

2009-01-0333

Application of Detached–Eddy Simulation for Automotive Aerodynamics Development

M. Islam and F. Decker
Audi AG, Wind–Tunnel Centre

E. de Villiers, A. Jackson and J. Gines
ICON Ltd.

T. Grahs and A. Gitt–Gehrke
Volkswagen AG

J. Comas i Font
SEAT S.A.

Copyright © 2009 SAE International

ABSTRACT

This paper presents a complete methodology for performing finite–volume–based detached–eddy simulation for the prediction of aerodynamic forces and detailed flow structures of passenger vehicles developed using the open–source CFD toolbox OpenFOAM®. The main components of the methodology consist of an automatic mesh generator, a setup and initialisation utility, a DES flow solver and analysis and post–processing routines.

Validation of the predictions is done on the basis of detailed comparisons to experimental wind–tunnel data. Results for lift and drag are found to compare favourably to the experiments, with some moderate discrepancies in predicted rear lift. Point surface–pressure measurements, oil–streak images and maps of total pressure in the flow field demonstrate the approach’s capabilities to predict the fine detail of complex flow regimes found in automotive aerodynamics.

Standard DES methods can cost an order of magnitude more than traditional methods, but optimisation and automation of mesh generation, setup and solution algorithms ensure quick turn–around times. Due to the fully parallel nature of these components, the entire process can be executed in a distributed fashion. Efficient solution algorithms provide exceptional accuracy when compared to Reynolds–averaged approaches without sacrificing stability, even when the flow exhibits high Courant numbers.

The proposed methodology is highly customisable, which allows for targeted developments to suit the individual needs of aerodynamics CFD. On the basis of the results presented here, the methodology is found to be appropriate and suitable for use in the industrial development process.

INTRODUCTION

The use of numerical simulations to predict the aerodynamic characteristics of road vehicles is now a standard practice in automotive development [1]. The current state–of–the–art makes use of either finite–volume solvers based on the Reynolds–averaged Navier–Stokes (RANS) equations or lattice–Boltzmann solvers, both in the context of commercial, proprietary computational fluid dynamics (CFD) software packages. Recent developments in the field of automotive aerodynamics have generated the need for accelerated development of the application of CFD therein. First, the steady increase in the number of different vehicle models brought to market has not been matched by a commensurate increase in wind–tunnel capacity, necessitating the use of alternatives to experimental testing, preferably CFD. Second, the continual improvement of vehicle aerodynamics, both in terms of drag coefficient as well as lift, requires the tools of aerodynamics development to perform at ever–increasing levels of accuracy; this applies to both wind–tunnel technology and CFD. Third, CFD methods must be able to keep pace with continually shortening development cycles, placing

The Engineering Meetings Board has approved this paper for publication. It has successfully completed SAE’s peer review process under the supervision of the session organizer. This process requires a minimum of three (3) reviews by industry experts.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of SAE.


ISSN 0148-7191

Positions and opinions advanced in this paper are those of the author(s) and not necessarily those of SAE. The author is solely responsible for the content of the paper.

SAE Customer Service: Tel: 877-606-7323 (inside USA and Canada)
Tel: 724-776-4970 (outside USA)
Fax: 724-776-0790
Email: CustomerService@sae.org

SAE Web Address: <http://www.sae.org>

Printed in USA



9-2009-01-0333

SAE International™

continually increasing demands on computational efficiency and robustness. Fourth, these increasing demands require the CFD software to be specifically tailored to the needs of the application, and therefore to contain considerable know-how derived from the application, as well as to be usable flexibly on a large scale.

In the present paper, we present a complete methodology for carrying out automotive aerodynamics CFD that addresses precisely these issues. The methodology is based on the open-source CFD toolbox OpenFOAM® [2] and is applied in an industrial context to a wide range of vehicles from the brands of the Volkswagen Group Audi, Volkswagen and SEAT. High numerical accuracy is achieved by applying detached-eddy simulations, as opposed to simulations based on the RANS or lattice-Boltzmann equations, which, in an industrial context, use $k-\epsilon$ or similar models for the integral length scales of turbulence. DES produces a much more accurate and detailed representation of the flow when compared to traditional simulation methods. Calculation costs are however significantly higher, such that the principal challenge in making such an approach viable is enhancing the efficiency of the solution process. Evidence indicates that simulations using RANS equations with $k-\epsilon$ turbulence modelling and steady-state solution algorithms can fail to capture all time and length scales important to vehicle aerodynamics [3]. Rapid turn-around times are achieved by full parallelisation of the complete simulation process, including mesh generation.

The method and results presented here form a summary of a focused, multi-year effort to develop and validate this particular technology, in which over 150 CFD simulations of vehicle aerodynamics were carried out on a palette of more than 20 different vehicles in order to optimise accuracy, performance and robustness.

The paper is structured as follows: In the next section, we present the numerical method used here, which is based on the finite-volume method for CFD. Subsequently, the method by which the computational mesh is generated is presented and discussed. The following section comprises an extensive validation study, in which the numerical simulations are compared to experimental data obtained in automotive wind tunnels used by the group. On the basis of these results, an assessment of the methodology presented here is made, followed by a discussion of intended future work, a summary and some general conclusions.

