CHALMERS

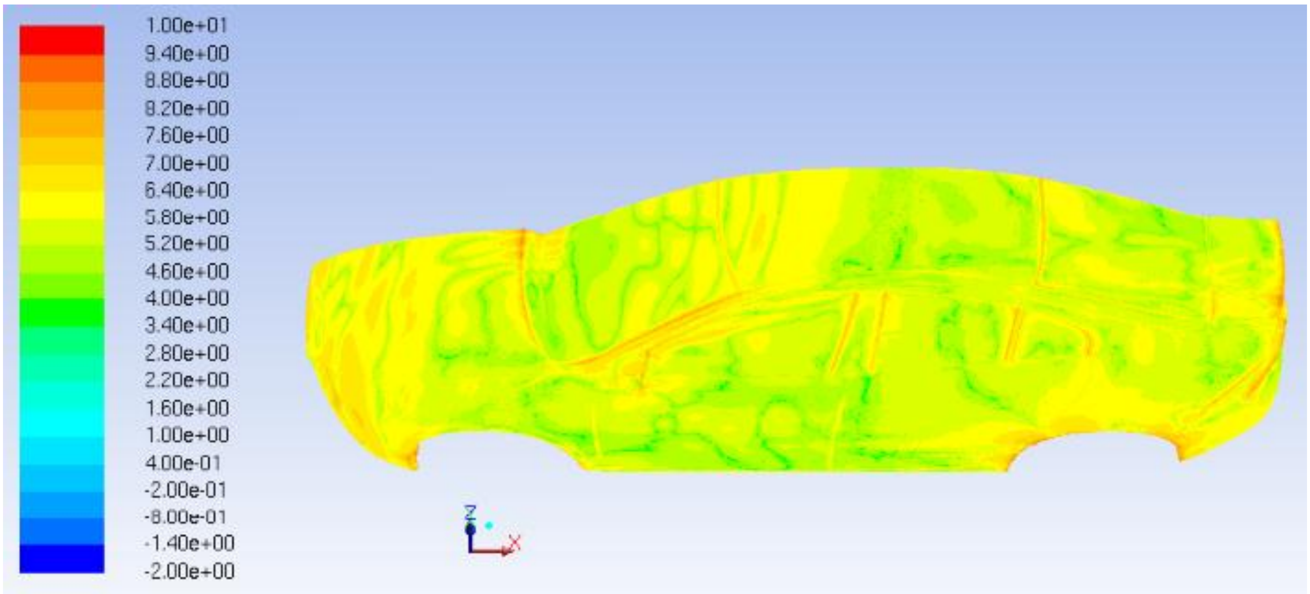

Aerodynamic Optimization of Ground Vehicles with the Use of Fluent's Adjoint Solver

Master's Thesis in the Masters Programme Automotive Engineering

JOHAN ZAYA

Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
Road Vehicle Aerodynamics
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2013
Master's thesis 2013:01

MASTER'S THESIS IN AUTOMOTIVE ENGINEERING

Aerodynamic Optimization of Ground Vehicles with the Use of Fluent's Adjoint Solver

JOHAN ZAYA

Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
Road Vehicle Aerodynamics
CHALMERS UNIVERSITY OF TECHNOLOGY

Göteborg, Sweden 2013

Aerodynamic Optimization of Ground Vehicles with the Use of Fluent's Adjoint Solver
JOHAN ZAYA

Master's Thesis 2013:01
ISSN 1652-8557
Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
Road Vehicle Aerodynamics
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: + 46 (0)31-772 1000

Cover: Plot of Shape Sensitivity Magnitude

Chalmers / Department of Applied Mechanics
Göteborg, Sweden 2013

Aerodynamic Optimization of Ground Vehicles with the Use of Fluent's Adjoint Solver
Master's Thesis in the *Automotive Engineering*
JOHAN ZAYA
Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
Road Vehicle Aerodynamics
Chalmers University of Technology

ABSTRACT

Like all other areas of the automotive development, environmental issues and fuel efficiency is one of the main driving forces for aerodynamic engineers where the aerodynamic drag force is the dominating resistance force at higher velocities. By improving the shape of a vehicle with respect to aerodynamic performance, the drag force can be reduced hence the fuel consumption can be reduced.

The use of Computational Fluid Dynamics is a widely used methodology to carry out simulations that describes the flow in and around a vehicle. With these simulations, aerodynamic engineers can gather information about the aerodynamic performance
Of the vehicle and make changes that can improve the performance. However, due to the many variables involved, this can become very computer demanding process and require a high number of design optimization cycles to finally reach valuable stage.

Recently, a new procedure used for optimization purposes, named the Adjoint Solver, has been the focus of researchers and engineers. The Fluent Adjoint solver compute derivatives of chosen engineering observation, such as drag, with respect to all inputs and provides a more direct guidance for optimal modifications to improve performance. The Adjoint Solver accomplishes to calculate the derivative data by running only one single computation, very similar to basic CFD computations, and by that providing valuable engineering insight that can both improve and reduce the number for design optimization cycles.

The main goal of this master thesis is to state if the Adjoint Solver is ready to be incorporated into Volvo Cars Aerodynamic development process.

This project has been carried out as an Master Thesis together with Volvo Cars and Chalmers University of Technology, in a close relationship with Ansys which are the developers of Fluent. Fluent's Adjoint Solver has been tested on its computation abilities, robustness, computer requirements and functionality. The tests are done on four different vehicle models provided by Volvo Cars and simulations are computed in wide variation of different case setups.

Based on the results from the simualtions, the conclusion is that Fluent Adjoint Solver is at the moment not at a stage where it is ready to be incorporated as a part of the development process. It is proven that the one can gain valuable engineering insight that surely can improve the development process, however, for external aerodynamics, the Adjoint Solver is not yet ready. Ansys will now continue to develop and improve the Adjoint Solver, with the issues discovered in this project in mind.

Key words: CFD, Aerodynamics, Optimization, Fluent, Adjoint Solver, Volvo Cars

Preface

This master thesis has been performed at the department of Applied Mechanics, Chalmers University of Technology, Sweden in collaboration with Volvo Cars during the spring of 2012 and presented in the autumn of 2012. Johan Zaya has been the student performing the master thesis with Simone Sebben as supervisor at Volvo Cars and Lennart Löfdahl as examiner at Chalmers.

Acknowledgments

I want to thank my supervisor at Volvo Cars, Simone Sebben, first of all for the opportunity of performing this exciting project and for providing all the support and help I needed throughout the whole process of this master thesis. Without her this project would not have been carried out. I want to thank Torbjörn Virdung from Ansys for all his help and valuable discussions regarding theAdjoint Solver. Without Torbjörn and Ansys, the results from this master thesis would not have been possible to be realized. I also want to thank my Examiner, Lennart Löfdahl for always being available whenever I had any questions regarding the project.

I want to thank the department of Applied Mechanics, division of Vehicle Engineering and Autonomous Systems for providing me with all of the necessary needs to carry out this project. I want to thank PhD student Alexey Vdovin for always being the extra helping hand during this project and also the rest of my colleagues at the Road Vehicle Aerodynamics Group.

Last I want to thank my family and friends for always giving much support and encouragement since day one of my studies. I'm glad for the opportunity, my passion is automotive and I'm looking forward to an exciting future.

Göteborg January 2013

Johan Zaya

Contents

1. Introduction

1.1. Background

One of the current challenges that CFD engineers face today in the automotive industry is the highly demanding work of design optimization and the need of incorporating a large number of computations into the engineering design cycle. The design optimization work has in the past decade been mainly through Design On Experiments (DOE) methods however, for aerodynamics, this is heavily computer demanding due to the large number of aerodynamic design variables involved. To move forward and further effectives the optimization work, the focus is now on Adjoint solver codes which is a new methodology used for optimization purposes in the automotive industry. The Adjoint Solver method computes the gradients, as in direction and magnitude, directly by solving a set of adjoint equations which makes it cost independent of the number of design variables.

1.2. Purpose

The purpose of this master thesis is to test the Adjoint solver code provided by Ansys in Fluent 14 referred as Fluent Adjoint Solver. The tests are done on a number of different car models and modifications to optimize the performance are carried out based on the adjoint solver results. The adjoint procedure will be tested for its robustness, computer requirements and easiness of use. Depending on the findings of this thesis, Fluent Adjoint Solver should be shortly incorporated into Volvo's aerodynamic engineering process or used as further development guide for Ansys in order to reach a stage where the procedure is fully ready to implement.

1.3. Goal

The goal of this master thesis is to finalize a conclusion weather if the Adjoint solver actually meets the expected benefits, where the results both improve the aerodynamic performance and reduce the development lead time by decreasing the amount of optimization design cycles needed. And with this conclusion state if Fluent Adjoint Solver is at a stage where it can be incorporated into Volvo's aerodynamic CFD engineering process.

1.4. Procedure

The methodology for this thesis will be based on the current Volvo AEDCAE01 Aerodynamic CFD procedure. The program ANSA is used to clean up the CAD model and prepare it if for meshing. Harpoon is the program mainly used for building up volume meshes in the Volvo procedure, however in this thesis both Harpoon and Tgrid will be used. Fluent will be used to perform the calculations and within in Fluent is the adjoint solver code, Fluent Adjoint Solver. The results are imported to Paraview for further analysis and visualization.

2. Theory

In the following chapter, theory covering the numeric simulations carried out using Computational Fluid Dynamics are briefly stated and explained. The theory behind the Adjoint solver is also briefly explained. For a more in depth explanation of CFD theory, the reader is guided to ref. [2][4] and for better understanding of Fluents Adjoint Solver ref. [5] is recommended.

2.1. Computation Fluid Dynamics (CFD)

Computation Fluid Dynamics, referred as CFD, is the use of computer simulations to compute and analyze the dynamics of a fluid and the method is used in various areas such as ventilation, internal combustion engines and vehicle aerodynamics as in this project. CFD simulations are carried out by dividing the physical domain into small finite volume elements and numerically solved by the governing equations that describe the behavior of the flow. The governing equations are derived from the laws of conservation, which become very complex, almost impossible to solve analytically and therefore requires numerical simulations. There are many advantages with CFD, such as reduction of cost and lead time when developing new design and is a widely used method in the industry, however due to the complexity of the physics, there still are difficulties. Large and trustworthy results needs a high amount of computer capacity and are only as good as the operator and the physics embedded.

Fluid dynamics is the study of fluids in motion, e.g. airflow around a vehicle, and the physics can be described by the conservation laws of mass, momentum and energy [1].

Conservation of mass is expressed by the continuity equation, which states that the amount of mass flow that enters a control volume must be equal to the amount leaving it.

Continuity equation:

$$\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}+\frac{\partial w}{\partial z}=0 \qquad (1)$$

Conservation of momentum (Newton's Second Law of Motion) is expressed by the Navier Stokes equations and is used to obtain a relation between pressure, momentum and viscous forces for a Newtonian fluid.

Conservation of momentum equation:

$$\rho g_x-\frac{\partial p}{\partial x}+\mu\left(\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2}+\frac{\partial^2}{\partial z^2}\right)=\rho\frac{\partial u}{\partial t}$$

$$\rho g_y-\frac{\partial p}{\partial y}+\mu\left(\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2}+\frac{\partial^2}{\partial z^2}\right)=\rho\frac{\partial v}{\partial t} \qquad (2)$$

$$\rho g_z-\frac{\partial p}{\partial z}+\mu\left(\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2}+\frac{\partial^2}{\partial z^2}\right)=\rho\frac{\partial w}{\partial t}$$

Conservation of energy (First Law of Thermodynamics) is expressed by the energy equation stating that the total amount if energy within a system stays constant.

$$\textbf{\textit{Conservation of momentum:}}\quad \rho c_p \frac{dT}{dt} = k\,\nabla^2 T + \Phi \tag{3}$$

As mentioned above, the governing equations become very complex. In order to numerically solve them, the differential equations are discretized into large system of algebraic equations in order to solve them. The general approach is to assume incompressible and isothermal flow where the Mach number under 0.3 and a constant temperature since vehicles travel at a relatively low speed. The governing equations are therefore simplified and the energy equation can be neglected.

