

AUTOMATIC OPTIMISATION OF AUTOMOTIVE DESIGNS USING MESH MORPHING AND CFD

¹Francisco Campos*, ²Paolo Geremia, ³Evangelos Skaperdas

¹ICON, UK, ²ESTECO, UK, ³BETA CAE Systems S.A., Greece

KEYWORDS – aerodynamics, optimisation, modeFRONTIER, ANSA, FLUENT

ABSTRACT - The growing access to extensive computational resources and the introduction of parallelisation methods has facilitated the integration of CFD analysis and optimisation techniques as a way to reduce lead-times from concept to market. This paper describes an example of an automotive application where response surface optimisation methods have been successfully applied to fluid related design problems using a combination of ANSA, FLUENT and the multi-objective optimisation product modeFRONTIER, as part of a proof-of-concept exercise.

TECHNICAL PAPER -

1. INTRODUCTION

The traditional CFD design process is often based on a “trial-and-error” type approach. Starting from an initial geometry, CAD changes are introduced manually based on results from a limited number of design iterations and CFD analyses. The process is usually complex and time-consuming, thus limiting the number of configurations that are verified and reducing the chances of finding an optimum solution.

The introduction of a mathematical framework to find an optimum design through the use of the latest optimisation techniques “standardises” the design procedure and eliminates most of the drawbacks found in the traditional approach. In this sense, designers are exploiting the benefits of parallel execution of CFD analyses by taking full advantage of the most sophisticated IT frameworks available. More precisely, distributed computing based on high-performing CPU technology allows an extensive use of optimisation tools for designs involving large-sized computational fluid domains. As a result, large numbers of configurations can be verified and analysed fully automatically to find an optimum.

In order to reduce the product development time and satisfy the growing design requirements to stay competitive in the market, new design approaches have been introduced. Consequently, designers are giving more and more importance to the quality of their work, complying with the principle of finding the best solution with the minimum effort.

Design parameterisation plays an important role in this context, since the definition of design parameters leads to a very efficient modification of the shape of the fluid domain geometry. Following this approach, the designer can investigate the effects of geometry parameterisation on model performance and find out which is the optimal shape that meets all the requirements needed. This task is normally accomplished by using optimisation software which allows the parametric model to be automatically adjusted in order to achieve the optimal solution. The techniques currently employed for this purpose include a parametric CAD model, topological optimisation and mesh morphing approach.

Parameterisation of CAD models in the optimisation loop is based on the fact that a finite number of design variables can define the geometrical shape of the model by using parametric curves and surfaces, following Bezier and NURBS theory [1]. This approach is widely used by engineers since most of the CAD tools available in the market feature parametric definition of the geometrical model. However, to integrate this process into the CFD optimisation loop requires a fully automatic mesh generation tool, leading to additional time and difficulty in the process. This is especially true in cases involving complex geometries which could result in poor quality meshes, thus affecting the reliability of the CFD analysis stage.

An alternative approach is that of Topological optimisation. The approach is based on the idea that cells with back flow can be retrieved and discarded, thus isolating the “healthy” cells which feature no flow recirculation problems. The overall space available for the fluid can then be shaped by switching off recursively the “bad” cells. The main advantage is that just a single CFD model is required and, most importantly, no parameterisation is required.

Finally, the mesh morphing approach allows for the model shape to be modified with no need for parameterisation of the CAD model or automatic mesh generation. An existing mesh is deformed based on predefined actions, such as stretching or contraction of nodes, thus saving computational time during each optimisation loop. However, a significant number of mesh changing parameters are needed to perform an efficient mesh modification while avoiding issues related to excessive mesh deformation.

In this work, the latter approach was applied in the optimisation of the rear part of a car body using a parametric mesh morphing model. The benefits of employing a Response Surface Methodology were also demonstrated.

2. RESPONSE SURFACE METHODOLOGY

Engineering design optimisation problems typically boil-down to performing the optimisation of a baseline system by evaluating the lowest number of analyses while extracting as much information as possible out of the design space. This approach requires a suitable sampling of the design space by using the most advanced Design of Experiments techniques available.

The traditional approach is to test one factor at a time (OFAT). The first factor is moved while the other factors are held constant, then the next factor is examined, and so on. Design of Experiments (DOE) provides a strong and universal framework to design and analyse all comparative experiments. OFAT usually requires many runs to obtain sufficient information and this is generally prohibitive. The DOE approach is in direct contrast to OFAT because it considers all factors simultaneously. Design of experiments uses the best factor settings to obtain a certain amount of information.

The DOE approach can also be applied for CFD optimisation. A limited number of evaluations exploring different design configurations can be performed by employing a “high-fidelity” CFD model. This is particularly useful when confronted by models with considerable mesh sizes and extensive running times.