NUMERICAL METHOD

The basis of large-eddy simulation (LES) is that the large scales of motion, which contain most of the energy, do most of the transporting and are affected the strongest by the boundary conditions are calculated directly, while the small scales are represented by a model. In finite-volume practice, this separation of scales is achieved by applying a spatial averaging filter to the Navier-Stokes equa-

tions [4],

$$\nabla \cdot \bar{\mathbf{u}} = 0 \quad (1)$$

$$\frac{\partial \bar{\mathbf{u}}}{\partial t} + \nabla \cdot \bar{\mathbf{u}} \bar{\mathbf{u}} = -\frac{1}{\rho} \nabla \bar{p} + \nabla \cdot \nu (\nabla \bar{\mathbf{u}} + \nabla \bar{\mathbf{u}}^T) + \boldsymbol{\tau} \quad (2)$$

where $\bar{\mathbf{u}}$ and $\bar{p}$ represent the filtered velocity and pressure, respectively, and $\boldsymbol{\tau}$ is the sub-grid scale (SGS) stress that results from the filtering of non-linear terms and has to be modelled. The SGS stresses represent a much smaller part of the turbulent energy spectrum than the RANS turbulent energy, so that the accuracy of the stress model may be less critical to the accuracy of the overall solution than in RANS computations.

For LES to be considered accurate, some highly limiting restrictions have to be adhered to with respect to the resolution of the mesh near solid boundaries. Specifically, the near-wall resolution must be of the order of $y^+ = 1$ with other dimensions similarly small. These restrictions make LES impractical on aerodynamic geometries due to excessively large mesh sizes required. The detached-eddy simulation (DES) approach largely circumvents this problem by sacrificing some of LES' inherent accuracy in the near-wall region. DES refers to an approach whereby unsteady RANS turbulence modelling and mesh spacing is used in the boundary layer, while LES is employed in the core and separated regions of the flow. In the near-wall regions, the RANS turbulence model, which has been calibrated in thin shear-layer flows, has complete control over the solution. In the LES region, the turbulence model changes to an SGS formulation. DES can be considered a good approximation, as long as the scale of the eddies in the boundary layer is much smaller than the eddies in the bulk flow. This situation is almost always the case for external aerodynamic flows.

For the the current investigation, the DES formulation proposed by Spalart is used [5]. The model is based on the Spalart-Allmaras (S-A) one-equation eddy-viscosity model [6] normally used in RANS calculations. To accomplish the transition from near-wall RANS-based simulation to LES treatment of the interior flow in one formulation, the near-wall distance y_w is replaced by $\tilde{d}$, defined as

$$\tilde{d} = \min(y_w, C_{DES} \Delta), \quad (3)$$

which acts as the S-A RANS model for $y_w \ll \Delta$ and a SGS model for $\Delta \ll y_w$. C_{DES} is an empirical constant calibrated to a value of 0.65 using the known decay rate of isotropic turbulence [7]. The current investigation uses a version of the standard S-A DES model that has been modified to allow for its use in conjunction with wall functions and polyhedral meshes [8]. Here, Δ is defined as the cube root of the cell volume, not the maximum cell dimension more traditionally used.

To minimise numerical dissipation, which can easily be more diffusive than the SGS model, second-order accurate energy-conserving numerical schemes are used [9, 10]:

- fully implicit backward differencing in time (Gear's method);
- non-oscillatory minimally blended central-differencing scheme for momentum convection;
- a second-order TVD/NVD scheme for SGS viscosity convection;
- explicit non-orthogonality correction.

To maintain overall stability and efficiency, an addition provision has been added to reduce the order of accuracy locally in regions of high Courant number. The local blending mechanism blends first-order upwind differencing with the central schemes used for the momentum convection to increase stability in regions where small cell size or poor grid quality might normally lead to instability. In this way, stability of the solution is significantly increased without an appreciable reduction in overall accuracy.

MESH GENERATION

With the requirement for fast prototyping of complex geometries, the automatic generation of high-quality meshes has become a critical factor in the CFD design process. Traditional manual/sequential mesh generation can create severe bottlenecks in the CFD simulation process. For this reason, a new automatic parallel mesh generation capability has been developed using the functionality available in the OpenFOAM toolbox. The key features of the new mesh-generation tool are:

- mesh generation in parallel;
- fully automatic operation;
- generation of hexahedral and split-hexahedral elements with an overall body-fitted structure;
- local patch-wise surface (including curvature) refinement and volume refinement away from the surface;
- surface-layer mesh generation on a patch-wise basis;
- surface features in the geometry automatically preserved in the mesh;
- a guarantee that the generated mesh satisfies user-defined mesh quality checks;
- no principal requirement to perform any cleaning up of CAD surface data;
- developed using the open-source toolbox OpenFOAM.