Further more, since stochastic, three-dimensional and time dependent, fluid flows are almost always turbulent and will experience fluctuations and therefore Reynolds decomposition method is used. The velocity and pressure are split into two parts, an average and a fluctuating part stated as

$$\bar{u} = \frac{1}{T}\int_0^T u\,dt \tag{4}$$

$$\begin{aligned} p &= \bar{p} + p' \\ u &= \bar{u} + u' \\ v &= \bar{v} + v' \\ w &= \bar{w} + w' \end{aligned} \tag{5}$$

This is called Reynolds-Averaged Navier-Stokes (RANS) equations where the additional terms are included in the continuity and momentum equations and time averaged.

$$\frac{\partial \bar{u}}{\partial x} + \frac{\partial \bar{v}}{\partial y} + \frac{\partial \bar{w}}{\partial z} = 0 \tag{6}$$

$$\rho g_x - \frac{\partial \bar{p}}{\partial x} + \frac{\partial}{\partial x}\left(\mu \frac{\partial \bar{u}}{\partial x} - p\overline{u'^2}\right) + \frac{\partial}{\partial y}\left(\mu \frac{\partial \bar{u}}{\partial y} - p\overline{u'v'}\right) + \frac{\partial}{\partial z}\left(\mu \frac{\partial \bar{u}}{\partial z} - p\overline{u'w'}\right) = \rho \frac{d\bar{u}}{dt} \tag{7}$$

$$\left\{ \begin{array}{c} \rho = density \\ t = time \\ p = pressure \\ \mu = static\ viscosity \\ g = gravity\ constant \\ c_p = specific\ heat\ capacity \\ T = temperature \\ k = coefficient\ of\ thermal\ conductivity \\ \Phi = viscous\ dissipatio \end{array} \right\}$$

2.2. Adjoint Solver Theory

The Adjoint solver is a CFD optimization tool developed to further help engineers in their process of developing and improving the performance of a fluid system. This by providing detailed sensitivity data about the system being simulated hence locating the main problem regions.

The Adjoint Solver method is still relatively new, however the use of an adjoint solver is getting highly popular and the availability is increasing. There are several CFD code providers that are implementing Adjoint Solvers as extend to the conventional fluid flow solvers, though it must be said that the method and the codes are still under development and not always ready to be fully implemented. Car manufacturers are testing adjoint solvers to see if the process is at a stage where it can be implemented within their own development procedure and depending on what type of code that is being used, the results vary. The code being used for this project is Fluent Adjoint Solver provided by Ansys within Fluent 14.0.

The key factor of an adjoint solver is the gathering of sensitivity data, this data provides specific information about the fluid system which in conventional solver is very difficult. In standard flow calculations, the user supplies a number of attributes together with the system mesh such as material properties, physics models and boundary conditions. When the simulation is successfully converged, detailed data about the flow state for the system is provided. This data is then analyzed, the post processing shows the performance of the system and by that one can state what changes should be carried out. However, the post processing analyses does not state what effect the change will have on the performance. In order to state the effect due to the change, the user must carry out a second simulation which then again is post processed and so on. This would be a so-called Design Of Experiments procedure, where the changes become an iterative process of experiments based on the knowledge gained from the post processing. This can sometimes be an ineffective way of working.

The effect of modification depends in how sensitive the flow is to the particular parameter that is being modified. This is a derivative data that can become very large with respect to the input data required and the output data that is produced. Determining this derivative data is the remarkable feature of the Adjoint Solver. The Adjoint solver accomplishes to calculate the derivatives of a chosen observation of a simulated system by running only a one single computation. With one single computation, the Adjoint Solver calculates the derivatives with respect to all of the inputs for the system in order to measure the performance, such as lift or drag, or total pressure drop through a system. This could be the derivative of drag with respect to the shape of a vehicle, or the total pressure drop with respect to the shape of the flow path.

Understanding these sensitivities becomes a very valuable set of data where it can both improve the efficiency of the development process and the performance of the system. The adjoint data identifies areas of high and low sensitivities and by that, the engineer can focus on the regions with the biggest impact on the performance and be guided to do the right changes to the geometric shape of the system. Besides stating the sensitive regions and parameters, the Adjoint Solver quantifies the effect of the modification carried out on the shape and estimate a value of how it will improve the performance.

The mathematical background of adjoint computations is of course very heavy and can be briefly looked in to by reading the Fluent Adjoint Solver User Guide [5] However with the formula in figure 1, the key idea can be explained. The formula simply states the value of a chosen observation with respect to all sensitivity coefficients and node displacement. This can be compared to manually re-calculating flow and drag on a surface where each point has been moved. N points will require N flow calculations which really isn't a realistic to perform in real case where are the number of points is vastly more than the amount of calculations an engineer has the ability to perform. With this simple way of describing the procedure, one can understand the powerful ability of Adjoint computations.

Shape sensitivity: Sensitivity of the observed value with respect to (boundary) grid node locations

$$\delta(Drag) = \sum_{mesh} \underline{w}^{n} . \delta \underline{x}^{n}$$

Node displacement

Shape sensitivity coefficients: Vector field defined on mesh nodes

Figure 1: Shape Sensitivity formula, Ref.[6]

With Fluent Adjoint Solver there's also one other very valuable feature which is the ability to morph the geometry based on the sensitivity data gathered from the adjoint computations. By Ansys, this is described as an extra step to the adjoint computations, which very conveniently is to the output and post-processing of derivative sensitivities. This means that modifications on the geometry are carried out directly within Fluent Adjoint solver with respect to the sensitivity data and the morphing set up controls. The procedure is a Bernstein polynomial-based mesh morph, which can smoothly modify all mesh types where the user chooses a rectangular control volume box to that defines the region wished to be morphed.

3. Methodology

In this chapter, the methodology carried out for this master thesis will be explained. The methodology is divided into two steps, step one is conventional flow computations and step two is the Adjoint Solver computations. The goal is to compare methods, then state the benefits and usability of the Adjoint Solver. First, the workflow of the adjoint procedure is described, followed by a description of the vehicle models used and the different testing configurations.

3.1. Conventional flow calculations - Volvo Procedure

Conventional flow computations are pretty standard in the industry where the procedure of different manufacturers doesn't vary a lot. The procedure is pretty much the same when it comes to work flow, a CAD model of the vehicle is imported to a pre-processer where the surfaces and geometry are improved if needed. The CAD model is then meshed, both surface and volume mesh for external aerodynamics and furthermore also an interior mesh for cooling and likewise. The mesh is very important, a fine mesh with small cell sizes will result in a better representation of the flow, however finer meshes need more computer capacity to compute so the mesh quality is many times dependent on the availability of computer capacity. The method used for the basic fluid flow calculations in this project is according to the Volvo Cars AEDCAE01 Aerodynamic CFD procedure. The CAD model is imported into ANSA, meshed in Harpoon and the code used for flow calculations is Fluent. Since the aim of this project is to state if the Fluent Adjoint Solver is ready to implement in to Volvo's procedure, the method of work has to fit Volvo's aerodynamics way of work, and therefore the methodology is according to the Volvo procedure. However the Adjoint Solver may cause some limitations and by that other approaches must be tested. This will be described further in the configurations section.

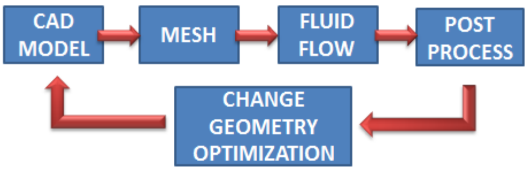

Figure 2: DOE optimization cycle

3.2. Fluent Adjoint Solver

At the moment there are few Adjoint solver codes provided in industry, however they are all still under development and not always fully ready for implementation. Mostly it is pilot versions and previews used for testing and further development. The Adjoint solver provided by Ansys in the Fluent code is one of the adjoint codes that has come the furthest in its development. The focus in this project is the Fluent Adjoint Solver code.

The Adjoint Solver in Fluent is provided as an extend to the conventional solver and is a simple extra step with similar workflow as for standard flow calculations.
Step 1; is to compute a fluid flow simulation of the system, which is done no differently than standard computations. The fluid flow simulation is done accordingly to the Volvo Aerodynamic CFD procedure, however only to an extend where the adjoint solver actually works due to its current limitations.
Step 2; is to load the adjoint solver addon module and choose to observe drag, lift or pressure. For this project, the idea is to use the Adjoint Solver to optimize the shape with respect to drag.
Step 3; is to set up Adjoint Solver controls and convergence criteria. This is a very vital part of the testing procedure, Ansys has already stated that there are convergence difficulties with the adjoint solver and therefore much focus is put on controls. This will also be further described in the configurations section.
Step 4; is just like standard simulations, initialize the adjoint solution and iterate to convergence.
Step 5; is to post process the adjoint solution, check if it converged and the go one to the essential part of the adjoint solver, extracting the sensitivity data. The post-processing of sensitivities is done directly in Fluent by using standard tools such as contours, vectors, xy-plots etc. Via the contours tool, there a number of different fields which the sensitivities can be analyzed, e.g. Magnitude of Sensitivity to Body Forces, Sensitivity to Body Force XYZ-Component and Sensitivity to Flow Blockage. As stated before, the idea is to optimize the shape with respect to drag and with that, the given drag sensitivity data is analyzed by the log10(Shape Sensitivity Magnitude) field. This field gives a log10 plot of the sensitivity magnitude with respect to shape, which allows the importance of the surfaces in a domain to be ranked more easily on how they affect the drag. This provides valuable engineering insight where the most critical regions are displayed and identified and by that one is given guidance to shape modifications that will improve the result of the chosen observable. The guidance to shape modifications cannot only be given by the sensitivity magnitude plots, there are also plots of vectors or node normal sensitivity plots displaying both magnitude and direction on how a modifications should be carried out. As stated before, the key factor of the Adjoint Solver is the gathering of sensitivities which becomes a tool that can improve the optimization design cycle and further effectives' DOE by decreasing number of design iterations where one can directly find what areas to focus on. One further step, that even more can improve the work efficiency, is the ability to carry out shape modifications on the system by directly morphing in Fluent. Within Fluent there is a morphing tool that is

integrated with the Adjoint Solver.

Step 6; is to use the morphing tool, where modifications are carried out by setting up morphing boxes around the whole model or just around specific regions depending on what is wished to be modified. The actual amount of change within the morphing boxes is defined by setting up a number of controls points in each coordinate of the box and stating a value of allowed modification percentage. Prior to the shape modification within the morphing boxes, an estimated value of the affect on the chosen observable is given with respect to the sensitivity data and morphing set up. Once the model has been modified, the next step is a second standard fluid flow to evaluate the actual change in performance. In appendix 1, Fluent Adjoint Solver procedure demonstrated in steps together with the user interface.

