

USING OPENFOAM AND ANSA FOR ROAD AND RACE CAR CFD

Robert Lewis*, Andrew Mosedale, Ivor Annetts

TotalSim Ltd, UK

KEYWORDS –

aerodynamics, optimisation, RANS, DES

ABSTRACT –

Optimisation of the glass-house of a small car has been carried out in a parallel study of RANS and DES numerical methods. A response surface method based on a Kriging analysis using Latin hypercube sampling has been used to carry out the optimisation. The parameters to be optimised were front and rear screen rake, a-pillar angle and roof slope. The RANS-based optimisation worked well, although it was found that the baseline was already near minimum drag for the parameters. The roof slope was found to be the dominant factor in changing the simulated drag. The DES approach suggested a similar direction of optimisation to the RANS cases and a low-drag configuration was found. However the data was too noisy to effectively complete the optimisation, with the convergence being dominated by oscillations in the separation of the front wheel-wakes. This forced the simulations to be run for long periods to extract the underlying statistical average behaviour, making it more expensive. Further work is needed to understand how and when DES can be applied in such cases.

TECHNICAL PAPER -

1. INTRODUCTION

Traditional design optimisation has been a resource-intensive activity, with variable results. The rapid increase in computing power and growth of CFD methods has seen a radical change in the ability to control the model parameters using mesh deformation software and run many different iterations of a design, thereby increasing the potential for performance improvement. Advances in optimisation methods have also helped by reducing the number of runs required to accurately resolve the behaviour of a given design space.

Another advantage of the CFD optimisation approach is that it can be automated. Having created a parametrically 'morphable' mesh and set the design parameters, many runs can be carried out to test and refine optimal solutions in a relatively short space of time with little user input and limited by computational power. Automated optimisation is one of the current challenges in CFD now becoming possible with increases in resources.

Most industrial CFD uses Reynolds Averaged Navier-Stokes (RANS) solvers which are generally robust and reliable but have limitations, particularly in resolving separated, turbulent flow. An improvement on this is Detached Eddy Simulation (DES) where the flow away from the surface is hopefully better resolved. DES solvers are currently significantly more expensive than RANS methods, in part because they solve for transient flow and have to be run for more iterations, but are seen to be an increasingly necessary step in understanding flows of engineering interest.

This research aims to demonstrate using DES solvers as a basis for optimisation. A comparison is made with RANS-based optimisation for a generic small car geometry, both using response-surface methodology to determine the difference in the optimum predicted and the individual effect of each shape change parameter.

2. MODEL DETAILS

Creating the Mesh

The model used is that of a typical small car. It is a basic model with no engineering underbody detail. Figure 1 shows the baseline geometry tested. The objective of these simulations was to study the effect on the overall drag of the vehicle of altering the shape of the 'glass-house'. A water-tight STL object was quickly created from the original basic geometry representation using a combination of 3-matic and ANSA. This was used to generate the volume mesh using snappyHexMesh (OpenFOAM) which allows different refinement regions to be set. Figure 2 shows the detail captured in the surface mesh. Wake boxes were also defined to set a specific level of refinement around the car as shown in Figure 3. The volume mesh consists of approximately 4.8 million cells. The domain extends approximately 10 car lengths upstream of the vehicle and 20 downstream. It is 10 car lengths high and 5 on either side.

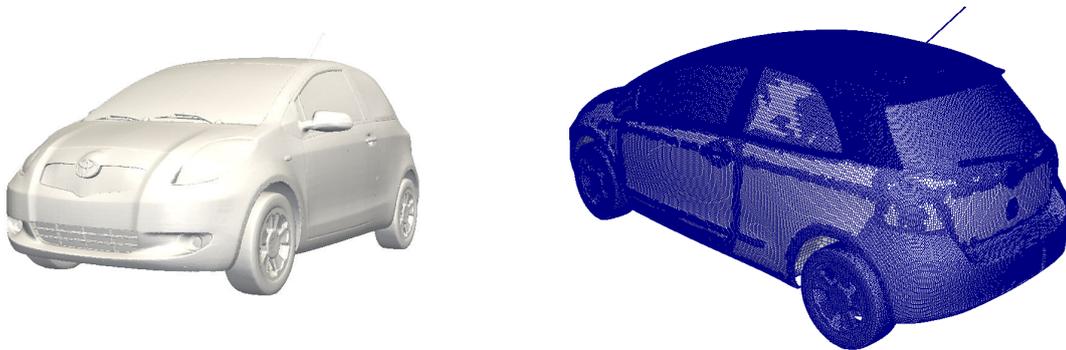


Figure 1 – Baseline geometry

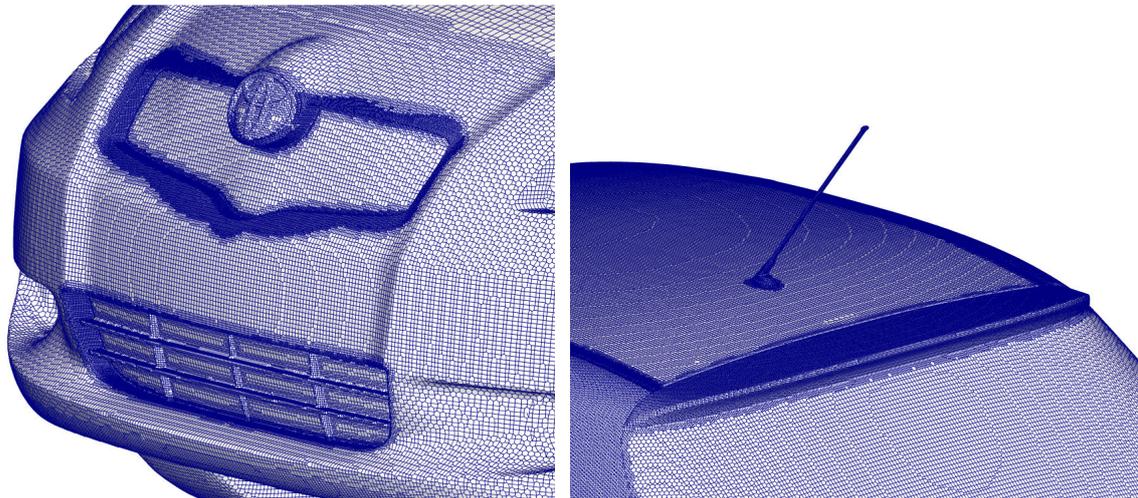


Figure 2 – Surface Mesh detail

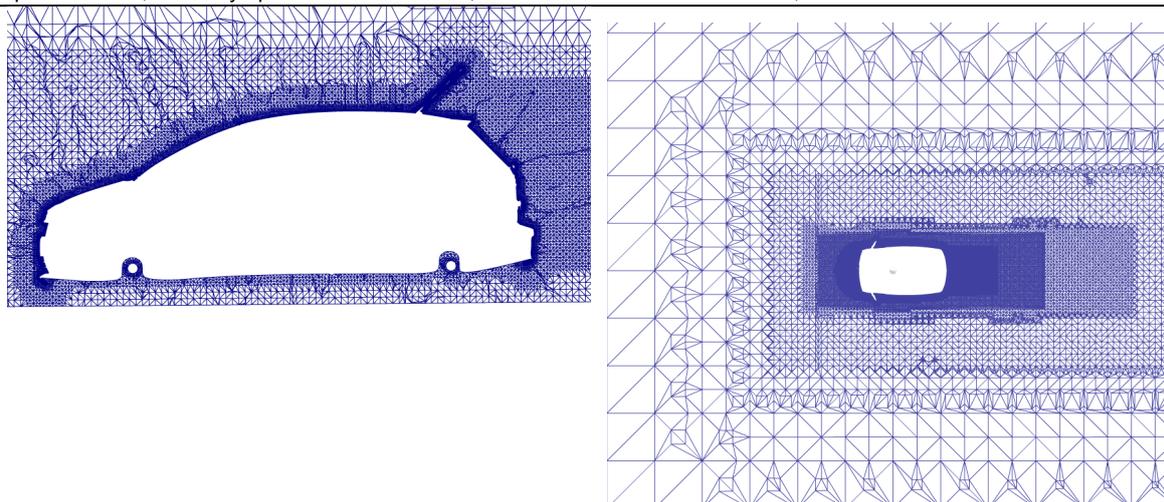


Figure 3 – Volume Mesh (Hex and Poly cells – appear triangulated in visualisation)

Case Setup

The flow was assumed to be isothermal and incompressible, using the SST k-omega turbulence model for RANS. The inlet flow was set to 26.82m/s (60mph) and at 1° yaw coming from the driver's right-hand side to simulate realistic driving conditions. Running with yaw is also important numerically when simulating a symmetric model as the flow structures can otherwise jump between alternate solutions. The rear and left boundaries were specified as pressure outlets. The ground plane was also moving at 26.82m/s and the wheels were rotating walls with a rotating reference frame around the spokes. The top boundary was set as a symmetry plane. The car body was defined as no-slip wall. The engine-bay was sealed to remove the need to simulate the under-hood flow and reduce model size.