The mesh-generation algorithm has been designed to keep any user interaction in the CFD process to a minimum through the use of intelligent algorithms and optimised parameter defaults. It is possible with these newly

developed tools to go directly from CAD to a mesh with tens of millions of cells in hours rather than days (a typical 30 million cell mesh can be generated in under 1 hour on 64 processors).

In order to achieve accurate and robust flow solutions, it is essential that the mesh being generated can satisfy a set of mesh quality checks. At all stages in the mesh-generation process, face- and cell-based mesh quality is monitored to ensure that the final mesh that is generated satisfies a set of user-defined cell- and face-based quality criteria. Currently, the mesh-quality checks that are monitored during the generation phase are orthogonality, face pyramid volumes, face areas, face skewness, face interpolation weights, cell volume ratio, face twist and cell determinant. If at any stage in the process mesh quality cannot be guaranteed, the mesh reverts back locally to a previously valid error-free mesh. In the following sub-sections, a more detailed description of the mesh-generation procedure is presented.

STAGE 1: CASTELLATED MESH GENERATION Initially, a Cartesian base mesh with cells of near unit aspect ratio is generated. The mesh extends throughout the entire solution domain. This base mesh is now refined based on user defined patch-wise surface-refinement levels. The refinement engine that is used is based on a fast octree surface intersection checking routine. Figure 1




Figure 1: Initial surface refinement stage

shows details of how an initial base mesh (with a refinement level of 0 everywhere) is refined in the neighbourhood of a surface which requires a user defined refinement level of 2. Additional surface refinement can also be performed based on the local surface curvature.

As well as surface refinement, volume-refinement strategies are used to refine the mesh away from the surface, where mesh resolution is required e.g. wake regions.

Once the refinement phase is complete, an intersection check is done to determine which parts of the mesh should be discarded. This is accomplished by first determining if there is a surface between neighbouring cell centres using an octree-based search. If a surface is detected, the mesh between the neighbouring cells is locally decoupled. As a final stage a topological walk is performed to agglomerate all connected elements with the remainder being deleted from the mesh.

An example of the type of intermediate castellated surface mesh is shown in Figure 2, using an Audi A6 as a test case, which will be the focus of extensive results presented in this paper. The different colours show the parallel surface mesh decomposition for a mesh generated on 8 processors.




Figure 2: Visualisation of castellated surface mesh for Audi A6 test case

STAGE 2: BOUNDARY-RECOVERED MESH In the second stage of the mesh-generation process, the surface geometry is recovered in the CFD mesh. The surface faces of the castellated mesh, generated at the end of stage 1, are moved (snapped) to the surface. This is achieved by constraining the faces to be planar and iteratively projecting them onto the surface using a technique to preserve the underlying surface feature edges in the reconstructed surface mesh. Figure 3 shows a basic schematic of how the boundary faces are projected onto the underlying surface.





Figure 3: Schematic illustration of projecting faces to surface during boundary recovery

As an additional step, at the end of this boundary-recovery stage, co-planar boundary faces from the same cell are merged, provided this leads to an improvement in the overall mesh quality. The resulting surface mesh generated after these snapping and merging operations is shown in Figures 4 and 5 on the Audi A6 test case. At every stage during the mesh-generation process cell- and

face-based mesh quality are monitored; if a step leads to mesh elements that do not satisfy some user-defined quality criteria, the mesh is locally moved back to a previously generated valid state.

CAD



Surface mesh





Figure 4: Original CAD surface and resulting surface mesh for Audi A6 test case

CAD



Surface mesh




Figure 5: Original CAD surface and resulting surface mesh of vehicle underbody for Audi A6 test case

STAGE 3: BOUNDARY-LAYER ADDITION With the need for improved viscous-flow modelling, an optional final stage in the mesh generation process is to add layers of cells next to the surface. This can be performed on a patch-wise basis. The base mesh, generated after stage 2, is projected off the surface and a layer of cells is generated in the void that is created. As in the previous stages of the mesh-generation process, if layer addition results in violation of any of the mesh quality criteria, surface layers are not generated in this region.

Figure 6 shows the mesh on processor boundaries at the




Figure 6: Finite-volume mesh on processor boundaries, showing surface layers generated on rear of an Audi A6

rear of the Audi A6 test case. This shows the surface layers that have been generated in this region. The layer mesh is terminated at concave and convex edges if the resulting mesh quality in the layer mesh cannot be maintained around these edges.

An overview of the final finite-volume mesh that has been generated on the Audi A6 case is shown in Figure 7. This shows the refinement of the mesh in the interior cells away from surface which is the result of the volume-refinement strategy employed during stage 1.




Figure 7: Visualisation of finite-volume mesh of an Audi A6

VALIDATION

The validation of the simulation methodology is based on data obtained from wind-tunnel experiments of the aerodynamics of a wide range of vehicles produced by the Volkswagen Group. The data are obtained from measurements carried out in the Audi aeroacoustic wind-tunnel and the Volkswagen climatic wind tunnel. Here, we present detailed validation results for a generic vehicle model and a production vehicle, followed by an overview of some of the results obtained for a number of other vehicles with widely differing shapes, including compact hatchbacks, sports coupés, saloons and super-sports cars. This very wide range of validation cases is chosen

as a test of the general applicability and robustness of the methodology, an essential feature for use in the actual vehicle-development process.