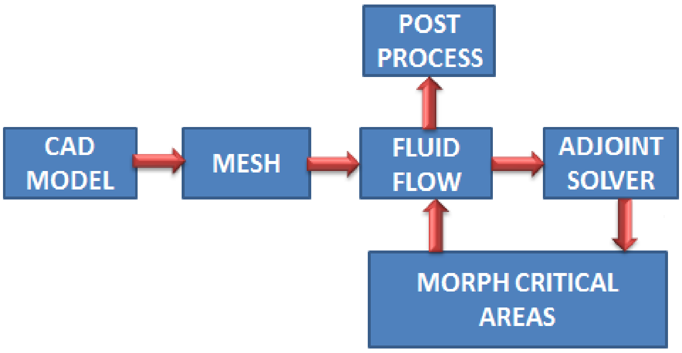

Figure 3: Adjoint optimization cycle

3.3. Vehicles for simulation

The models being used for this project is provided by Volvo Cars. Initially, the goal was to test the adjoint procedure on a variation of Volvo models such as a sedan, a estate and a SUV together with simplified concept models. The idea for this is to make sure that the Adjoint Solver is tested on a wide variation of vehicle shapes, both with detailed and simplified underbody and wheels and state how the results would correlate to each other with respect to different shapes, mesh sizes, instabilities and convergence difficulties. Further on, the goal was also to test the adjoint procedure on specific areas that are more difficult to analyze and time consuming with conventional flow solver, areas such as the underbody and around the exhaust system which can be a big factor on drag. A similar study done by Volkswagen [9], even though it already had a quite optimized aerodynamic shape, showed that the most potential for improvement was found on the underbody. However, it was quite early in this project recognized that testing all these models would not be possible within the time limit. This is due to the many difficulties encountered with running Adjoint Solver computations. This will be explained more in detail in the results chapter.

Due to the limitations, the vehicle models used for Adjoint solver simulations is are four different models, where three of them are very simplified concept models from Volvo referred as VRAK (Volvo Research for Aerodynamic Knowledge) models and are only used for testing purposes such as this project. The fourth is a more detailed Volvo S60 sedan model. Even though the models used for testing is limited to only four, it still is a good variation of shapes and model quality, which can give a good conclusion about the current Fluent Adjoint Solver usability.

Vehicle model 1, seen in figure 4, is a VRAK estate model with a very simplified surface without any major details such as front grill or headlights, no side mirrors and closed wheels. Since it is already stated that the Adjoint Solver has convergence difficulties, this model is a good base for testing. Estate models tend to have more vortices' in the rear due to the poor aerodynamic shape, which causes instabilities in the flow field. Testing with these type of instabilities is a vital part of this project. At the same time, this model is compared to the other VRAK models and the more detailed Volvo S60 model.

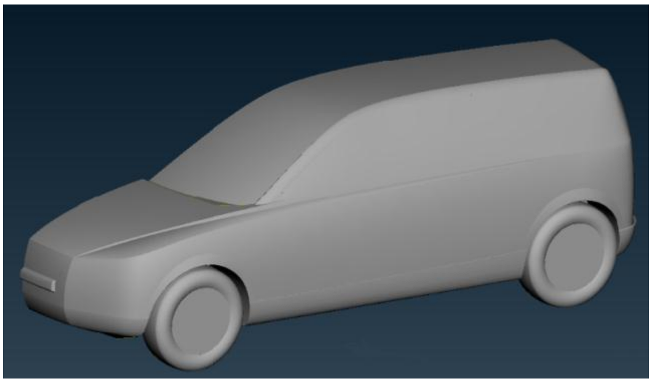

Figure 4: Model 1 VRAK_BASE

Vehicle model 2 and 3 are almost identical hatchbacks where the difference is only a small rear wing as seen in figure 5 & 6. They are just as simplified as the vehicle model 1, therefore comparing the three VRAK is a good way to evaluate convergence problems with respect to each other.

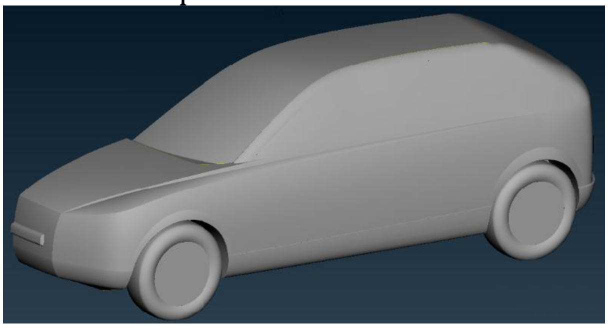

Figure 5: Model 2 VRAK_FAST

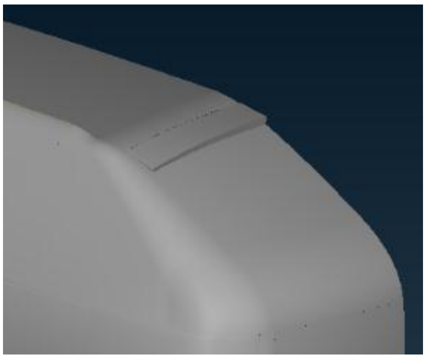

Figure 6: Model 3 VRAK_FASTW, rear wing view

Vehicle model 4 is a sedan Volvo S60 which is quite close to the final version of the production car. This model is more detailed than the other models, with headlights, detailed windshields, side mirrors etc. However, still in some area it has been simplified such as closed wheels, no grill opening or door handles. This model is a good comparison to the other models, where it is more detailed and already quite far in its aerodynamic optimization cycle.

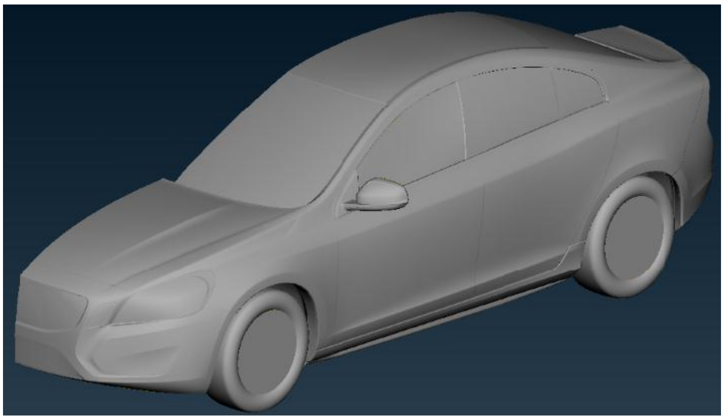

Figure 7: Model 4 Volvo S60

3.4. Configurations for testing

The methods for testing the Adjoint Solver will be described in this section. As previously explained, the aim is to state if Fluent's Adjoint solver shortly can be incorporated into Volvo's aerodynamic CFD procedure and if so, will it provide the expected benefits that the industry is looking for. In order to reach such a conclusion, the Adjoint solver must be tested in every aspect of its procedure. Not only in term of it's functionality and computing ability, but also for it's robustness, computer requirements and easiness of use.

The testing methods are divided into three main parts. The first part is to test the Adjoint solver's ability to handle different type's of mesh and mesh quality, the second part is testing the different control setting for running a an adjoint solution and finally overview the post-processing of sensitivity data and the morphing feature.

3.4.1. Mesh abilities:

The method carried out in order to test the Adjoint solver, with respect to the mesh it is solving, is also divided into several steps. The Adjoint solver will be tested on its' ability to handle different volume mesh builders, then test to see if there is any mesh size limitations and also test to what degree the mesh quality will affect the adjoint solution. The importance of a good quality mesh cannot be overemphasized, it's is must in order to get a good representation of the flow and provide trustworthy results from a computation that has converged. A computation that has not converged is more or less useless and since Ansys already stated that Fluent Adjoint solver has convergence difficulties, this becomes a big focus within this project.

In the AEDCAE01 Volvo procedure, Harpoon is used as a mesh generator with specific guidelines on how to build up the mesh such as size of wind tunnel, refinement boxes and so on. With these guidelines, one is expected to generate a high quality mesh that will give good representation of the flow. Since the aim is to implement the Adjoint solver within Volvo procedure, the first step is of course setting up simulations according to the AEDCAE01 procedure and running the Adjoint solver process. Harpoon is a good mesh generator where it can handle very complex models and can efficiently generate meshes without major difficulties. However the down side is that it isn't the most accurate mesh generator, where other generators can provide better meshes when it comes to quality. There is a wide variety of mesh generators being used in industry, within this project the focus is on testing Harpoon, Tgrid and ANSA. Tgrid and ANSA are tested as mesh generators mainly due to that they already are known of and used before at Volvo with similarities to Harpoon which wouldn't mean a drastic change of current procedure if it is shown that Harpoon is excluded when it comes to adjoint computations. The testing is carried out by simply running simulations where a mesh is generated by all three types, Harpoon, Tgrid and ANSA. As described in the workflow section, an adjoint computation is carried out after running a standard fluid flow computation, so of course the mesh is re-generated for a new case.

To test if there is any mesh size limitations, different setups for e.g. cell size and refinement zone are carried out and compared after running simulations. Mesh size limitations can surely become a problem due the complex computations that are needed when running an adjoint simulation, and therefore the bigger the mesh size is the more demanding the computation will become regarding both computer requirements and sensitivity data gathering. The amount of volume cells is what is meant by mesh size and this is something that varies a lot between different mesh setups and models, it all depends on the complexity of the model and how accurate the models is wished to represented. A mesh of a Volvo model at its' latest setup with all components can easily overcome hundred million, however with the model used this project which are simplified the maximum mesh size is around 30 million cells.

The quality of the mesh is of course linked to the mesh size, without directly meaning that increasing the mesh size would improve the quality. A good quality mesh is a mesh with such accuracy that it will give a good representation of the flow and sometimes this means that the mesh needs to be finer by decreasing the cell sizes and introducing finer refinement zones and so on. This though means increasing the mesh size. If a mesh isn't fine enough, the simulation experience instabilities that the Adjoint Solver can’t handle where the computation doesn't converge or crash before reaching the set number of iterations. This is something crucial for Adjoint Solver, which is tested by generating various types of meshes with respect to quality, refinement and mesh size.

With the mesh size and quality in mind, together with the convergence difficulties in the, one further step is running simulations with a polyhedra mesh. By re-generating the TGrid Tetrahedral mesh to polyhedras, which very conveniently is done directly in Fluent when loading a case before running first fluid flow simulation, the mesh size decreases and therefore may become easier to compute for the Adjoint Solver. This decrease can be quite drastic, a mesh around 26 million cells can drop down to 5 million cells. This gives a chance to build finer meshes and still get a reasonably small size.

One other test, not really related to generating a mesh, is the division of parameters on a model referred as PID (Property Identification Division) division which simply is to divide the geometry into different regions. Each PID is treated as a separate zone in Fluent when setting boundary conditions. For standard fluid flow computations this quite straight forward where parts that differ from each other is separated into different regions without being to detailed. However when it comes to the morphing , this can be a problem since the morphing is based on the gathered sensitivity data. When morphing, a morphing box is set up with regions chosen to be morphed, it can be the whole car or just one PID region. The morphing is than carried out is based on the sensitivity data of that region and if the region isn't detailed enough, the morphing feature may struggle due to the vast sensitivity data within the region. A large PID region might end up with sensitivities that vary so much within the actual region, which will make it harder to morph smoothly and reach the desired modifications. Therefore a more detailed division of PID regions is also tested so see how that effects the morphing.

3.4.2. Adjoint Control Settings

The setup of adjoint control settings is similar to the boundary condition setup for the standard flow calculations, however with fewer controls that more or less are recommended to set as default. Once again the testing is about state the solving ability of the Adjoint solver with the convergence difficulties in mind, therefore even though it is recommended to run on default settings, the controls are tested in a variety of setups.

The control settings for running an adjoint simulation are few and are all included in the control box shown in figure 8. The main focus is on the Advancement controls and Under-Relaxation factors. In the Advancement controls one can choose to run with Preconditioning and Stabilize Scheme. Preconditioning is needed for most cases involving turbulent flow and Stabilize Scheme is for large cases and complex geometry. By the Advancement Controls and Under-Relaxation factors, the solution algorithm is set where higher values correspond to a more aggressive algorithm that is more likely to be unstable. The values seen in figure 8 are the default values set for the adjoint computations.

