependence on the cell types, the type of ground and wheel simulation, the type of rim, number of iterations and averaging, as well as simulations with unsteady or steady algorithms. The sensitivity maps may be partially inconclusive if the appropriate simulation conditions are not taken into account. Consequently, the parametric study

in section A.3 may serve as a guide in recognising automotive external aerodynamic sensitivity patterns bound by simulation conditions. Next, the sensitivity maps were unreliable in the regions of detached flow or highly rotational flow. Therefore, the credibility of the sensitivity maps should be based on a full understanding of the primal flow field and its simulation setup.
Despite the aforementioned key findings and limitations, the adjointOptimisationFoam algorithm is capable of large-scale automotive development and requires no additional third-party software to work, besides ANSA. The adjointOptimisationFoam algorithm has also proved its capability for a high Reynolds flow with porous zones and MRFs. As a result, it can simulate features similar to the simpleFoam algorithm. If familiar with the OpenFOAM terminology and simulation setup, the adjointOptimisationFoam may be directly implemented into the development strategy of Koenigsegg. The use of the adjoint method may be divided into two applications: (1) Use of sensitivity maps as a tool working along-side conventional CFD development, followed by manual morphing of areas with optimisation potential. (2) Steady optimisation of N cycles.

Sensitivity Maps

The sensitivity map have proved to identify potential optimisation areas which are not directly revealed during post-processing of C_p, C_f, velocity fields, or monitoring of aerodynamic forces. The sensitivity maps have proven to correlate with flow features observed during post-processing of the primal field. However, the area of separated or highly rotational flow causes the sensitivity maps to be of less use. Thus, the sensitivity map shall not stand by itself but should be compared to the primal field scalars and vectors, i.e. skin friction, mean pressure, and velocity. This should aid the engineer in getting the most out of the primal field and the sensitivity maps. Due to the limits of the sensitivity map, the process of manual shape optimisation/morphing cannot rely solely on the sensitivity map and needs critical oversight by an engineer. Conversely, the sensitivity maps may also be used as an indicator of areas with low or negligible optimisation potential regarding the function object. As an example, the sensitivity maps prove no optimisation potential in areas with zero or small pressure gradient, as expected.

Steady optimisation

The steady optimisation have proved that it is competent as an optimisation algorithm for external aerodynamic development of ground vehicles. Notable drag minimisation was achieved within four cycles of the FSMW model. Steady optimisation with the Spalart-Allmaras turbulence model was validated by running three independent simulations with different two-equation turbulence models. Conversely, the steady optimisation is still susceptible to mesh instability which appear when movements of the mesh exceed the quality criteria severely. Additionally, the engineer needs input from a sensitivity map or primal post-processing to place the control points in the first place. The engineer may have to redo the control point placement and adjust the max displacement and confinement variables in case of mesh instability. Different areas need different morphing treatment. During the steady optimisation, the sensitivity derivatives are calculated with respect to predefined control points. Deviance of sensitivity derivatives of control points vs. boundary nodes were proven and should be studied further. Moreover, noise from the sensitivity derivatives may be incorporated into the morphed geometry. This might be alleviated by changing the max morphing displacement of each cycle. The steady optimisation requires many CPU hours and much memory. For instance, a mesh of approximately 47 M cells ran on two Intel Xeon40 ran for 50 hours while completing approximately 3.5 optimisation cycles. This equates to 1, 143 CPU/hours per cycle. CPU/hours per cycle varies significantly with the mesh size, the number of primal and adjoint iterations, as well as the number of control points if using the Field Integral formulation. The great advantage of the steady optimisation is a semi-automated process that may introduce and inspire new alternative designs.

Value of Sensitivity Maps vs. Conventional Post-Processing

The sensitivity maps may be of little use if analysed without preceding knowledge of the primal flow. Vice versa, the conventional CFD development relies on the expertise of an engineer in the analysis of complex flow phenomena. The sensitivity derivatives successfully recognised the drag optimisation potential in all creases, which was also observed in skin friction surface plots of the FSMW model. Decrease of of the mirror of the FSMW and the Regera were identified by the sensitivity map, also revealed by monitoring the aerodynamic forces. Separated flow aft of the wheels, seen on velocity plots, created inconclusive regions in the sensitivity map. Rotational boundary conditions on the wheel arches revealed a wavy pattern. Closer inspection of this pattern revealed potential drag optimisation of the fenders by upstream covering of the wheels. Visual inspection of the primal field pressure, wall shear stress, and local velocity plots indicated no obvious drag optimisation potential in this area. Moreover, there must be smooth flow over the body with no detached flow. The sensitivity map can visualise optimisation potential of each crease. Next, pressure recovery over the rear end, clean separation, and a small base area is optimal. The sensitivity map propose the elongation of the trailing edge and some increase/decrease geometry slope to minimise and balance the wake. Therefore, the sensitivity map may partially guide the engineer in shaping the rear end. Hence, the sensitivity map can aid in awareness.

Sensitivity maps can be used to guide the engineer to regions of optimisation potential for a given function object, if used with consideration and with knowledge of the limitations. The additional time and effort of setting up the adjoint simulations may be neglected, if a standardised setup is used. The primal field still needs to be computed during both the conventional CFD procedure and for the adjoint method. The additional computational time of the adjoint solver may be limited to solving the adjoint equations to generate the sensitivity maps. However, the real temporal cost may be in producing a mesh that is sufficiently high quality to ensure convergence. At the same time, the increase of mesh quality is both beneficial for the conventional CFD procedure and the adjoint method. Steady optimisation by the adjoint method may be limited by the computational resources available to the development team, such as the cost of CPU/hours and memory use. Moreover, the organic shape deformation may be difficult to implement in actual production. Additionally, the challenges of sensitivity derivative noise may also impair the steady optimisation. If computational resources are available, the steady optimisation can certainly deform shapes that this author expects few engineers can do. Accordingly, the steady optimisation can be used as an inspiration followed by manual morphing performed by an experienced engineer. A final remark may be that both AUDI and Volkswagen have continuously been using the adjoint method, also based on the OpenFOAM platform, for more than a decade in which the adjoint method is an active discipline of the aerodynamic vehicle development [3, 4, 20].

Potential of the Adjoint Method

The potential of the adjoint method is not limited to external aerodynamics. The newly implemented adjointOptimisationFoam can handle various function objects, such as total pressure loss, for channel or duct flow optimisation. Additional function objects include minimization of mirror-induced noise by decreasing the turbulent viscosity squared, ν_t^2, or decreasing of aerodynamic moment on a vehicle [11]. More function objects can be used at the same time by weighting. The adjoint development efforts have primarily been on a steady-state simulation. This is because of the computational cost of solving unsteady adjoint PDEs. But, during the parametric study, changing from steady to unsteady state primal solver sensitivity proved to resemble the actual flow better. In this manner, moving towards an unsteady adjoint solver may be a leap towards high accuracy of the adjoint state variables. If computational resources are available, studies have shown gain in accuracy by changing from RANS-based turbulence modelling to (Delayed) Detached Eddy Simulation [3, 15, 20]. As for the adjointOptimisationFoam, the unsteady simulation of the primal PDEs can be extended to Detached Eddy Simulation, although only with the Spalart-Allmaras turbulence model.

5

Conclusion

The adjoint solver in the release of OpenFOAM v2006, adjointOptimisationFoam, was validated on a reference vehicle and successfully run on a Koenigsegg pre-production vehicle. However, limitations of the sensitivity maps were established in the areas of flow separation and highly rotational flow. Stability issues were confirmed in the adjoint formulations and mitigated in terms of filtering, zeroing, under-relaxation, and decreasing scheme order, to overcome divergence issues. Additionally, the mesh quality proved to have impeccable importance for the convergence of the adjoint solver. Encountered challenges when applying the adjoint method to optimise vehicle aerodynamics included:

- Convergence and numerical noise on results
- Accuracy in terms of discretization of the ATC terms, bounding, scheme order, and influence from the primal solution
- Steady state adjoint turbulence formulations
- Need of a higher quality mesh than for conventional CFD

Next, the parametric study proved that sensitivity derivatives related to the simulation setup. Changes of sensitivity derivatives were caused by type ground simulation, different rim configurations, type of cells, and simulating appropriate aerodynamic length and time scale in the primal flow. The findings of the parametric study may aid users of the adjointOptimisationFoam to create state-of-the-art sensitivity maps.

With established knowledge, the sensitivity map is not intended to substitute the conventional CFD development process, but it might aid in awareness during optimisation processes. Furthermore, the adjoint automated shape process, denoted steady optimisation, is inherently bound by initial mesh quality and requires a significant amount of CPU/hours and memory, compared to conventional CFD development and the computation of sensitivity maps. The steady optimisation performed drag reducing optimisation loops, in which a decrease of three drag counts were obtained within four shape optimisation cycles. There is need for further computation of cycles to understand the effectiveness of the steady optimisation. As a result, the adjoint method is a leap towards automation and guidance of aerodynamic optimisation. For the time being, there is still a need for an engineer to actively judge and justify in which areas a given geometry displays optimisation potential and how to morph the geometry. Time and effort of setting up an adjoint simulation may be neglected, if a standardised setup is used. The additional computational time of the adjoint solver may be limited to solving the adjoint equations to generate the sensitivity maps. Additional temporal cost may be due to the generation of a mesh of high quality that ensure convergence. Yet, a high quality mesh is beneficial for both the conventional CFD procedure and the adjoint method. A guide of recommendations were made to ease the implementation of the adjoint method in aerodynamics development and can be found in Appendix section A.5.

5.1 Future Work

This author has recommendations for future research within the field of the adjoint method applied on ground vehicles. For instance, in current version of the adjointOptimisationFoam, the adjoint solver accuracy might be increased in at least two ways: (1) Direct modelling of the viscous sublayer and logarithmic layer [22]. The enhanced wall modelling should aid in the accuracy of predicting when the flow separates. (2) Updating the primal solver to *Detached Eddy Simulation*. This requires computational resources beyond the limits of this thesis work. Delayed Detached Eddy Simulation of the primal solver can enhance the sensitivity map accuracy and identify transient flow features herein [20]. When running unsteady simulations, a full car simulation should be performed to capture the three-dimensional flow in the wake. Appropriate averaging of the temporally dependent state variables is crucial to create the best possible sensitivity map. As a result, unsteady full car simulation is expected to increase accuracy the sensitivity map in the entire region aft of the rear wheels.

This thesis could have continued the application of two-equation models for the adjoint solver. This would cause a decrease in scheme order, frozen turbulence assumption, and heavy filtering of the ATC term. However, it should be investigated how much these assumptions and decrease in solver precision actually impact the sensitivity maps.

Furthermore, the adjoint method may be extended to co-existing fields of automotive aerodynamics. Thus, the test of alternative function objects could be performed in terms of noise minimisation, external aerodynamic moment optimisation, and internal flow optimisation of total pressure loss.

Investigation of irregular patterns, observed on the sensitivity maps, should be studied with the purpose of creating smoother sensitivity maps and increasing the mesh morphing quality for steady optimisation. In this regard, the cell size should be investigated in terms of the adjoint distance variable which is used for calculating the sensitivity derivatives. Cell refinement in areas of curvature should be included herein. Additional investigation of which degree the turbulence state variables' are a source of sensitivity map noise should be conducted.

The hypothesis of sensitivity maps and steady optimisation being useful alongside conventional CFD development should be tested by engineers other than the author himself. This might include manual morphing to optimise for a given objective function by the engineer, first by the conventional CFD procedure, followed by access sensitivity maps and subsequent manual morphing once again. Running simulations of the two different scenarios may establish the delta values of which of the method proved best.

Finally, in terms of steady optimisation, a study of placement and density of control points should be performed in order to approach sensitivity derivatives derived with respect to boundary nodes.

Bibliography

[1] Ashton, N., Skaperdas, V., Sinclair O., and Van Noordt, W. First Automotive CFD Prediction Workshop: Cross-Plotting- Case2a/b. University of Oxford, AWSM, and BETA Simulation Solutions, 2019.

[2] Batten, P., Wang, J., and Peroomian, O. 1st Automotive Workshop, Oxford, UK. Metacomp Technologies, 2019.

[3] Blacha, T. and Islam, M. The Aerodynamic Development of the New Audi Q5. *SAE Int. J. Passeng. Cars - Mech. Syst*, 10(2):638–648, 2017. doi:10.4271/2017-01-1522.

[4] Blacha, T., Gregersen, M., Islam, M., and Bensler, H. Application of the Adjoint Method for Vehicle Aerodynamic Optimization. *SAE Technical Paper 2016-01-1615*, 2016. doi:10.4271/2016-01-1615.

[5] Calderon, J. C., and Zastawny, M. 1st Automotive CFD Prediction Workshop. Siemens Digital Industries Software, 2019.

[6] Chair of Aerodynamics and Fluid Mechanics. Drivaer model. URL: https://www.mw.tum.de/en/aer/research-groups/automotive/drivaer.

[7] Collin, C., Mack, S., Indinger, T., and Mueller, J. A Numerical and Experimental Evaluation of Open Jet Wind Tunnel Interferences using the DrivAer Reference Model. *SAE Int. J. Passeng. Cars - Mech. Syst.*, 9(2):657–679, 2016. doi:10.4271/2016-01-1597.