Numerical Methods

TotalSim has extensive experience in solving with OpenFOAM and has a robust and efficient algorithm for setting up and solving automotive cases with a modified RANS solver. The flow field is initialised with a potential solution and then run using a pressure-limited RANS solver for 500 iterations. The case is then typically run unlimited for another 2500 steps, with forces being averaged over the last 500 iterations. The whole process to mesh and solve takes approximately 3 hours across 12 dual-core CPUs.

The output of the RANS solution is then used to start a DES simulation with a version of pimpleFOAM using the SpalartAllmarasDDES turbulence modelling option. In order to collect useful statistics from this unsteady solution the case is run with a maximum Courant number of 5 for a total of 1 second of simulated time, with forces averaged over the last 0.5s. The Courant number is selected based on previous experience to advance the solution as quickly as possible within the constraints of the time-step not exceeding the eddy turnover time. The simulation time varies due to the adaptive time-stepping, taking approximately 14 hours with 12 dual-core CPUs.

3. BASELINE RESULTS

The main result of interest is the average drag force on the car. Figure 4 shows the convergence history for the RANS case. Determining the period to average over requires some care, particularly with DES cases. The flow is assumed to be statistically steady, however there is significant oscillation in the forces as seen in Figure 5. The flow is allowed time to settle and then the average force determined over several cycles. Initial tests used the average data from 0.2-0.3s however that is clearly not representative of the flow over a longer period. Table 1 shows the forces for the RANS and DES baseline runs. Although the

two results are similar, the DES case suggests lower drag. This is most likely due to the models identifying slightly different points of separation.

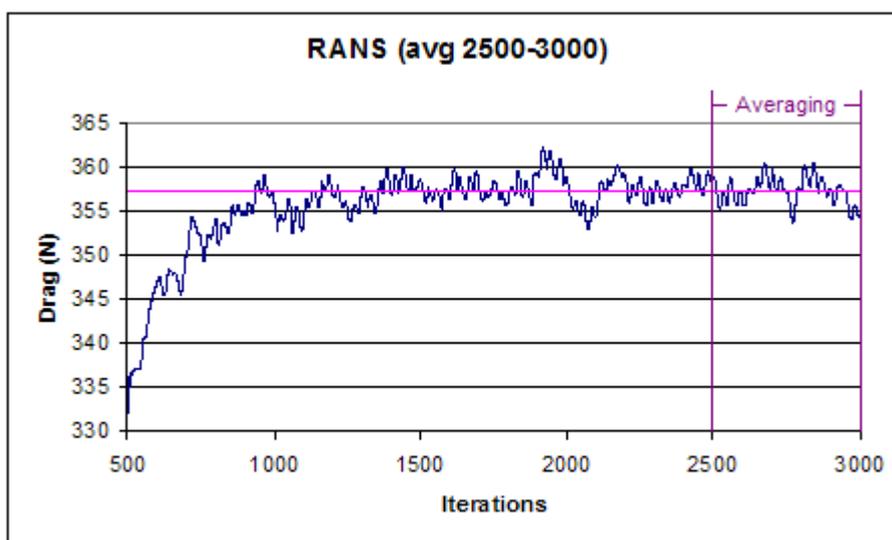


Figure 4 – Baseline drag force convergence for RANS simulation

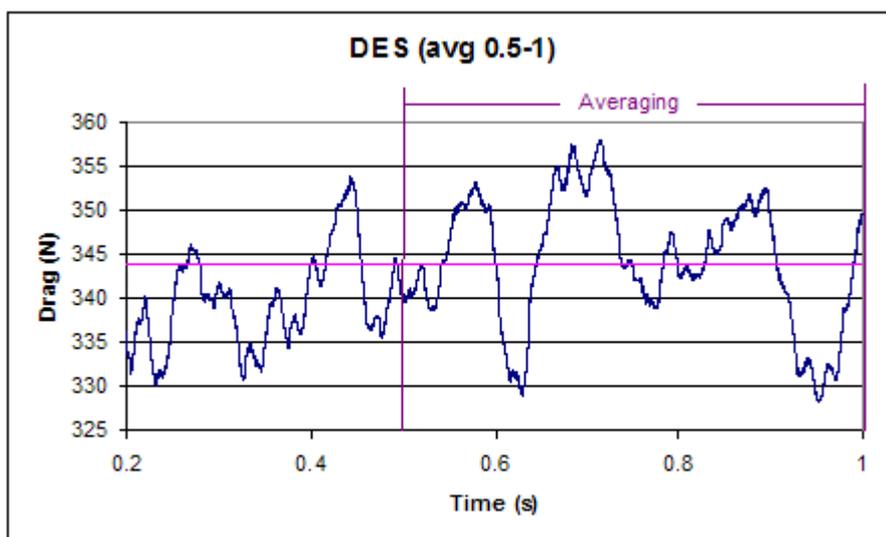


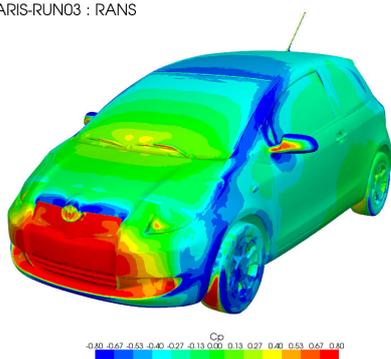
Figure 5 – Baseline drag force history in DES simulation

Method	Drag (N)
RANS	357.3
DES	344.2

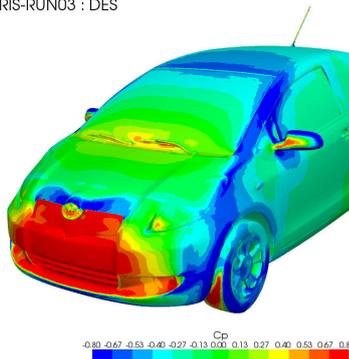
Table 1 – Baseline forces

Figure 6 shows the average pressure distribution for both methods. This reinforces the impression that both methods are essentially resolving the same flow features. It is perhaps more instructive to look at the total pressure losses in the wake of the car. The differences between the methods can perhaps be more clearly seen in Figures 7 and 8, with the turbulent flow behind the car being captured more realistically with the DES method. However, although it is possible to find qualifiable differences between the methods, it is not easy to relate that to the differences in drag.

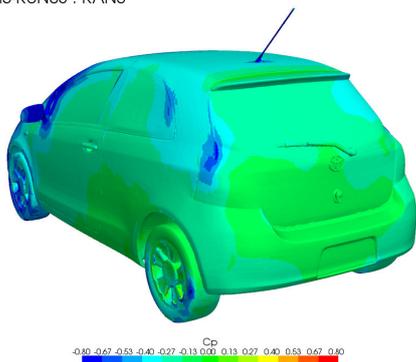
YARIS-RUN03 : RANS



YARIS-RUN03 : DES



YARIS-RUN03 : RANS



YARIS-RUN03 : DES

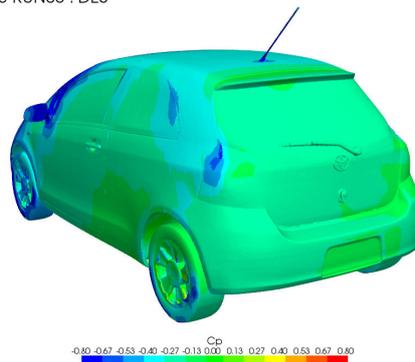
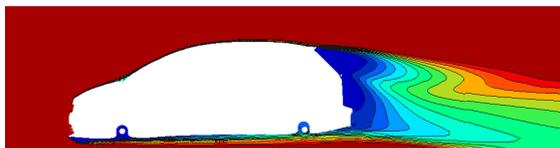
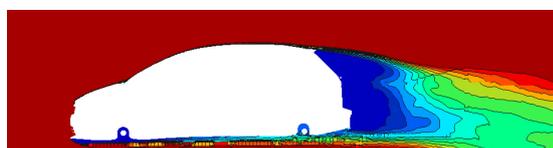