After the sampling phase of the design space is completed, statistical analysis can be applied to the database of results in order to assess the response of the system with respect to variation of the factors considered in the analysis. For this purpose, t-Student criterion as well as main effects analysis can be applied in order to investigate existing relationships between the output values (response variables) and the input variables (factors). As a result, the direct/inverse effect of each factor with respect to the response variables can be detected and quantified.

A further step in the design process can be made whenever the need for a better resolution of the response of the design space arises. In practice, the designer would typically like to understand well the behavior of the system under investigation. For this, a suitable fitting of the data available from the numerical computations is therefore needed. A technique known as Response Surfaces Methodology (RSM) is a widely employed approach for building meta-models to approximate the response from experimental data or analysis codes.

RSM is a collection of mathematical and statistical techniques useful for the modelling and analysis of problems in which a response of interest is influenced by several variables. These synthetic models, when found to be satisfactory, effectively replace the simulation in the next steps of optimisation or design space exploration. The important point to note is that these simplified models will calculate a response in near instantaneous time, with obvious advantages over an alternative such a full CFD analysis.

In this paper, the geometrical shape of the rear part of a passenger car body was parameterised according to a mesh-morphing related model, in order to analyse the response of the system with regards to the aerodynamics. The Latin Hypercube Sampling scheme was used to perform the preliminary exploration of the design space through the DOE approach. After the sequence of computations was completed, several RSM algorithms were considered in order to retrieve a proper mapping of the solution under examination. As a result, the most efficient designs were selected from the set of optimal solutions belonging to the so called Pareto frontier.

3. APPLICATION CASE: OPTIMISATION OF CAR AERODYNAMICS

The optimisation case in this example focuses on the external aerodynamics of the Audi A6 Avant shown in Figure 1. The aerodynamics engineer is responsible for enhancing the aerodynamic characteristics of the car so as to reduce drag, reduce wind noise, and prevent undesired lift forces at high speeds.



Figure 1. Audi A6-Avant (www.audi.com – AUDI AG © 2007)

The main objective of the simulations presented here was to find the optimal aerodynamic shape for the rear part of the car body. The optimisation problem was therefore defined by means of the following two objective functions: (i) minimisation of the drag coefficient, and (ii) minimisation of the car balance in terms of aerodynamic lift forces acting on the front and the rear wheels. The latter is defined as the absolute difference between front and rear lift (zero being the optimum).

The optimisation procedure consisted of generating a parametric mesh morphing model of the car geometry, performing a CFD analysis in batch mode and using the multi-objective optimisation product modeFRONTIER to find, screen and select the optimal design layout.

The work to perform the car optimisation can be summarised as follows:

- Generate parametric mesh morphing model in ANSA
- Setup CFD analysis in FLUENT
- Setup optimisation workflow within modeFRONTIER
- Perform statistical analysis
- Response Surface modelling
- Results assessment

3.1 MESH MORPHING WITH ANSA

A watertight model was created in ANSA v12.1.1 by merging a combination of geometrical data for the main aerodynamic surfaces of the vehicle and shell elements for the wrapped suspension and exhaust system.

The variable size surface mesh, shown in Figure 2, was restricted to around 330 thousand triangular elements, while the maximum Fluent skewness was kept below 0.6.

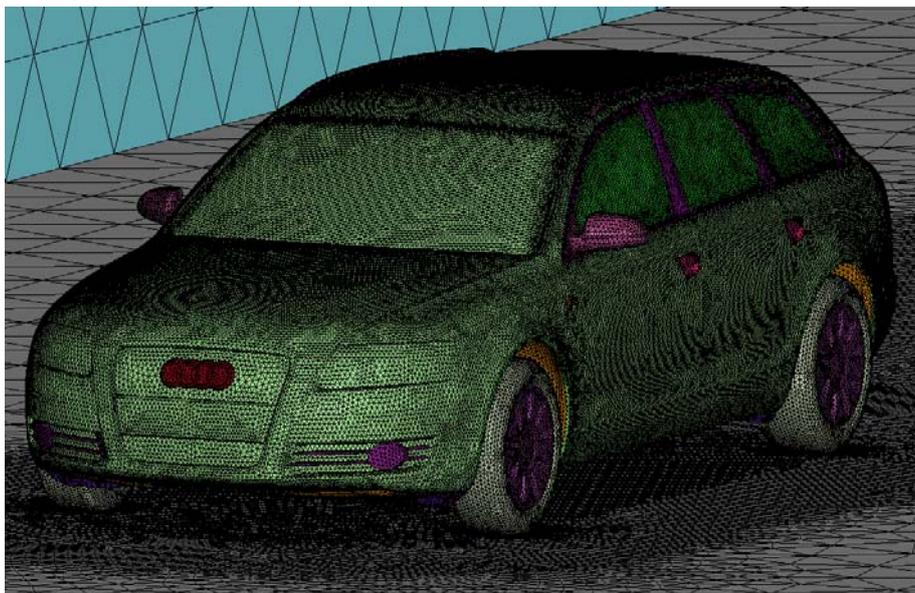


Figure 2. Surface mesh of vehicle and virtual wind tunnel box

The volume mesh consists of around 4.4 million elements. The Hexa-Interior algorithm of ANSA was used to generate variable size hexas in the main volume connected through pyramids and tetras with the surface mesh.