[8] Fangqing, L. A thorough description of how wall functions are implemented in openfoam. 2016. URL: http://www.tfd.chalmers.se/~hani/kurser/OS_CFD_2016.

[9] Fotiadis, G., Skaperdas, V., and Iordanidis, A. The Influence of Mesh Characteristics an Open Form Simulations are the Driver Model. 6th. BETA CAE International Conference, 2015.

[10] Frank, T., Gerlicher, B., and Alberto, J. DrivAer-Aerodynamic Investigations for a New Realistic Generic Car Model using ANSYS CFD. *Automotive Simulation World Congress*, 2013.

[11] Giannakoglou, K. C., Papoutsis-Kiachagias, E. M., and Gkaragkounis, K. T. User Manual: AdjointOptimisationFoam, an OpenFOAM-based optimisation tool. 2020. URL: https://www.openfoam.com/releases/openfoam-v2006/numerics.php.

[12] Giles, M. B., and Pierce, N. A. An Introduction to the Adjoint Approach to Design. *Flow, Turbulence and Combustion*, 65:393–415, 2000.

[13] Heft, A. I., Indinger, T., and Adams, N. A. Introduction of a New Realistic Generic Car Model for Aerodynamic Investigations. *SAE International*, 2012. doi:10.4271/2012-01-0168.

[14] Hucho, W. H. *Aerodynamics of Road Vehicles.* SAE International, 1998.

[15] Karpouzas, G. K, Papoutsis-Kiachagias, E. M., Schumacher, T., De Villiers, E., Giannakoglou, K. C., and Othmer, C. Adjoint Optimization for Vehicle External Aerodynamics. *International Journal of Automotive Engineering*, 7:1–7, 2016.

[16] Katz, J. *Race Car Aerodynamics: Designing for Speed*. Bentley Publishers, 2006.

[17] Kavvadias, I. S., Papoutsis-Kiachagias, E. M., and Giannakoglou, K. C. On the Proper Treatment of Grid Sensitivities in Continuous Adjoint Methods for Shape Optimization. *Journal of Computational Physics*, 301:1–18, 2015. doi:10.1016/j.jcp.2015.08.012.

[18] Kundu, P. K., Cohen, I. M., and Dowling, D. R. *Fluid Mechanics Sixth Edition*. Academic Press, 2016.

[19] OpenFOAM. User guide: Spalart-allmaras. 2020. URL: https://www.openfoam.com/documentation/guides/latest/doc/guide-turbulence-ras-spalart-allmaras.html.

[20] Othmer, C. Adjoint Methods for Car Aerodynamics. *International Journal of Automotive Engineering*, 4(6):1–23, 2014.

[21] Othmer, C., Papoutsis-Kiachagias, E. M., and Haliskos, K. CFD Optimization via sensitivity-based shape morphing. *4th ANSA μETA International Conference*, 2011.

[22] Papoutsis-Kiachagias, E. M., and Giannakoglou, K. C. Continuous Adjoint Methods for Turbulent Flows, Applied to Shape and Topology Optimization: Industrial Applications. *Arch Computat Methods Eng*, 23:255–299, 2016. doi:10.1007/s11831-014-9141-9.

[23] Papoutsis-Kiachagias, E. M., Asouti, V. G., Giannakoglou, K. C., Gkagkas, K., Shimokawa, S., and Itakura, E. Multi-Point Aerodynamic Shape Optimization of Cars based on Continuous Adjoint. *Structural and Multidisciplinary Optimization*, 59:675–694, 2019. doi:10.1007/s00158-018-2091-3.

[24] Spalart, P. and Allmaras, S. A One-Equation Turbulence Model for Aerodynamic Flows. *La Recherche Aérospaciale*, (1):5–21, 1994.

[25] Strangfeld, C., Wieser, D., Schmidt, H. J., Woszidlo, R., Nayeri, C. and Paschereit C. Experimental Study of Baseline Flow Characteristics for the Realistic Car Model DrivAer. *SAE International*, 2013. doi:10.4271/2013-01-1251.

[26] Sumer, B. M. *Lecture Notes on Turbulence*. Technical University of Denmark, DTU Mekanik, Section for Fluid Mechanics, Coastal & Maritime Engineering, 2013.

[27] Wieser, D., Schmidt, H. J., Mueller, S., Strangfeld, C., Nayeri, C., and Paschereit, C. Experimental Comparison of the Aerodynamic Behavior of Fastback and Notchback DrivAer Models. *SAE International*, 2014. doi:10.4271/2014-01-0613.

[28] Xu, J. and Li, M. Adjoint-based Sensitivity Analysis for External Aerodynamics. Master thesis, Chalmers University of Technology, 2017.

[29] Yazdani, R. Steady and Unsteady Numerical Analysis of the DrivAer Model. Master thesis, Chalmers University of Technology, 2015.

[30] Zymaris, A. S., Papadimitriou, D. I., Giannakoglou, K. C., and Othmer, C. Continuous adjoint approach to the Spalart-Allmaras turbulence model for incompressible flows. *Computers & Fluids*, 38:1528–1538, 2009. doi:10.1016/j.compfluid.2008.12.006.

A
Appendices

A.1 Discrete vs. Continuous Adjoint Solver

There are two types of adjoint formulations; the discrete and the continuous. The discrete adjoint formulation starts by discretizing the PDEs and subsequently differentiates these equations numerically, while the continuous adjoint formulation receives differentiated PDEs prior to the discretization. Both the discrete and continuous adjoint variants can provide sufficiently accurate derivatives for use in nowadays optimisation problems, if approaching infinite grid resolution [22]. OpenFOAM v2006 is built on the continuous adjoint method in the official release. This thesis is primarily concerned with the continuous approach. However, in an attempt to distinguish the two methods, a brief list of differences are provided below [12, 22].

- Sensitivity derivatives computed by the discrete adjoint method may match the reference values, i.e., from Finite Difference, to a higher degree of accuracy, since the discrete adjoint method takes into account the primal discretization scheme.
- As the mesh size increases this difference in accuracy diminishes and does not greatly affect the convergence of the optimisation algorithm.
- The continuous adjoint method is more forward to implement as the decretization scheme applied to the adjoint PDEs does not have to match the primal PDEs to compute accurate sensitivity derivatives.
- The continuous adjoint variant is less expensive in terms of CPU cost and memory requirements per iteration. This is due to modifications to the initial adjoint formulation which allows the computation of sensitivity derrivatives based exclusively on Surface Integrals and not Field Integrals. This reduces the cost and complexity of the continuous adjoint approach.
- It has been found that the surface-based method may compute sensitivities of the wrong sign in areas with geometrical singularities (i.e., at the trailing edge of an airfoil) in turbulent flow fields [17].
- In the continuous adjoint approach, the physical understanding of the adjoint system is more transparent, as closed-form expressions exist for the field adjoint equations, their boundary conditions and the sensitivity derivatives expression. Thus, the impact of each term can be quantified. If terms of negligible or small importance are identified, their computation can be avoided to simplify the method.

A.2 FSMW Spalart-Allmaras 140kmph Validation

Pressure Distribution

Figure A.1a shows the numerical distribution of the mean static pressure coefficient, C_p, of the FSMW. The experimental results are given as reference in Figure A.1b. As seen in Figure A.1a and Figure A.1b the base of the windshield inherits a stagnation area with a local high pressure. Subsequently, the flow accelerates towards the roof and over the A-pillar. The numerical pressure distribution over the windshield correlates well with the experimental measurements. Note the slightly different colour bar. It is observed that high-curvature convex regions of the bonnet, mirror, A-pillar, and wheels decrease pressure. Low pressure indicates suction. Low pressure at the A-pillar may be sources of vortex formation and recirculating flow.

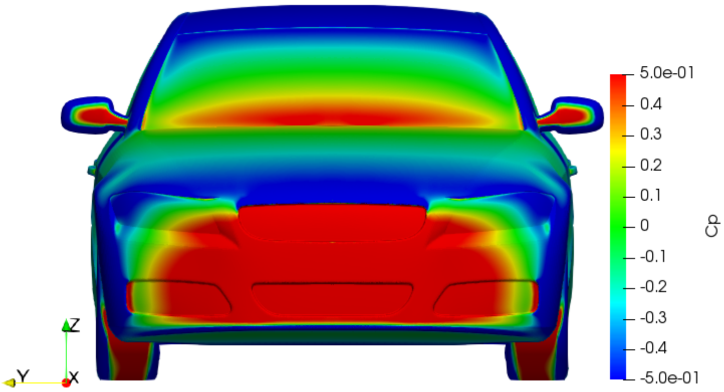

(a) Numerical pressure distribution front of the FSMW.

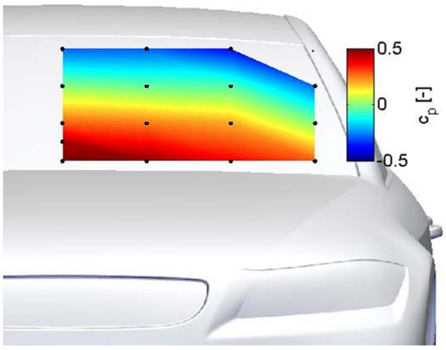

(b) Experimental surface pressure coefficient of wind shield from [13].

Figure A.1: Pressure distribution on front of FSMW.

Note that the change in pressure distribution in the windshield area near the mirror may differ $C_p \pm 0.05$ for the experimental values, since no mirror is presented at Figure A.1b [13]. The numerical mean static pressure distribution in Figure A.2a identify similar flow characteristics seen in the experimental pressure distribution in Figure A.2b. It is observed that the A-pillar vortex develops close to the root of the A-pillar and detaches when approaching the upper part of the driver side window.

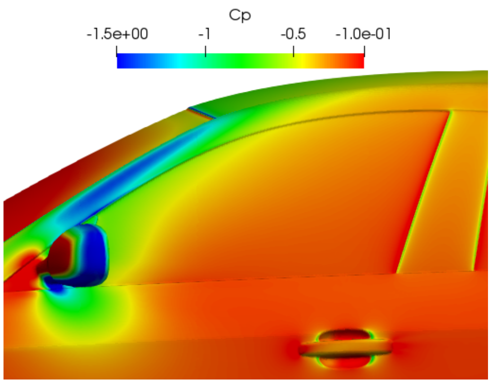

(a) Numerical pressure distribution side of the FSMW.

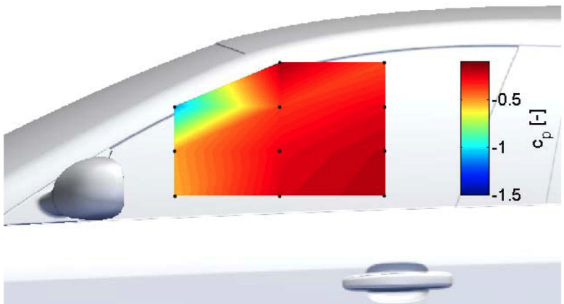

(b) Experimental surface pressure coefficient of driver side window from [13].

Figure A.2: Pressure distribution on side of FSMW.

Furthermore, the flow is the separated in the wake of the mirror. The pressure distribution on the side of the FSMW is not greatly affected by the mirror wake flow.
Both Figure A.3a and Figure A.3b show a low pressure zone over the C-pillar. Moreover, a near zero static pressure zone at the lower part of the rear window and at the upper surface of the trunk may indicate a reattachment of a separation bubble. The sharp trailing edge of the roof causes a pressure drop. Some pressure recovery is observed over the rear window, with the existing adverse pressure gradient present.

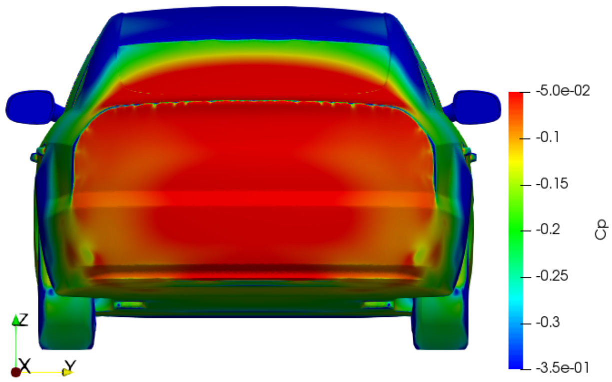

(a) Numerical pressure distribution rear of the FSMW.

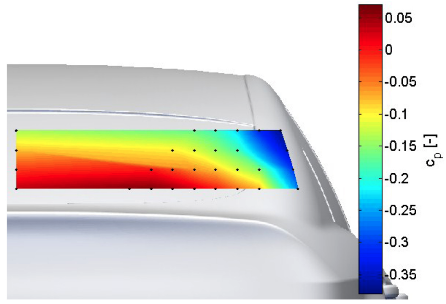

(b) Experimental pressure distribution of rear from [13].

Figure A.3: Pressure distribution on rear of FSMW.

Static pressure increase and occurs on the upper trunk surface. This is true for both the numerical simulation and experimental results. The low pressure at the upper region of the C-pillars inherits accelerated flow due to the model geometry. Pressure recovery over the rear window are confirmed by Figure A.4b from [27]. Figure A.4b also confirms a sudden pressure drop due to the trailing edge of the roof.