Figure 6 – Time-averaged C_p distribution for both numerical methods

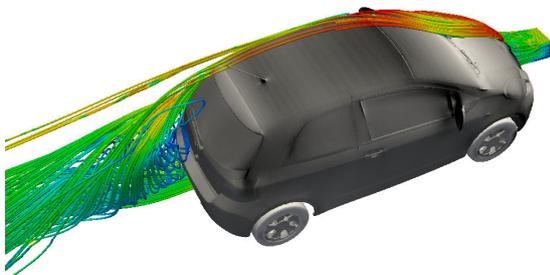


YARIS-RUN03 : RANS

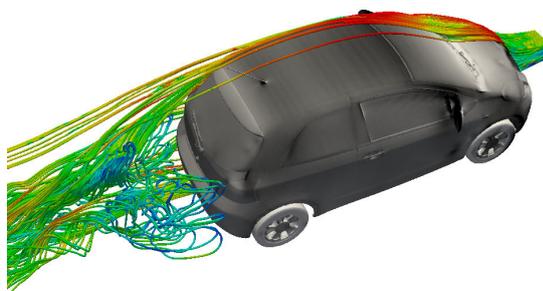


YARIS-RUN03 : DES

Figure 7 – Wake slices of time-averaged total pressure for the baseline runs



YARIS-RUN03 : RANS



YARIS-RUN03 : DES

Figure 8 – Streamlines based on the velocity vector

4. OPTIMISATION METHODOLOGY

Sculptor Setup

The optimisation of the glass-house requires parameter-controlled morphing of the geometry. This is done in *Sculptor*TM through the creation of a net as shown in Figure 9. This is similar to the parametric mesh deformation available in ANSA. The control points allow the model to be deformed with precision, specifically controlling the front and rear screen angles, roof slope and a-pillars. These four parameters were chosen as governing the aerodynamic properties of the glass-house geometry. They can be varied in combination, allowing a wide design space to experiment within.

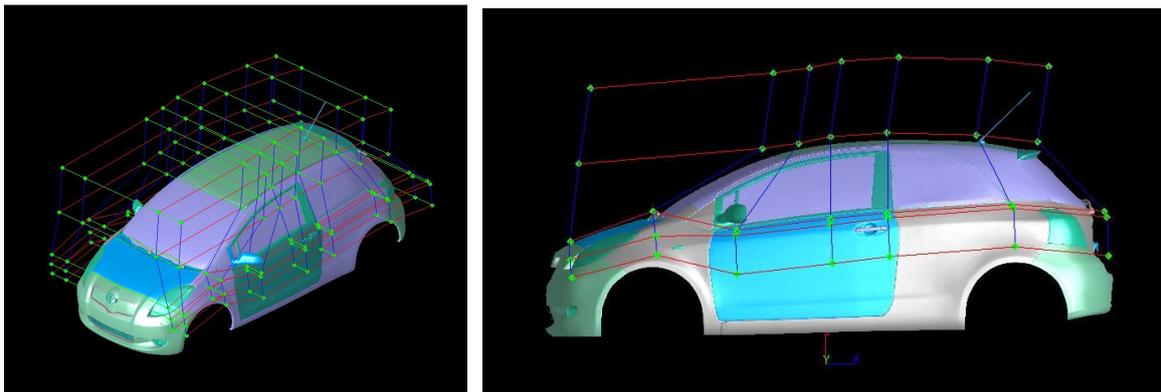


Figure 9 – Sculptor ASD volume

Typically it is desirable to allow the parameters to vary as much as possible; however there are non-aerodynamic constraints which need to be imposed. In this case the roof height and total internal volume needed to be considered. Optimising based on a single objective such as drag reduction could easily lead to impractical solutions, in this case most likely flattening the glass-house entirely. Therefore limits to the design space need to be imposed. Similarly care needs to be taken in setting the morphing parameters to isolate the desired effect. For this case it was decided that maximum roof height would remain fixed so as not to compromise the headroom in the cabin. Figures 10a and 10b show the range allowed for each parameter, defining the limits of the design space.

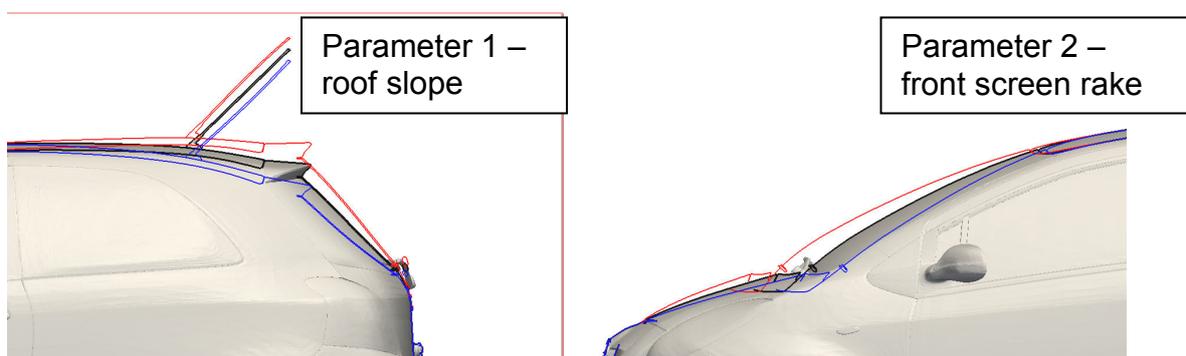


Figure 10a – Limits of design space

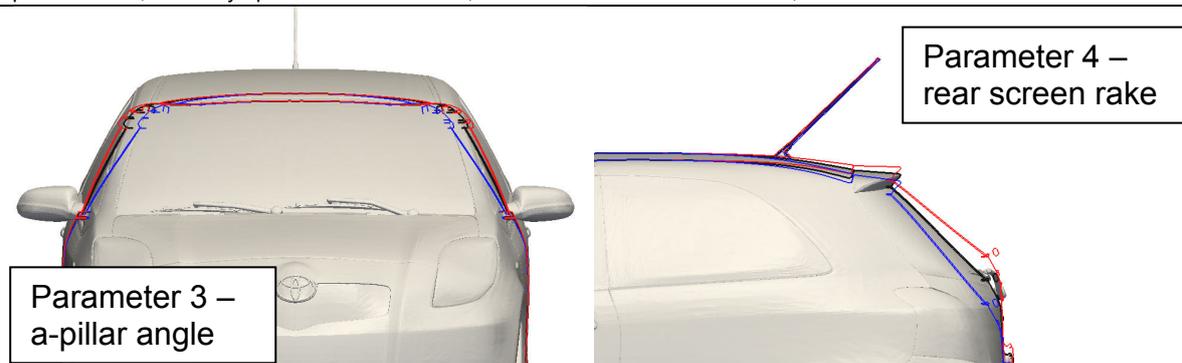


Figure 10b – Limits of design space

Design of Experiments Analysis

Having determined the parameters to be optimised, some method is needed to carry out the actual optimisation. As has been mentioned previously, the aim is to create a response surface. In order to construct the initial surface, data has to be obtained for sufficient sample points for the response to be reliable. Traditionally each parameter would be varied independently of the others, but this requires a large number of runs and is inefficient. Here, the sample runs are determined by fulfilling the requirements of an optimised Latin hypercube. This is to ensure the chosen number of runs properly represent the entire design space though the parameters are varying simultaneously which allows the maximum information to be obtained from the smallest sample.

For this set of four parameters, the initial Latin hypercube sampling consisted of 15 runs. The parameter combinations for each run are fed into the Sculptor batch process to generate the appropriately morphed surface. For consistency, each run is meshed to the same quality as the baseline run before the simulation is carried out. This whole process is automated having set up the case with the baseline run. The advantage of CFD is clear here. Although the DoE approach reduces the total number of runs, the process still requires significant resources that would prove impractical to conduct physically in a wind tunnel. As before with the baseline case; having generated the results using the RANS method, DES is run for each case. This provides two sets of drag results for each sample point. Figure 11 compares the measured drag for both methods on each run. The DES method consistently reports lower drag, but approximates the same trend as the RANS results. It is also worth noting that the RANS result for the baseline run already gives the second lowest drag figure, suggesting there may not be much room for improvement – possibly because the original car was optimised with a RANS method when it was designed.