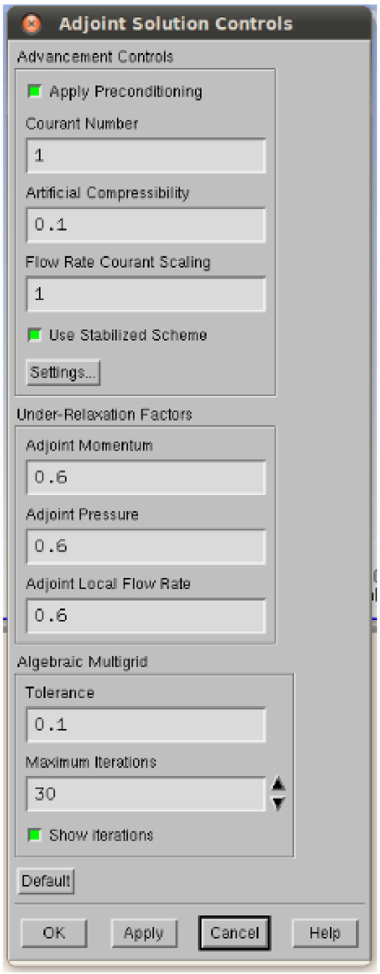

Figure 8: Adjoint Control Setting, default set

The testing is carried out by running simulations with different adjoint control setups on all four vehicle models with various types of meshes. By this, it is stated how adjoint control settings correlate with the complexity of the models and mesh generation, once again it is vital to investigate and find the best suitable setup for a stable computation.

3.4.3. Fluent Adjoint Solver Post-Processing and Morphing

As previously explained, the key factor of the Adjoint solver is the gathering of derivative sensitivity data and with that the ability to modify the geometry by mesh morphing with respect to the sensitivities and the morph set up.

As described in the workflow section, the post-processing of sensitivity data is carried out directly within fluent after an adjoint simulation is computed. There are several fields that can be used to post process the computation data, however the best suitable and mainly used field is the log10Shape Sensitivity Magnitude field which ranks the sensitivities of the chosen observable, which is drag in this project, with a map across the entire surface geometry as seen in figure 9.

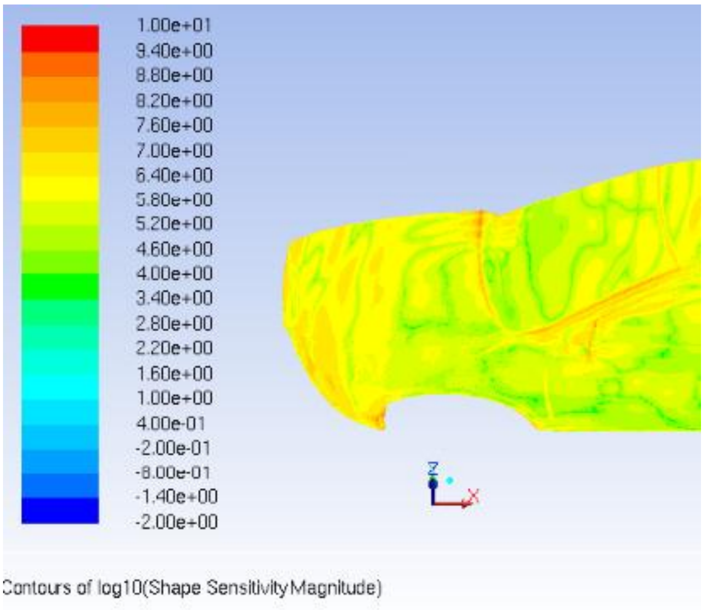

Figure 9: Plot of Shape Sensitivity Magnitude

Post-processing the derivative shape sensitivity data provides valuable engineering insight and is used as a very intelligent and efficient guide to design modifications that will improve the performance with respect to the drag. However, the extra step of actually modifying the geometry by morphing what really makes the Fluent Adjoint solver design optimization cycle an efficient process. Directly within Fluent together with the post-processing fields, the geometry is modified and can be sent to a second fluid flow computation. This takes away several time consuming steps of re-modeling the geometry in CAD, cleaning the model, generating a new mesh and so on. Surely, without the morphing feature, the adjoint sensitivity data is still of high value, however the modifications still has to be manually carried out within CAD or by conventional morphing in a pre-processor such as ANSA. Therefore the morphing is a very valuable feature and is tested in every aspect of it functionality. The actual procedure of morphing with Fluent Adjoint Solver is in detail described in appendix 1. Briefly explained, a morphing box is built around the region wished to be modified and within that control box a number of control points is set. The control points are distributed uniformly in each coordinate and the spacing, depending on the length of the morphing box, define the scale of the modification together with the scale factor defines the movement of the control point. An example is seen in figure 10, where the length of the morphing box is 1 meter with one hundred control points, with a scale factor of 1 it means that the actual modification of the control point in the surface is one centimeter.

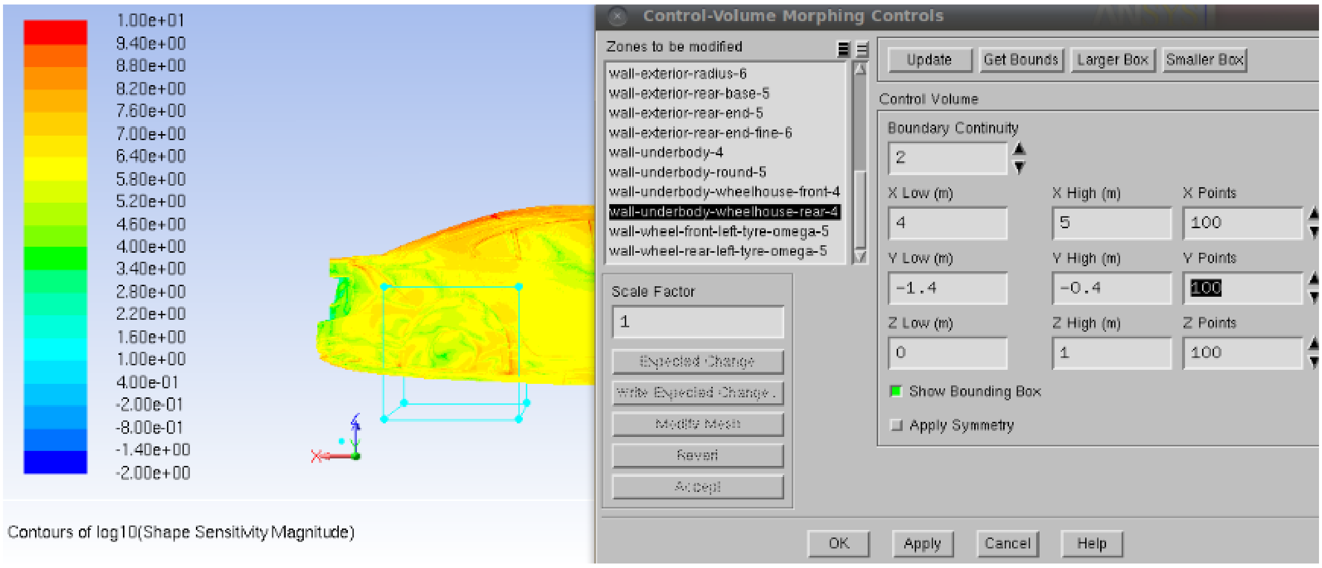

Figure 10: Morphing Setup

Testing of the morphing feature is carried out to get a grasp of the whole functionality and the ability to accurately handle a modification. The morphing feature is tested on all four vehicles models with different setups in both morphing and adjoint controls such as described in previous sections.

4. Results and discussion

In this chapter, a summary of all the results gathered from the many simulations that's been computed is presented. More than 50 adjoint simulations have been computed, with the many variations described in the methodology chapter. All these simulations showed the same tendencies, therefore the results from these simulations are not described in detail, it is described as a summary with discussion. However, one case that has gone through the whole adjoint procedure will be described in detail as an overview to get good understanding of the results and functionality.

4.1. Summary of results and discussion

As just explained, the results from the many adjoint simulations are consistent with the same tendencies with respect to adjoint sensitivity data, instabilities and convergence difficulties. From these results, a number of clear points can be established and are described below.

- **Fluent Adjoint Solver is very mesh dependent**
- **Cannot handle a mesh generated from Harpoon**
- **In most cases, cannot handle a mesh size with over 20 million cells**
- **Very unstable iterations, cannot reach convergence**

Since the future goal is to incorporate Fluent Adjoint Solver into Volvo's aerodynamic development procedure, the Volvo AEDCAE01 procedure is used where the mesh is generated in Harpoon. Therefore the first simulations were carried out with a Harpoon mesh and directly this was an issue. Basic fluid flow simulations, setup and computed according to the AEDCAE01 procedure, was never an issue. However going forward and running an adjoint simulation was never fully successful, Fluent Adjoint Solver never manage to complete a full adjoint computation where the simulation was set up with a mesh generated in Harpoon. All car models where tested with a variation of mesh size, still not a single one adjoint simulation was successfully computed, the adjoint computation either crashed after running a maximum of 700-800 iterations or simply did not start running at all. Therefore one conclusion is that the Adjoint Solver cannot handle a mesh generated in Harpoon. The reason for this is at the moment not fully clear, this has been discussed with Ansys which is the developer of Fluent, however the answer is not certain yet and will be investigated. One thought of reason is the mesh quality, as explained in the methodology chapter, Harpoon is sometimes criticized and not seen as one of the better mesh generators when it comes to mesh quality. This of course is a direct effect on the AEDCAE01 procedure, where Harpoon at the moment doesn't work when it comes to adjoint computations and therefore if Volvo wishes to use Fluent Adjoint Solver, other programs must be used to generate the mesh. As mentioned in the methodology chapter, Tgrid and ANSA, was also used to generate meshes and tested with the Adjoint Solver. Using Tgrid to generate a mesh wouldn't be a drastic change to Volvo's way of work where the program is already known and used in some cases. ANSA however is rarely used as a mesh generator since it can be quite complex and time consuming, especially when working with complex geometry. With this in mind, a few simulations were carried out with a ANSA mesh however the focus was on running simulations with meshes generated in Tgrid. The change from Harppon to Tgrid or ANSA didn't show any significant change when running basic fluid flow computations, however when running adjoint computations it was shown that the Adjoint Solver could handle meshes from both Tgrid and ANSA, still though not all the way to convergence.

This was a milestone in the project, where finally it was possible to analyze sensitivity data from the adjoint computations and go through with the whole adjoint process. This though ended up in discovering much greater issues with the Adjoint Solver, which is the maximum mesh size of 20 million cells and convergence difficulties. To be able to reach convergence, the mesh quality has to be very good so that the flow is well represented with no major instabilities, however since the mathematical background of the computations is very complex and quite computer demanding, the adjoint solver cannot handle large mesh size. The output sensitivity data and the effect of change with respect to all inputs simply becomes too much, which is better described by looking at the shape sensitivity formula presented in figure 11, where matrix x increases vastly. This leads to a catch 22, where the Adjoint solver cannot handle relatively large mesh sizes however it requires a fine mesh with very good quality.

$$\delta(Drag) = \sum_{mesh} \underline{w}^n . \delta \underline{x}^n$$

Figure 11: Shape Sensitivity Formula, Ref.[6]

The idea for solving this issue was to change the tetrahedral meshes generated in Tgrid to polyhedral meshes. The mesh for all four vehicle models were between 23-30 millon cells when generated in Tgrid, when transformed to a polyhedral number decreased to between 5-7 million cells. This meant that much finer meshes could be simulated in the Adjoint Solver and hopefully result in less instabilities and be able to reach convergence. However this still was not possible, even though the meshed was no much finer when running with polyhedral meshes, the adjoint solver could not successfully reach convergence with any of the many simulations. Even with a very fine mesh of 35 million cells generated in Tgrid on with the Volvo S60 sedan vehicle model which is considered to cause less instabilities than an estate model, the simulation still did not converge.