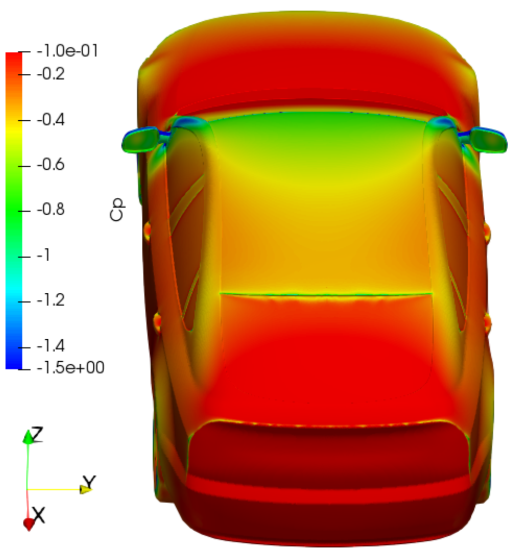

(a) Numerical pressure distribution rear of the FSMW.

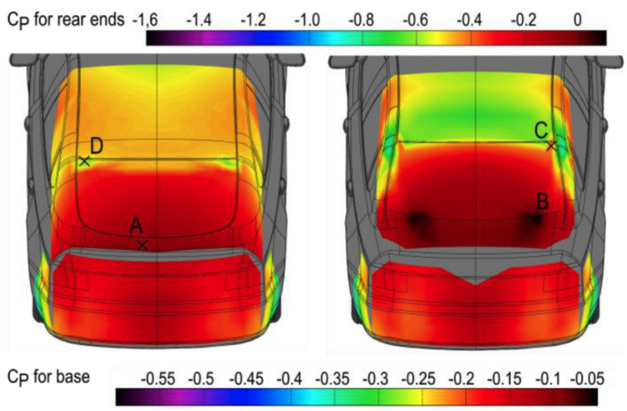

(b) Experimental surface pressure coefficient of 2/3 rear view for the fastback (left) and notchback (right) from [27]. Pressure distribution of notchback may be ignored.

Figure A.4: Pressure distribution of FSMW.

Pressure recovery is observed as the flow moves over the roof for both numerical and experimental studies. The base pressure is relatively high and uniform for both experimental and numerical studies, which is good for drag reduction. Clean separation at the edges of the base is not necessarily achieved due to the wake flow of the rear wheels. Note that the pressure distribution over the underbody is not presented as it is simplified to be flat and smooth, and will not be covered during the adjoint simulation. The underbody shows a nearly uniform negative pressure distribution, with regions of lower pressure between the front wheels.

Vorticity

Figure A.5a and Figure A.5b shows the pressure distribution along with the normalized vorticity. The pressure recovery starts mid-roof, in Figure A.4b and Figure A.5. This pressure recovery is seen from increasing pressure coefficient over the roof, backlight, and the rear trunk surface.

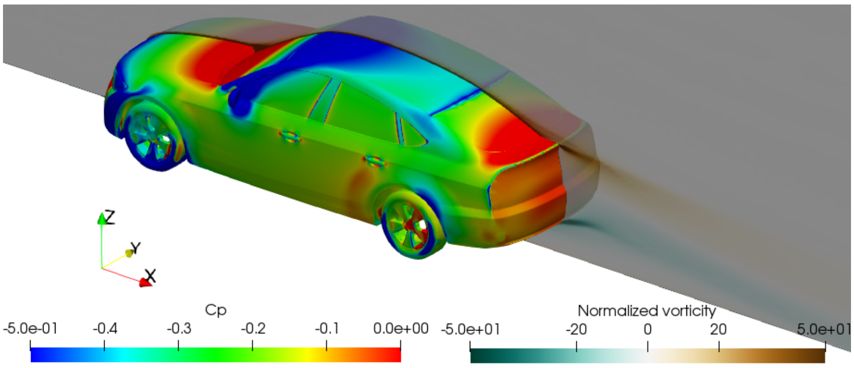

(a) Numerical pressure distribution, velocity field and normalised vorticity FSMW.

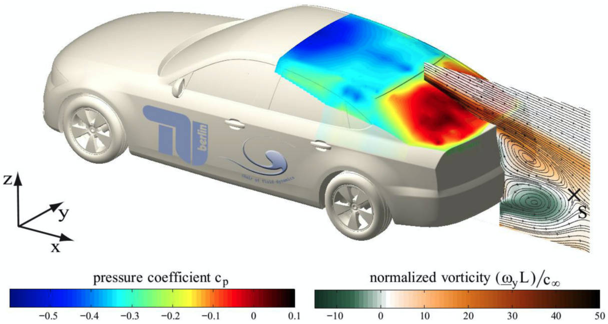

(b) Experimental pressure distribution, velocity field and normalised vorticity from [27].

Figure A.5: Pressure distribution, velocity field and normalised vorticity.

The flow over the backlight remains partially attached. A shear layer is generated in the boundary layer, identified by the dark brown area over the rear window. The shear layer convects downstream, thickens, followed by a separation at the trailing edge of the trunk. Two free shear layers emerge, visualised by the brown and green regions, which roll up and form vortices. A free stagnation point, denoted S at Figure A.5b, is observed where the two counter-rotating vortices meet. The upper and lower shear layer merge and form a common wake profile downstream of this point. Note that the numerical and experimental counter-rotating vortices deviate in size lengthwise. Thus, it appears that the wake region of the numerical simulation is longer than the experimental results. This may be due to different Reynolds numbers with the experimental Reynolds number being 2.8M [27] and the numerical 12.2M. Alternatively, one can question the capability of the steady state simulation, which may be insufficient in resolving the complex 3D flow. A time-averaged unsteady simulation may inherit compute a smaller wake than a steady state simulation. Moreover, it should be noted that the numerical simulation does not take into account the spanwise velocity through the symmetry plane, due to the nature that it is a symmetry car simulation. Both the numerical and experimental flow solutions depict mean flow. The actual instantaneous structures appear less organized, imbalanced, and disruptive.

A.3 Parametric Study of FSMW Sensitivity Maps

A parametric study of the FSMW sensitivity maps were carried out to study the sensitivity and reliability of the adjoint method when changing simulation parameters. All simulations in this parametric study were carried out with the Spalart-Allmaras turbulence model both for the primal and adjoint simulations. All simulations are of the model FSMW with smooth rims, apart from simulations F and S with open five spoke rims.

	Cells		Road and Wheel BCs	Iterations Primal / Adjoint	TI / Eddy viscosity ratio	Steady / Unsteady
	Surface	Interior				
A	Prism	Hexa	Stationary	5,000 / 5,000	1% / 1.5	Steady
B	Prism	Tetra	Stationary	5,000 / 5,000	1% / 1.5	Steady
C	Hexa	Hexa	Stationary	5,000 / 5,000	1% / 1.5	Steady
D	Prism	Hexa	Stationary	10,000 / 10,000	1% / 1.5	Steady
E	Prism	Hexa	GS,RWBC	10,000 / 10,000	1% / 1.5	Steady
F	Prism	Hexa	GS,RWBC,MRF	10,000 / 10,000	1% / 1.5	Steady
N	Prism	Hexa	Stationary	15,000 / 15,000	1% / 1.5	Steady
Q	Prism	Hexa	GS,RWBC,MRF	5,000 / 5,000	1% / 2	Steady
S	Prism	Hexa	GS,RWBC,MRF	10,000 / 10,000	1% / 1.5	Unsteady

Table A.1: Table of parametric study of the sensitivity maps of the FSMW model. *GS (Ground Simulation), **RWBC (Rotating Wall Boundary Condition), ***MRF (Moving Reference Frame).

The study is made with a target of providing a more holistic view of the sensitivity map.

A.3.1 (A) Hexa Cell Interior vs. (B) Tetra Cell Interior

This study was carried out to see the effect of changing interior cells from hexa to tetra or vice versa on the sensitivity maps. The model in question is the FSMW symmetry model, with smooth rims. Each sensitivity map is virtually cut over the symmetry plane (XZ). It should be confirmed that both simulations are identical in all parameters, including local cell length and volume refinement. Mesh A (hexa interior) has 20,791,135 cells, while mesh B (tetra interior) has 25,307,606 cells. Changes are observed at the leading and trailing edges of the FSMW Figure A.6.

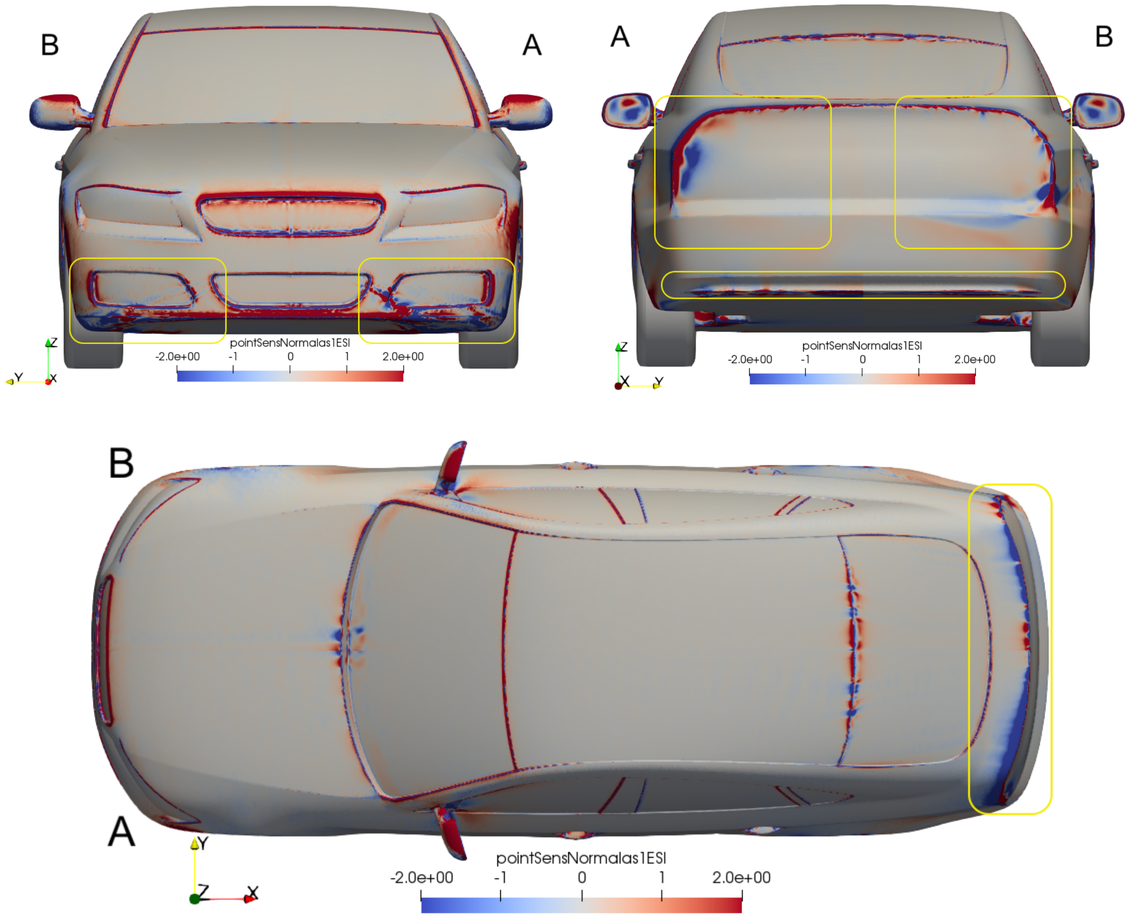

Figure A.6: Sensitivity maps front, rear, and top of (A) Hexa Cell Interior vs. (B) Tetra Cell Interior.

Changes are observed on front and rear wheel arches are observed in Figure A.7. Opposite sign of the sensitivity derivative is observed in lower region of the fenders. A significantly different sensitivity map is observed on the side trailing edge. Thus, changes to the sensitivity maps are correlated to changes of the interior cell type. Tetra cells have a larger flux to volume ratio, which might explain the significant changes over the trailing top and side edge, respectively in Figure A.6 and Figure A.7.

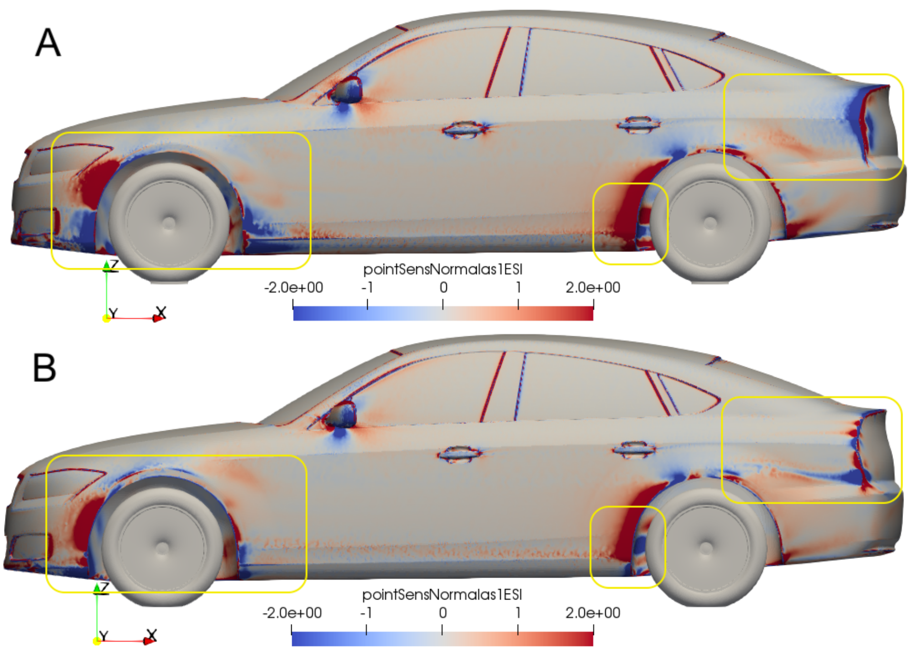

Figure A.7: Side sensitivity maps. (A) Hexa Cell Interior vs. (B) Tetra Cell Interior.