Refinement boxes were placed in the wake of the vehicle and wing mirrors, thus allowing control of the volume element size, as shown in Figure 3. An average Fluent tetra skewness of 0.27 was achieved, with max skewness kept below 0.95.

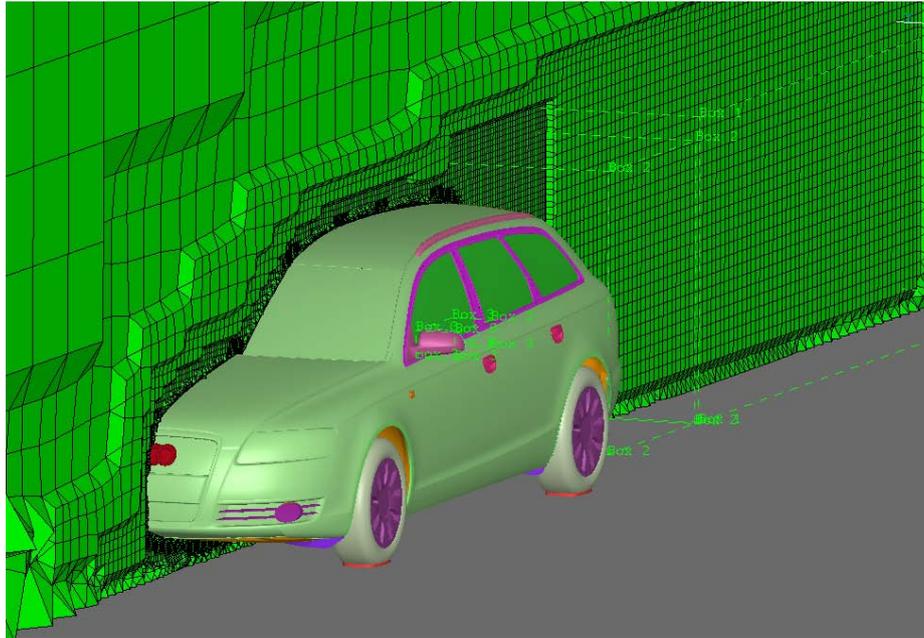


Figure 3. Hexa-Interior volume mesh with refinement Boxes in the vehicle and mirror wakes

After the mesh was built, morphing boxes were created, split and fit exactly onto the main features of the vehicle, as shown in Figure 4. The boxes allow precise control of morphing operations to examine various shapes parametrically. Morphing parameters were specified in predefined degrees of freedom so that the position of the rear spoiler as well as the inclination of the D-pillar could be altered. A total of four design parameters were set to influence the shape of the rear of the car, namely: rear spoiler position in X, Y and Z, and D-pillar angle.