The computational meshes are generated using triangulated CAD geometries prepared to yield water-tight closed surfaces. The computational domain is simply a large rectangular box of sufficient dimensions to avoid any blockage effects or interference of the walls. The floor of the domain is treated with a combination of slip and no-slip wall boundary conditions: The no-slip wall begins at a position upstream of the vehicle corresponding to the distance required to generate the experimentally determined boundary-layer thickness on the basis of the 1/7th-power law. The roof and sides of the domain are given symmetry-plane boundary conditions, the inlet a fixed-value boundary condition for velocity and SGS viscosity, and the outlet a standard fixed-pressure boundary condition. In the work reported here, the CFD simulations are typically run for a total physical time of around 2.5 seconds in order to flush any start-up transients out of the domain and to allow the flow field to reach a quasi-steady state. The flow field is initialised with a potential-flow solution.

VOLKSWAGEN RED-MODEL VALIDATION CASE A wealth of experimental data have been collected for a generic vehicle model — referred to as the “Red” model due to its colour — used at Volkswagen for wind-tunnel calibration and CFD validation, and shown in Fig. 8. Due




Figure 8: Volkswagen Red model

to its geometric simplicity but the simultaneous complexity of its aerodynamics, this model lends itself well to the validation study undertaken here. In particular, the rounded tail end of the vehicle presents a significant challenge to CFD due to the pressure-gradient induced separation and vortex generation that occurs there. The experimental data are obtained from measurements carried out in the Volkswagen climatic wind tunnel. The CFD data are obtained from a simulation done on a mesh with approximately 26 million cells and run for a total of 2.5 seconds physical time; the flow-field data reported below are obtained from a time-average of the solution over the last 0.25 seconds of the run.

	c_D [-]	c_{Lf} [-]	c_{Lr} [-]
Experiment	0.249	-0.052	0.128
Simulation	0.265	-0.048	0.118

Table 1: Comparison of predicted and measured aerodynamic coefficients for Volkswagen Red-model test case

The most fundamental set of data for the validation exercise is the aerodynamic coefficients, shown for the simulation and experiment in Table 1. On the whole, the accuracy of the predicted force coefficients can be seen to be very good. The predicted drag coefficient c_D deviates from the experimental value by an overprediction of 6.4% or 16 counts. Prediction of front-lift coefficient c_{Lf} is excellent, with an error of only 2 counts. Considering the difficulty of accurately predicting the flow over the rear upper surface of the vehicle, the underprediction of rear-lift coefficient c_{Lr} by 10 counts or 7.8% is also a very good result.

Integral values of the aerodynamic force coefficients are important overall measures of the accuracy of the flow predictions, but can in practice be misleading due to cancellation of errors. For this reason, it is important to consider also the flow structures predicted by the simulation. First, a comparison of the static-pressure coefficients, c_p , on the vehicle surface is informative. c_p is defined by

$$c_p = \frac{p - p_\infty}{\frac{1}{2}\rho_\infty U_\infty^2}, \quad (4)$$

where p is the static pressure, p_∞ is the free-stream static pressure, ρ_∞ is the free-stream density and U_∞ is the free-stream velocity. This comparison is shown in




Figure 9: Comparison of experimental and predicted static-pressure coefficients c_p for Volkswagen Red model — rear view

Fig. 9, where the difference between the predicted and measured values of c_p is visualised at the discrete points shown with the given colour scale. These points correspond to the numerous pressure taps embedded in the surface of the wind-tunnel model. The figure indicates that a discrepancy exists between the predicted and mea-

sured static-pressure distribution on the rear upper surface of the vehicle, in particular in the area of maximum curvature, at the transition from the side walls to the rear surface. As will be seen shortly, in this region the flow is partially attached, partially separated; clearly there are some differences in the predicted flow structures there that result in the deviation from the experimentally determined pressure. As the base pressure of the vehicle is well-matched, it is likely that the x -component of the underpredicted pressure on the upper surface is a significant cause of the overpredicted drag coefficient. Underprediction of the pressure is also observed at the lower edge of the bumper in the transition to the underbody. Note that the apparent high overprediction of c_p at the A-pillar is due to erroneous positioning of the probe locations in the simulation and should therefore not be taken into consideration.




Figure 10: Comparison of experimental and predicted static-pressure coefficients c_p for Volkswagen Red model — underbody view

Examination of the pressure distribution on the underbody, shown in Fig. 10, also provides further information. It appears that the largest discrepancies in the simulation are found in the vicinity of the wheels, where complex vortical structures and regions of separated flow are found. Particularly between and around the front wheels, the pressure is underpredicted. With the front-lift coefficient is predicted very well, it can be found that this underprediction in local pressure is compensated by a slight underprediction on the vehicle's bonnet, not shown here. Moderate underprediction of pressure is also seen on the portion of the underbody behind the wheels, partially compensating in rear lift for the underprediction of pressure on the upper surface.