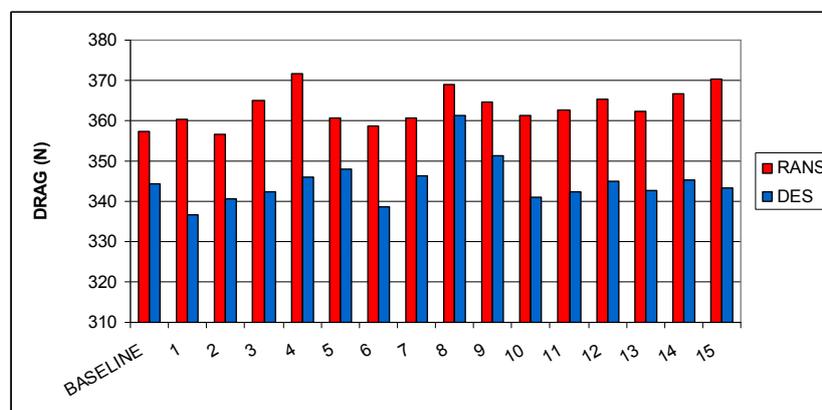


Figure 11 – Drag comparison between RANS and DES for initial experiments

Response Surface Modelling

A Kriging analysis (1) is carried out on the data which generates a response surface which can be searched for minima. The response surface can be used as a surrogate model, once it has been trained with the simulation data. However it is advisable in most cases to simulate the identified optimal parameters to see how the results compare with the predicted minimum drag.

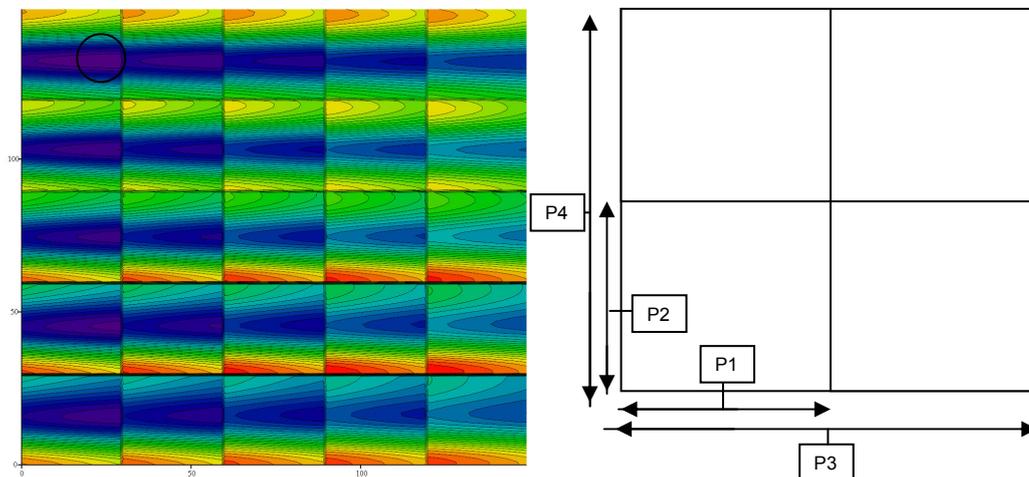


Figure 12 – RANS response surface

The response surface based on the fifteen RANS samples and baseline result is shown in Figure 12. It is only a representation of the four-dimensional surface however it is possible to see the located minima (purple) to be tested. The analysis also affords a lot of information on the sensitivity of the drag to each parameter which can be useful in refining the study. There is a clear dominant parameter in this figure – the roof slope parameter, represented on the y-axis of the sub-cells. Figure 13 shows the response surface generated using the DES data, and demonstrates the effect of allowing the simulation to run longer. The analysis can cope with a certain amount of noise, but the underlying data has to be representative for a sensible result. In this case, the initial test runs were not averaged over a long enough period and the time-step was too high for the flow to converge to a statistically steady oscillation, resulting in a poor response surface. Running the DES cases for a longer period with a smaller time-step limit gives a better response; which has a different view of the significance of each parameter to the RANS-derived results.

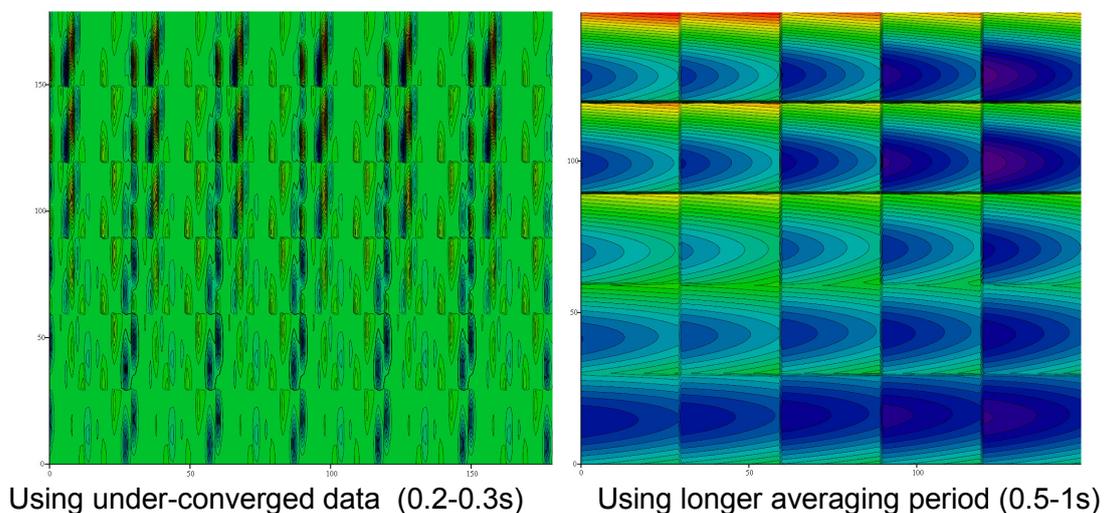


Figure 13 – DES response surface

The next step is to refine the response surfaces. It is possible to expand the initial DoE sampling to obtain more data points, however it is advantageous to use parameters in the region of the minima as identified in the first optimisation sweep. This effectively guides the optimisation process to refine the response surface in the region of interest. Each method identified two minima to explore. These sets of parameters were run and the results are shown in Table 2.

RUN	P1	P2	P3	P4	Predicted Drag (N)	Method	Simulated Drag (N)
RANS 1 st Minimum	171.6	-25.0	-200.0	130.1	355.6	RANS	359.3
						DES	341.1
RANS 2 nd Minimum	311.4	7.1	-200.0	-77.5	355.7	RANS	361.8
						DES	342.1
DES 1 st Minimum	-500.0	33.0	500.0	-150.0	334.9	RANS	366.1
						DES	344.2
DES 2 nd Minimum	-500.0	-38.9	500.0	130.1	339.6	RANS	359.9
						DES	336.7

Table 2 – Forces from first optimisation sweep with parameter values

The results for both methods are consistent in which minimum is more effective for each optimisation. The first minimum listed from the RANS optimisation and the second minimum from the DES study most closely match the expected minimum drag. These minima have very similar slope (P2) and rear screen (P4) angles, the parameters governing the shape of the rear of the car, but interestingly they point in different directions for the front of the car (P1 and P3). It is also notable that the DES method finds a minimum which holds in RANS space, whereas the RANS-derived minimum does not perform particularly well in the DES solution. Adding these results helps refine the response surface in the regions of interest, as shown in Figure 14. There is no drastic shift from the previous results but the additional data does help identify the specific configuration needed to attain minimum drag. A new minimum was found from each of these surfaces and run, with the results adding another level of refinement.

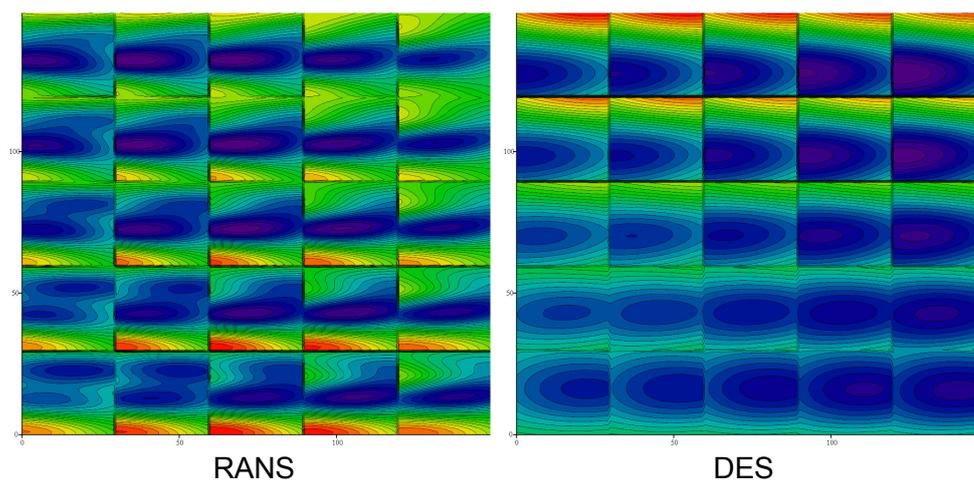


Figure 14 – Refined response surfaces from the first optimisation sweep

Converging to an optimum solution

Determining when a response surface has been sufficiently refined depends on a number of factors which are often not well defined. Figure 15 shows the evolution of the response surface for each method through each sweep of the optimisation. The initial data comes

from the baseline run and fifteen DoE sample runs. The second iteration includes four more runs, two based on the initial RANS data and two based on the initial DES data. The third iteration adds one run derived from each method and the fourth adds the final ‘optimised’ run for each method.