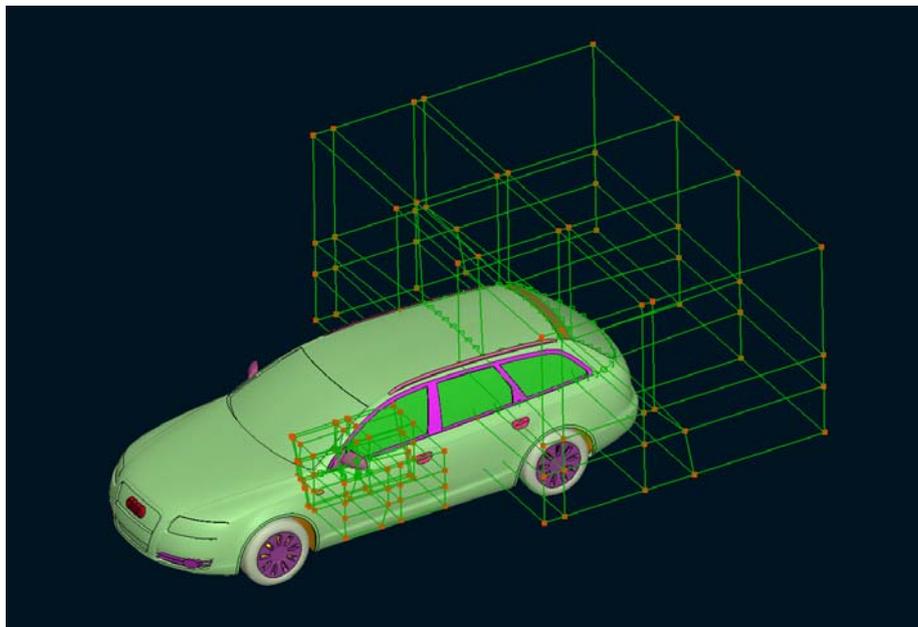


Figure 4. Mesh morphing model within ANSA

A preliminary feasibility study was performed in order to determine a valid range for each of the design parameters described above. Any parameter set causing mesh generation difficulties or excessive mesh deformation was detected and the values for the range of each parameter readjusted accordingly. This task was accomplished in modeFRONTIER by systematically exploring all of the possible combinations of the design parameter sets within the bounding values.

Finally, all mesh morphing operations were automated in modeFRONTIER using the ANSA batch running capabilities. A session file was created to read the morphing parameters from a text file containing the morphing definitions. The process was repeated for each design variation by changing accordingly the values listed in the text file.

3.2 CFD MODELLING WITH FLUENT

The CFD analysis stage was performed in FLUENT 6.3. The air flow was computed as an incompressible-subsonic turbulent gas (i.e. constant density). The Spalart-Allmaras one-equation turbulence model, specifically intended for aerodynamic calculations, was applied in every simulation.

A normal inlet velocity boundary condition was specified at the entry of the virtual wind tunnel with a value of 19.44 m/s (or 70 kph driving speed). The prescribed incoming air temperature was set to 293 K, while the turbulence intensity and length scale were set to 5% and 0.4 m, respectively. Similarly, the static pressure was fixed at 0 Pa relative to a reference pressure of 101,325 Pa at the tunnel outlet. The ground and wheels were assumed to be stationary, while all surfaces representing the car body were defined as adiabatic and no-slip wall boundaries. Finally, the tunnel walls were modelled as symmetry planes.

The SIMPLE solution procedure for pressure-velocity coupling was employed in the calculations with the AMG solver. The default differencing schemes were utilised for all flow variables. An upper limit of 500 iterations was specified in the event that a case would not fully converge.

FLUENT was run in batch mode using a journal file to allow modeFRONTIER to perform the analyses and extract all relevant information fully automatically. The latter was achieved by means of post-processing monitors for drag, lift and moment coefficients.

3.3 modeFRONTIER SETUP

The optimisation software employed in this work was modeFRONTIER. This package allowed the process integration of ANSA and FLUENT with the design study strategy employed. This was achieved by reading a specific set of morphing parameters as input variables for ANSA, and retrieving drag coefficient and aerodynamic lift forces values as output variables from FLUENT.

For every step of the optimisation process ANSA and FLUENT were run in batch mode in order to evaluate each specific design configuration. Furthermore, modeFRONTIER parallelisation capabilities were exploited to allow concurrent FLUENT analyses to be performed in a multi-processor cluster machine.

The whole process, including the data and logic flows, were defined under modeFRONTIER's workflow environment, as shown in Figure 5.