The ability of a CFD method to capture the three-dimensionality of the flow is also of paramount importance. Further insight into the aerodynamics of the model and the accuracy of the CFD solution is therefore gained by examining the surface-flow topology, obtained by generating oil-streak patterns on the model's surface. Figure 11 shows a comparison of experimental and predicted surface-flow topology on the rear upper surface of the model. As mentioned, this region of flow is particularly interesting due to the pressure-gradient induced separation arising from the strong curvature of the surface. Evidence of the flow's complexity is immediately evident, as




Figure 11: Comparison of experimental and predicted surface-flow topology in rear region of Volkswagen Red model

is the CFD method's ability to predict its large-scale topology very well. Matched very well are the attached flow coming off the top corner of the side and roof and extending diagonally across the rear surface, the separation of the flow at the edge between the side and the rear surface up to the point of transition to attached flow, and the small zone of separated flow at the centre of the rear surface. All the salient features of the flow structure are captured with high fidelity. This result is somewhat in contrast with the discrepancies observed in the predicted static pressures in this region and merits further investigation.

In general, it can be seen that the predictions made by the DES methodology when applied to this case are of very high quality, given its complexity. Both the integral force coefficients and the flow structures are predicted well, notwithstanding the discrepancies observed.

AUDI A6 VALIDATION CASE A further analysis of the simulation results is carried out using experimental data for the Audi A6, for which particularly detailed data are

available. The data used for the validation of this test case are obtained from experiments carried out in the Audi aeroacoustic wind tunnel [11, 12]. Although this facility is equipped with technology for carrying out aerodynamic testing with full ground simulation (*i.e.* boundary-layer suction upstream of the vehicle, rotating wheels and a moving belt between the wheels), test results reported on here are for the more conventional configuration of static ground and wheels. Furthermore, the vehicle is measured in its "mock-up" configuration, *i.e.* with all inlets for cooling air flow closed.

The simulation results reported on here are obtained from a DES run on a finite-volume mesh of approximately 18 million cells, where the flow field and aerodynamic forces are time-averaged over the last 0.25 seconds of the simulation.

The aerodynamic coefficients for the simulation and experiment are shown in Table 2. The predicted coefficients

	c_D [-]	c_{Lf} [-]	c_{Lr} [-]
Experiment	0.271	0.068	0.116
Simulation	0.267	0.070	0.142

Table 2: Comparison of predicted and measured aerodynamic coefficients for Audi A6 test case

can be seen to compare very favourably with the experimental data. The drag coefficient c_D is predicted to within 4 counts, or less than 1.5% of the experimental value, better than the level of accuracy typically required of CFD simulations for external aerodynamics of road vehicles. The relative error of the front-lift prediction is also excellent, at less than 3% with 2 counts. However, the rear-lift coefficient, c_{Lr} , is overpredicted considerably, by 26 counts or 22%.

Some clues explaining this discrepancy can be found by examining the distribution of the static-pressure coefficient c_p along the vehicle's centreline, shown in Fig. 12. The figure shows a comparison of c_p measured at discrete points on the centreline with the continuous c_p distribution there extracted from the CFD results. Across the bonnet, the static pressure can be seen to be slightly overpredicted, with a switch to underprediction as the flow approaches the base of the windscreen. Along the windscreen and roof of the vehicle, the CFD results compare very well to the experiments. At the end of the roof, a spike in the predicted pressure is generated by the presence of a rooftop antenna there. As the flow progresses along the rear window and boot lid, static pressure is recovered. However, it can be seen that the predicted pressure recovery is lower than in the experiments, yielding a possible explanation for the overprediction in rear lift for this case.

For the underside of the vehicle, the comparison is shown in Fig. 13. Note that the erratic character of the c_p distribution in the simulation is due to the geometric complexity




Figure 12: Comparison of experimental and predicted static-pressure coefficients along centreline of upper vehicle surface




Figure 13: Comparison of experimental and predicted static-pressure coefficients along centreline of lower vehicle surface

of the underbody components and the local maxima and minima in pressure that are generated by the small-scale flow structures there. In general, the predicted distribution is seen to follow the experimental one well, although between the two front measurement points and between $x = 1$ m and $x = 2$ m larger discrepancies are found. With regard to the overprediction of rear lift, no obvious overprediction of the pressure in the rear part of the underbody is seen, providing further reinforcement of the hypothesis that this discrepancy might be attributable to insufficient pressure recovery on the rear window and boot lid. Further confirmation of this can be found in the off-centre pressure distribution, not shown here.

Figure 14 shows a comparison of the experimental and predicted flow topology near the surface of the vehicle in the rear C-pillar region. The images generated from the simulation results show the surface streaklines generated on the basis of the wall shear stress, while the colour




Figure 14: Comparison of experimental and predicted surface-flow topology in C-pillar region of Audi A6

scale represents the near-wall velocity, ranging from low (blue) to high (red). Particularly for saloon vehicles, the flow structures in this region are known to be subtle and of considerable importance for the overall aerodynamics of the vehicle. Here, the predictions made by the DES method can be seen to be quite close to the structures observed in the oil-streak visualisation of the experiments. In particular, the fact that the flow remains attached in the transition from the base of the C-pillar to the boot lid and rear window, as well as the predicted size of the separation bubble at the base of the rear window, provides further evidence of the good fidelity of the turbulence modelling of this method.