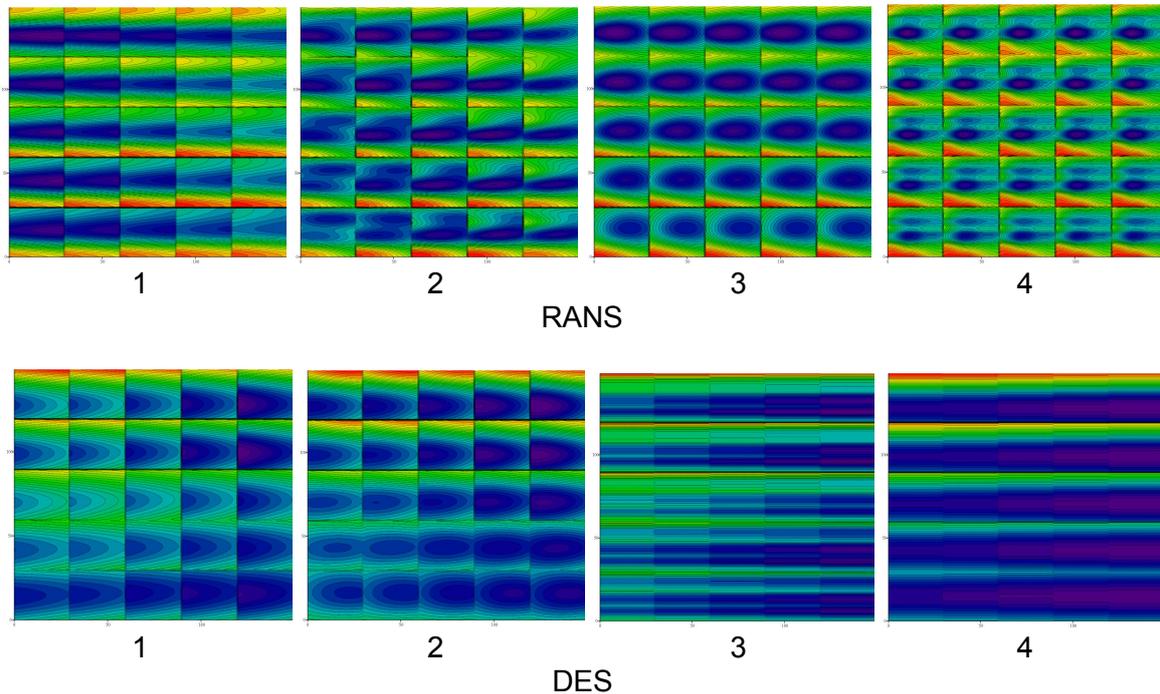


Figure 15 – Response Surface Refinement

While the RANS-based response surface refinement appears well-behaved, the DES-based surface does not. The analysis identifies any parameters to which the drag response is insensitive, resulting in zero variation on one of the axes. This is to be expected as a priori knowledge of the relative influence of each parameter on the resulting flow would negate the need for costly simulations. However, the inability to determine a response for a given parameter may also indicate further refinement is needed. The figures for sensitivity, normalised by the most sensitive parameter at each point, are shown in Figures 16a and 16b. This reveals something not immediately obvious from the plots of the response surface – the RANS optimisation has also found an insensitive parameter. It is also intriguing to note that the parameter identified as insensitive by the DES runs is the second-most important parameter in the RANS optimisation – the front screen rake.

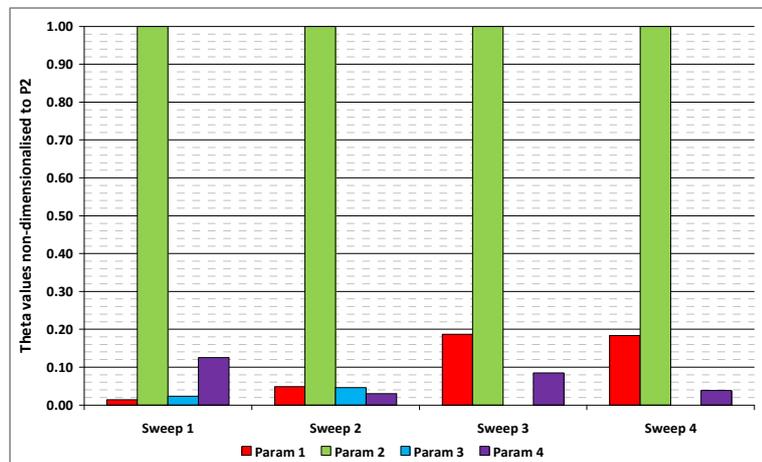


Figure 16a – Parameter sensitivity in RANS

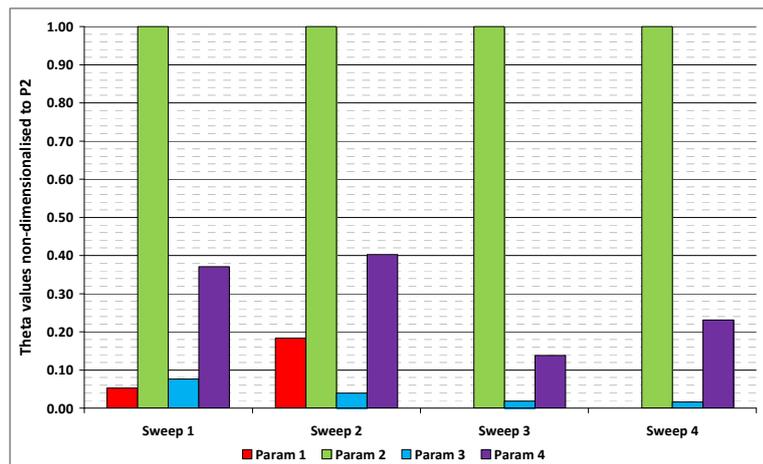


Figure 16b – Parameter sensitivity in DES

Another point to note is the degree to which the roof slope parameter (P2) dominates the response; this may be a factor in swamping the measurable effect of the other parameters accurately. The optimisation method assumes that the response of the system varies predictably and can be measured accurately but this is often not the case. The analysis gives a noise factor which indicates how well the data fits the predicted surface. As can be seen in Figure 17, changing the data set can lead the analysis to interpret some of the data as noise in generating a 'best fit' response surface. This applies to both methods, although the DES case is arguably worse in this respect which is consistent with the uncertainties in determining the average forces. Perhaps the main conclusion to draw is that interpreting the results of these methods is not straight-forward.

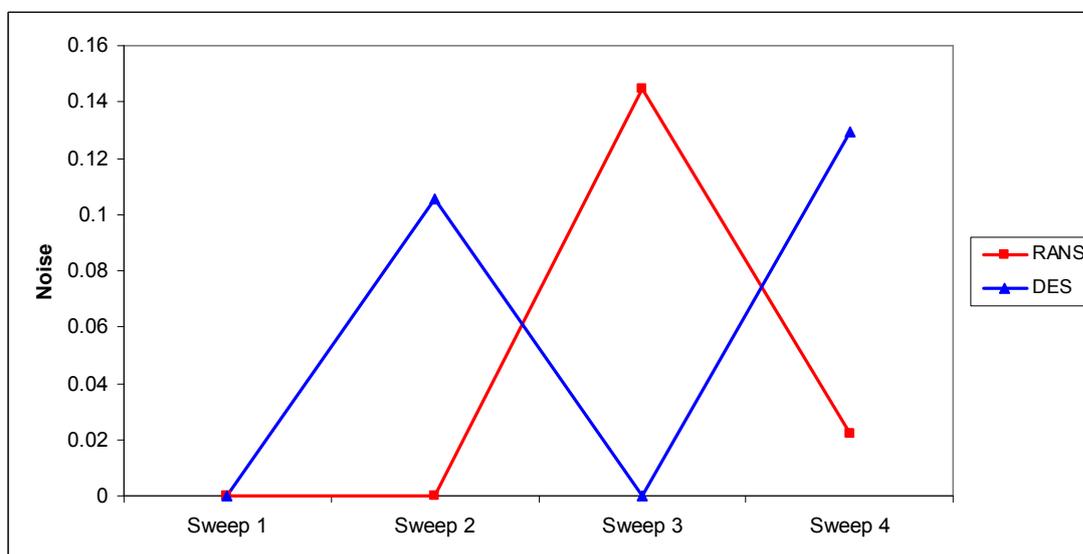


Figure 17 – Evolution of calculated noise

While these tools can be very powerful in understanding the response behaviour of the system, in practical terms the aim of the refinement is to locate the optimum configuration. Often this needs only be defined as an approximate region of the design space, with appropriate appreciation of the relative sensitivity and uncertainty of each parameter. The most important factor therefore is the object of the optimisation, in this case the drag. Figure 18 shows how the simulated drag varies with each run compared to baseline.