A separate study, performed by the author of this thesis reveals changes to force coefficients during modelling of tetra and hexa cells according to respective turbulence models. *Spalart-Allmaras*, *Realizable* k-ϵ, k-ϵ, and k-ω *SST*, are respectively denoted by *SA*, *RKE*, *kE*, and *SST* in Figure A.8. Most turbulence models drag coefficients are within the experimental values, except for the k-ω *SST* with hexa interior and *Realizable* k-ϵ with tetra interior. Both the *Realizable* and *SST* turbulence models have a large deviation of min and max values recorded within the averaging interval. This may be due to less numerical damping of the *Realizable* and *SST* turbulence models compared to the $Spalart-Allmaras$ and k-ϵ turbulence models. The large max and min deviations from the mean, indicating that the actual flow is transient. The Spalart-Allmaras is observed to be the turbulence model least affected by the interior cell type, and corresponds well with all experimental drag coefficient values. The Spalart-Allmaras turbulence model predicts lift coefficients with less accuracy, as observed in Figure A.8b and Figure A.8c. Generally, all turbulence models predict lift coefficients with some deviation from the experimental values. This trend too is observed in [2, 9]. The deviation from the experimental values could be measurement related, since the lift values are small in magnitude. Finally, Spalart-Allmaras predict less C_d, C_{lf}, and C_{lr} than *Realizable* k-ϵ, in both hexa and tetra interior meshes. This trend is consistent for all turbulence models.

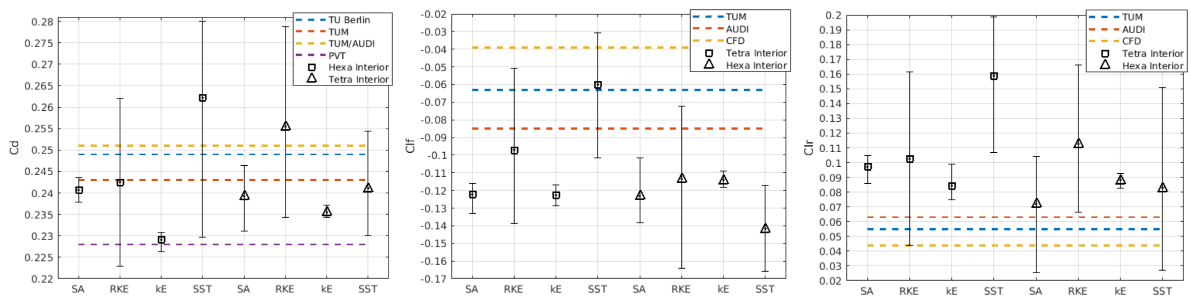

(a) C_d (b) C_{lf} (c) C_{lr}

Figure A.8: Primal force coefficients. Averaged last 3,000 out of 10,000 iterations.

A.3.2 (A) Prism Cell Surface vs. (C) Hexa Cell Surface

This study was carried out to see the effect of changing surface cells (boundary layer cells) from hexahedra to prisms or vice versa on the sensitivity maps. The model in question is the FSMW symmetry model, with smooth rims. Each sensitivity map is virtually cut over the symmetry plane (XZ). Minor changes are observed at the cowl area, trailing edge of the roof, and mirror housing in Figure A.9 and Figure A.10. Additionally, the trailing edge of the trunk has different distribution of inwards (blue) wall-normal displacement. Different irregular displacement zones under the doors are observed in Figure A.10. The rear fender, in Figure A.10, shows significantly different wall normal displacement.

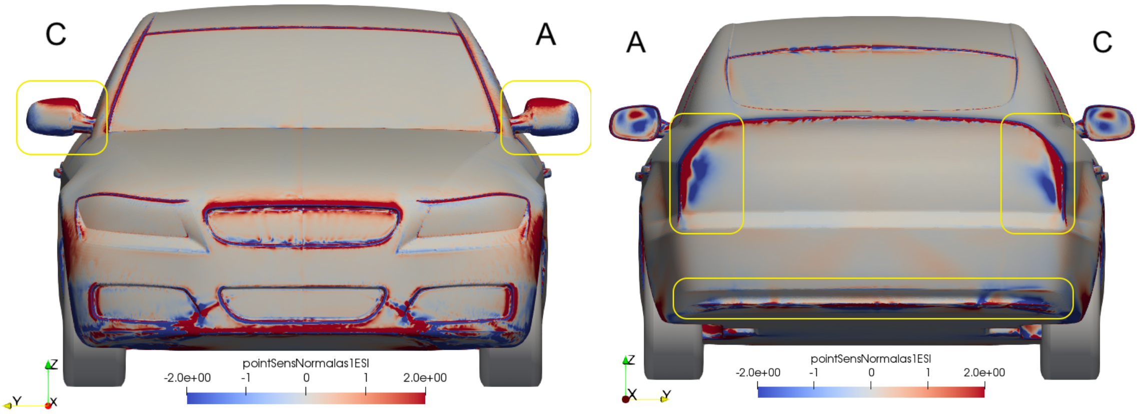

Figure A.9: Sensitivity maps of front and rear of (A) Prism Cell Surface vs. (C) Hexa Cell Surface.

It should be mentioned that both mesh A (Prism cells) and mesh C (Hexa cells) have equal cell height and length, and volume refinement. Due to the geometry of the cells, mesh A has a total of 20,781,135 cells while mesh C only has 16,055,019 cells. In addition, prism cells are considered advantageous when modelling non-planar surfaces, to keep a constant cell height. Thus, the sensitivity map of mesh A is considered more accurate.

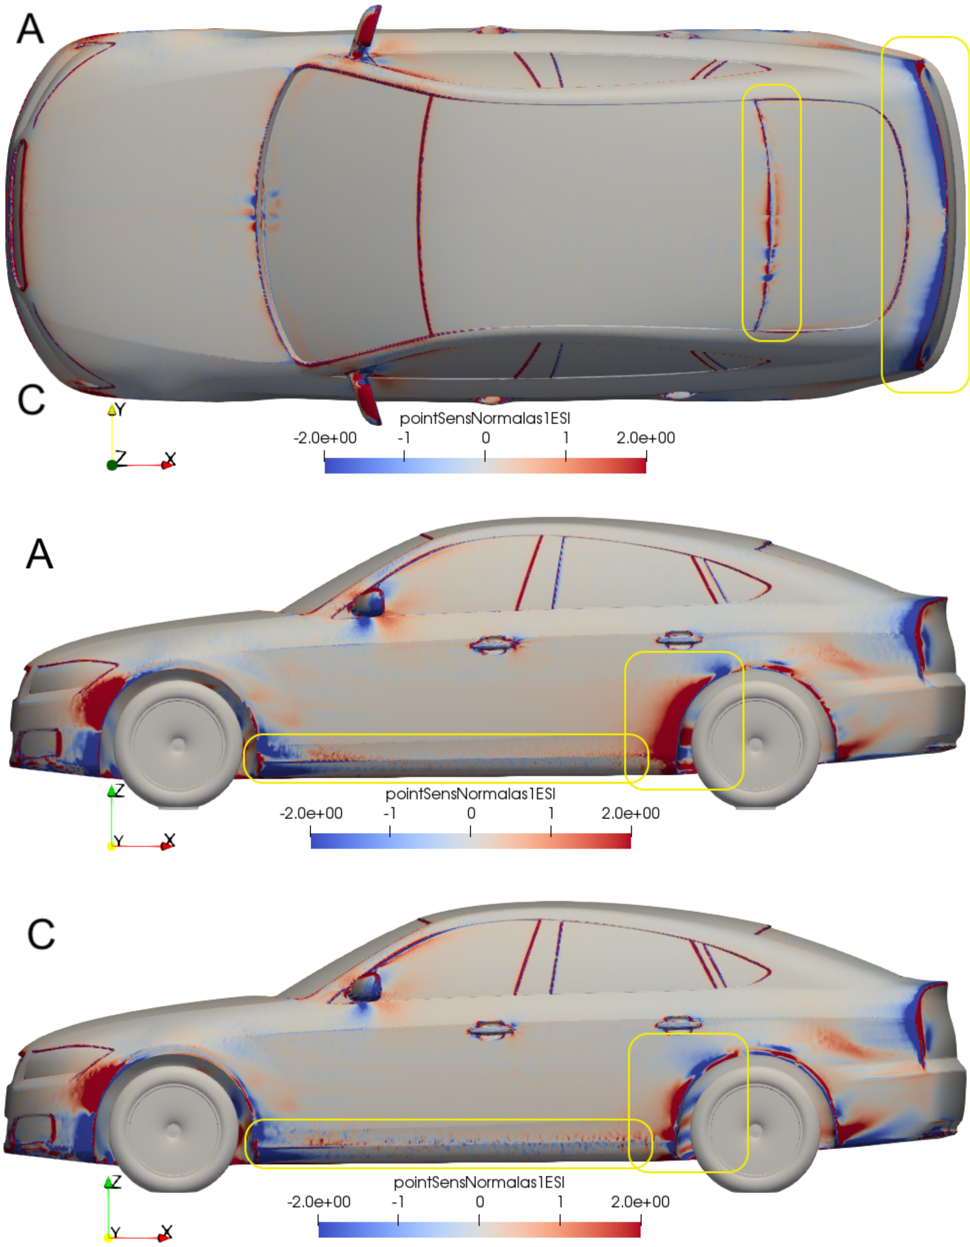

Figure A.10: Sensitivity maps top and sides of (A) Prism surface vs. (C) Hexa surface.

A.3.3 (A) 5,000 Iterations, (D) 10,000 Iterations, (N) 15,000 Iterations

This study was carried out to see the effect of varying the number of iterations of primal and adjoint simulations. The model in question is the FSMW symmetry model, with smooth rims. Note that the continuity equation is run continuously over both the primal and adjoint simulations. The primal residuals converge at around 5,000 iterations for all simulations. However, the adjoint residuals continue to decrease. This may explain the changes in sensitivity maps despite of a seemingly converged primal flow. Averaging is carried out differently, namely, in the primal simulation, only when convergence is achieved. Thus, the primal simulation A was averaged from 2,000-5,000 iterations, primal simulation D was averaged from 5,000 to 10,000 iterations, and primal simulation N was averaged from 10,000 to 15,000 iterations.

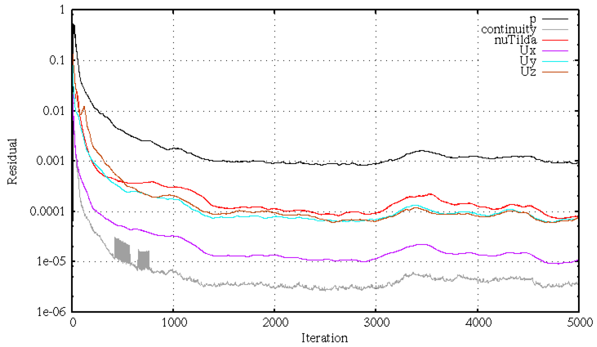

(a) Primal residuals simulation A.

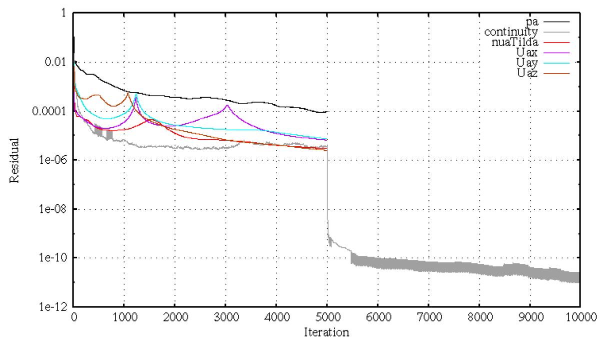

(b) Adjoint residuals simulation A.

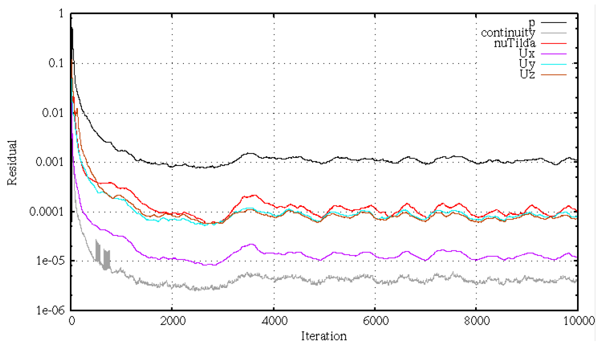

(c) Primal residuals simulation D.