Figure 15: Comparison of experimental and predicted surface-flow topology around rear end of Audi A6

A further detail of note is separation of the flow and the vortex generation in the vicinity of the tail light. Figure 15 shows the comparison of simulation with experiment in this region. Here also, attachment and separation of the flow in the appropriate regions is observed to be predicted largely correctly, even if some minor differences in the directionality of the flow in the vicinity of the longitudinal vortex can be seen.

Another region of particular interest in the vehicle's aero-




Figure 16: Comparison of experimental and predicted surface-flow topology around A-pillar and side window of Audi A6

dynamics is the A-pillar and wing mirror. In Figure 16, a comparison of the surface-flow structures in this region is shown. Predicted well is the wake of the mirror at the vehicle's surface, straddling downstream of the mirror the lower region of the window and the upper door shoulder. The flow immediately across the A-pillar is also predicted well. Some discrepancies are observed at the reattachment line generated on the upper section of the side window by the A-pillar vortex, indicating that the vortex size is slightly underpredicted. Possibly related to this is the fact that the simulation predicts a stronger upward orientation of the flow on the upper section of the window, whereas the oil-streak images indicate flow more parallel to the shoulder of the door. An investigation of the influence of mesh resolution of the prediction of these detailed flow structures may lead to further insight.

Prediction of the size and strength of vortices generated by the vehicle can be analysed well by means of total-pressure measurements. Figure 17 shows plots of the




Figure 17: Comparison of experimental and predicted total-pressure coefficient in wake plane of Audi A6

total-pressure coefficient, c_{pt} , defined as

$$c_{pt} = \frac{(p + \frac{1}{2}\rho U^2) - p_{\infty}}{\frac{1}{2}\rho_{\infty} U_{\infty}^2}, \quad (5)$$

in a plane at $x = 3.9$ m in the vehicle's coordinate system, roughly 10 cm downstream of the rear bumper. In the simulation result, two off-centre local minima in c_{pt} can be observed just above the large wake region. These correspond to the A-pillar vortices, and can be observed to be somewhat stronger than in the experiment. The figure also indicates that the wake generated by the front wheels, the outer region close to the floor, is underpredicted by the CFD simulation. Predicted well are the general structure of the wake, including the degree of blockage of the flow through the underbody. This lends further weight to the hypothesis that the overprediction in rear lift is primarily due to insufficient pressure recovery on the rear window and boot lid.

The flow-field images shown in the previous figures are derived from the time-averaged flow field of the last 0.25 seconds of the simulation. Figure 18 shows a plot of the instantaneous total-pressure coefficient at the end of the simulation. In comparison to the flow field shown in Fig. 17, the instantaneous field clearly exhibits the small-scale turbulent structures resolved by an LES formulation, where only the structures smaller than the filter (cell) size are modelled. All the essential macroscopic features of




Figure 18: Instantaneous total-pressure coefficient in wake plane of Audi A6

the flow field are present — the vortices generated by the A-pillar and wing mirror, the underbody flow, the wakes of the wheels — but the fine structures are not damped out.




Figure 19: Visualisation of SEAT Ibiza, Volkswagen Passat and Audi R8 aerodynamics validation cases showing near-wall velocities and wake turbulence structures

VALIDATION OVER WIDE VEHICLE RANGE The results of the detailed comparison of flow structures, surface pressures and integral force coefficients for the Volkswagen Red model and Audi A6 generally show very good agreement between simulation and experiment for important characteristics of vehicle aerodynamics. In order to test the robustness and range of applicability of the simulation method, a large test suite of vehicles — spanning a wide range of basic vehicle shapes — was created on which the method was used. Here, the objective was to ascertain to what degree the fundamentally different aerodynamics of different vehicle shapes — such as sub-compact hatchbacks, sports coupés, squireback vehicles and super-sports cars — can be captured by a single, consistent simulation approach. Successful demonstration of this capability is essential if the methodology is to be applied reliably in a real vehicle-development process.