With these results in mind, it is safe to say that the Adjoint Solver at the moment isn't ready yet for commercial use, it simply needs more development time. Surely Ansys have done their own testing of Fluent Adjoint Solver, however the models used were very simplified and not in the same level when It comes to complexity or mesh size as the one used in this project.

Though one must keep in mind that the focus of this project has been external vehicle aerodynamics, specifically drag performance. The Adjoint solver may also be used in other areas, such as pressure drop through a system which can be a much simpler case to set up and compute and therefore the results can have a more positive outcome.

- **User should follow recommendations for adjoint controls setup**

As described in the methodology chapter, the Adjoint Solver is tested with a numerous of different adjoint control setting. Since convergence was never reached at any of the simulations, a vast number of different adjoint control set ups where tested with all four vehicle models to see if any specific setting could actually improve the computation residuals and be able to converge. However nothing was successful, throughout the simulations it was shown that running with control settings as recommended gave the best residuals, best residuals with respect to each other. A computation with residuals that has not convergence should not be taken in consideration in normal cases.

- **Morphing feature is highly valuable and very effective**
- **However very unstable due to no convergence and mesh size limitations**

As previously explained, the key factor of the Adjoint Solver is gathering the derivative sensitivity data and the morphing feature is more of an extra feature within Fluent that very conveniently is linked with adjoint results. Of coursethe sensitivitiy data provided by the Adjoint solver is valuable and provide great engineering insight, however it is the morphing feature that takes the procedure an extra step and making the Fluent Adjoint Solver procedure a very effective way of work and quite unique in it's way where one can within Fluent post-process the sensitivity data and directly carry out modifications by morphing the mode. Within the morphing set up, the effect on performance by the desired modification is estimated and with that one can evaluate if the modification is worthy making or should be reconsidered. The sensitivity data together with the morphing feature becomes a very effective way to carry out modifications that can improve the performance, where the design optimization cycle leads to better results with much less number of cycles when comparing to the conventional DOE methods used in industry.

However since the morphing modifications were carried out with computations that didn't converge, this was very unstable and could not always handle to modify the mesh since the sensitivity data is unstable with too much variations in the model. This can be described by viewing figure 12, where the sensitivity plots on the vehicle are unevenly distributed and do not correlate to realistic results. Another issue that kept occurring was that the mesh collapsed due to negative cell volume when trying to morphing. This happens since the mesh size is limited which means that the cell size in the volume mesh is relatively large and can not handle the modification. If the mesh was much finer, the cell sizes would be smaller and therefore easier to modify.

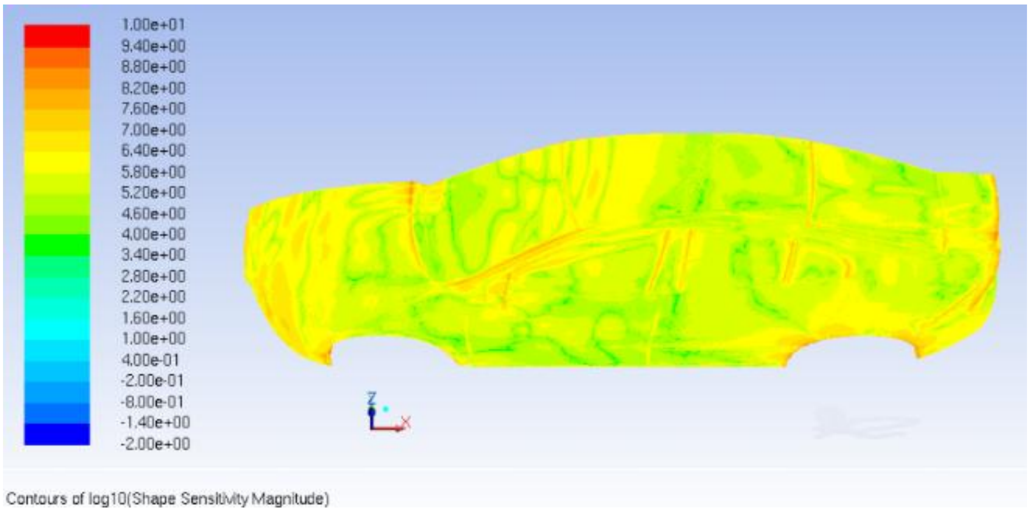

Figure 12: Plot of Shape Sensitivity Magnitude

With all of the mentioned above in mind, still the Adjoint Solver has proven it self to be a very valuable tool for engineers where the sensitivity data is a great guidance to improve the aerodynamic performance of a vehicle. With the sensitivity data, engineers can point out and focus on the most critical areas and directly a clear scope of what and how the model should be modified with respect to the chosen observable. Evan though the computations in this project did not converge and therefore are not trustworthy, still the results are interesting where the most critical regions could clearly be pointed out and further on morphed. The effect of the modifications carried out based on the adjoint results was expected to be greater, however the Adjoint Solver is at he moment simply isn't ready and need more development time with focus on the issues that are described in this report. One must have in mind that this is the first official release of Fluent Adjoint Solver by Ansys and is still under development.

4.2. Case follow-through: VRAK_BASE_X5

In this section a case setup with the VRAK_BASE estate model that has gone through the whole adjoint procedure optimization cycle will be described in detail.

4.2.1. 1:st Fluid Flow computation

First of all, the division of PID's were improved where the regions are divided in detail, which was proven to ease the morphing ability, is seen in figure 13 and the reason for this will be described in the morphing process.

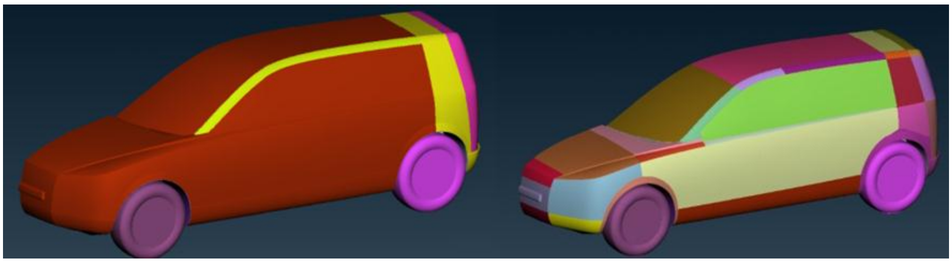

Figure 13: Comparison of PID Division

The surface mesh and the wind tunnel was built up in ANSA, and then imported into Tgrid to generate the volume mesh. The volume mesh generated in Tgrid was considered to be a good quality mesh with a mesh size of 26 million cells. Since a mesh of 26 million cells is larger than what the Adjoint Solver can handle, the mesh was transformed into a polyhedral mesh when imported into Fluent. The difference between the meshes can be seen in figure 14 and 15.

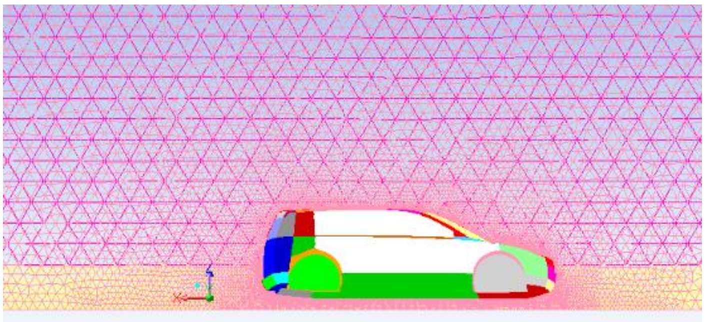

Figure 14: Tetrahedral Mesh, 26 million cells

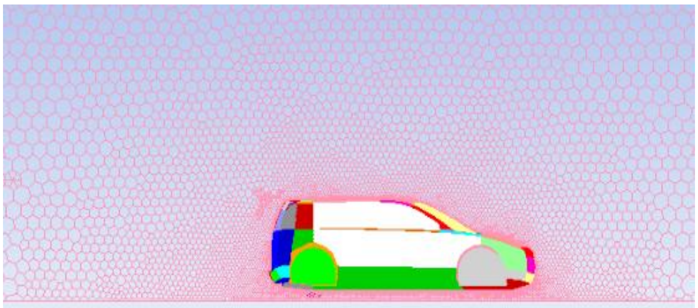

Figure 15: Polyhedral Mesh, 6 million cells

The boundary conditions for the basic fluid flow were like all other cases set up according to the Volvo AEDCAE01 procedure and ran for 3000 iterations. As seen in figure 16, the computations converged perfectly fine and the results can be seen in figure 17, where total force resulted in 256 N and the Cd to 0.264.

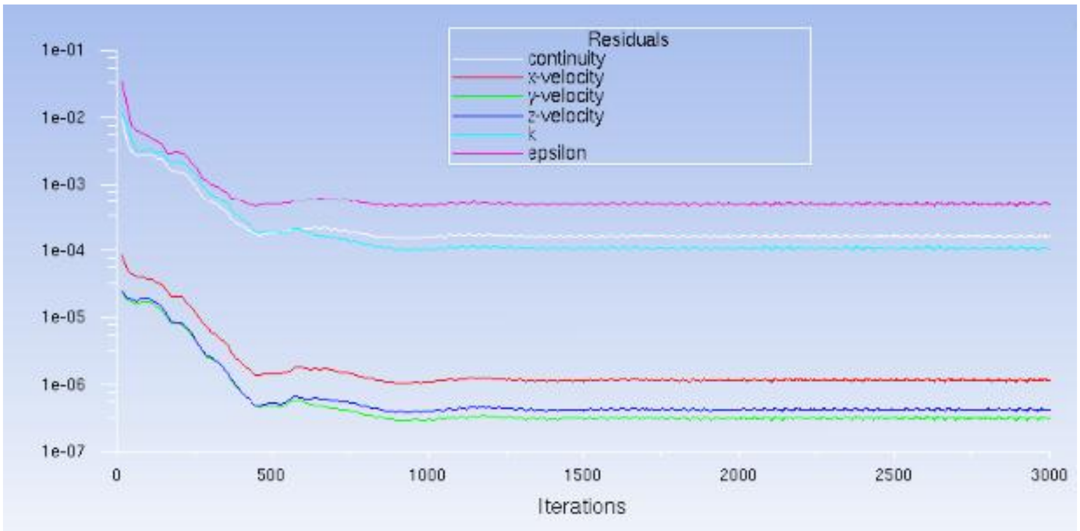

Figure 16 Residuals of 1:st Fluid Flow simulation

	Force x-dir (N)	Cd*A	Cd	%	Force z-dir (N)	Cl*A	Cl
Exterior	211.01	0.454	0.217	82.4%	202.79	0.436	0.209
wheelhouse 4	31.01	0.067	0.032	12.1%	-107.90	-0.232	-0.111
Wheel front	8.14	0.018	0.008	3.2%	-1.49	-0.003	-0.002
Wheel rear	6.06	0.013	0.006	2.4%	-1.68	-0.004	-0.002
Powertrain	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Eng bay	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Suspension	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Coolpack	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Miscellaneous	0.00	0.000	0.000		0.00	0.000	0.000
TOTAL	256.210	0.551	0.264	100%	91.717	0.197	0.094
TOTAL (incl. cooling drag)	$\boxed{256.210}$	0.551	$\boxed{0.264}$	0%			

Figure 17 Results of 1:st Flow Flow simulation

4.2.2. Adjoint computation

The following step was running an adjoint computation, where the adjoint controls where set up according to the recommendations and can be seen in figure 19. Besides being the recommended control settings, is was proven by previous simulations that a control setup like as seen in figure 19 gives the best residual results. Since the VRAK_base model is simplified and the mesh size is relatively small, Stabilize scheme is not needed. However the under-relaxation factors are decreased from 0.6 set as default, to 0.2 in order for a less aggressive solution to better stabilize the computation. The adjoint simulation is set for 2000 iterations and as seen in figure 18, the residuals did not converge.