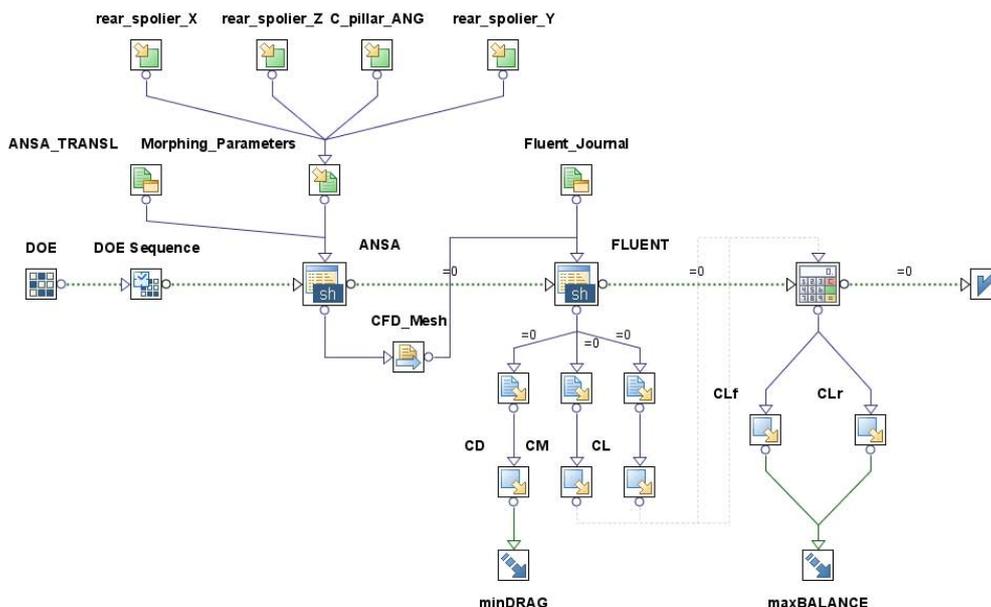


Figure 5. modeFRONTIER Workflow definition

3.4 DESIGN OF EXPERIMENTS AND STATISTICAL ANALYSIS

The Latin Hypercube Sampling (LHS) preliminary exploration sequence was employed [3]. Given a multivariate statistical distribution, LHS technique samples the joint probability distribution (i.e. the statistical distribution of all the variables) independently generating the values of the marginal probability distributions (i.e. the statistical distribution of each single variable). LHS arranges the values of the marginal probability distributions in order to sample the multivariate space in a relatively uniform way. Consequently, the correlation index between the design parameters was drastically reduced to almost nil, as shown in green in Figure 6, thus confirming the fact that LHS distributes efficiently the points around the design space in this particular case.

Following this approach, 40 designs were created and then evaluated using FLUENT according to a uniform distribution applied to each of the four input variables considered.

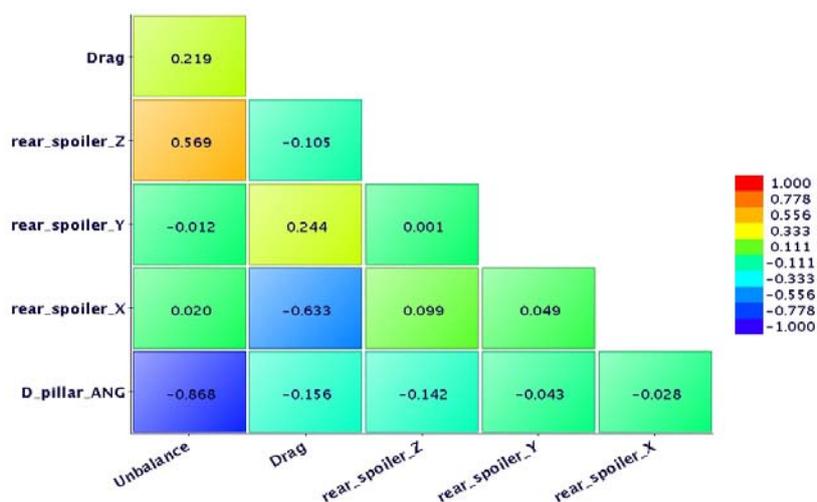


Figure 6. Correlation Matrix of the variables

Figure 6 also depicts the correlation index between the output variables (i.e. drag and unbalance) and the input variables. Thus, the measure of the linear association between each couple of variables was automatically computed and displayed.

Advanced statistical tools were then exploited in order to acquire more information about dependencies between input and output variables. For this purpose, the Student Chart was created for both the output variables in order to assign a precise measure of the influence of each factor (i.e. the input variables). This tool is based on the t-Student test considering the significance of the effect of each factor on a specific response variable. As a result, any direct/inverse effect can be detected and then quantified (see Figure 7).

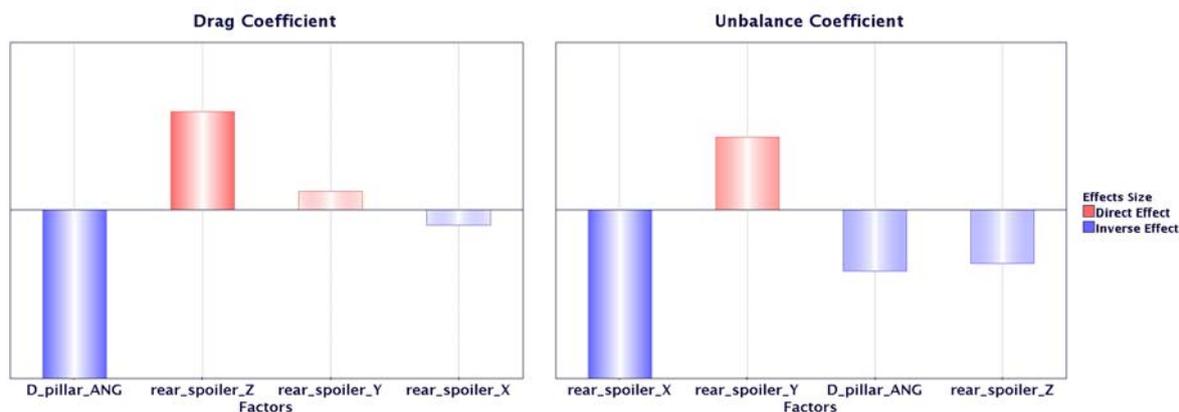


Figure 7. Student Chart – Drag Coefficient (left) and Unbalance (right)

The Student Chart and the Correlation Matrix both reveal a strong inverse relationship between the D-pillar angle and the drag, and also between the rear spoiler variation in X and the unbalance. Similarly, the rear spoiler variation in Z and the rear spoiler variation in Y have significant direct effects on drag and balance, respectively.