Table 3 presents an overview of the accuracy of the predicted drag and lift coefficients for a selection of vehicles from the test suite used in the validation activities. The

	Δc_D [-]	Δc_{Lf} [-]	Δc_{Lr} [-]
SEAT Ibiza	0.018	-0.017	0.045
SEAT Leon	0.021	-0.005	0.030
VW Golf	0.003	0.034	0.024
VW Passat	0.011	-0.033	0.035
VW New Beetle	0.016	0.001	0.030
Audi A3	0.007	-0.018	0.034
Audi A5	0.011	-0.036	0.031
Audi A6	-0.004	0.002	0.026
Audi Q5	-0.001	-0.006	0.047
Audi TT	-0.001	-0.006	0.051
Audi R8	0.022	0.021	-0.012

Table 3: Comparison of predicted aerodynamic coefficients with experiment for range of vehicles

table shows that, on the whole, good accuracy is obtained for the prediction of drag coefficient over the whole range of vehicles. A few notable exceptions can be identified, however. Both SEAT vehicles suffer from a significant overprediction in drag that is not observed in the other squareback vehicles in the suite. Closer investigation of these particular cases is required to ascertain the source of the discrepancy, including a verification of the exact correspondence of all the details of the simulated geometry to that of the vehicle tested in the wind tunnel. At 16 counts, the error in the prediction for the Volkswagen New Beetle is also quite high. This result can likely be attributed to the fact that the aerodynamics of this particular vehicle are dominated by pressure-gradient driven separation at the round rear surface of the vehicle, always a particular challenge for any CFD methods that employ wall modelling. Furthermore, the size of the error is consistent with the results obtained for the Volkswagen Red model, which has a similar fundamental flow structure in the rear. The overprediction of the drag coefficient of the Audi R8 by 22 counts is also too high to be considered acceptable. Differentiating this vehicle from the others is

the particularly low ground clearance and the flow in the underbody diffusers. Investigations as to the source of the discrepancies here are still underway.

With regard to front lift, a clear trend is difficult to identify, except perhaps a general tendency for underprediction. Here also the Audi R8 does not fit this trend, with the 21-count overprediction likely related to the overprediction observed in drag coefficient. The significant overprediction of the front lift for the Volkswagen Golf merits further investigation.

The clearest trend to be observed in the deviation from the experiments is found in the rear-lift coefficient. Here, again with the notable exception of the Audi R8, the trend for overprediction is obvious. Generally, an overprediction of around 30 counts is found in most vehicles. Worst among these vehicles is the Audi TT, which, with its strongly curved rear surface, is a particularly challenging case for CFD with wall modelling. Although the overprediction of rear lift yields a result that errs on the side of caution, greater accuracy is required, and this necessitates a deeper analysis of this issue.

ASSESSMENT OF OVERALL APPROACH

The validation of the simulation results demonstrates that the methodology presented here is able — in spite of the discrepancies observed — to deliver predictions with very good accuracy. Given the broad scope of the present investigation, the overall accuracy is, in the experience of the present authors, considered at least on par with that of other CFD codes and is sufficient to be useful for actual production-vehicle development. Further criteria, some of which are mentioned in the introduction of the paper, can be considered to complete the assessment.

After accuracy, of utmost importance is the time-to-solution, as any simulation methodology not able to keep pace with the rapid evolutionary steps within a modern vehicle design cycle cannot be successfully integrated into the development process. Due to the fully parallelised operation of each of the steps in the overall simulation process and the good parallel performance of the CFD solver, given sufficient computing resources, turn-around times of less than 24 hours can be obtained, beginning from a triangulated surface mesh and ending with a quasi-steady DES solution. For example, the Audi A6 case, which can be considered to be representative, required 0.5 hours clock time for mesh generation, 5 minutes for case initialisation and 28 hours for the flow simulation on 192 cores of a 3.3 GHz Xeon Linux cluster with an Infini-band interconnect and 2 GB RAM per core.

Conformity with the CAD process is also an important attribute of this methodology. In this study, water-tight triangulated CAD surfaces were used to represent the vehicle geometry, a standard approach in current automotive CFD application. However, this can easily be extended to surface data with reduced topological and quality re-

quirements, thereby enabling even tighter integration with vehicle body and component design.

Other than the obvious advantage of the absence of license fees, the open-source software approach offers several advantages over conventional proprietary software for this application. As discussed in the introductory section of this paper, the continually increasing demands on the performance of CFD software require higher and higher degrees of tailoring for specialised applications. For the work reported on here, the fundamental customisability of the OpenFOAM toolbox was found to be exceptionally advantageous and powerful.

Related to this is the flexibility with which alternative and/or advanced modelling approaches or numerics can be implemented and applied. Practitioners of advanced CFD will know the importance of investigating the possibilities offered by modelling and numerics developed specifically for their flow phenomena. Furthermore, where details of modelling and numerics in proprietary CFD codes are usually hidden, in the open-source code used here they are — by definition — fully transparent. To the extent that these details are important for determining accuracy and assessing sources of inaccuracy, the openness of the software is of considerable advantage to the user. Again here, the general open-source framework, and the object-oriented C++ design of the code in particular, enabled the generation of the high-quality results reported in this paper.

Based on these conclusions, we find the methodology presented here to be appropriate and suitable for use in the industrial development process.

FUTURE WORK

As shown in the validation section, there currently appears to be a consistent tendency of the chosen methodology to overpredict rear lift, likely due to insufficient pressure recovery on the rear upper surface of the vehicle. Moreover, super-sports cars with low ground clearance and smooth underbodies suffer from a clear overprediction in drag. Both topics are to be analysed in more detail through a combination of additional numerical and experimental investigations.