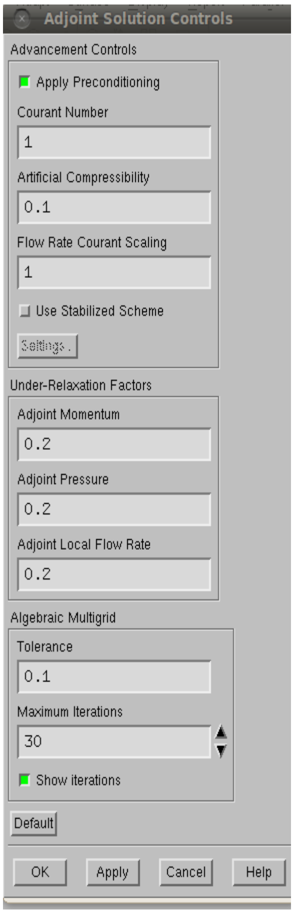

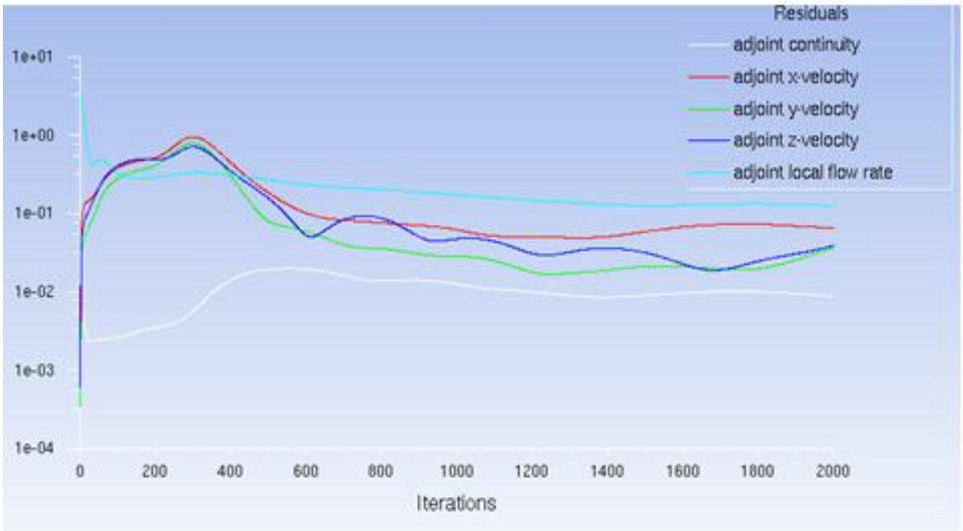

Figure 18: Residuals for Adjoint Solver simulation

Figure 19: Control Settings for Adjoint Solver simulation

4.2.3. Post-processing and morphing

The residuals did not converge, however still the the model was taken through the whole process. After running an adjoint computation, the sensitivity data is directly post-processed within fluent. The Shape Sensitivity Magnitude is plotted all over the vehicle model, where the red color indicates the most sensitive regions and therefore what to focus when carrying out modifications. When comparing to other models, as seen in previously showed figure 12, the sensitivity plot more evenly plotted indicating better residuals and less instabilities.

In this case, the VRAK_BASE model was morphed in three separate regions which are the front end of the hood, the a-pillar and the upper rear ending as seen in the figures above where the lines indicate the morphing boxes. The morphing boxes are built around the PID regions, and since the PID regions were divided in more detail, the morphing is only carried out with respect to the sensitivity data within that region and not affected by surrounding regions as when having very large PID regions. In Figures 20 the morphing boxes for this case is shown.

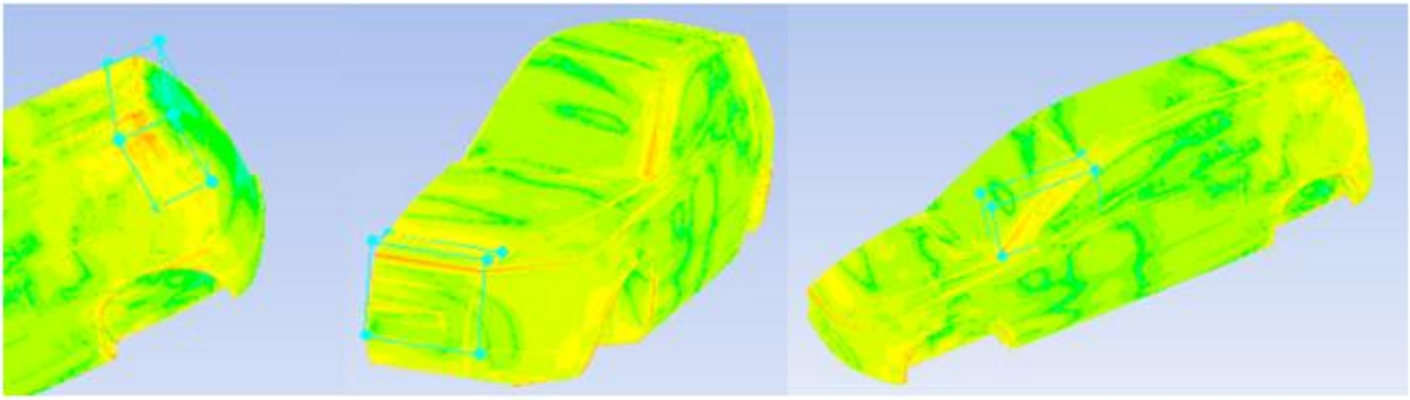

Figure 20: Morphing box setup for most critical regions

In Figure 21, the whole morphing setup is shown together with plots of the Shape Sensitivity Magnitude. Expected change for the modification carried out on the upper rear end region is a decrease of about 7.7 N. All three regions where individually morphed and the total number of expected change added up to a decrease of about 12 N, which surely is a good loss of drag force when knowing that the total amount of drag force was 256 N. The scale factor for the morphing was set to 1.5 which is more than recommended when carrying out an modification, it is rater recommended to carry out several steps with lower scale facto, e.g. 0.5 in order to get geometry modifications within reasonable amount. Morphing with a scaling factor of 1.5 was done to get clear modifications on the geometry and view how it really affects the result.

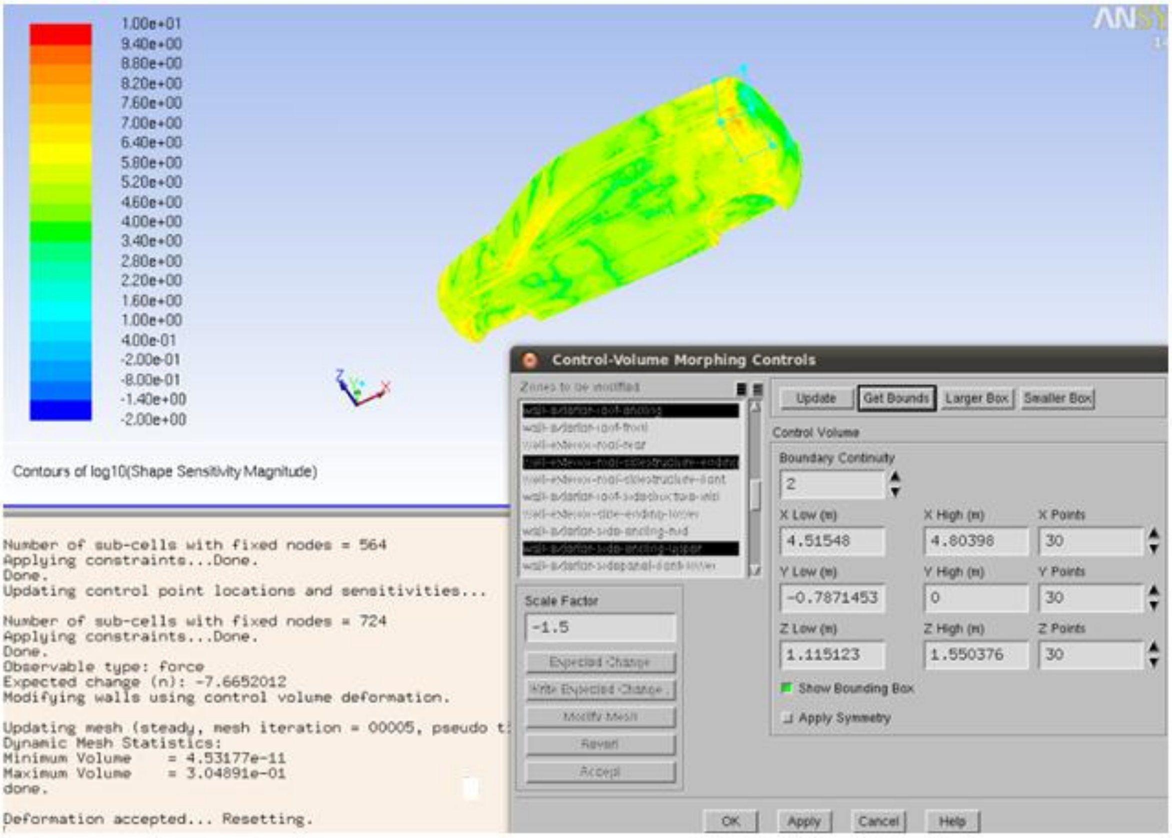

Figure 21: Morphing setup for case VRAK_BASE_X5

A better approach to morph and improve the geometry is to set a morphing box around the whole car. However this was not really possible since the morphing feature at the moment has a limitation of 99 control points in each coordinate, so a morphing box around the whole vehicle would mean that the relation ship between controls points and the actual length scale of the model would not correlate, the distance between the controls points becomes to large and therefore cannot properly morph the geometry. Another issue is that the morphing feature is quite computer demanding, initializing a morph with on large region can be very time consuming even though using a powerful industry computer.

The difference in geometry can be seen in the figures below, where the yellow model is before morphing and the purple model is after being morphed. The changes may be small, however they are still quite clear. There is a small change on the front end of the hood and the a-pillar where the shape is slightly raised. On the top rear end though, the change is quite drastic as seen in the figure where the top of the c-pillar is a given a weird shape which is due the high scale factor number.

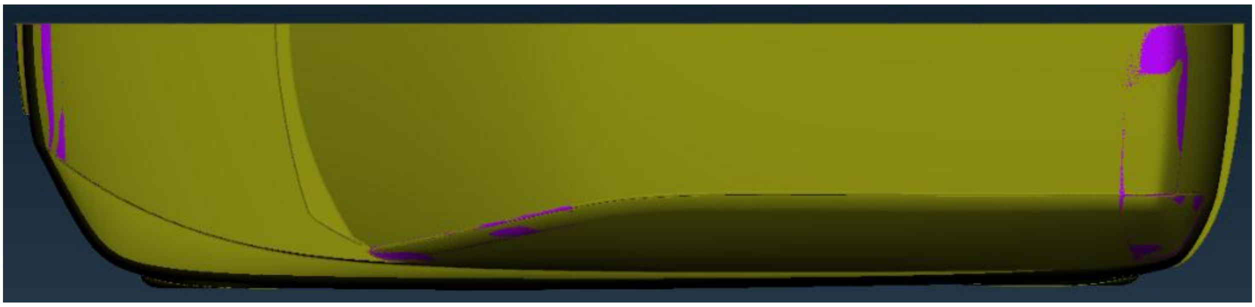

Figure 22: VRAK_BASE_X5 geometry before morphing

Figure 23: VRAK_BASE_X5 geometry after morphing

4.2.4. 2:nd Fluid Flow computation

After the VRAK_base model was morphed and saved with modifications made on the geometry, the model was sent to a second fluid flow simulation, which is the end of the adjoint optimization cycle. Figure 24 shows that the simulation converged just as the first fluid flow simulation, however the results was not as expected. The expected change was estimated to around a loss of 12 N to the drag force, however the actual number was only a loss of 2 N and 2 drag count as seen in figure 25. The main reason for this is of course the convergence issue, the morphing is based on sensitivity from a simulation that didn't converge and therefore one can not expect to get accurate results.