3.5 RESPONSE SURFACE MODELLING

The RSM tool featured several algorithms designed for data interpolation. Different techniques were employed to produce a proper fit to the data available from the 40 CFD simulations. Three state-of-the-art algorithms were considered, namely (i) Multivariate Polynomial Interpolation, (ii) Neural Networks and (iii) Radial Basis Functions.

The Multivariate Polynomial Interpolation is based on the Singular Value Decomposition (SVD) algorithm. It is considered to be a very powerful technique for solving a linear system of equations with particular benefits when dealing with numerically difficult situations [4].

Neural Networks (NN) is a very efficient and powerful interpolation tool [5]. Inspired by the brain structure and functions, NN can learn from training data. In this way NN can model any generic non-linear relationship between input and output variables.

Radial Basis Functions (RBF) is another efficient tool for multivariate scattered data interpolation [6]. Scattered data means that the training points do not need to be sampled on a regular grid and in practise this provides a truly meshless method. Since RBF are interpolating response surfaces they pass exactly through training points, whereas both the Polynomial and the NN methods are pure approximating algorithms, since they perform a surface fitting throughout the set of the points available for the training.

In order to evaluate the accuracy of both the Polynomial and NN surfaces the residual quantities were considered. The residual quantities are the differences in terms of relative or absolute error between the value of the points belonging to the training set and the corresponding value predicted by the response surface. More precisely, the mean and the maximum values were considered for each response surface created. Furthermore, the regression index was exploited as a measure of the quality of the response surface. This criterion measured the amount of the overall interpolation error between effective and predicted values of the points belonging to the training set. A regression index value close to 1 indicates a good synthetic model.

For the RBF functions, the “leave-one-out error” criterion was considered. In turn, each point belonging to the training set is excluded from the training procedure. The value predicted in the excluded point by the so created surface is then compared to the known value. A smaller value on average indicates that the response surface is better trained on the whole dataset.

Following this approach, first, second and third-order Polynomial Interpolation functions and a fully automatic NN were exploited. Regarding the RBF surfaces five different radial functions were applied: Gaussians (G), Duchon’s Polyharmonic Splines (PS), Hardy’s Multi Quadrics (MQ), Inverse Multi Quadrics (IMQ), and Wendland’s Compactly Supported C2 (W2). This list represents a complete set of state-of-the-art and widely used radial functions that can be found in literature [6].

Analysis of the RSM interpolation errors on all the nine response surfaces algorithms revealed that the best performer was the RBF algorithm for both of the output variables considered. Furthermore, the Gaussian model was selected for the unbalance, while the Duchon’s Polyharmonic Splines was applied to the drag coefficient.

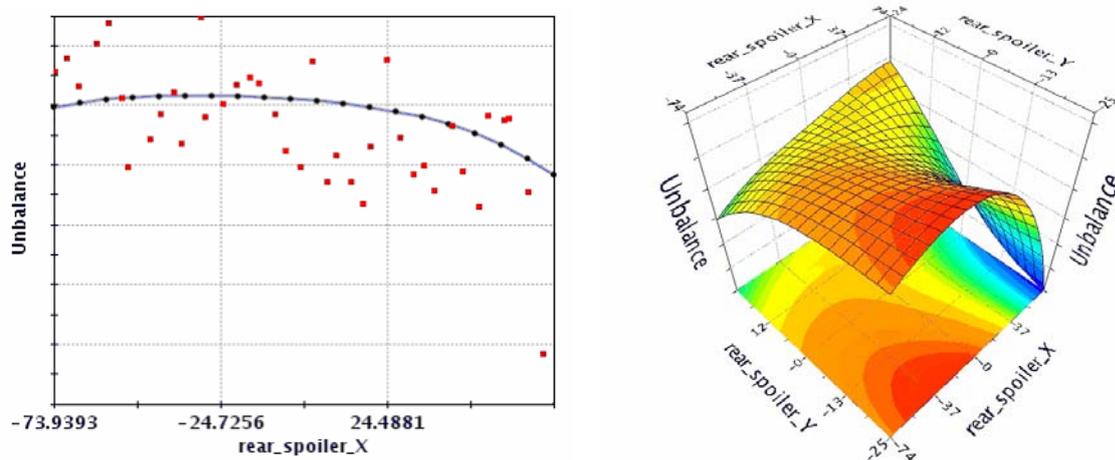


Figure 8. RBF surface for Unbalance

Two-dimensional plot: vs. rear_spoiler_X (left)

Three-dimensional contour plot: vs. rear_spoiler_X and rear_spoiler_Y (right)

Figure 8 shows the matching for the RBF is sufficiently accurate. The smoothness of the resulting function also provides a good fitting of the design space with respect to the points considered in the training set.