The work to date presented in this paper has focused on external aerodynamics of vehicles in a mock-up configuration, *i.e.* without the flow of air into and out of the engine bay for cooling purposes. The influence of cooling-air flow on vehicle aerodynamics can be very significant and very complex, and is therefore of considerable interest in the development process. Work is therefore underway to validate the prediction of vehicle aerodynamics that includes cooling flow. Figure 20 shows a visualisation of an underhood simulation of the Audi A6 discussed in detail previously in the aerodynamics validation section. As a first basic step, the simulation is performed on the basis of the steady-state RANS-based incompressible isothermal




Figure 20: Visualisation of velocity magnitude from underhood simulation of Audi A6

solver in OpenFOAM using the S-A turbulence model, using the identical mesh-generation process described previously. Mass-flow rate of air across the radiator components is used as a plausibility check of the results, with the accuracy of the predicted value found to be within 3.6% of the experimental value. Given the excellent accuracy of this prediction, the next steps include detailed analysis of the aerodynamic predictions with both RANS and DES solvers.

More generally, the scalability and flexibility of this overall approach makes optimisation-based design, incorporating genetic-evolution or DOE algorithms, a more realistic undertaking than up to now.

SUMMARY AND CONCLUSIONS

In this paper, a complete methodology was presented with which DES can be applied productively within the development process for vehicle aerodynamics. The approach is based on the open-source CFD toolbox OpenFOAM and comprises the automatic generation of a body-fitted computational mesh, the initialisation of the flow field and boundary conditions and a DES flow solver, all fully parallel.

The method is applied to a very wide range of test cases for validation, from a generic vehicle model to production vehicles of various types. Both the detailed analyses of the simulation results, as well as the general overview of the results, show that many important aspects of vehicle aerodynamics are predicted well, ranging from integral drag and lift coefficients to details of the flow topology. Some moderate discrepancies in the simulations are found, most notably the consistent overprediction of rear lift, and these merit further investigation, but they are not deemed to be fundamental.

The degree of automation and parallelisation inherent in the methodology, combined with some of the key characteristics of open-source software, yield a powerful and flexible new approach to advanced industrial CFD for vehicle aerodynamics.

ACKNOWLEDGEMENTS

The authors wish to thank H.G. Weller, M. Janssens (OpenCFD Ltd), J. Betz, K. Gruber, B. Jacquet, S. Kolpatzik, N. Lindener, D. Patnaik, T. Schütz, K. Zens (Audi AG), F. Purschke, G. Rapin, H. Schmidt, R. Sundermeier (Volkswagen AG), X. Agustin, R. Mari (SEAT S.A.), C. Frick (science + computing AG), M. Renner, M. Schröter (GNS Systems GmbH), F. Campos and S. Weston (Icon Ltd) for their valuable support of and contributions to the work presented here.

OpenFOAM is a registered trademark of OpenCFD Ltd.

REFERENCES

- [1] S.R. Ahmed and L. Lüthmann. Numerische Verfahren. In W.H. Hucho, editor, *Aerodynamik des Automobils*, Wiesbaden, 2005. Vieweg-Verlag.
- [2] <http://www.opencfd.co.uk>.
- [3] M. Islam. Simulation of Unsteady Flows. In J. Wiedemann and W.H. Hucho, editors, *Progress in Vehicle Aerodynamics – Unsteady Aerodynamics*, Stuttgart, 2004.
- [4] J. Smagorinsky. General circulation experiments with primitive equations. *Mon. Weather Review*, 1963.
- [5] P.R. Spalart, W.-H. Jou, M. Strelet, and S.R. Allmaras. Comments on the feasibility of LES for wings and on a hybrid RANS/LES approach. *Advances in LES/DNS*, 1998.
- [6] P.R. Spalart and S.R. Allmaras. One-equation turbulence model for aerodynamic flows. *Proc. AIAA ASM*, 1992.
- [7] M. Shur, P.R. Spalart, M. Strelets, and A. Travin. Detached Eddy Simulation of an Airfoil at High Angle of Attack. In W. Rodi and D. Laurence, editors, *Fourth International Symposium on Engineering – Turbulence Modelling and Experiment*, Corsica, 1999. Elsevier.
- [8] E. de Villiers. *The Potential of Large Eddy Simulation for the Modeling of Wall Bounded Flows*. PhD thesis, Imperial College, London, UK, 2006.
- [9] H. Jasak. *Error Analysis and Estimation for the Finite Volume Method with Applications to Fluid Flows*. PhD thesis, Imperial College of Science, Technology and Medicine, London, UK, 1996.

- [10] W.C. Gear. *Numerical Initial Value Problems in Ordinary Differential Equations*. Prentice–Hall, Englewood Cliffs, N.J., 1971.
- [11] G. Wickern and N. Lindener. The Audi Aeroacoustic Wind Tunnel: Final Design and First Operational Experience. *SAE Paper 2000-01-0868*, 2000.
- [12] G. Wickern, W. von Heesen, and S. Wallmann. Wind Tunnel Pulsations and their Active Suppression. *SAE Paper 2000-01-0869*, 2000.