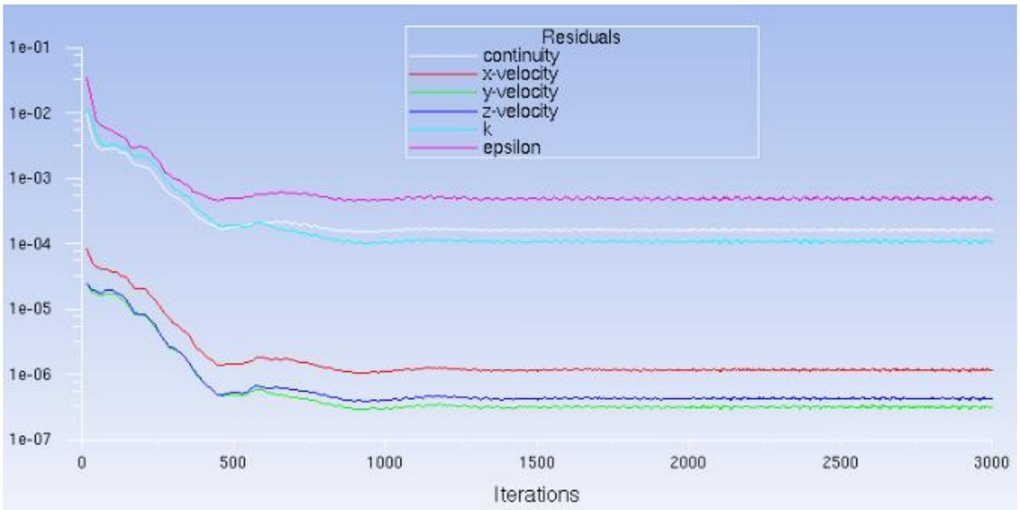

Figure 24 Residuals for 2:nd Fluid Flow simulation

	Force x-dir (N)	Cd*A	Cd	%	Force z-dir (N)	Cl*A	Cl
Exterior	209.30	0.450	0.215	82.3%	210.46	0.453	0.217
wheelhouse 4	30.78	0.066	0.032	12.1%	-108.08	-0.232	-0.111
Wheel front	8.31	0.018	0.009	3.3%	-1.26	-0.003	-0.001
Wheel rear	6.07	0.013	0.006	2.4%	-1.76	-0.004	-0.002
Powertrain	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Eng bay	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Suspension	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Coolpack	0.00	0.000	0.000	0.0%	0.00	0.000	0.000
Miscellaneous	0.00	0.000	0.000		0.00	0.000	0.000
TOTAL	254.456	0.547	0.262	100%	99.349	0.214	0.102
TOTAL (incl. cooling drag)	254.456	0.547	0.262	0%			

Figure 25: Results of 2:nd Fluid Flow simulation

5. Conclusions

- **Fluent Adjoint Solver is at the moment not ready for external aerodynamics.**
- **Fluent Adjoint Solver can provide valuable engineering insight.**
- **Fluent Adjoint Solver can both improve and decrease the lead development time.**

The main conclusion gained from this project is that Fluent Adjoint Solver is at the moment not fully ready for external aerodynamics and therefore not ready to be incorporated into Volvo's CFD Aerodynamic development process. However it is still quite clear that when ready, the Adjoint Solver surely will be a great tool to further improve the development process and reduce the overall lead time. The derivative sensitivity data provides valuable engineering insight that will help engineers get a better grasp the performance, what to focus on and how it should be carried out.

6. Future Work

- **Fluent Adjoint Solver is still under development, release with Fluent version 14.5 and 15 will be more robust**
- **The main focus will be improving convergence abilities & mesh size**
- **Provide better shape sensitivity insight**
- **ANSA has received the VRAK models used for this project**
- **Ansys has received VRAK models used for this project ans will be alble to test it before a new Adjoint Solver is released.**

The future work is on Ansys as they are the providers of Fluent Adjoint Solver. The Adjoint Solver is still under development and is promised to be more robust in future release where the main focus is on stable computations and mesh size in order to reach convergence, which have been the main issues in this project. Further Ansys will look into improving the post-processing functions of the derivative sensitivities by providing different set of tools to analyze the data in order get a better understanding. The morphing feature will also be overviewed and improved in the areas needed such as number of control points and collapsed mesh cells. During this project, ANSA has also been contacted regarding the sensitivity data output. Since the morphing feature is at the moment not robust enough, ANSA is working to develop a function that morph geometry within the pre-processer that modifies the geometry according to the sensitivity data. Furthermore Ansys is provided with the models used in this project so that the developers themselves can try similar simulations and when the Adjoint Solvers reaches a stage where it can handle all type of simulations, Volvo is informed and the question of incorporating the adjoint procedure into Volvo's Aerodynamic development can again be raised.

References

[1]. White FM. Fluid Mechanics. 6th ed. McGraw-Hill; 2008..

[2]. H.K. Versteeg & W. Malasekera. An Introduction to Computational Fluid Dynamics, The finite Volume Method 2nd edition. Pearson Education Limited; 2007.

[3]. Bernard RH. Road Vehicle Aerodynamic Design – An introduction 2nd ed. MechAero Publishing; 2011.

[4]. Wolf-Heinrich Hucho, editor. Aerodynamics of Road Vehicles Fourth Edition. SAE International; 1998.

[5]. ANSYS. Fluent Adjoint Solver User's Guide. Fluent Inc.; 2011.

[6]. Gilles Egenspieler. ANSYS Fluent Adjoint Solver Presentation. Fluent Inc.; 2012.

[7]. Volvo Cars. AEDCAE01: Aerodynamic performance assessment of a car exterior with closed front; 2009.

[8]. Evangelos M. Papoutsis-Kiachagias. CFD optimization via sensitivity-based shape morphing. National Technical University of Athens. Greece; YEAR?

[9]. Carsten Othmer, CFD topology and shape optimization with adjoint methods. Volkswagen AG; 2006.

Appendix 1

The procedure and controls used for running an adjoint simulation will is described in this appendix. As stated in the report, the method is quite straight forward and easy to set up.

1. Load the Adjoint Module: The Adjoint Solver is available in Fluent version 14 and is loaded by typing the following, as seen in figure 26:

Define > models > addon-module > Enter module number: 6

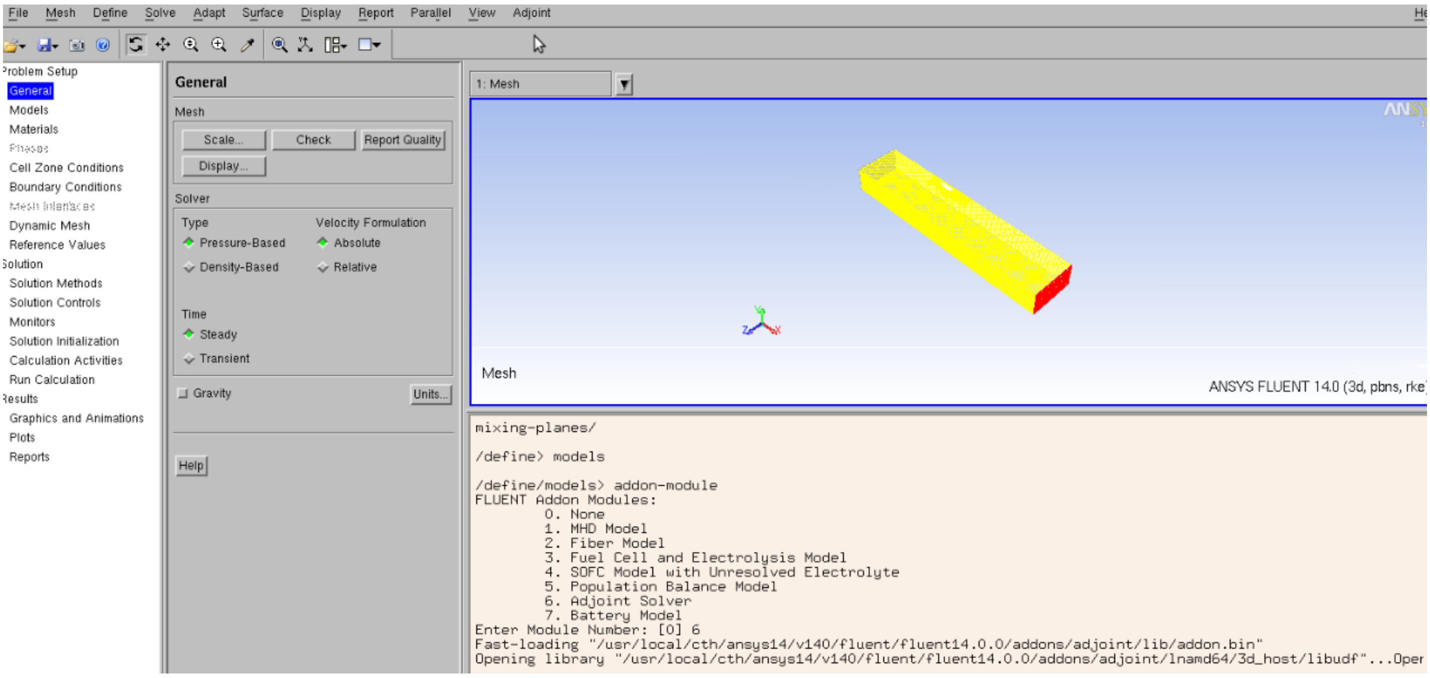

Figure 26 Step 1. Load Adjoint Module

2. Set up Adjoint Controls: Second step is to choose what to observe within Adjoint Observables tool box as seen in figure 27 where drag is chosen. Followed by setting the Adjoint solution controls and then finally set a number of iterations.