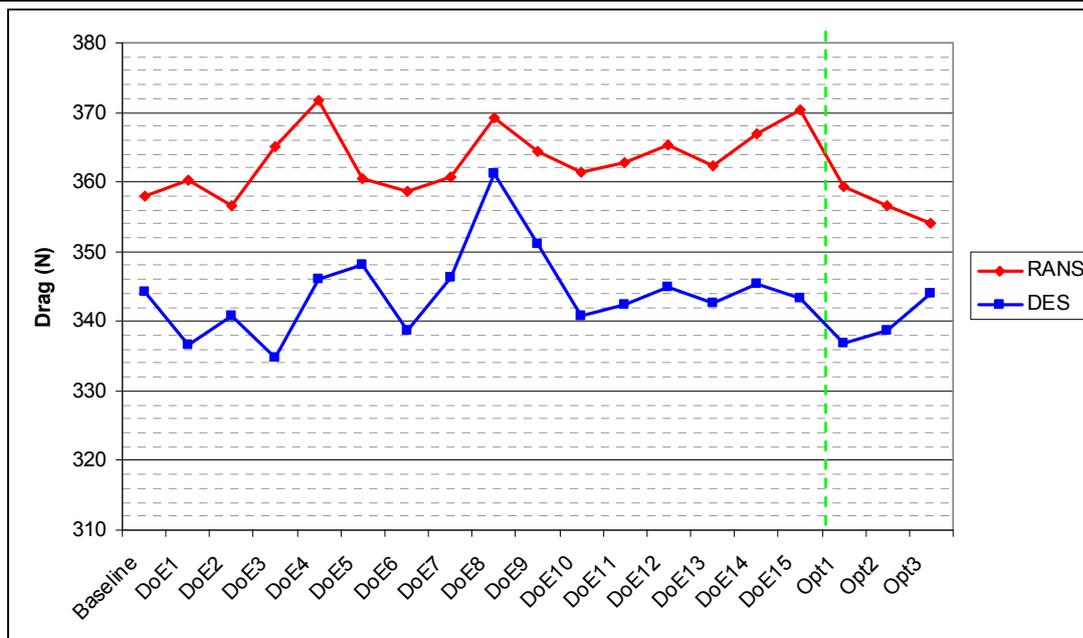


Figure 18 – Simulated drag evolution for each optimisation

The RANS refinement shows a clear trend of decreasing drag in the final three points – the runs identified as minima from the preceding data. This is the expected behaviour, and delivers a 4N drag reduction on the baseline case. The DES trend by contrast does not improve with refinement, although all three optimum runs are below the baseline. This is not surprising given the factors discussed previously and considering that the variation seen is within the standard deviation of the drag during the course of each simulation. This should not affect the overall direction of the optimisation, and indeed the response surface is largely valid and further refinement runs in the region of the minimum would likely improve the trend. The limitation is more in relying on the measured data from the single final run as proof of the optimisation. This is a good example of the value of understanding the detail of the analysis so as to be able to determine how well-posed the initial problem is.

Design Compromises

As has been mentioned previously, there is usually more than one competing objective in design. One such consideration is the trade-off in cabin volume caused by the drag optimisation. The response surface methodology provides a framework to study such issues by comparing the response for different variables when probing the same point in the design space. This is illustrated in Figure 19, with a Pareto plot showing the relationship between internal volume and drag. This is a projection of the response surface which delineates the feasible design space in terms of the desired objective functions, providing a visual representation of the trade-off between competing objectives. In this case it suggests that the lowest drag configurations do not require an overall reduction in volume from the baseline. In fact, there is room to increase the internal volume of the car while reducing drag. In practise this is an invaluable tool for selecting the best compromise, and sometimes for finding where there is room for gains in all areas.

Another use of the Pareto plot is to compare the performance of different models. It has been seen that the RANS and DES methods do not always agree on what a low-drag configuration is. Figure 20 shows the joint behaviour of the methods. Most of the runs from the DoE inhabit the middle of the space, with the optimised runs clustered in the low-drag corner. However the extent of the viable space is now revealed, along with interesting features. One point to note is that a low-drag run in RANS will also be low-drag in the DES simulation, but the converse is not true. A prudent designer would use this plot to select a design which was robust to the method used and gave low drag for both by selecting a point

in the bottom-left corner. Again, this is all that is needed in practical terms when a solution exists that satisfies or exceeds all given criteria simultaneously.