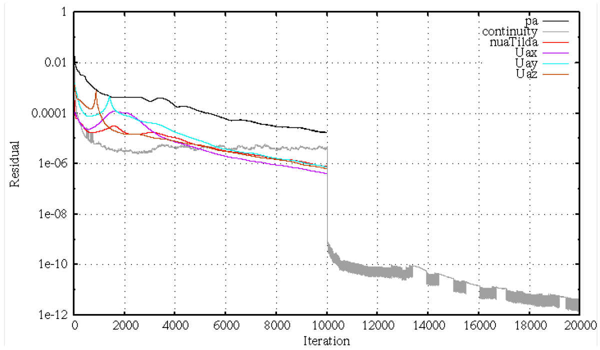

(d) Adjoint residuals simulation D.

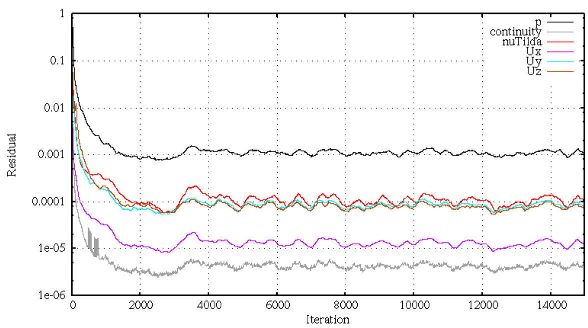

(e) Primal residuals simulation N.

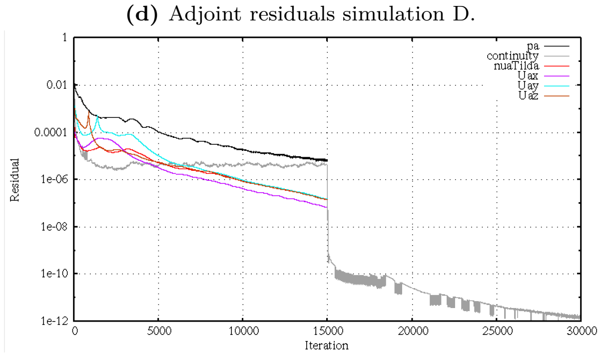

(f) Adjoint residuals simulation N.

Figure A.11: Residuals vs. number of iterations for primal and adjoint solutions.

Each sensitivity map is virtually cut over the symmetry plane (XZ). All simulation setups A, N, and D are completely identical, besides the number of iterations and averaging of the primal flow. Table A.1 provides further details of the setup. Simulation A has run

5,000 primal iterations and 5,000 adjoint iterations. Simulation D has run 10,000 primal iterations and 10,000 adjoint iterations. Finally, simulation N ran 15,000 primal iterations and 15,000 adjoint iterations. It is observed in Figure A.12, that self-similar sensitivity maps are obtained for the front of the car. However, the sensitivity maps exhibit different patterns on the rear end in Figure A.12. Simulation D and N are forming nearly the same pattern, however, with differences in the region of the trailing edge side/rear corner. Thus, the adjoint flow might have converged in front of the car, but not downstream of the car.

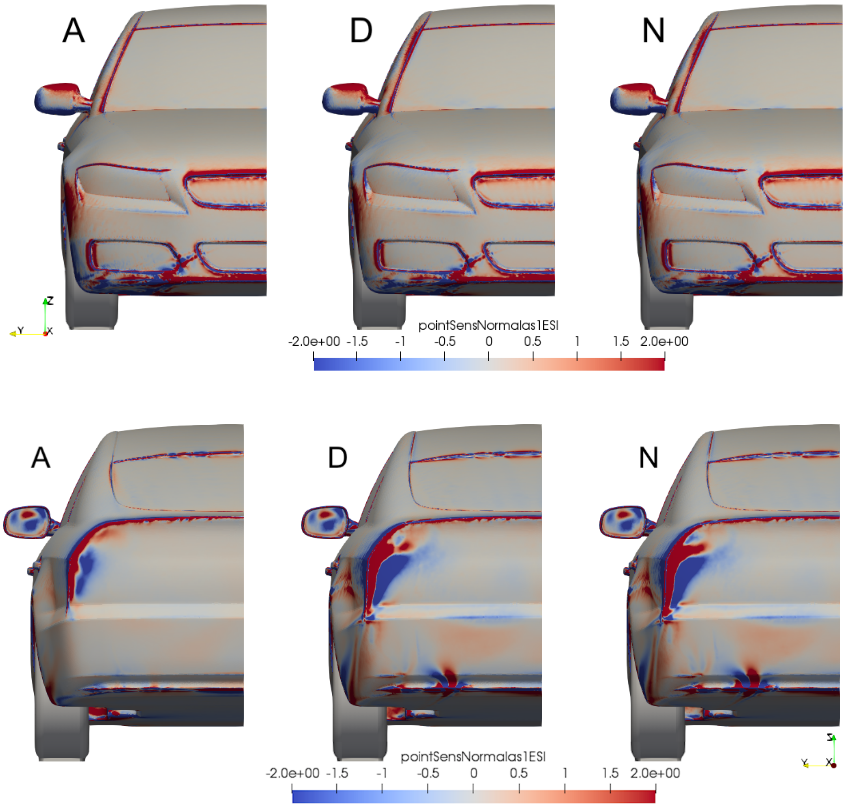

Figure A.12: Sensitivity maps front and rear of (A) 5,000 Iterations, (D) 10,000 Iterations, (N) 15,000 Iterations.

Next, changes in sensitivity maps between the different simulations are observed at the trailing trunk edge and the crease between the trailing edge of the roof and the leading edge of the rear window. In addition, change in gradient magnitude is observed over the A-pillar.
Significant changes are observed between the sensitivity maps downstream of the rear wheel. This is most like due to wheel wake flow interacting with the sensitivity derivatives. This wake may not have converged during simulation A. However, the region of the rear fender upstream of the rear wheel has developed self-similar sensitivity maps. The impact of iteration vs. wake development of the wheels is also observed aft of the front wheel. However, change to the front fender upstream of the front wheel is observed over the sensitivity maps. For this particular simulation setup, it is evident that the leading fender regions should be displaced wall-normal outwards to cover the wheels.

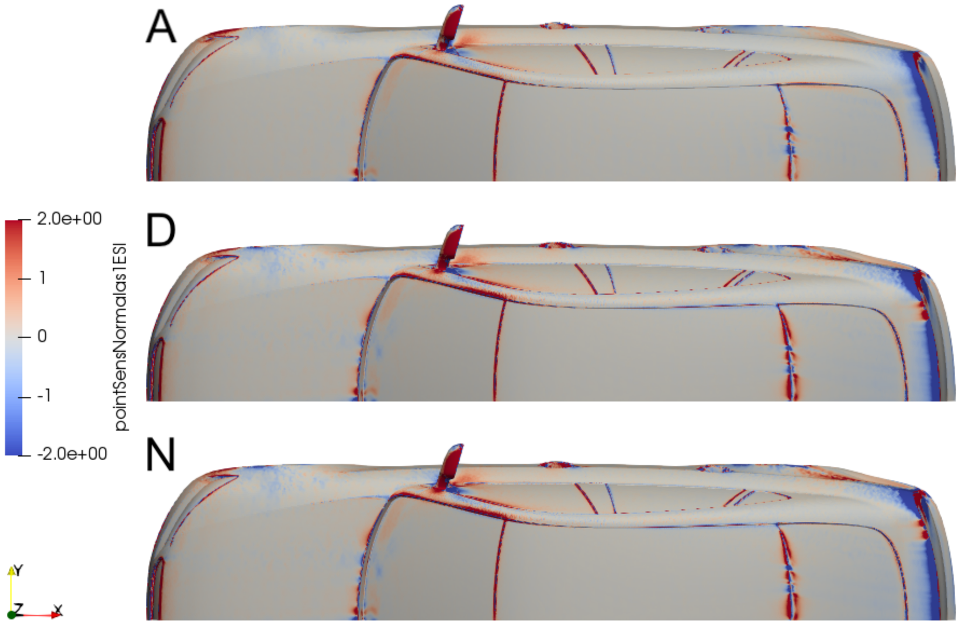

Figure A.13: Top sensitivity maps of (A) 5,000 Iterations, (D) 10,000 Iterations, (N) 15,000 Iterations.

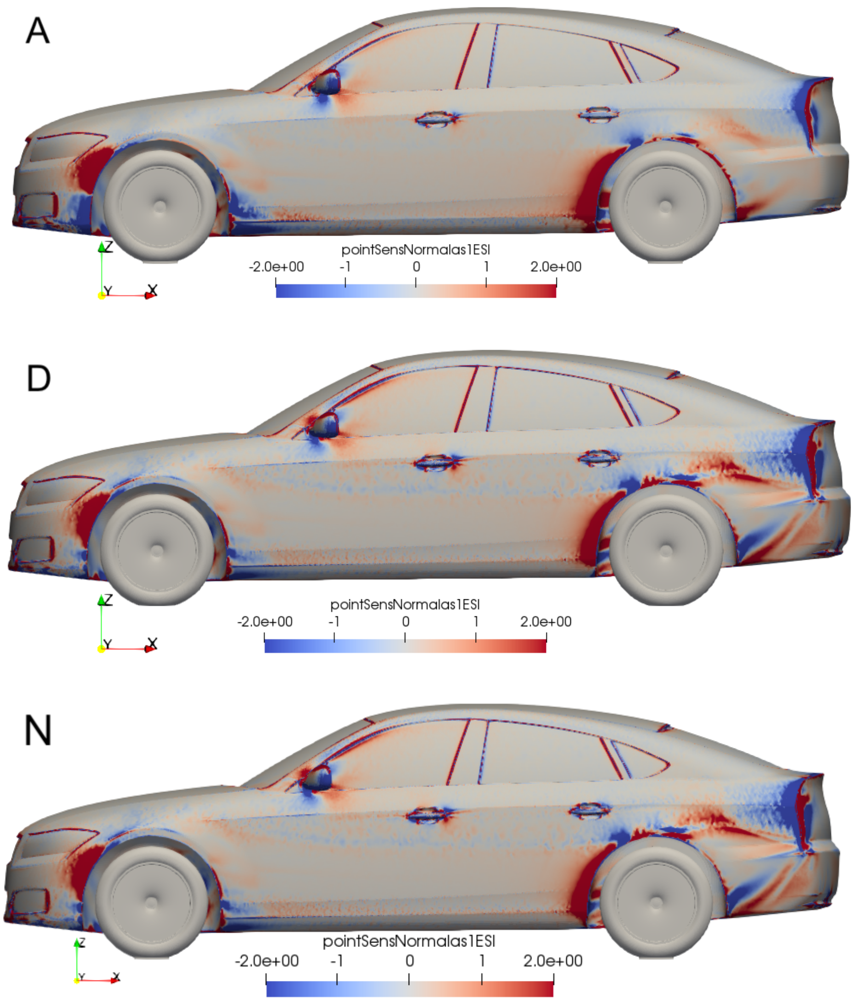

Figure A.14: Sensitivity maps sides of (A) 5,000 Iterations, (D) 10,000 Iterations, (N) 15,000 Iterations.

A.3.4 (A) vs. (Q) Changes of Turbulence Properties and BC's

This study investigate the effect on sensitivity maps when changing turbulence intensity, eddy viscosity ratio, and boundary conditions over the symmetry plane. The model in question is the FSMW symmetry model with smooth rims. Each sensitivity map is virtually cut over the symmetry plane (XZ). Both simulations were identical in setup and mesh, apart from the aforementioned boundary conditions, and run with 5,000 primal iterations (last 3,000 iterations averaged) and 5,000 adjoint iterations. As observed on Figure A.15 and Figure A.16, the symmetry plane boundary condition has an impact on the sensitivity map. Simulation (A) has been run with symmetry boundary conditions at the symmetry plane, while Simulation(Q) has been run with zero gradient and slip boundary condition. This creates minor changes to the trailing trunk edge, cowl area near symmetry plane and the trailing roof / leading rear window edge in Figure A.16.