3.6 OPTIMISATION PHASE

After the response surfaces were selected for the two output variables, it was possible to re-use the created meta-models instead of using the CFD solver in order to evaluate the performance of the car aerodynamics. Since the design evaluation using the RSM is very fast, a massive optimisation of the design space was carried out using the response surfaces created. The NSGA-II optimisation algorithm [7] was used for the present work as it is based on a fast and elitist multi-objective evolutionary algorithm. At the end of the optimisation phase, the best designs (by definition those belonging to the Pareto frontier) were retrieved.

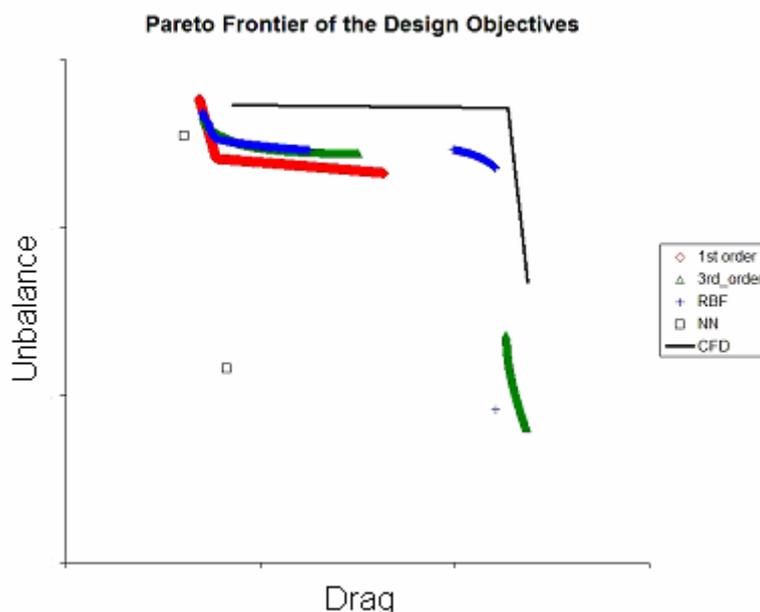


Figure 9. Scatter Chart of the Pareto frontier for the CFD and the RSM results

For each RSM created, an optimisation run was performed and then the Pareto frontier was detected. Apart from the second-order Polynomial Function, all the other response surfaces showed acceptable agreement with the Pareto curve provided by the CFD runs (as shown in Figure 9). Moreover, all of the RBF functions showed a similar behavior in terms of results obtained.

The plot of the Pareto frontier highlights once again the better accuracy of the RBF model in fitting the dataset available. This response surface model was successfully employed to find, screen and select the layout of the optimal shape of the car for enhanced aerodynamics performance.

The use of RSM led to a drastic reduction in analysis time. A full optimisation that would have required hundreds of FLUENT computations (maintaining the same optimisation algorithm design settings) was replaced with only 40 design evaluations using the RSM approach proposed in this work.

4. CONCLUSIONS

The fast and efficient mesh-morphing tools available in ANSA, as well as parallel executions of FLUENT solver capabilities, were successfully coupled with modeFRONTIER in order to manage the whole design process, as well as the RSM strategy applied, as part of a proof-of-concept exercise.

The Latin Hypercube Sampling scheme was used to perform the preliminary exploration of the design space through the DOE approach. After the sequence of computations was completed, several RSM algorithms were considered in order to retrieve a proper mapping of the solution under examination. As a result, the most efficient solutions were selected from among the set of the not dominated solutions belonging to the Pareto frontier.

In order to select the best solution, the results from the RBF response surface were considered. Both performance objective goals were successfully achieved. The optimisation of the RBF response surface model allowed a fast comparison of the results for drag and unbalance for both the baseline design and the optimised configuration (see Figure 10). Overall, the drag coefficient was reduced -2.7%, while the unbalance was reduced -8.3%.