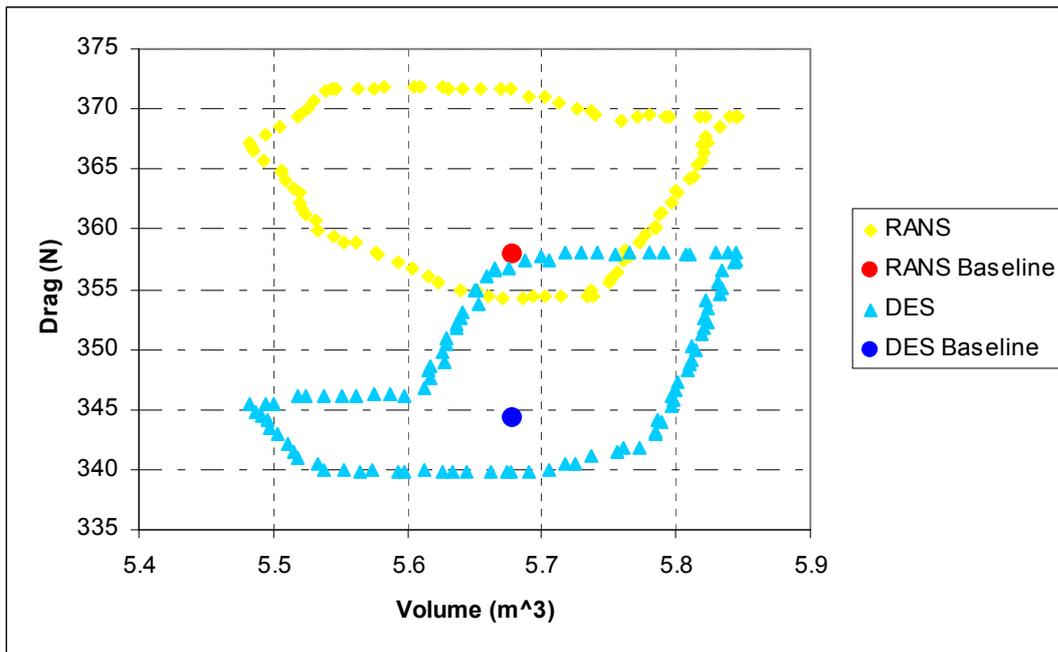


Figure 19 – Pareto efficiency of drag versus cabin volume

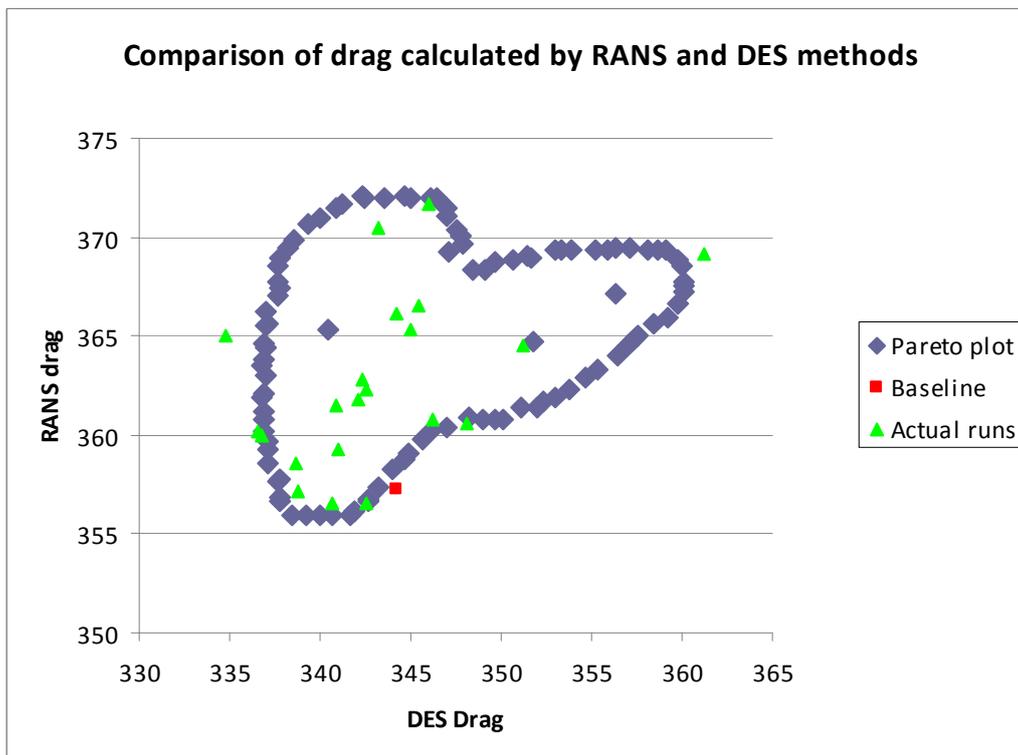


Figure 20 – Calculated drag method comparison

One final point of interest is the differing locations in the design space of the maximum drag for each method. While maximising drag is rarely a desirable objective, it is valid in terms of

the method comparison to consider what differences there are in the resolved flow field for these configurations. Looking at the flow from one of the optimisation runs located in this area of the design space offers some explanation for the difference. Figure 21 shows an oilflow-type representation of the surface flow for each method for OPT-RUN4. The rear of the car is where most differences are manifested; in this case the RANS solution is showing clear vortical flow which induces drag. The DES solution does not contain the same coherent structures which would account for the differences in the response surfaces. In terms of the design space, this run represents a roof with significant slope down towards the rear of the car. This appears to be feeding a large vortex in the RANS solution, effectively nullifying the spoiler. It would appear the more accurate eddy capturing of the DES method causes this vortex to break down into turbulent flow, which is a preferable regime in this region of the car. This is consistent with the wake slices seen for the baseline runs in Figure 6, with the DES solution being notably less regular and structured in the wake.

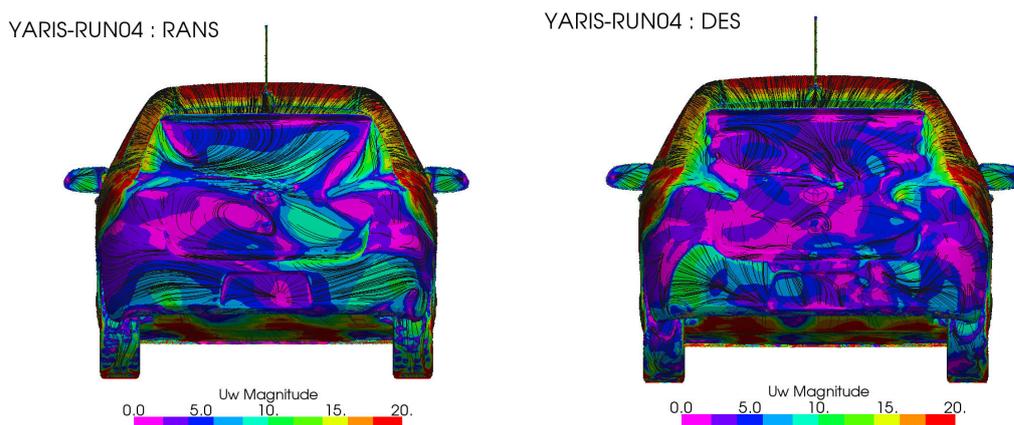


Figure 21 – Near-wall velocity for a high-drag run

Looking at the instantaneous wake for this high-drag run demonstrates the different nature of the methods. Figure 22 shows the more physical flow captured by the DES solution.

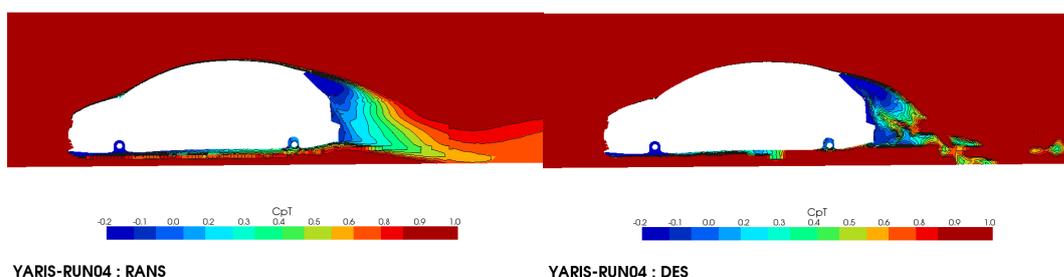


Figure 22 – Instantaneous wake flow for a high-drag run

5. OPTIMISED RESULTS

This final section will look at what has changed physically from the baseline to the optimum configuration to better understand how changing the parameters affects the flow. The lowest drag case for each method will be used. The geometry for these runs is shown in Figure 21. There are significant similarities in the shape of the car. The front screen angle is changed by the same amount from both RANS and DES optimisation. Also, although the combination of rear screen and roof slope parameters gives a different rear end, crucially the effective rake of the rear is the very similar for both methods compared to the baseline position.

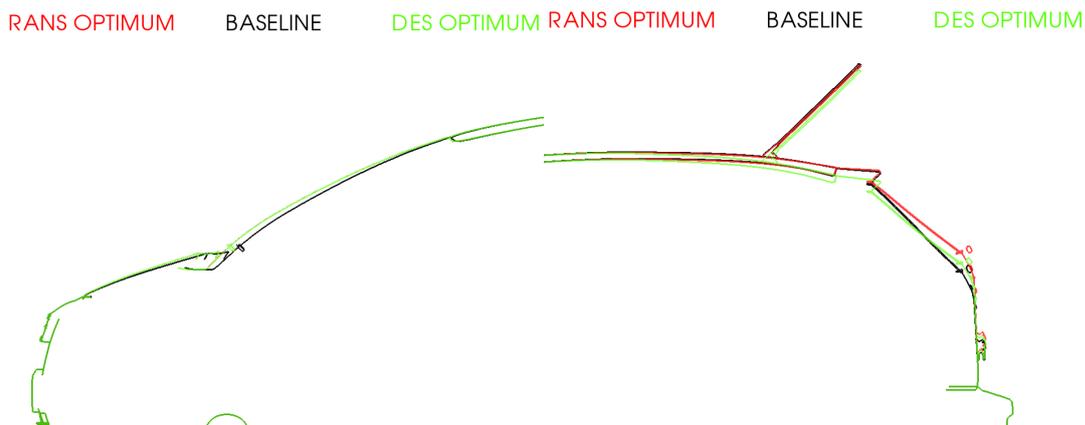


Figure 21 – Geometry changes in optimum configurations

	RANS	DES
Baseline	357.3	344.2
Optimum	353.9	334.7

Table 3 – Drag reduction with optimum runs over baseline for each method

The calculated drag figures for these runs are shown in Table 3. The gain is relatively small, making it difficult to determine the changes in flow responsible. Figure 22 shows how the drag accumulates over the length of the body.