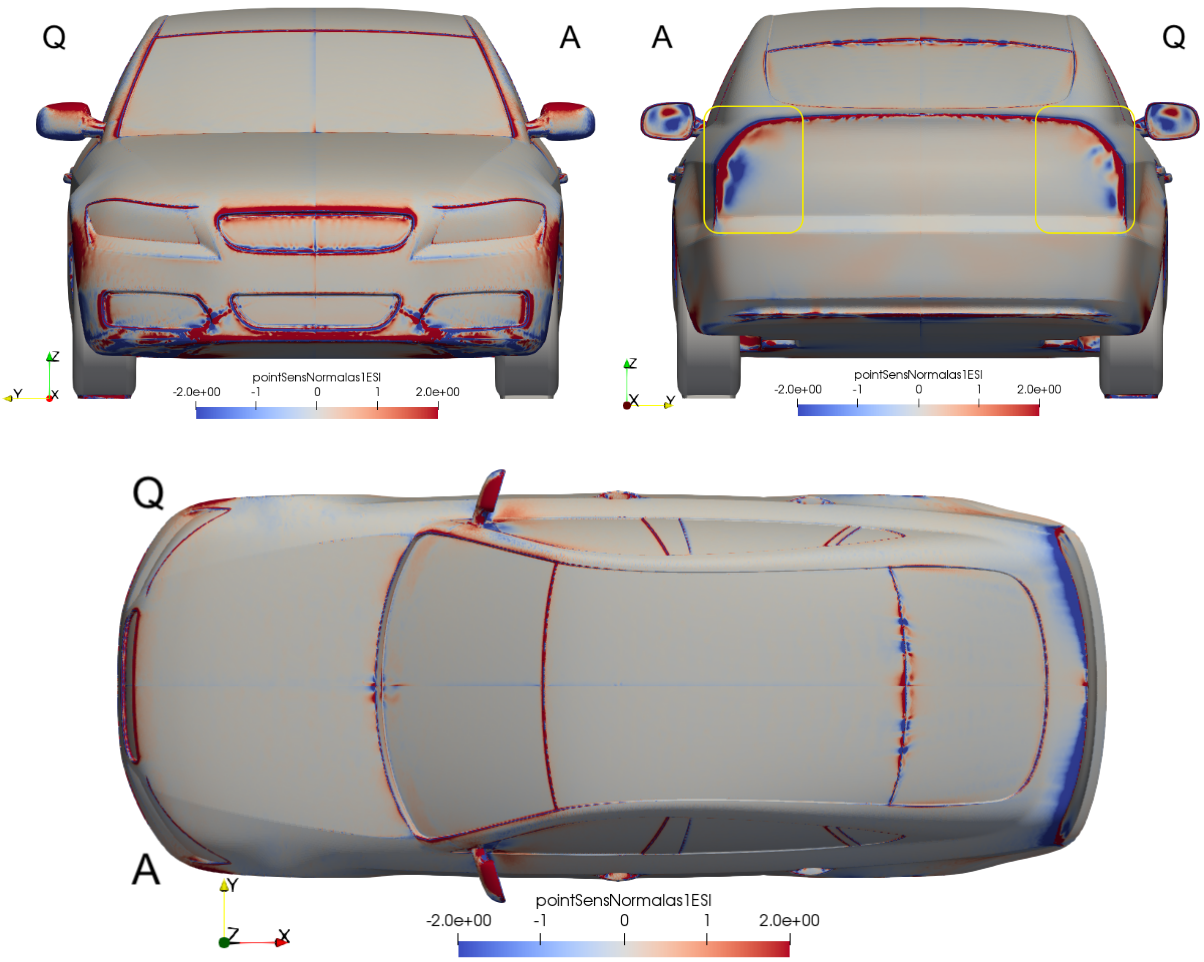

Figure A.15: Sensitivity Maps of front, rear, and top of (A) vs. (Q) with changes of Turbulence Properties and BCs.

The change in turbulence properties changes the sensitivity maps, although to a small degree. Simulation A has a eddy viscosity ratio of 1.5, while simulation Q has eddy viscosity ratio of 2. Changes in the sensitivity map on the trailing trunk/side edge of the car are observed both on Figure A.15 and Figure A.16. The higher eddy viscosity ratio in simulation Q is observed to cause changes in the sensitivity map, since higher turbulence stresses yield the modelling of higher Reynolds stresses. The effect of increasing Reynolds stresses in simulation Q can also be observed in Figure A.16 just behind the front wheel at the lower part of the side and in front of the rear wheel on the lower part of the rear fender.

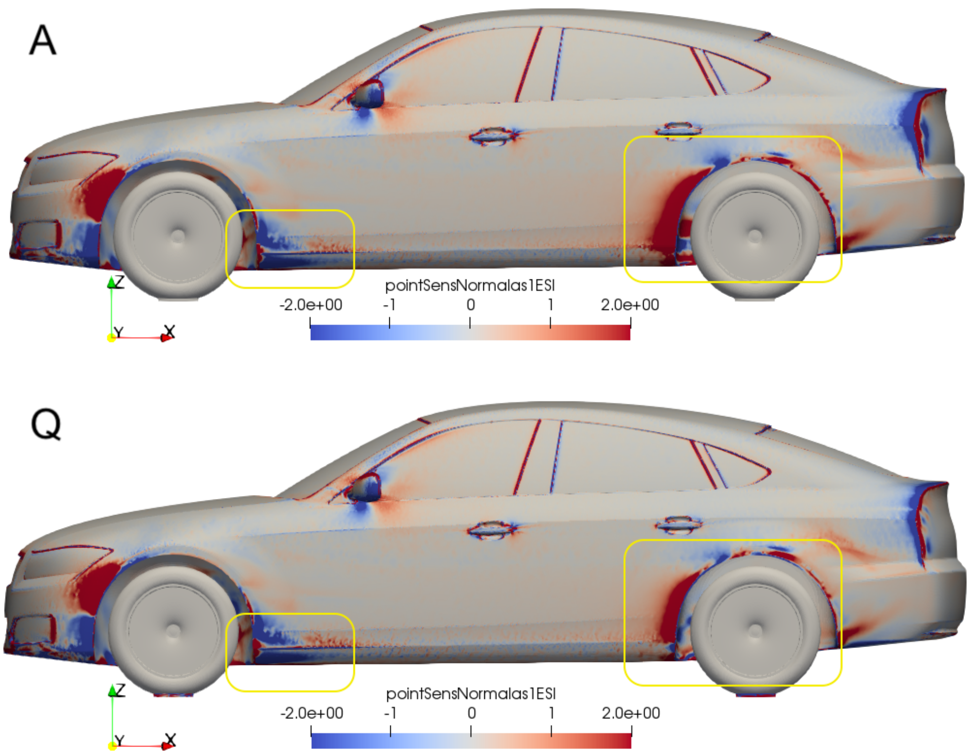

Figure A.16: Sensitivity Maps of sides of (A) vs. (Q) with changes of Turbulence Properties and BCs.

A.3.5 (E) Moving Ground and Rotating Wheels vs. (D) Stationary Road and Wheels

This study was carried out to see the effect of changing to ground simulation on the sensitivity maps. The model in question is the FSMW symmetry model with smooth rims. Each sensitivity map is virtually cut over the symmetry plane (XZ). This study utilises the same mesh and identical solver setup apart from the boundary conditions at the wheels and road. As expected, no greater changes are observed upstream of the wheels on the front of the FSMW in Figure A.17. However, small changes to the front bumper lip and in the region of the lower side grille are observed, primarily due to the moving ground. Next, significant changes to the sensitivity maps are observed on the rear end in Figure A.17. Here, the wake of the vehicle is changing shape depending on the ground simulation and rotational wheels or not. The more realistic case E, with moving ground simulation and rotating wheels, proves lesser areas of gradient on the upper region of the rear end. However, changes are recommended to the trailing edge of the bottom of the car. Skepticism is advised in such regions where the sensitivity maps suggest organic shape deformations.

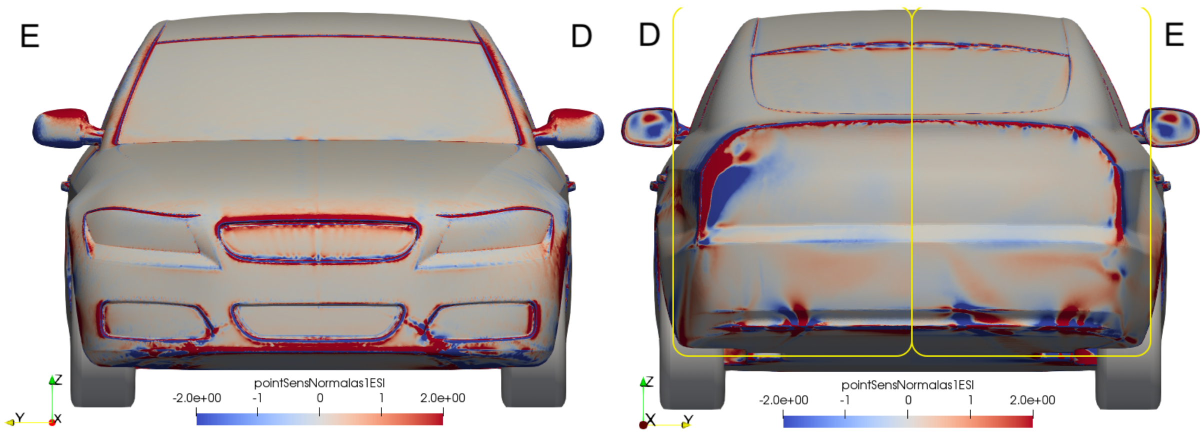

Figure A.17: Sensitivity maps of front and rear (E) Moving Ground and Rotating Wheels vs. (D) Stationary Road and Wheels.

The trailing trunk edge sensitivity map is significantly changed by the ground simulation and rotating wheels in Figure A.18. This is due to a change of pressure distribution in the wake of the vehicle. The asymmetric wake flow properties are not achieved with the symmetry FSMW flow simulations. It is to expect that a full car simulation may have a great impact on sensitivity maps as well. Next, the sensitivity map of simulation D is counter intuitive, as wall-normal inwards displacement is recommended on the lower part of the rear window, despite the possible flow separation and strong adverse pressure gradient in this region.
The sensitivity maps are affected significantly by the rotating wall boundary conditions on the wheels. Simulation E with rotating wheel boundary conditions performs less well due to the expected pressure fluctuation in the vicinity of the rear wheel. Thus, a wavy sensitivity map pattern is observed on the rear wheel arch. However, the sensitivity maps still recommend wall normal outwards displacement covering upstream of the rear wheel, as simulation D. Boat-tailing of the trailing side of the car is also recommended, although in different regions of the trailing region of the side. Again, simulation E shows a wavy

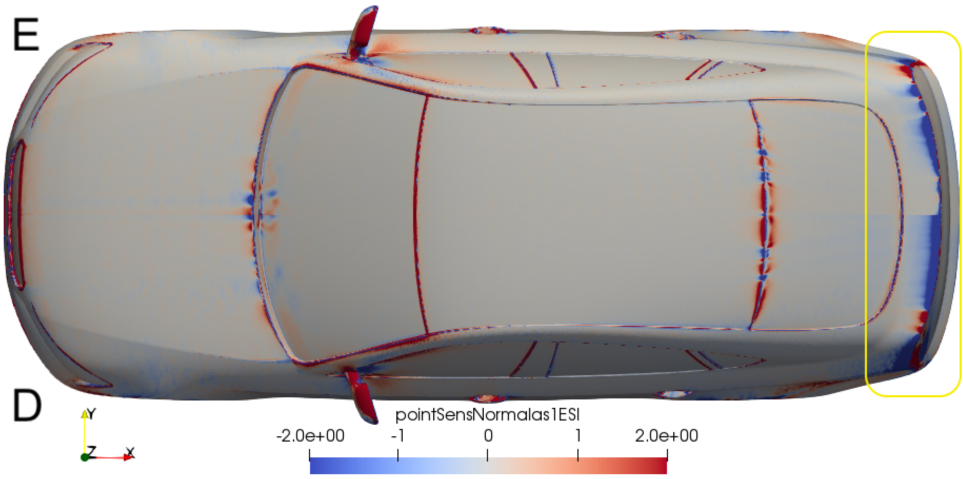

Figure A.18: Sensitivity maps of top (E) Moving Ground and Rotating Wheels vs. (D) Stationary Road and Wheels.

pattern in the sensitivity map, presumably due to pressure fluctuations induced by the rotating wheels. A wavy pattern upstream of the front wheel is observed too. Simulation D recommends a general wall-normal outwards displacement to cover the front wheel, while Simulation E recommends both wall-normal inwards and outwards displacement. Thus, an engineer is needed to justify where to morph. Finally, pressure fluctuations are observed on the mid-section and low section of the doors. This pattern was observed on the skin-friction plots, in regions of low/high skin friction.

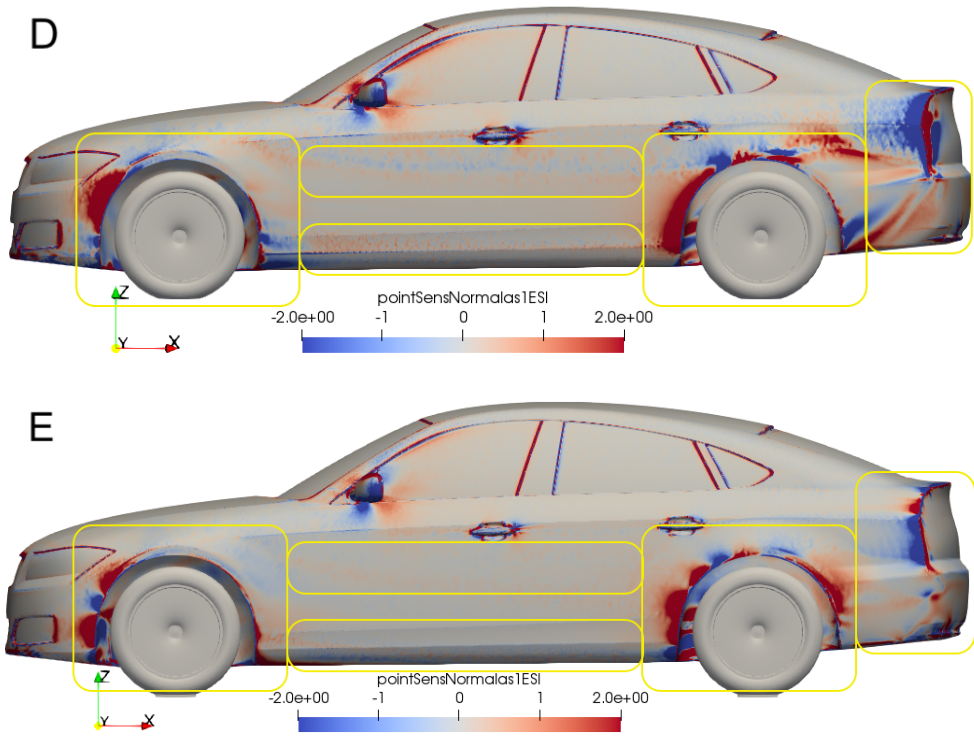

Figure A.19: Sensitivity maps of sides (E) Moving Ground and Rotating Wheels vs. (D) Stationary Road and Wheels.