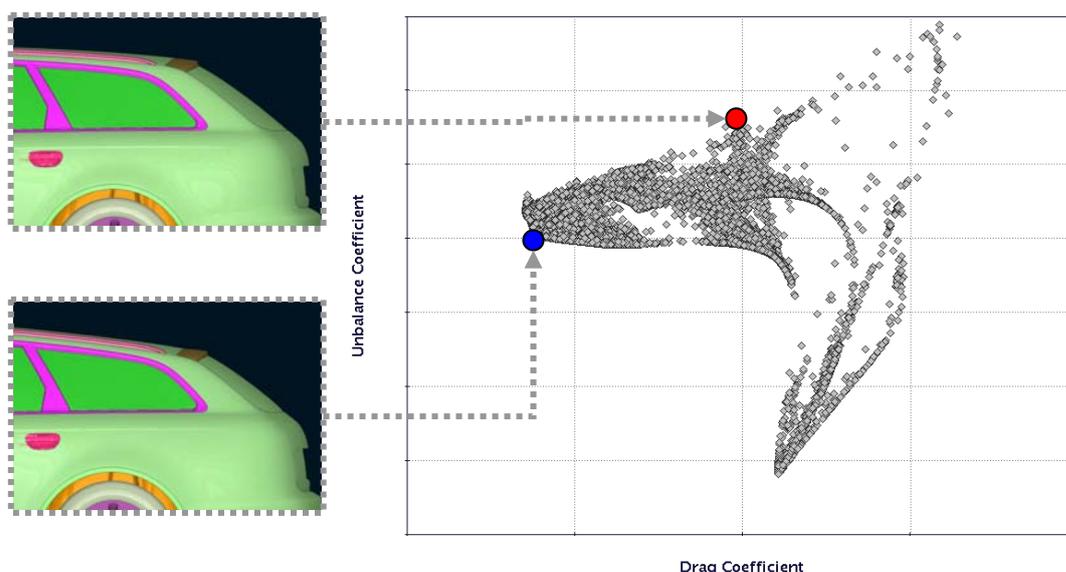


Figure 10. Scatter Chart of the Design Objectives using the RBF model (red dot – baseline design; blue dot– optimised design)

The solutions presented above are restricted by the accuracy of the CFD methods employed. For example, it is well known that RANS turbulence models are sometime unable to predict the on-set of separation, or even whether flow separation occurs at all [8]. This type of situations tends to generate misleading results which could affect the outcome of the optimisation exercise. Consequently, a good knowledge of the limitations of the numerical methods in use is fundamental when selecting the best design from the Pareto frontier. This knowledge is reinforced by extensive experience gathered from widespread wind tunnel testing and systematic validation of selected aerodynamic CFD calculations.

REFERENCES

- (1) Haslinger, J., and Mäkinen, R. A. E., Introduction to Shape Optimization, and Computation, SIAM, ISBN: 0898715369
- (2) Campos, F., Geremia P., Weston, S., and Islam, M., Automatic Optimisation Of Aerodynamic Designs Using CFD-Based Methods, Proceedings of 6th MIRA International Vehicle Aerodynamics Conference
- (3) modeFRONTIER v3 User manual – www.esteco.com
- (4) Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P., 1992, Numerical Recipes in C. The Art of Scientific Computing, 2nd ed., Cambridge University Press

- (5) Haykin, S., Neural Networks – A Comprehensive Foundation, Prentice Hall International Editions, 1999.
- (6) Buhmann, Martin D., 2003, Radial Basis Functions: Theory and Implementations, Cambridge University Press
- (7) Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. 2000, A Fast and Elitist Multi-Objective Genetic Algorithm-NSGA-II, KanGAL Report Number 2000001
- (8) Menter, F. Test Case 9.4: Flow around a simplified car body (Ahmed body). Comments (F. Menter), 10th joint ERCOFTAC (SIG-15) -IAHR-QNET/CFD Workshop on Refined Turbulence Modelling, October 2002, viewed on 14/11/2005, <http://labo.univ-poitiers.fr/informations-lea/Workshop-Ercoftac-2002/case9.4/Case9.4_comments.pdf>

ACKNOWLEDGEMENTS

The authors would like to thank Audi AG for providing the data and permission to use the A6 Avant for this work.

CONTACT

For information regarding this paper, please contact s.weston@icon-cg.co.uk



www.icon-cg.co.uk

Icon is an independent provider of confidential CAE simulation services and has had a long association with industry by delivering CFD related projects at both OEM and supplier level since 1992. Close links with academic institutions, and continuous involvement with clients in various industries, has allowed Icon to establish a wide base of knowledge and offer technology transfer to its clients.



www.esteco-uk.com

modeFRONTIER is a multi-objective design optimisation software package developed by ESTECO. It is widely used in automotive and aerospace engineering to find, screen and select the optimal design in an organised fashion. modeFRONTIER can drive your FLUENT tools to explore the design space and iterate towards these improved and/or optimal designs.



www.beta-cae.gr

BETA CAE Systems S.A. is a private engineering software company specialised in the development of state of the art CAE pre- and post-processing software systems.

The company's flagship products, ANSA pre-processor and μETA post-processor package, currently hold a leading position in the CAE software market worldwide.