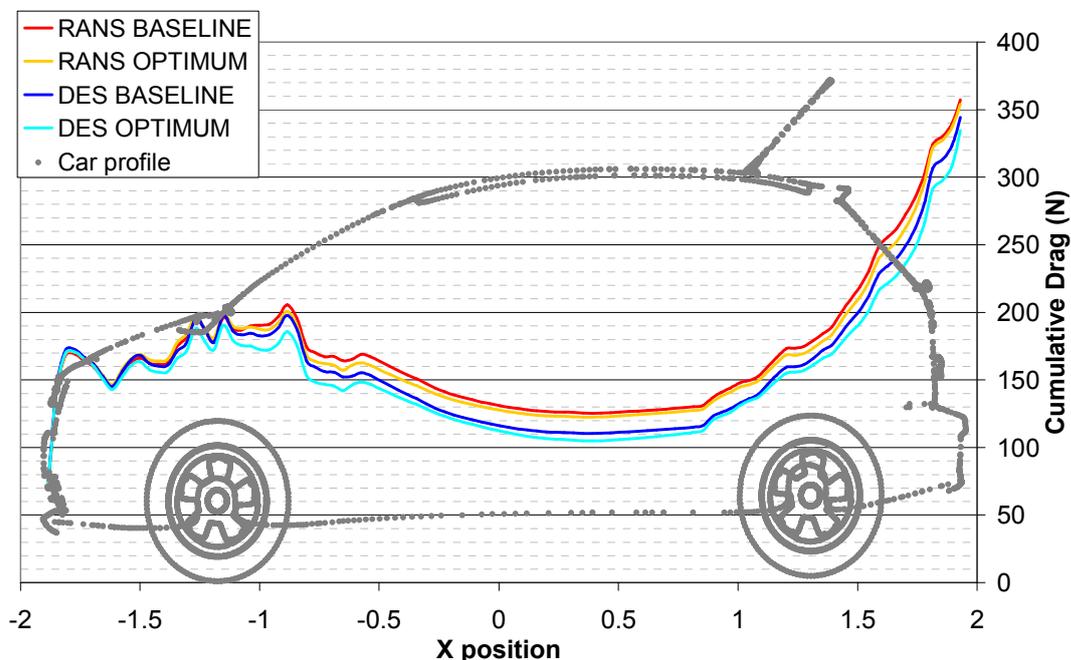


Figure 22 – Drag accumulation over the length of the car

There are several things to note about this plot. The first is that the differences in drag manifest themselves near the front of the car. For the RANS runs, the drag reduction due to the change in front-screen angle in the optimum is clear from the gap between the lines in the figure, however these gains are lost by the time the flow reaches the middle of the roof.

Figure 23 shows the pressure distribution, with a smaller stagnation region at the base of the front screen. The gap opens up again over the rear part of the roof where the slope angle has changed, as expected, but the gains are once more lost as the flow negotiates the rear bumper, leaving only the small net gain. A similar analysis can be made for the DES runs. This suggests that more parameters may be needed to allow the optimisation process to find a sympathetic shape for the very rear of the car to maintain the gains found further forward.

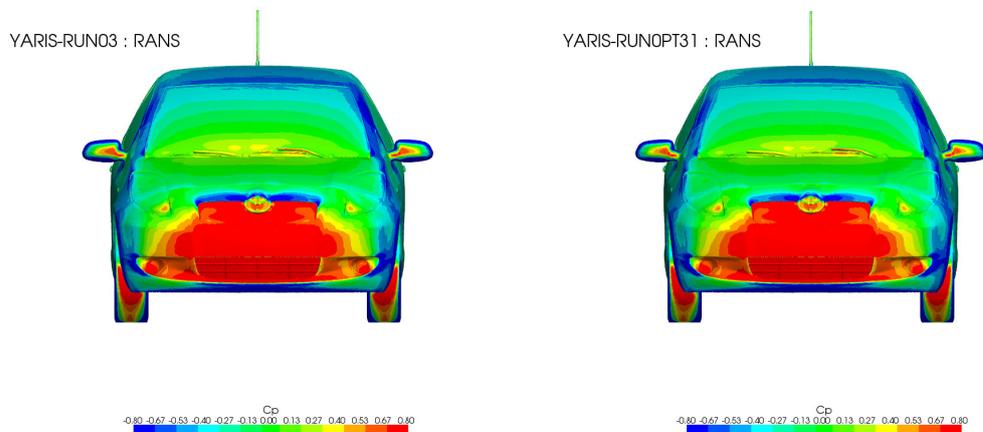


Figure 23 – Affect on flow stagnation of changing front screen angle in the Opt run

Careful consideration of where the differences in drag occur also helps in understanding the differences between the RANS and DES methods. Figure 24 shows the instantaneous velocity of the flow in a slice through the wheels. Wheel wakes can often be a significant factor in overall drag and some of the differences seen in Figure 22 can be attributed to the front wheels.

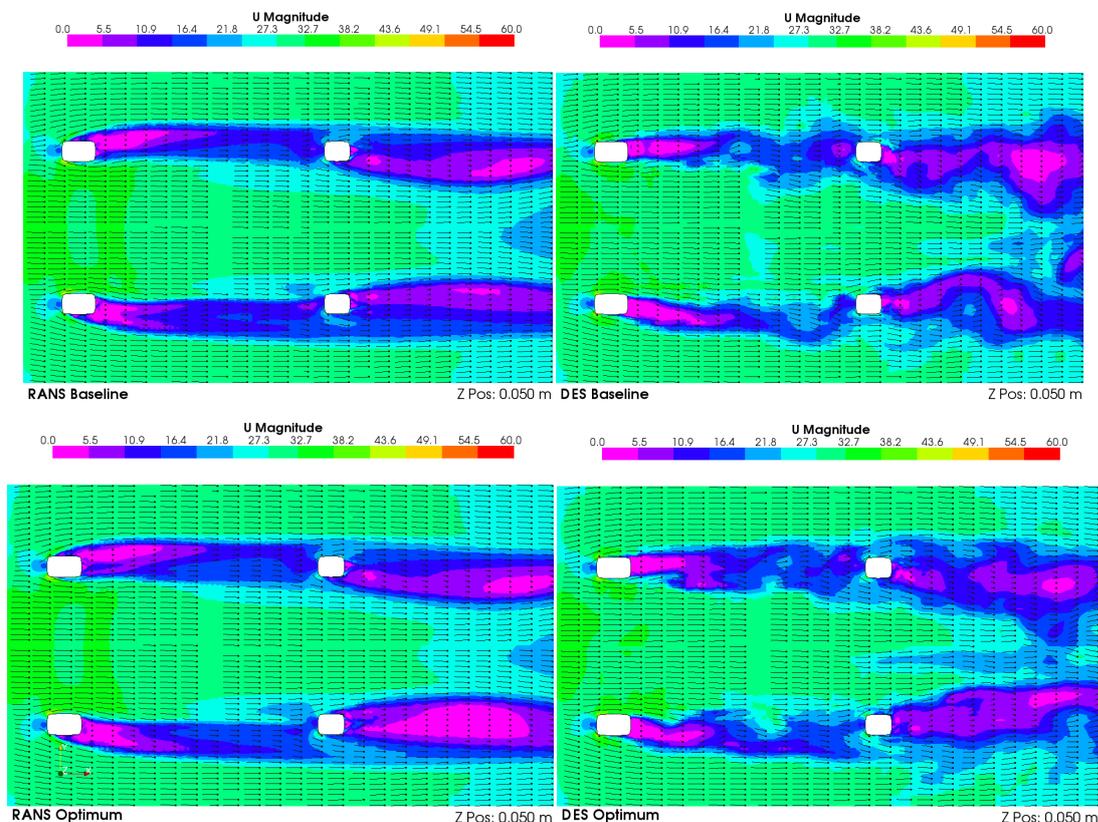


Figure 24 – Near ground flow velocity showing wheel wakes

As expected, the flow appears quite different for the two methods, and there are also differences in the optimum runs despite the changes to the geometry only being made to the top of the car. The most significant feature is the separation point from the front wheels. The RANS cases separate from the front of the tyre, leading to sizeable outboard wakes. By contrast the DES cases remain attached longer, with smaller wakes as a result. This would explain the discrepancy in total drag between the methods. It should be noted that this is an instantaneous view; however it is reasonable to assume that the DES case is more likely to deviate over time, and in fact it may be the oscillation of this or another separation point which dominates the convergence and generates the noise in the analysis. More work needs to be done in assessing the influence of the wheels on the overall flow.

6. CONCLUSIONS

In conclusion, a basic model of a small car has been simulated and optimised with both RANS and DES numerical methods, based on parametric morphing of the glass-house and using response surface methodology.

- The optimisation process works well with RANS-derived data, finding a 1% reduction in drag.
- The roof slope is the most dominant parameter in the study, further drag reductions could be found by replacing the insensitive parameters.
- The DES approach requires significantly more run-time to gather statistical information resulting in a 500% increase in computing time over the RANS method. The DES solutions reported lower drag than the RANS for each run, most likely due to the separation point on the front wheels.
- Reducing drag in both RANS and DES is possible and does not need to compromise the overall cabin volume; however the region of maximum drag in the design space is different for each method.
- DES can be used for optimisation but involves greater levels of uncertainty than RANS.

REFERENCES

- (1) Jones, D. R., A Taxonomy of Global Optimization Methods Based on Response Surfaces, J. Glob. Opt. **21** p345-381, 2001

ACKNOWLEDGMENTS

OpenFOAM is produced by OpenCFD Ltd. WWW.OPENCDF.CO.UK OpenFOAM is a registered trademark of OpenCFD.