A.3.6 (E) Smooth vs. (F) Open Rims

This study was carried out to see the effect of changing rim type on the sensitivity maps. The model in question is the FSMW symmetry model, with smooth rims (E) and open rims (F). Each sensitivity map is virtually cut over the symmetry plane (XZ). This study utilises the same mesh and identical solver setup apart from the boundary conditions and mesh of wheels. More specifically, the smooth rims have a rotating wall boundary condition applied to both rims and tires. The open rims have rotating wall boundary conditions applied to the tires and rims, while the pockets of fluid in between spokes are modelled by a moving reference frame. Both simulations have the ground simulation active. Apparent changes on the sensitivity map are observed upstream of the wheels at Figure A.20a. However, clear differences in the rear wheel wake and interaction with the base of the wake create a much different sensitivity map on the rear end in Figure A.20b. It should be mentioned that different fine wheel refinements may also be the cause of the differences observed in the sensitivity maps. Namely, simulation E has a cell count of 20,791,135 cells. Simulation F has a cell count of 38,529,671 cells.

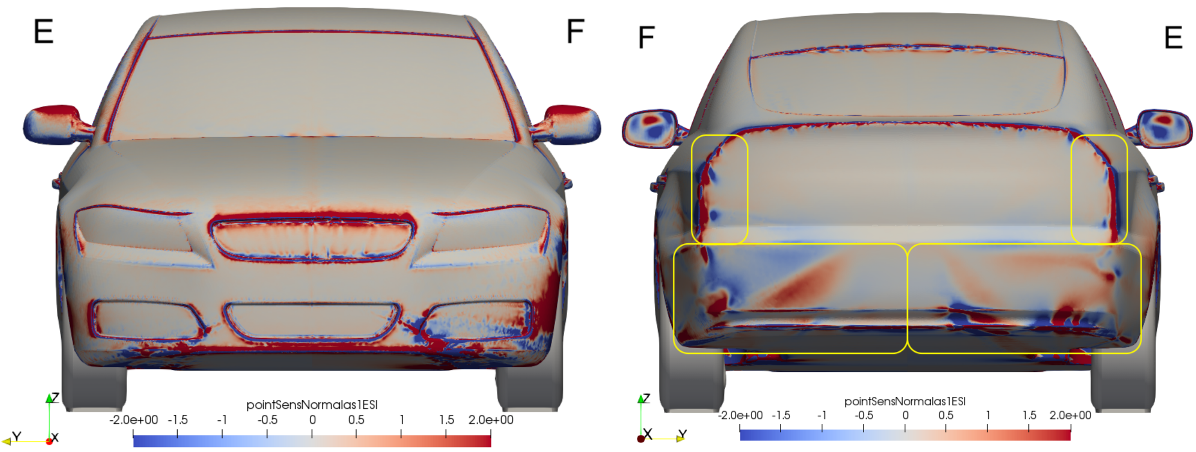

(a) (b)

Figure A.20: Sensitivity maps of front and rear (E) Smooth vs. (F) Open Rims.

Next, small changes in the creases at the cowl area and at the roof trailing edge are observed in Figure A.21. Conversely, these changes are minor and might be due to pressure fluctuations or slightly different generation of mesh. More noticeable difference is observed at the trailing trunk edge. Interestingly, the change of rear wheel wakes by the rims induces pressure changes on the base of the vehicle, presumably when merging with the vehicle wake.

Pressure changes are observed at the wheel arches. The magnitude of the gradient is larger for the open rim setup (F), compared to (E). I.e., the wall normal outwards displacement of the front bumper in the region ahead of the front wheel indicates the need for better covering the front wheel. The wavy patterns on the wheel arches are continuously caused by the rotating wall boundary conditions. The trailing side edge of the FSMW is also dominated by a wavy pattern, although wall-normal inwards displacement is expected. Additionally, the green-house and mirror is left unchanged by the rim change.

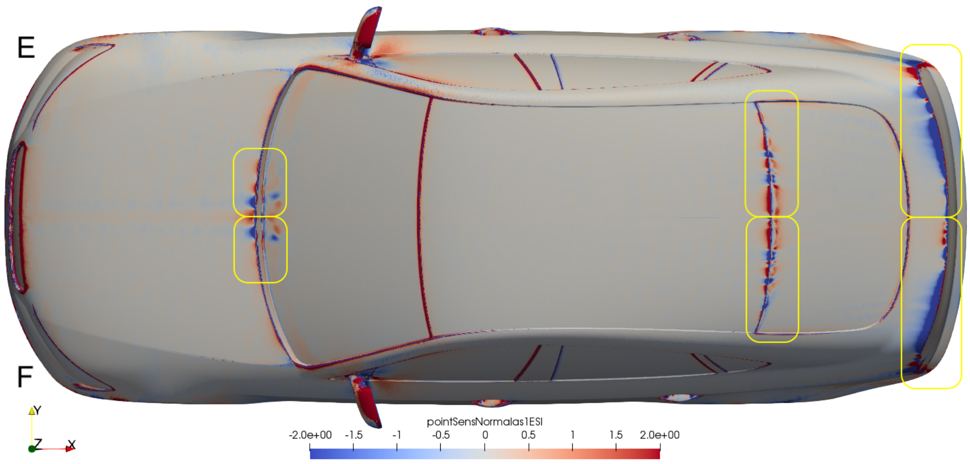

Figure A.21: Sensitivity maps of top (E) Smooth vs. (F) Open Rims.

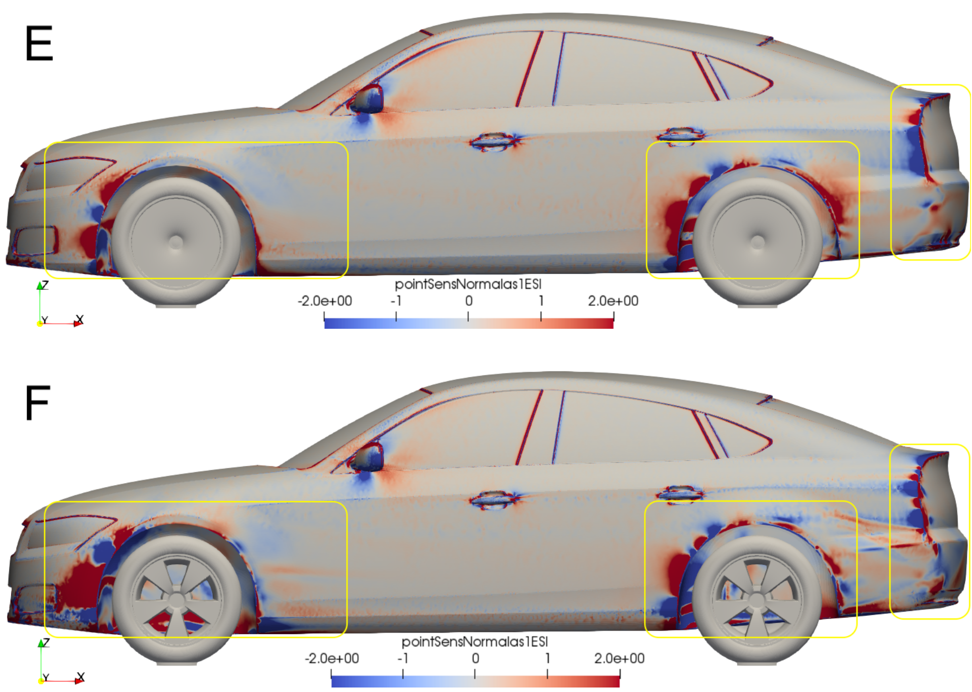

(a)

Figure A.22: Sensitivity maps of sides (E) Smooth vs. (F) Open Rims

A.3.7 (S) Unsteady vs. (F) Steady

This study was carried out to see the effect of changing from steady to unsteady simulation simulation on the sensitivity maps. The model in question is the FSMW symmetry model with open rims. Each sensitivity map is virtually cut over the symmetry plane (XZ). This study utilises the same mesh and identical solver setup apart from running simpleFoam (F) and pisoFoam (S). Both the steady and unsteady simulation has been averaged over the last 5,000 iterations of the primal flow. This equals 2 seconds of flow averaging for the unsteady flow. The unsteady simulation may have the advantage of being more accurate, as the unsteady primal simulation is run based on a fully converged steady state primal solution. The change to an unsteady solver has not induced major changes to the sensitivity maps in front of the car. The irregular pattern of the surface gradient of the lower region of the front bumper, seen in simulation F Figure A.23, has been smeared out in Simulation S, and less wall normal displacement in this region is identified. Less strong sensitivity gradients are also observed at the rear in Figure A.23 for simulation S. The sensitivity pattern for the steady state wake, in sensitivity map F, is not represented in the sensitivity map S. However, a more irregular pattern on the rear sensitivity map S has been formed.

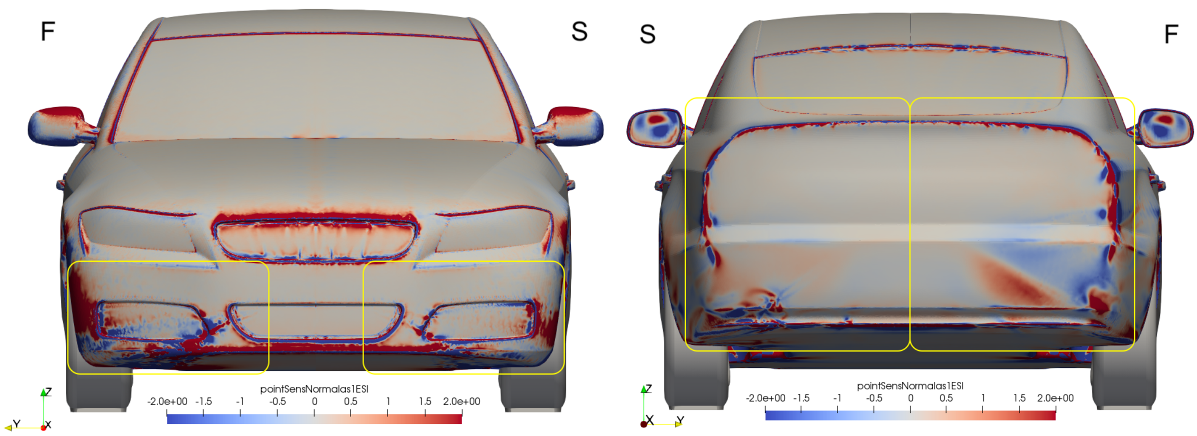

Figure A.23: Sensitivity maps of front and rear of (S) Unsteady vs. (F) Steady.

Changes between unsteady (S) and steady (F) solver on the top surfaces of the sensitivity maps are observed on the trailing trunk edge in Figure A.24. Here, more wall-normal outwards displacement at the center section of the trailing trunk edge, whilst less wall-normal inwards displacement is recommended at the outer/side region of the trailing trunk edge. The rear window leading edge/crease is also affected by the change of the solver, to a lesser degree.
Distinguisement of side sensitivity maps in the region aft of the rear wheel is observed in Figure A.25. The sensitivity map S shows the contours of unsteady flow properties, formed by an unsteady rear wheel wake. The region of the rear fender in front of the rear wheel is also affected, where the upper rear wheel arch is displaying opposite sign, when comparing F with S. This is also the case in the lower front wheel arch aft of the front wheel. Major changes to the sensitivity maps, when comparing F and S are observed on the front fender upstream of the front wheel in Figure A.25. Here, the unsteady simulation creates a wavy pattern, compared to the steady. To cover the front wheel by a wavy front, the fender

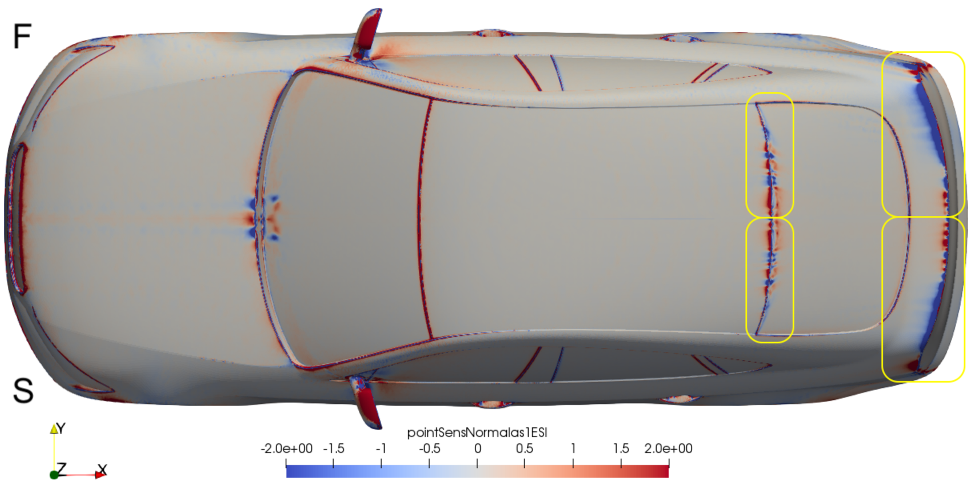

Figure A.24: Sensitivity maps top of (S) Unsteady vs. (F) Steady.

should be viewed with some skepticism. Note that the rotating wheel boundary condition was responsible for the wavy pattern. Thus, the unsteady simulation may enhance this wavy pattern further. Finally, the sensitivity map of the mirror housing is kept constant throughout all parametric studies. This proves that that the mirror is a component suited for automatic morphing, as it is relatively unaffected by the changes of various simulation parameters.