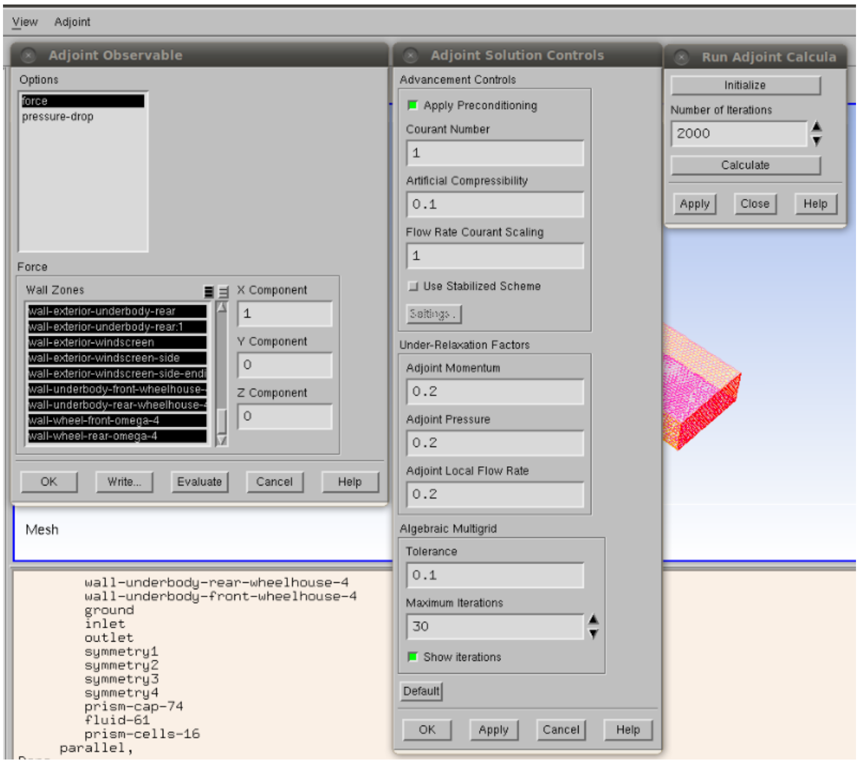

Figure 27: Step 2. Set up Adjoint Controls

3. Residuals: Step 3 is to check the residuals just as conventional fluid flow simulations.

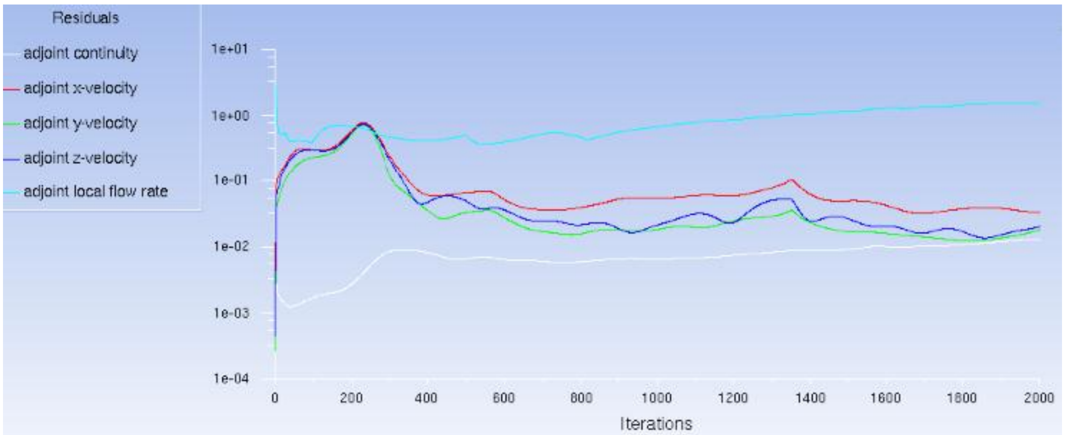

Figure 28: Step. Check resiudals

4. Post-process: The post-processing is done directly within Fluent by and can be viewed with a number of various tools. Recommended is to use the log10(Shape Sensitivity Magnitude) field within Contours, which plots the sensitivities over the entire vehicle model as seen in the figure 30.

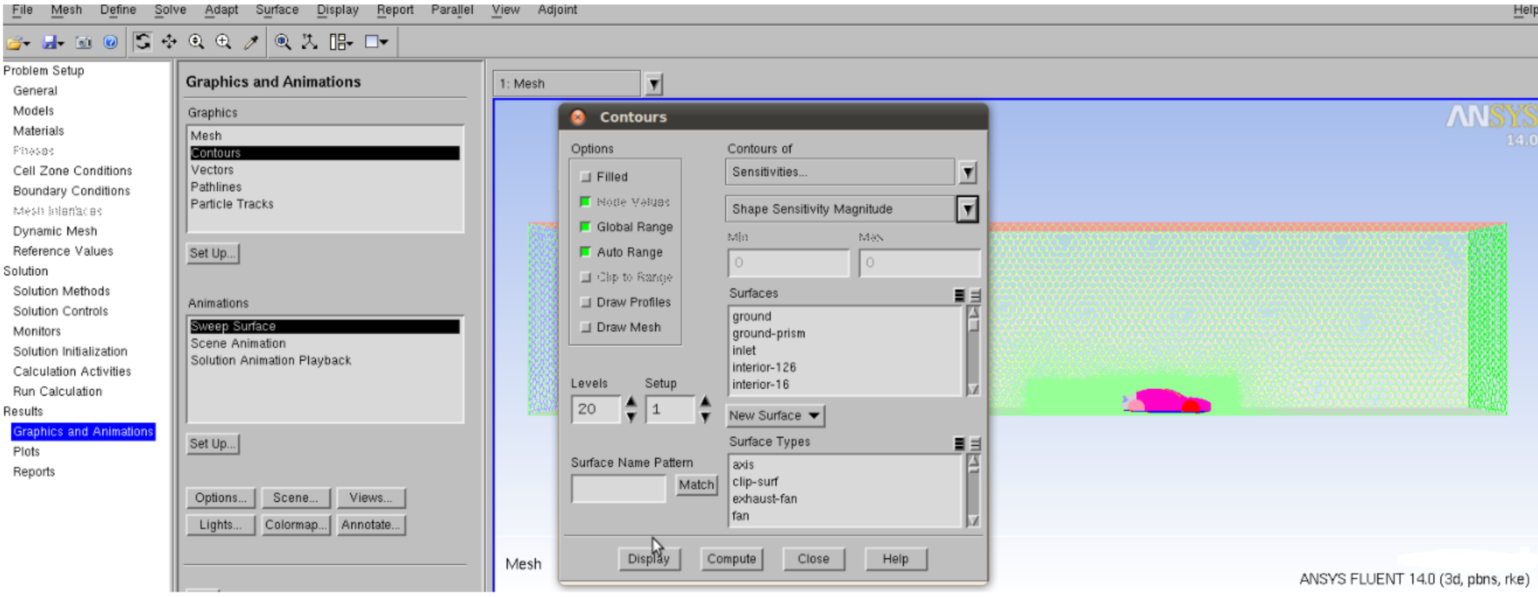

Figure 29: Step 4. Choose Sensitivty Plots

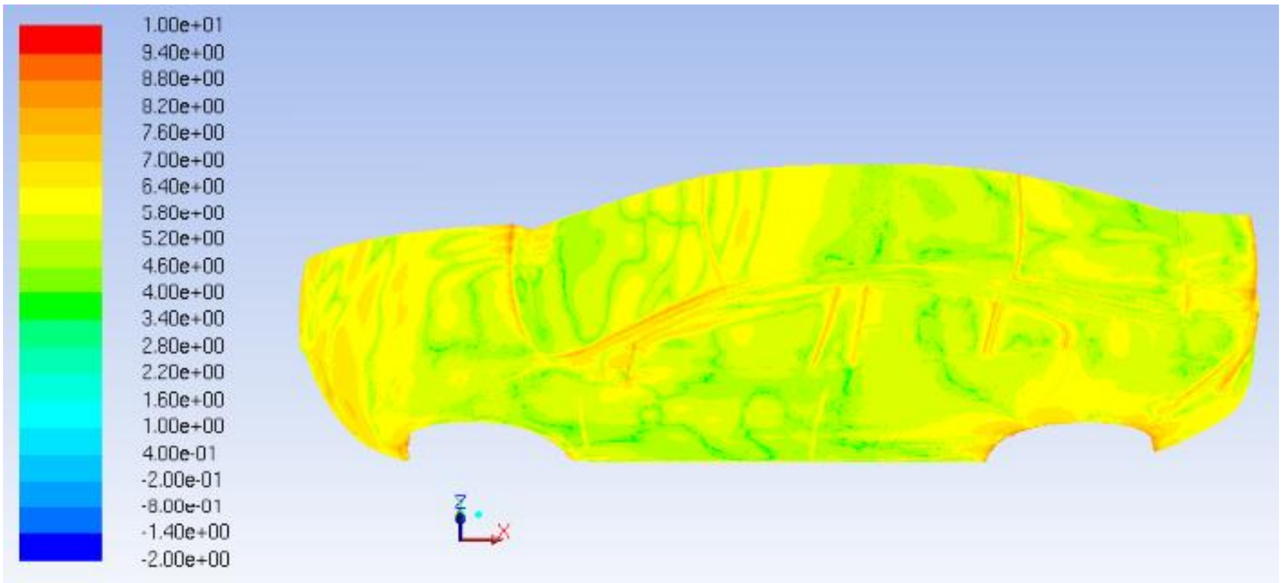

Figure 30: Plot of Log10(Shape Sensitivity Magnitude)

5. Morphing: The last step before running a second fluid flow simulation, is to modify the most critical areas by using the morphing feature within Fluent. Morphing boxes are set around the regions wished to be morphed and modified with respect to the sensitivity data and the morphing controls as described in the morphing methodology section in the report.

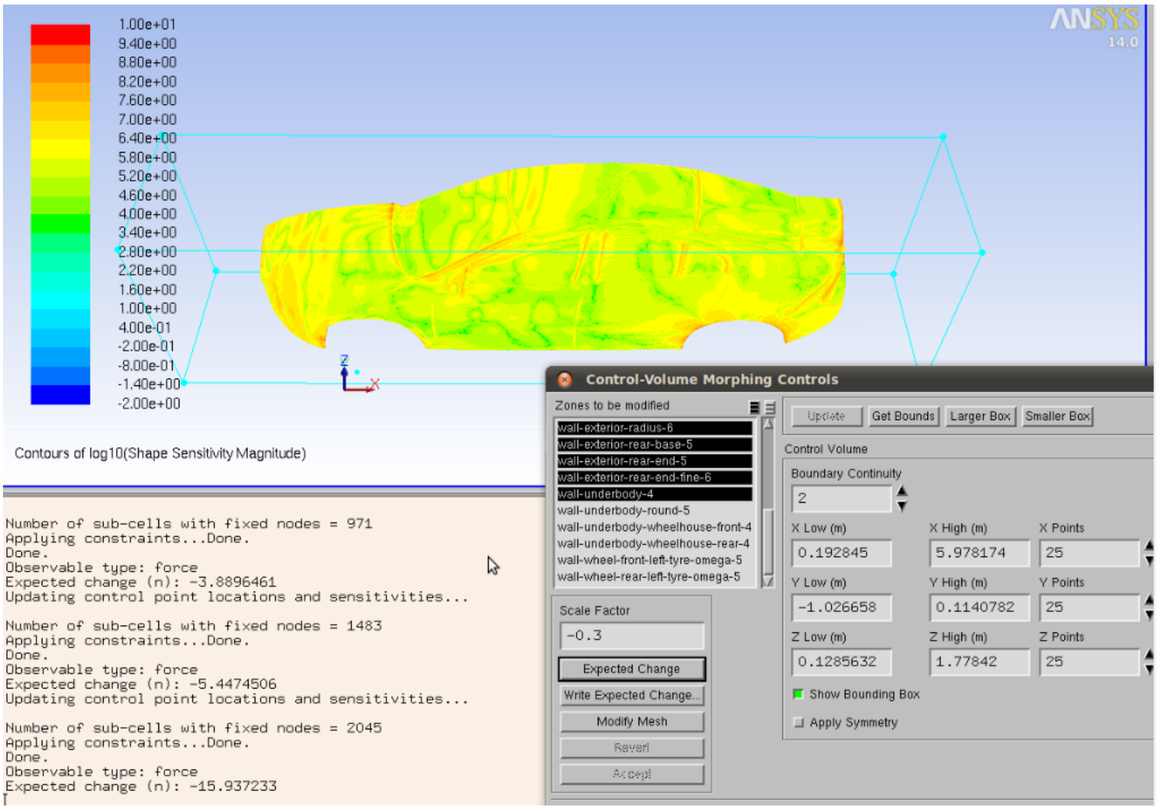

Figure 31: Step 5. Morphing controls