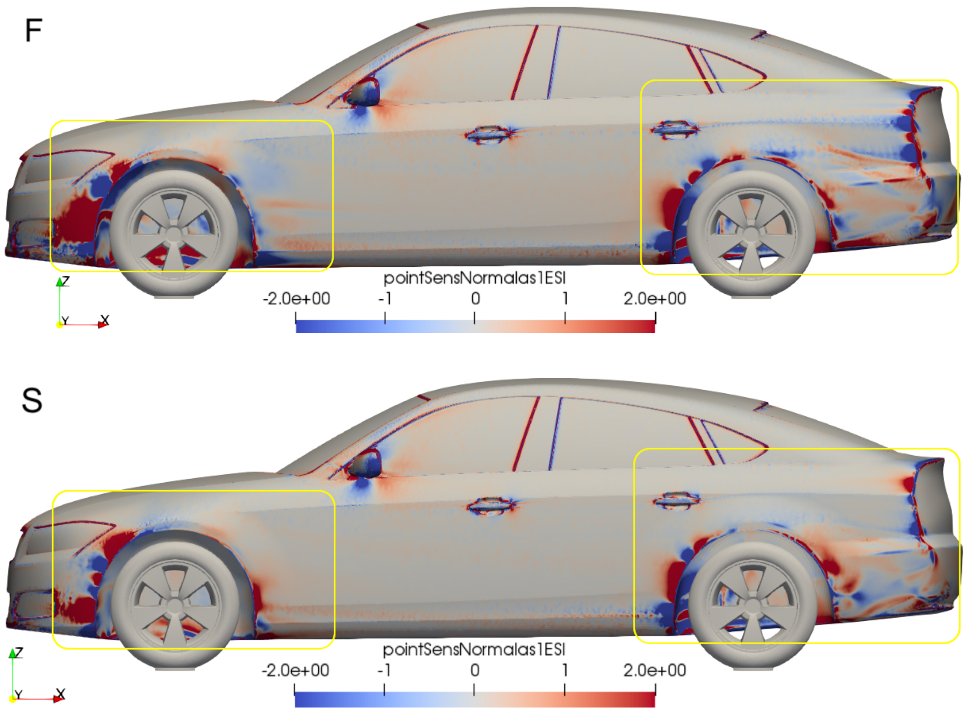

Figure A.25: Sensitivity maps sides of (S) Unsteady vs. (F) Steady.

A.4 Flow Properties and Sensitivity Maps

Figure A.26 shows the relation and interaction of the sensitivity map and vortical structures. Note that Figure A.26a and Figure A.26b is derived from the same primal and adjoint simulation (F). The colour bar of the sensitivity map has been scaled to show properties of the flow. The reader is advised to focus on the section of the doors, where irregular colours are observed on both the sensitivity map and Q surface plot. The patterns show some correlation with vortical structures which may cause disturbances/irregularities to the sensitivity map. This proves that the sensitivity maps are results of the flow.

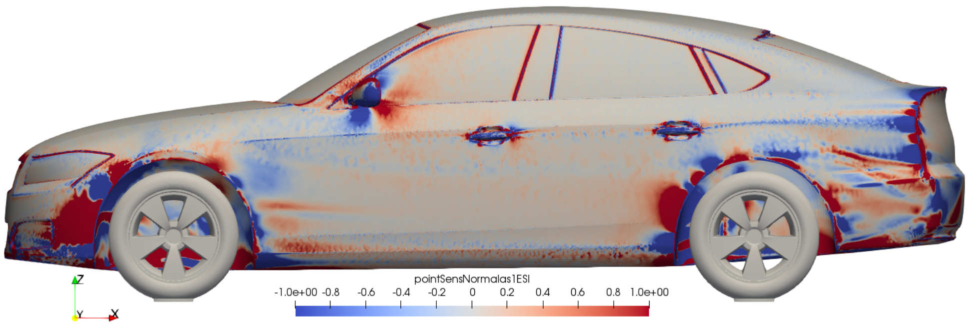

(a) Sensitivity map side of the FSMW.

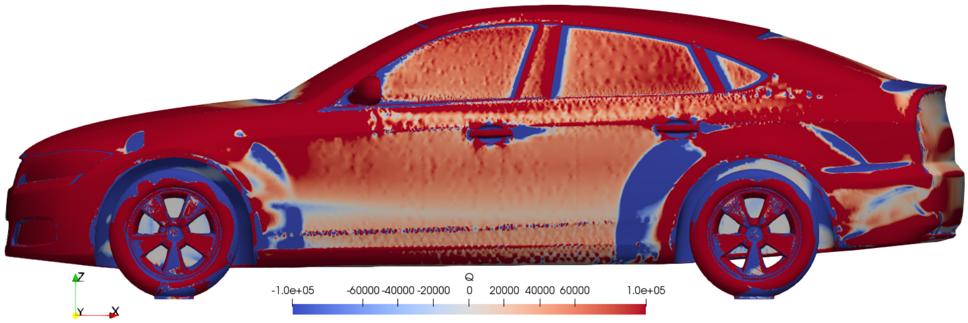

(b) Q-criteria map upper side of the FSMW.

Figure A.26: Sensitivity map and Q-criteria map of FSMW primal solution.

A.5 Recommendations for adjointOptimisationFoam

A general guide when using adjointOptimisationFoam was made by the wish Koenigsegg to sum up the recommendations stated in the report. The recommendations can be separated into three steps of the full process simulations when using adjointOptimisation-Foam, respectively pre-processing in ANSA, solving in OpenFOAM, and post-processing in ParaView.

Pre-processing

- Import highest possible CAD file into ANSA (preferably *.CATIA, *.STEP, or *.IGES in that order). Avoid *.STL files.
- If not familiar with ANSA, perform as many geometrical modifications in CATIA.
- Simplify geometry as much as possible to achieve the mesh quality necessary for running adjointOptimisationFoam without mesh related divergence.
- Mesh should include cells with:
  - Non-orthogonality below 60.
  - Internal skewness below 4.
  - Cell determinant above 1E-8.
  - Prism cells for resolving boundary layer.
  - Growth ratio within prism layer and in contact to surrounding hexa cells not exceeding 1.3.
  - Appropriate prism layer height.
  - Additional refinement added where prism layers are collapsed (refinement cell length in order of prism cell length and height).
  - Added refinement at leading and trailing edges to resolve abrupt pressure changes(refinement cell length in order of prism cell length and height).
  - Added refinement in highly concave/convex and self-intersecting geometry (refinement cell length in order of prism cell length and height).
- Most issues of adjoint instability arose in prims layer cells. Additional care should be taken.

Solver Settings

Recommended solver settings are given according to OpenFoam case folders of *0*, *constant*, and *system*. Start by reading thoroughly through the adjointOptimisationFoam User Manual [11]. Additional comments are given below.

0

- Initialise adjoint variables with zero internal values.
- Apply adjoint boundary conditions according to the adjointOptimisationFoam User Manual [11].

constant

- Apply dicts of MRFProperties as usual.

system

- Apply porous zones under fvOptions as usual.
- Initially, keep all schemes as second order. Adjoint variables may be decreased to first order.
- Use under-relaxation as User Manual [11] and decrease for stability.
- Use bounding of all adjoint variables. Check residuals after a preliminary run and remove unused bounding of adjoint variables.
- Use zeroing of patches and walls in the zeroATCPatchTypes for industrial cases.
- Use ATCModel as standard for a compromise of accuracy and stability for industrial cases.
- Avoid using extraConvection unless necessary, as this a numerical diffusive filter, compromising accuracy.
- Use nSmooth under the ATC model by using a Laplacian-like filter nSmooth times.
- Run sufficiently many primal iterations, with appropriate averaging to ensure a fully converged field is fed to the adjoint solver. Lack of iterations and averaging will appear in sensitivity derivatives.
- Monitor forces and force coefficient and residuals to observe convergence in the primal solution. Monitor residuals and more importantly inspect adjoint field variables to check convergence in the adjoint solution.
- Calculate sensitivity derivatives with respect to boundary points, not faces, when calculating sensitivity maps.
- Use the Enhanced Surface Integral formulation for calculating sensitivity maps.
- Use Field Integrals formulation when morphing with the volumetric B-Splines, denoted *volumetricBSplinesFI*.
- Use max displacement value in the order of a few millimeters to filter out noisy sensitivity derivatives, ensure consistent mesh quality, and stability.

Post-Processing

- Change colour-bar of sensitivity derivatives to understand the recommended wall-normal displacement of the sensitivity-map
- Refer highlighted areas on the sensitivity maps to the primal field distribution of pressure, wall shear stress, and velocity.
- Sensitivity derivatives may change in magnitude and sign when calculating w.r.t. control points contra boundary nodes.
- Keep healthy skepticism of sensitivity maps.

A.6 Terminology

Figure A.27 and Figure A.28 are created to aid the reader in identifying the regions that are analysed and discussed. The terminology figures should be useful for the chapter of Results and Discussion. Terminology may differ depending on British or American English. Terminology is used for this thesis only, as different terms might be used academically or in the industry. The *green house* is defining all windows. The *base* or *rear base* refers to the rear end with fully separated flow. A *count* should be defined as 0.001 times of a force coefficient. For instance, one drag count is $C_d = 0.001$.

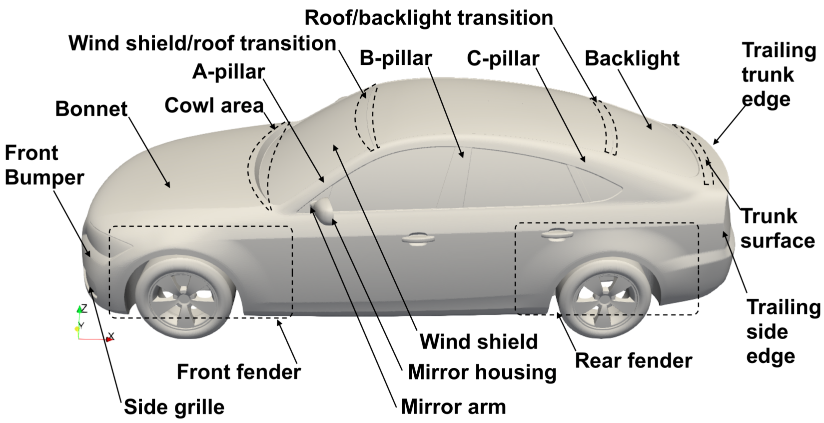

Figure A.27: Terminology used for the describing areas of the FSMW.

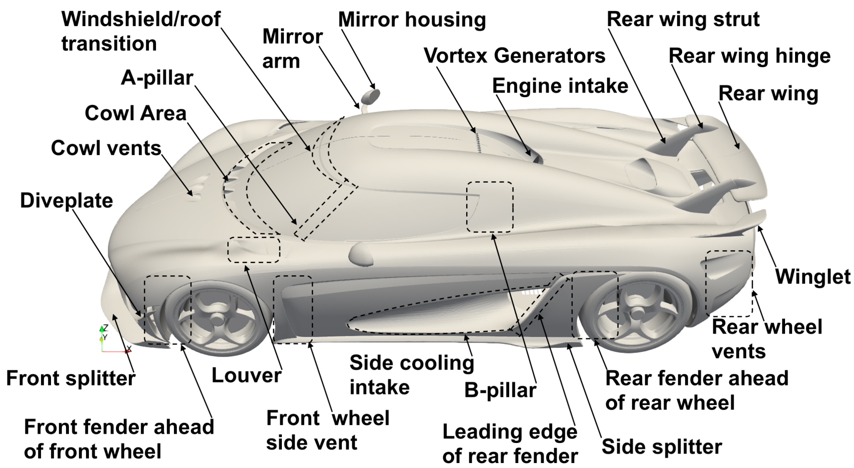

Figure A.28: Terminology used for describing areas of the Koenigsegg Regera.

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY