Imperial College London

Department of Aeronautics

Application of spectral/hp element methods to high-order simulation of industrial automotive geometries

Walid M. A. HAMBLI

August 2022

Supervised by Professor Spencer J. SHERWIN

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Aeronautics of Imperial College London and the Diploma of Imperial College London

Declaration of originality

The work presented hereafter is based on research carried out by the author at the Department of Aeronautics of Imperial College London and it is all the author's own work except where otherwise acknowledged. No part of the present work has been submitted elsewhere for another degree or qualification.

Walid M.A. HAMBLI

August 2022

Copyright Declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Flow predictions around cars is a challenge due to massively separated flow and complex flow structures generated. These flow features are usually poorly predicted by present industrial computational fluid dynamics (CFD) codes based on a low fidelity Reynolds averaged Navier Stokes (RANS) approach simulating the mean effects of turbulence. On the other hand, high fidelity approaches resolve turbulent scales but require many more degrees of freedom than classical techniques for an accurate solution. Previous applications have shown that the coupling of the spectral/hp element method and implicit large eddy simulation (iLES) turbulence treatment could be a potential candidate to perform high-fidelity simulations. This work aims at transferring the spectral/hp element technology to the automotive industry in which high Reynolds numbers and complex geometries are typical. Recent developments in stabilisation techniques such as the discontinuous Galerkin kernel spectral vanishing viscosity (SVV) and high-order meshing capabilities open the possibility of the application of the spectral/hp element method to complex cases. The technology is first implemented to an industrial case proposed by McLaren Automotive Limited (MLA) at a realistic Reynolds number of 2,3 million based on the front wheel diameter and is compared to a RANS numerical development tool. Differences in terms of vortical structures arrangement, principally due to the front wheel wake are highlighted. In parallel, a workflow is developed to systematically address similar complex cases.

The interaction between h-refinement, related to the size of the elements of the mesh, and

p-refinement, corresponding to the polynomial expansion order, is investigated on the SAE notchback body. Two different hp-refinement strategies with similar numbers of degrees of freedom are employed, the first one with a fine mesh and a third-order accurate polynomial expansion and the second one with a coarse mesh and a fifth-order accurate polynomial expansion. Results show that a minimum level of h-refinement is necessary to capture flow features and that p-refinement can subsequently be used to improve their resolution.

The final part focuses on wheel rotation modelling. Scale-resolving techniques are intrinsically unsteady and therefore require sophisticated techniques to correctly model rotating wheels. A procedure, built upon an immersed boundary method (IBM) called the smoothed profile method (SPM), is developed to model complex three-dimensional rotating geometries, in particular rim spokes. It is finally applied to an isolated rotating wheel case and results are compared to the moving wall (MW) and the moving reference frame (MRF) modelling techniques. It is concluded that the SPM is in better qualitative agreement with experimental results present in the literature than the two other modelling strategies.

بسم الله الرحمن الرحيم والصلاة والسلام على رسول الله.

اللَّهُمَّ انْفَعْنِي بِمَا عَلَّمْتَنِي وَعَلِّمْنِي مَا يَنْفَعُنِي وَزِدْنِي عِلْمًا الْحَمْدُ لِلَّهِ عَلَى كُلِّ حَالٍ.

'In the name of Allah, The Most Gracious, The Most Merciful and blessings and peace be upon his Messenger.
O Allah! Benefit me with that which You have taught me and teach me that which will benefit me and all praise is due to Allah in every condition.'

Sunan Tirmidhi, Hadith: 3599

Acknowledgements

I would like to firstly thank my supervisor, Professor Spencer Sherwin for his invaluable guidance during this almost four-year journey. His precious advice and assistance, not only technical but also on a human level, have been of great help in achieving this research project. I will always be grateful to him.
I am also extremely thankful to McLaren Automotive Limited for sponsoring my project, in particular Arthur Slight, Gary Ahlin and Roberto Della Ratta Rinaldi for their supervision, for their input and for integrating me into the team. My special thanks also go to Robin Algoo for his friendship and precious technical advice.
I would also like to thank my colleagues at Imperial College London, James Slaughter and especially Filipe Buscariolo. Thank you for your friendship and advice.

I must acknowledge Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service, using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council for the computing hours.

Finally, I am deeply and forever grateful, firstly to my parents Ammar and Leila, my wife Sabrina, my sister Meriam and my brothers Khaled and Mohamed and also my uncles, Yacine and Mehdi, for always supporting and encouraging me. I love you.

Contents

1 Introduction **11**

1.1 High performance car aerodynamics 11

1.2 Wheel aerodynamics 14

1.3 Computational fluid dynamics for automotive development 15

1.4 Motivation and layout of the thesis 16

2 Introduction to the spectral/hp element method **18**

2.1 Spectral/hp methods 18

2.1.1 Presentation in one dimension 18

2.1.2 Extension to multi-dimensions 25

2.2 High-order meshing 29

2.2.1 Geometry definition 30

2.2.2 Linear mesh generation 31

2.2.3 Linear mesh deformation into a high-order mesh 31

2.3 Velocity correction scheme 33

2.4 Stabilisation techniques 37

2.4.1 Implicit Large Eddy Simulation: Spectral vanishing viscosity (SVV) . 37

2.4.2 Dealiasing 42

2.4.3 Conclusion 43

3 McLaren Automotive geometry: P14 **44**

3.1 Literature review 44

3.1.1 Previous applications of the spectral/hp element method for automotive simulations 44

3.1.2 Objectives 49

3.2 Case presentation and workflow 51

3.2.1 Geometry 51

3.2.2 Mesh generation 55

3.2.3 Numerical settings 62

3.3 Results and analysis 67

3.3.1 Decomposition of aerodynamic forces 68

3.3.2 Surface pressure distribution 72

3.3.3 Three-dimensional flow structures 76

3.3.4 Streamwise evolution of flow structures 78

3.3.5 Unsteady flow structures 86

3.4 Conclusion 88

4 Influence of hp refinements: SAE body **91**

4.1 Literature review 91

4.1.1 Automotive bluff bodies 91

4.1.2 Notchback flow topology 93

4.1.3 The notchback SAE body 95

4.1.4 Previous studies on the SAE notchback body 96

4.1.5 Objectives 101

4.2 Simulation settings 102

4.2.1 Computational domain 102

4.3 Mesh 103

4.3.1 Numerical parameters . . . 109
4.4 Results and discussion . . . 111
4.4.1 Forces . . . 111
4.4.2 Pressure distribution . . . 112
4.4.3 Flow visualisation in the notch region . . . 116
4.4.4 Three-dimensional flow structures . . . 122
4.4.5 Discussion . . . 131
4.5 Conclusion . . . 134

5 Rotating wheel modelling: Smoothed Profile Method 136
5.1 Literature review . . . 136
5.1.1 Isolated wheel aerodynamics . . . 137
5.1.2 Rotating wheel modelling techniques . . . 151
5.1.3 Smoothed profile method . . . 162
5.1.4 Objectives . . . 171
5.2 Concentration field generation . . . 173
5.2.1 Distance from a mesh point to the immersed body surface . . . 173
5.2.2 Second condition to determine the location of the mesh point relative to the immersed body . . . 175
5.2.3 Mesh resolution in the transition interface . . . 176
5.3 Concentration field rotation . . . 178
5.3.1 Strategy . . . 178
5.3.2 Interpolation . . . 180
5.4 Isolated wheel simulation . . . 183
5.4.1 Geometry and mesh . . . 183
5.4.2 Numerical settings . . . 188
5.4.3 Results . . . 190

5.5 Conclusion 201

6 Conclusion and further work **204**

6.1 Application to the P14 geometry 205

6.2 hp-refinement strategy 206

6.3 Rotating wheel modelling 206

6.4 Further work 207

A Mesh metrics **210**

A.1 P14 210

A.2 SAE body 211

A.3 DrivAer wheel 211

Chapter 1

Introduction

1.1 High performance car aerodynamics

Growing concerns over global warming have led to strict regulations on exhaust emissions for road car manufacturers. Current discussions about the Euro 7 norm on ground vehicles should reduce by more than 50% nitrogen oxides (NO_x) emissions of diesel vehicles in 2025. Aerodynamic drag accounts for approximately 75% of the total resistance at $100\,km/h$ according to Hucho [1]), hence very important efforts on road cars aerodynamics optimisation to reduce aerodynamic drag to improve fuel efficiency. Since the first models, car bodies have evolved towards more aerodynamic shapes with reduced drag as can be seen in figure 1.1. Even if very low drag bodies like the Eiffel body (1912) on the bottom left of figure 1.1 have been proposed very early, logarithmic scaling for drag coefficient values on the plot, however, indicates that improvements have become very difficult to achieve solely by optimising the shape of cars. Moreover, the actual booming market of Sports Utility Vehicles (SUVs) with inherently high drag bodies clearly shows that various other criteria such as fashion play an important role in commercial road car design and can be opposed to aerodynamic efficiency. Yet, the trend towards aerodynamic drag optimisation has gained a lot of attention very re-

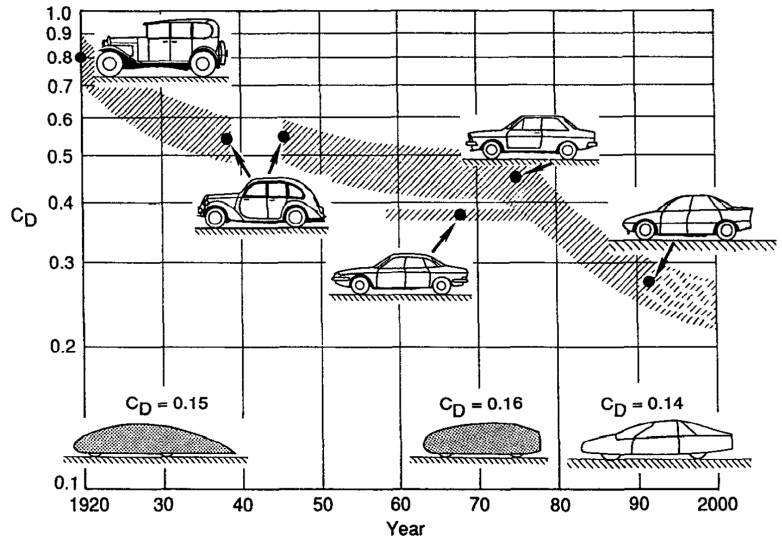


Figure 1.1: Road car drag evolution from 1920s to 2000s (after Hucho [1])

cently due to the development of electric vehicles as gains directly translate into increments in maximum range.

Aerodynamics plays a major role in both road car and race car design even if aerodynamic optimisation targets differ between the two cases. When it comes to race cars, weight and aerodynamics are the most efficient levers to improve performance on the track. In Formula One, Nakagawa et al. [2] have shown that one week of wind tunnel testing is equivalent in terms of impact on lap time reduction as diminishing the mass of the vehicle by 10 kilograms. Given the very strict regulations on the mass of vehicles and the fact that they are originally designed underweighted, aerodynamics turns out to be the most efficient way to improve performance. Even if early designers focused on reducing aerodynamic drag to increase maximum speed, they started following a totally different path from the 1970s. Indeed they took advantage of the fact that race cars operate on the ground to create additional negative lift, also called downforce, sacrificing drag counts in the process. This turns out to be more beneficial on the track. By placing inverted wings near the floor, the interaction with the track creates large suction under the wing increasing grip and allowing higher cornering speeds.

This phenomenon is called ground effect. The plot by Katz on figure 1.2 clearly shows that in the 1990s, race cars designed to create downforce could withstand more than twice higher values of lateral acceleration compared to low downforce race cars without loosing grip.

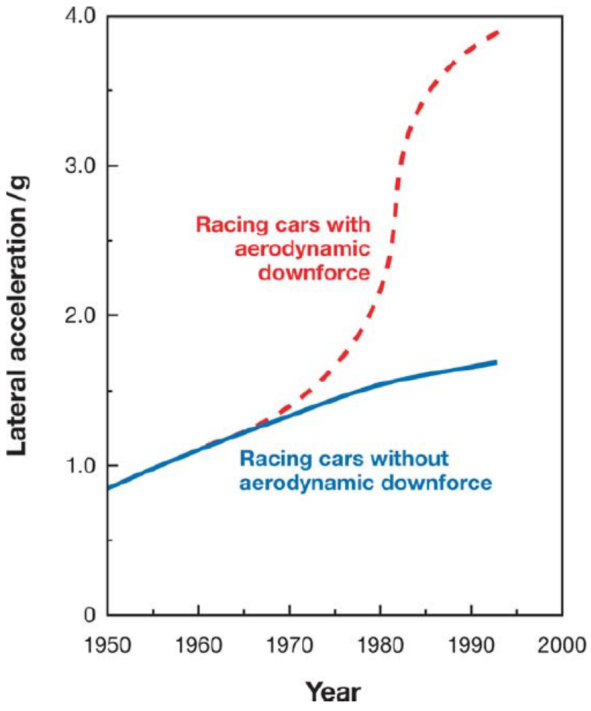


Figure 1.2: Evolution of maximum lateral acceleration of race cars with and without high downforce target

High performance cars are born from the ambition to adapt performances of race cars to road legal cars. Therefore, unlike usual commercial cars, the focus does not lay on aerodynamic drag but rather on maximising downforce and aerodynamic efficiency. They are often equipped with body parts and appendices directly inspired from race cars such as diffusers and wings to create additional downforce as can be seen on the rear view of the McLaren Senna on figure 1.3. These downforce generation parts are strongly impacted by flow structures originating from the wheels. Moreover the engine is often placed in a rear position that makes engine air intakes very dependent on the flow downstream from the front wheels. Therefore, understanding the flow in wheelhouses is of primary importance as it strongly

impacts the aerodynamic behaviour of the whole car.



Figure 1.3: McLaren Senna road legal car released in December 2017

1.2 Wheel aerodynamics

Wheels are responsible for 25% to 40% of the total drag ([3], [4], [5]) of road cars. Still, this is only in the seventies that designers have started focusing on the effect of wheels on car aerodynamics. First measurement campaigns focused on isolated wheels. In 1974, Fackrell and Harvey [6] measured the surface pressure on a rotating wheel on a moving belt in wind tunnel which enabled them to identify a pressure jump at the contact patch location, which is the junction between the wheel and the floor. Flow visualisation using smoke also showed the importance of wheel rotation on flow structures generated by the wheel, with most wind tunnel measurement campaigns on cars being carried out with static wheels by manufacturers at the time. Wheels have been placed on moving belts, deformable tyres with grooves have been used to take into account deformation of the contact patch and wheel shoulders and rims have been added to study wheel hub flow. Experimental and numerical studies, such as Wäschle [7], show that very strong variations in forces are witnessed on cars when wheels are rotated, mainly due to the wake interaction between the wheels and the rear of the car. More recent studies also focus on unsteadiness and highlight vortex shedding originating from the wheel impacting time evolution of forces on the vehicle or passenger comfort by exciting

motion-sickness frequencies in particular cases as observed by Aultman et al. [8].

1.3 Computational fluid dynamics for automotive development

First studies on the airflow around cars have been carried out either in real conditions on road or later in controlled environment in wind tunnel allowing more easily replicable test conditions. It is only in the last twenty to thirty years with the increase in available computational power that computational fluid dynamics (CFD) has become a major tool for car development. Although experimental visualisation techniques have successfully been applied to isolated wheels, it is much more challenging to implement them in the case of shrouded wheels due to space restrictions imposed by the wheelhouse. Hence the relevance of CFD when it comes to study the flow around shrouded wheels. Early CFD works, such as the simulation of Fackrell's wheel by Axon et al. in 1998 [9], focused on isolated wheels to develop and validate methodologies to replicate accurately wheels complex operating conditions such as rotation and contact patch deformation.

Even if simulation techniques have constantly improved, most industrial methodologies are based on low fidelity Reynolds averaged Navier Stokes (RANS) approach in which the whole turbulence spectrum is modelled. This method is widely used due to its ability to provide fairly accurate results in relatively short time and robustness. Yet, this type of modelling frequently fails at predicting massively separated turbulent flows which are typical around road cars, for instance in wheel houses as shown by Krajnovic [10]. RANS is thus usually employed to indicate trends rather than providing precise descriptions of the flow. Therefore, more accurate techniques must be employed to more finely understand the airflow around cars. Turbulence resolving techniques have recently gained in popularity and usually perform better in presence of large flow separations. Although they are increasingly employed in

development processes, it remains limited due to the complexity of geometries involved and high Reynolds numbers characteristic of automotive industrial applications imposing very fine mesh resolutions to resolve small turbulent scales. Moreover, they require low diffusive numerical techniques in order to preserve resolved small scales from artificial dissipation. High-order methods are a class of numerical schemes relying on more elaborate discrete decompositions of continuous physical quantities and can provide both lower numerical diffusion and improved scalability on modern high performance computing (HPC) architectures due to compact discretisation. Yet, they require more efforts to implement and their inherently low numerical dissipation compared to their low-fidelity counterpart is detrimental in terms of robustness.

1.4 Motivation and layout of the thesis

The application of high-order methods to industrial cases remains very limited principally due to associated complexity and demand in computational resources. The enhanced accuracy is a clear asset to perform further aerodynamic optimisation on road cars in cases where low-fidelity techniques and expensive wind tunnel testing campaigns fail at providing answers. High-fidelity simulation techniques are foreseen as powerful tools on modern and next-generation HPC facilities. Yet, development is still required to transfer them to the automotive industry since numerical methods and several steps of the simulation workflow differ from traditional techniques. The goal of the work is to evaluate the applicability and performances of Nektar++ [11], a simulation tool leveraging the spectral/hp element method belonging to the high-order framework, on industrial automotive cases. In the first part, the spectral/hp element framework on which Nektar++ is based, is introduced. The solver is then applied to an adapted commercial car proposed by McLaren Automotive and compared to a commercial RANS tool. The study helps identifying two crucial points to address in

order to facilitate the application of the tool to automotive development processes. The first point is high-order meshing. It is considered as crucial on complex industrial geometries and the impact of the interaction between geometrical refinement and high-order approximation on flow field predictions is investigated in the case of the generic SAE notchback body [12]. The second key point is wheel modelling. Rotation of different wheel parts, tyre and rim, is a real challenge to mimic numerically. An immersed boundary method called smoothed profile method (SPM) is adapted to tackle complex three-dimensional geometries and assessed in the case of an isolated rotating wheel in contact with a moving ground.

Chapter 2

Introduction to the spectral/hp element method

2.1 Spectral/hp methods

2.1.1 Presentation in one dimension

Nektar++ is based on the spectral/hp element framework which will briefly be developed on the case of the one-dimensional Helmholtz equation using the reference work of Karniadakis and Sherwin ([13]). The method is on the frontier between spectral element (SE) and finite element (FEM) methods and operates on the weak form of partial differential equations (PDE). Let Ω be the domain $\Omega = [0, L]$ where $L > 0$. The one-dimensional (1D) Helmholtz equation can be written

$$\frac{\partial^2 u}{\partial x^2} - \lambda u + f = 0 \tag{2.1}$$

where $\lambda > 0$, f is a source function and under the boundary conditions

$$u(0) = u_D \qquad \text{and} \qquad \frac{\partial u}{\partial x}(L) = u_N. \tag{2.2}$$

Let us now introduce, under sufficient regularity hypothesis, the inner product on Ω

$$(f, g) = \int_\Omega fg \; dx. \tag{2.3}$$

Applying the inner product to equation 2.1 with a test function v

$$\int_\Omega \frac{\partial^2 u}{\partial x^2} v \; dx - \int_\Omega \lambda uv \; dx + \int_\Omega fv \; dx = 0$$

which gives after integrating by parts

$$\int_\Omega \frac{\partial u}{\partial x}\frac{\partial v}{\partial x} \; dx + \int_\Omega \lambda uv \; dx = \int_\Omega fv \; dx + \left[\frac{\partial u}{\partial x} v\right]_\Omega . \tag{2.4}$$

The last equation is usually called the weak form of the equation 2.1. A well adapted functional space for u and v ensuring that the quantities above are well defined and that a solution can be found is the Sobolev space $H^1(\Omega)$. Moreover test functions are set to be homogeneous on Dirichlet boundaries. Therefore, useful spaces to introduce are $\mathcal{U} = \{u \in H^1(\Omega) |\, u(0) = u_D\}$ and $\mathcal{V} = H^1_0(\Omega) = \{v \in H^1(\Omega) |\, v(0) = 0\}$. $\mathcal{U}$ and $\mathcal{V}$ potentially contain an infinite number of functions, therefore the problem in equation 2.4 cannot be directly solved by a computer. Subspaces $\mathcal{U}^\delta \subset \mathcal{U}$ and $\mathcal{V}^\delta \subset \mathcal{V}$ with a finite number of functions must be extracted. The discrete problem can be rewritten: find $u^\delta \in \mathcal{U}^\delta$ such that

$$a(u^\delta, v^\delta) = l(v^\delta) \qquad \forall v^\delta \in \mathcal{V}^\delta \tag{2.5}$$

It can be noticed that the Neumann boundary condition is implied by the second term on the right-hand side of the weak form 2.4. Yet, the Dirichlet boundary condition must be explicitly enforced. This can be achieved by breaking down the discrete solution as follows

$$u^\delta = u^H + u^D \tag{2.6}$$

where $u^H \in \mathcal{V}^\delta$ and $u^D \in \mathcal{U}^\delta$ is a lift to ensure the Dirichlet boundary condition. Once the lift is set, the weak problem becomes: find a trial function $u^H \in \mathcal{V}^\delta$ such that

$$a(u^H, v) = l(v) \qquad \forall v \in \mathcal{V}^\delta \tag{2.7}$$

where

$$a(u^H, v) = \int_\Omega \frac{\partial u^H}{\partial x}\frac{\partial v}{\partial x}\, dx + \int_\Omega \lambda u^H v\, dx \qquad \text{and} \qquad l(v) = \int_\Omega f v\, dx + u_N v - \int_\Omega \frac{\partial u^D}{\partial x}\frac{\partial v}{\partial x}\, dx$$

Trial and test functions are now both in the same space as in the Galerkin method. In the spectral/hp element framework, the finite dimension spaces are obtained through subdivision of the domain in several *elements* of typical size h. Then trial and test functions are both approximated in each of them by a polynomial of order p.

2.1.1.1 h-type decomposition

The h-type decomposition corresponds to the partition of the domain into several subdomains as is usually the case in FEM. For instance, in 1D the domain Ω can be decomposed in three sub-domains $\{\Omega_i\}_{1\leq i\leq 3}$ with respect to the four degrees of freedom (DoF) $0 = x_0 < x_1 < x_2 < x_3 = L$ as on figure 2.1. Here, the finite space of functions to approximate the solutions (or *shape functions* in FEM) are the linear hat functions $\{\Phi_i\}_{(0\leq i\leq 3)}$.

This is called the *global expansion basis*. A discrete solution u^δ to problem 2.5 can be decomposed as follows

$$u^\delta = \sum_{i=0}^{3} \hat{u}_i \Phi_i$$

where $\{\hat{u}_i\}_{0\leq i\leq 3}$ are coefficients which are determined by resolving the linear system derived

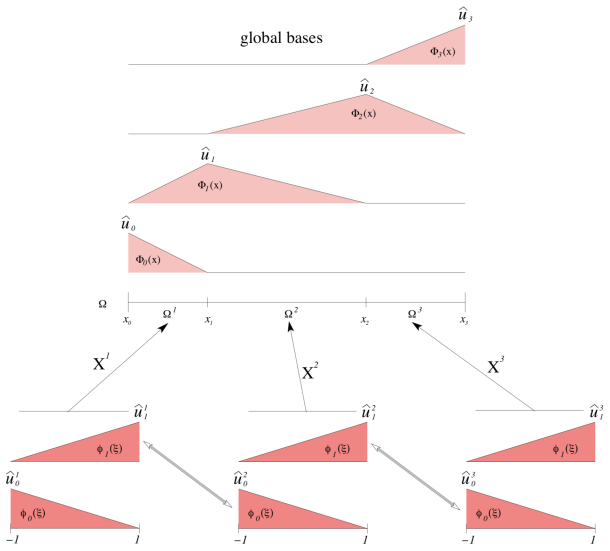


Figure 2.1: h-type decomposition applied to the domain Ω

from equation 2.5. Indeed, by introducing the *Laplacian matrix*

$$\boldsymbol{L} = \left(\int_\Omega \frac{\partial \Phi_i}{\partial x} \frac{\partial \Phi_j}{\partial x} \, dx \right)_{0 \leq i,j \leq 3}$$

the *mass matrix*

$$\boldsymbol{M} = \left(\int_\Omega \Phi_i \Phi_j \, dx \right)_{0 \leq i,j \leq 3}$$

the right-hand side vector

$$\boldsymbol{F} = \left(\int_\Omega f \Phi_i \, dx - u_N \Phi_i - \int_\Omega \frac{\partial \Phi_i}{\partial x} u_D \, dx \right)_{0 \leq i \leq 3}$$

and the trial function coefficients vector

$$\boldsymbol{U} = (\hat{u}_i)_{0 \leq i \leq 3}$$

the linear system to solve is

$$\boldsymbol{V}^T(\boldsymbol{L}+\boldsymbol{M})\boldsymbol{U} = \boldsymbol{V}^T\boldsymbol{F}, \quad \forall \boldsymbol{V} \in \mathbb{R}^4$$

since $\hat{u}_0$ is fixed by the Dirichlet boundary condition. As the matrix equation is true for all V, the final equation to solve is

$$(\boldsymbol{L}+\boldsymbol{M})\boldsymbol{U} = \boldsymbol{F}. \tag{2.8}$$

The global expansion basis itself is built from the *local expansion basis* $\{\phi_0, \phi_1\}$ defined on the standard element $\Omega_{st} = [-1, 1]$ which is then mapped to the elements $\{\Omega_i\}_{1\leq i\leq 3}$. In this case, the coordinate mapping χ_i for the element Ω_i is

$$x = \chi_i(\xi) = \frac{1-\xi}{2}x_{i-1} + \frac{1+\xi}{2}x_i, \; \forall \xi \in \Omega_{st}.$$

This mapping can easily be inversed

$$\xi = \chi_i^{-1}(x) = 2\frac{x - x_{i-1}}{x_i - x_{i-1}} - 1, \; \forall x \in \Omega_i.$$

The global basis functions Φ_0 and Φ_1 are defined using the local basis

$$\Phi_0(x) = \phi_0\left[\chi_1^{-1}(x)\right], \; \forall x \in \Omega_0$$

$$\Phi_1(x) = \begin{cases} \phi_1\left[\chi_1^{-1}(x)\right] & \text{if } x \in \Omega_1 \\ \phi_0\left[\chi_2^{-1}(x)\right] & \text{if } x \in \Omega_2 \\ 0 & else. \end{cases}$$

Therefore, an other way of decomposing u^δ is

$$u^\delta = \sum_{i=1}^{3}\sum_{p=0}^{1} \hat{u}_i^p \phi_p(\xi) = \sum_{i=1}^{3}\sum_{p=0}^{1} \hat{u}_i^p \phi_p\left[\chi_i^{-1}(x)\right]$$

where $\{\hat{u}_i^p\}$ are the coefficients corresponding to the decomposition in the local basis. Therefore, generalising this expansions to a case where Ω is split in N domains with P local basis functions, the discrete solution u^δ can be written

$$u^\delta = \sum_{n=0}^{N_{dof}-1} \hat{u}_n \Phi_n = \sum_{n=1}^{N}\sum_{p=0}^{P} \hat{u}_n^p \phi_p\left[\chi_n^{-1}(x)\right]. \tag{2.9}$$

An advantage of the spectral/hp element method is that many computations can be operated on a standard element and transposed to each element of the physical space through a mapping.

2.1.1.2 p-type decomposition

In the previous example, the local expansion basis is formed of linear functions but in practice higher order polynomials can be used. Numerical methods to solve large linear systems such as equation 2.8 operate better on well-conditioned matrix. It would therefore be judicious to choose an orthogonal polynomial basis with respect to the scalar product $(\cdot,\cdot)$ introduced in equation 2.3. Another interesting feature for the polynomial basis is to be hierarchical, e.g. with increasing polynomial order, to adapt high order modes representation to the polynomial order. Expansions are usually classified as nodal or modal. Nodal expansions are attached to mesh nodes and the coefficients physically represent the value of the approximated solution at the nodes. The most famous example of nodal expansion are the Lagrange polynomials. Yet, this type of expansion is not hierarchical unlike modal expansions which intrinsically fulfil this property. The Jacobi polynomials $P_p^{\alpha,\beta}$ is a modal expansion defined on the interval

$[-1, 1]$ as the solution of the Sturm-Liouville problem

$$\frac{d}{dx}\left[(1-x)^{1+\alpha}(1+x)^{1+\beta}\frac{d}{dx}P_p^{\alpha,\beta}\right] = \lambda_p(1-x)^{\alpha}(1+x)^{\beta}P_p^{\alpha,\beta}$$

where $\lambda_p = -p(\alpha+\beta+p+1)$. Jacobi polynomials are also orthogonal with respect to the weighted scalar product

$$\int_{-1}^{1} P_p^{\alpha,\beta}(x)P_q^{\alpha,\beta}(x)(1-x)^{\alpha}(1-x)^{\beta}dx = C\delta_{pq}$$

where δ_{pq} is the Kroenecker symbol and C a scalar depending on α, β and p. It can be mentioned that Legendre polynomials are obtained with $(\alpha, \beta) = (0, 0)$ and Chebyshev polynomials with $(\alpha, \beta) = (-\frac{1}{2}, -\frac{1}{2})$.

We noted before that in FEM, test and trial functions must be in H^1. This implies global continuity of solutions. Still, the value at the nodes, especially at the boundaries of the subdomains is not directly prescribed with modal expansions. Hence a feature called *boundary interior decomposition* of the polynomial basis introduced in spectral/hp. Two linear modes are added to control the value of the approximation at the boundaries of the subdomain.

$$\psi_p(\xi) = \begin{cases} \frac{1-\xi}{2} & \text{if } p = 0 \\ \frac{1-\xi}{2}\frac{1+\xi}{2}P_{p-1}^{1,1}(\xi) & \text{if } p \in \{1, ..., P-1\} \\ \frac{1+\xi}{2} & \text{if } p = P. \end{cases} \tag{2.10}$$

The choice of $P_{p-1}^{1,1}$ is selected in [13] for numerical efficiency reasons. A graphical representation of this local expansion basis in the case $P = 5$ can be seen on figure 2.2. The framework introduced in this section in one dimension is extended to two and three dimensions in the next section.

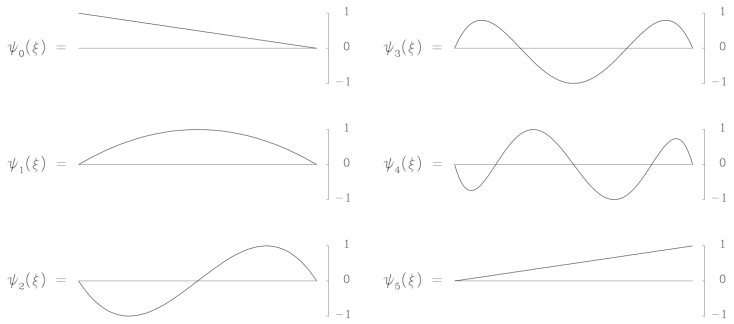


Figure 2.2: Boundary interior mode decomposition for $P = 5$ from [13]

2.1.2 Extension to multi-dimensions

2.1.2.1 Standard elements and expansion basis

The spectral/hp element framework has been extensively presented in one dimension. The derivation of the weak form of the problem is very similar in two and three dimensions, therefore this section will principally focus on the construction of the standard elements and expansion basis in multi-dimensions. Similarly to the one-dimensional development, let us first introduce the simplest standard elements in two and three dimensions, in terms of extension from the one dimensional case, which are quadrangles and hexahedrons:

$$Q_{std}^{2D} = \{(\xi_1, \xi_2) \mid -1 \leq \xi_1, \xi_2 \leq 1\} \tag{2.11}$$

$$Q_{std}^{3D} = \{(\xi_1, \xi_2, \xi_3) \mid -1 \leq \xi_1, \xi_2, \xi_3 \leq 1\}. \tag{2.12}$$

The construction of the basis in these elements is straightforward; it is expressed as a tensor product of the one dimensional base:

$$\psi_{p,q}(\xi_1, \xi_2) = \psi_p(\xi_1)\psi_q(\xi_2), \quad \text{for } p \leq P_1,\ q \leq P_2 \tag{2.13}$$

in two dimensions and

$$\psi_{p,q,r} = \psi_p(\xi_1)\psi_q(\xi_2)\psi_r(\xi_3), \quad \text{for } p \leq P_1,\ q \leq P_2,\ r \leq P_3 \tag{2.14}$$

in three dimensions where ψ is the basis introduced in the previous section and P_1, P_2, P_3 are the polynomial order in each dimension.

When it comes to complex geometries, unstructured meshes are often preferred as they offer greater flexibility. Nektar++ usually works with unstructured meshes including triangles and tetrahedrons which treatment slightly differs from quadrangles and hexahedrons. Standard region and the corresponding expansion bases for triangles and tetrahedron have been introduced by Sherwin and Karniadakis [14]. Figure 2.3 shows how the standard triangle region T_{std} with coordinates (η_1, η_2) is mapped from the standard quadrangle region by collapsing one of the edges of the quadrangle with the following transformation:

$$\eta_1 = 2\left(\frac{1-\xi_1}{1-\xi_2}\right) - 1,\ \eta_2 = \xi_2 \tag{2.15}$$

$$T_{std} = \{(\eta_1, \eta_2) \mid -1 \leq \eta_1, \eta_2 \leq 1\}. \tag{2.16}$$

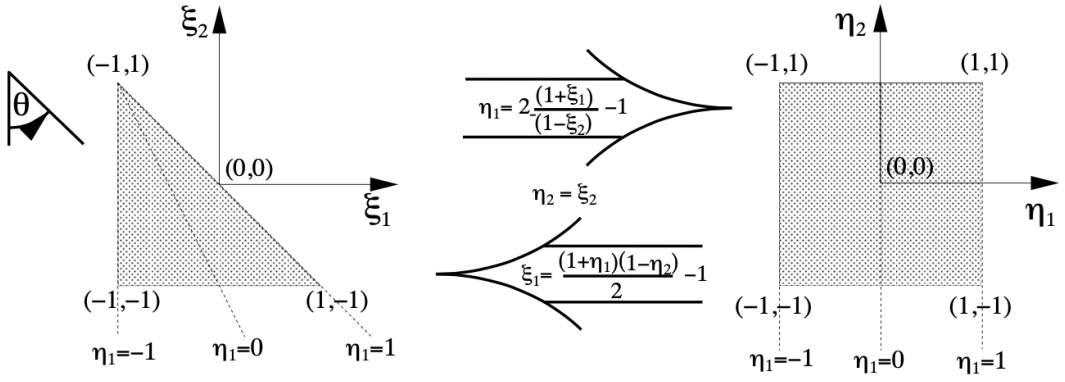


Figure 2.3: Mapping from the standard quadrangle to the standard triangle (from [13])

Still, the expansion basis for the triangle is more complex to construct than for the quadrangle. It is rather formed as a product of a one dimensional tensor $\psi_p(\eta_1(\xi_1, \xi_2))$ by the modified tensor $\psi_{pq}(\eta_2(\xi_2))$ presented on figure 2.4. The mapping from the standard hexahedron to the standard tetrahedron can similarly be defined and is fully described by Karniadakis and Sherwin [13] as well as the corresponding expansion basis. The mapping for high order elements will now be introduced.

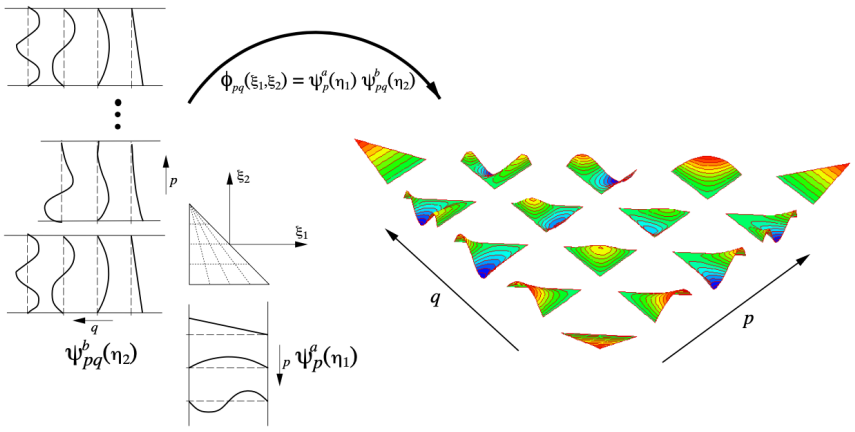


Figure 2.4: Construction of a fourth order ($P = 4$) triangular expansion using the product of two modified principal functions $\psi_p(\eta_1(\xi_1, \xi_2))$ and $\psi_{pq}(\eta_2(\xi_2))$ (from [13])

2.1.2.2 Element mapping

The operations in the spectral/hp element method are defined on standard elements introduced in the previous section. A mapping is then applied to transpose them to more complex elements of high-order unstructured meshes. The first type of transformation considered here is from the standard element into a straight-sided deformed element. The expression proposed for quadrangle mapping by Karniadakis and Sherwin [13] from the standard square region $(\xi_1, \xi_2) \in [-1, 1]$ from the previous paragraph is for $i = 1, 2$

$$\chi(\xi_1,\xi_2) = x_i^A\frac{(1-\xi_1)(1-\xi_2)}{4}+x_i^B\frac{(1+\xi_1)(1-\xi_2)}{4}+x_i^D\frac{(1-\xi_1)(1+\xi_2)}{4}+x_i^C\frac{(1+\xi_1)(1+\xi_2)}{4}. \tag{2.17}$$

The mapping can be built for triangles using the collapsed vertices parametric mapping from equation 2.15

$$\chi(\xi_1,\xi_2) = x_i^A\frac{(-\xi_1-\xi_2)}{2} + x_i^B\frac{(1+\xi_1)}{2} + x_i^C\frac{(1+\xi_2)}{2} \tag{2.18}$$

where $i = 1, 2$.

The second type of mapping presented here consists in the deformation of the edges of standard elements into curvilinear edges. The deformation of standard regions was constructed using only the linear vertex modes. Equation 2.17 can be generalised in the form

$$\chi_i(\xi_1,\xi_2) = \sum_{p=0}^{p=P_1}\sum_{q=0}^{q=P_2}\hat{x}_{pq}^i\psi_p(\xi_1)\psi_q(\xi_2) \tag{2.19}$$

where ψ_i is the hierarchical modal basis from equation 2.10 and all the coefficients aside from those corresponding of linear vertex modes are set to 0. Curvature of edges is incorporated in the previous expansion by allowing non-zero high-order modes expansion coefficients on the vertices and interior nodes. Taking the example of the quadrangle, if the location of the four vertices and the shape functions associated to the curved edges f^A, f^B, f^C and f^D are known and continuous at the vertices, then Gordon and Hall [15] have shown that the blending function could be written as a linear combination for $i = 1, 2$

$$\begin{aligned}\chi_i(\xi_1,\xi_2) = f^A(\xi_1)\frac{1-\xi_2}{2} + f^C(\xi_1)\frac{1+\xi_2}{2} + f^D(\xi_2)\frac{1-\xi_1}{2} + f^B(\xi_2)\frac{1+\xi_1}{2}\\ -f^A(-1)\frac{1-\xi_1}{2}\frac{1-\xi_2}{2} - f^A(1)\frac{1+\xi_1}{2}\frac{1-\xi_2}{2}\\ -f^C(-1)\frac{1-\xi_1}{2}\frac{1+\xi_2}{2} - f^C(1)\frac{1+\xi_1}{2}\frac{1+\xi_2}{2}.\end{aligned} \tag{2.20}$$

Both mappings are illustrated on figure 2.5.

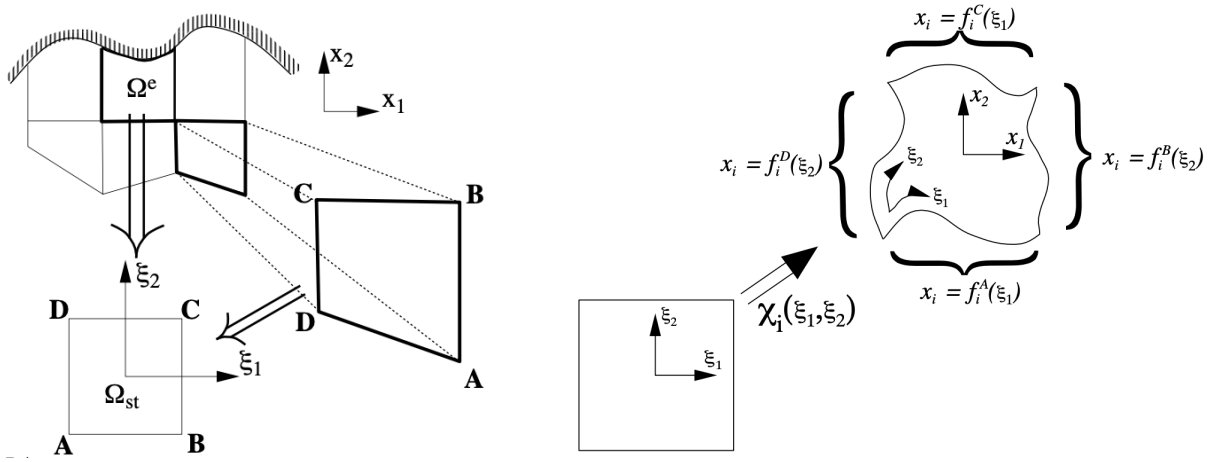


Figure 2.5: Illustration of two-dimensional element mappings for quadrangles by Karniadakis and Sherwin [13] (straight-sided deformation on the left and edges curvature on the right)

2.2 High-order meshing

Mesh generation is a crucial step as it has a huge impact on results as it is shown with the comparison of the simulations results of an automotive bluff body with different meshes in part 3, hence the need of a reliable meshing pipeline. The spectral/hp element framework developed in the previous section is well suited for simple domains. In order to be applied in an industrial process, the meshing tool needs to be both flexible and robust. In this framework, flexibility means that parameters such as mesh structure and refinements can be finely tuned and robustness comes from the ability to generate meshes for complex geometries with satisfactory mesh quality metrics such as skewness and elements validity as defined by Knupp [16] or smooth volume changes between adjacent elements. High-order meshing techniques have been extended to unstructured meshes by Sherwin and Peiro [17]. Indeed, polynomial approximation within each element introduced by the high-order scheme results in meshes significantly coarser compared to those designed for low-order techniques. However

the decrease in the number of elements results in poor approximation of complex surfaces. The idea in Nektar++ is to employ high-order curved elements in replacement of low-order straight elements in order to generate body-conforming grids despite lower element count. The resulting high-order generation process using NekMesh, the meshing tool embedded within Nektar++, based on the works of Turner [18], Moxey et al. [19] and Lombard et al. [20] is presented in this section and goes from geometry definition to the generation of the linear base mesh to its transformation into the final high-order mesh and which has been used generate the meshes for the three cases presented in the study.

2.2.1 Geometry definition

The meshing process starts with the definition of the geometry which has to be as close as possible to the original part as small discrepancies can lead to severe differences in flow detachment locations resulting in flow misprediction around bluff bodies such as cars. Complex geometries can rarely be defined analytically, therefore geometries are obtained with Computer Aided Design (CAD)commercial tools. As explained in section 2.2.3, a projection of the linear mesh upon the geometry definition is performed by NekMesh to produce the high-order mesh and requires a very sharp definition of the geometry. Unlike commercial meshing tools in which small inconsistencies in CAD representation can be healed a posteriori with surface wrapping techniques for instance, Turner et al. explain in [21] that invalid elements generated during the projection step upon the body comes from poor CAD design with porous or self-intersecting surfaces. Therefore geometry checking features embedded within Commercial CAD tools indicating and correcting errors have frequently been used to create thorough CAD definitions. The second consequence of the projection step is that the whole domain has to be included in the CAD definition. Boundary surfaces, the inlet and the outlet of the domain have all to be included in the CAD geometry. The linear mesh can then be generated using a CAD surface satisfying these two criteria.

2.2.2 Linear mesh generation

In the approach proposed by Turner [18], the high-order mesh is generated by deforming a linear mesh to adapt it to the boundary surfaces. Although NekMesh has the capability to independently create a high-order mesh from a geometry definition, the complexity of the geometries studied in this work requires a commercial meshing tool to generate a linear mesh. The CFD software StarCCM+ [22] v13.02 has been used in this work to generate linear meshes. The software enables to conveniently set mesh topology and h-refinements of the linear mesh by enforcing surface or volumetric refinements. The mesh has to respect a precise structure specified by Moxey et al. in [19] so that NekMesh can operate the conversion to a high-order mesh. Only two types of three-dimensional elements can be used in the computational domain: tetrahedrons and prisms. Prisms are used for the boundary layer discretisation while tetrahedrons are used in the rest of the computational domain.
The linear meshing process in StarCCM+ unfolds as follows. First the CAD definition is imported and a surface mesher triangulates the geometry definition following the various surface refinements imposed. Then the computational domain is discretised into tetrahedrons whose size is dictated by volumetric refinements settings. Finally, layer of prismatic elements of specified thickness is generated at the interface between the surface mesh consisting in triangles and the tetrahedral volume mesh. It is important in the process, especially for projection, that only one layer of prismatic cells is generated and splitting is performed directly by NekMesh during the high-order conversion process.

2.2.3 Linear mesh deformation into a high-order mesh

High-order meshes are usually coarser than their linear counterparts for a given geometry due to p-refinement increasing the total number of degrees of freedom compared low-order methods. Thus, there often remains discrepancies between the linear mesh resulting from the

previous step and the actual geometry. High-order meshing alleviates the issue by projecting the layer of linear prismatic elements interfacing with the surface mesh upon the CAD surface, creating in the process high-order elements with curved boundaries and a body-conforming mesh. First, NekMesh triangulates the CAD geometry. For each linear surface mesh node, the algorithm then searches for the closest node in the CAD triangulation. If the node is not on the CAD surface and its distance to the closest vertex of the CAD triangulation lies above the set node adjustment limit, the edges attached to it will not be curved. Else, the surface mesh vertex will be moved to coincide with the triangulated corresponding vertex in the CAD triangulation. Once all the surface mesh nodes have been processed, high-order nodes are inserted on linear mesh elements and projected onto the CAD surface, resulting in high-order curved prism elements. The algorithm determining the placement of high-order nodes along the linear mesh elements can be based on spring energy optimisation with the mesh being assimilated to a spring network system, which improves the quality of the generated high-order meshes. More information about the optimisation process can be found in Sherwin and Peiro [17]. The number of high-order elements inserted on linear surface mesh elements is set by the user and determines the order of the mesh. A higher order mesh allows one to capture more accurately complex geometrical features, but implies higher order quadrature rules to integrate equations in order to avoid geometrical aliasing as will be shown later in section 2.4.2.

The next step is a bad element detection routine to detect faulty elements generated by the projection. It is based on the quality metrics defined by Johnen et al. [23] for high-order triangular surface elements and Gargallo-Peiro et al. [24] for high-order tetrahedral elements and quantifying the distortion compared to a corresponding ideal reference element, that is the Jacobian of the shape functions applied in the mapping. The ideal reference element corresponds to the high-order element without curvature along the edges. Negative Jacobian values indicate invalid elements which are then transformed back into their original

linear shape. The high-order conversion process ends by splitting the layer of prismatic cells into several layers with an iso-parametric approach with a geometric growth rate between successive layers preventing from self-intersections as described by Moxey et al. [19] and illustrated on figure 2.6. It is important to note that the process requires both the linear mesh and the CAD definition.

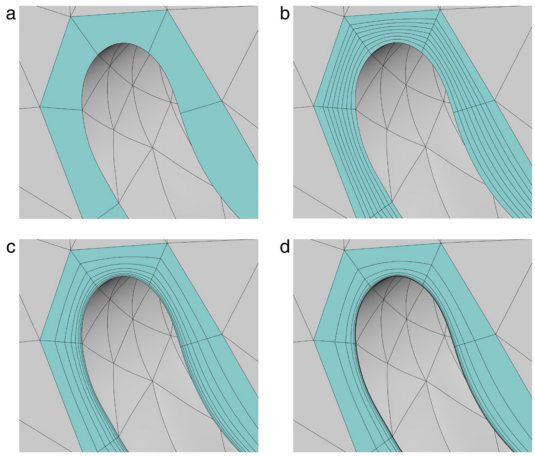


Figure 2.6: Splitting of a macro-boundary layer into eight sub-layers with different geometric growth rates. (a) Macro boundary-layer. (b) Growth rate of 1. (c) Growth rate of 1.5. (d) Growth rate of 2. From Moxey et al. [19]

2.3 Velocity correction scheme

The velocity correction scheme (VCS) was first introduced by Karniadakis et al. [25] and stability and convergence were analysed later by Guermond and Shen [25]. Let us start by writing the Navier-Stokes equations in the following form:

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0 \tag{2.21}$$

$$\frac{\partial \boldsymbol{u}}{\partial t} = \boldsymbol{N}(\boldsymbol{u}) - \nabla p + \nu \boldsymbol{L}(\boldsymbol{u}) \tag{2.22}$$

with

$$\boldsymbol{N}(\boldsymbol{u}) = -(\boldsymbol{u}\cdot\nabla)\boldsymbol{u} \qquad \text{and} \qquad \boldsymbol{L}(\boldsymbol{u}) = \nabla^2\boldsymbol{u} \tag{2.23}$$

where $\boldsymbol{u}$ is the velocity field, p is the pressure divided by the density and ν is the kinematic viscosity. The idea underpinning this method is to integrate the non-linear advection term $\boldsymbol{N}(\boldsymbol{u})$ explicitly 2.21 on $\boldsymbol{u}$ using Adams-Bashforth discretisation. In a second step, the pressure Poisson equation enforcing the divergence-free condition 2.21 on $\boldsymbol{u}$ is solved and finally this step is followed by implicit integration of a Helmholtz problem on $\boldsymbol{u}$ to implicitly treat the viscous term $\boldsymbol{L}(\boldsymbol{u})$ using Adams-Moulton discretisation. Implicit integration of the diffusive terms prevents one from having too tight time step restrictions to ensure stability for the diffusion equation.

Let us assume the Navier-Stokes equations need to be solved in the domain Ω with the following boundary conditions

$$\boldsymbol{u} = \boldsymbol{u}_D \text{ on } \Gamma_D \tag{2.24}$$

$$\frac{\partial \boldsymbol{u}}{\partial \boldsymbol{n}} = \boldsymbol{u}_N \text{ on } \Gamma_N \tag{2.25}$$

where Γ_D (resp. Γ_N) is the boundary on which the Dirichlet (resp. Neumann) boundary condition $\boldsymbol{u}_D$ (resp. $\boldsymbol{u}_N$) is applied. In the first step, non-linear convective terms are treated explicitly and the time derivative is approximated by a backward differentiation formula

$$\frac{\gamma_0\boldsymbol{u}^{n+1} - \sum_{q=0}^{J_{i-1}} \alpha_q \boldsymbol{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J_e-1} \beta_q \boldsymbol{N}(\boldsymbol{u}^{n-q}) - \nabla\bar{p}^{n+1} + \nu\boldsymbol{L}(\boldsymbol{u}^{n+1}) \tag{2.26}$$

where J_e (resp. J_i) is the explicit (resp. implicit) order of integration and α, β and γ are the integration coefficients of the stiffly-stable scheme. $\bar{p}^{n+1}$ is the pressure field ensuring

incompressibility at t^{n+1} and defined as:

$$\nabla \bar{p}^{n+1} = \frac{1}{\Delta t} \int_{t_n}^{t^{n+1}} \nabla p \, dt \tag{2.27}$$

Let us now introduce:

$$\boldsymbol{u}^+ = \sum_{q=0}^{J_{i-1}} \alpha_q \boldsymbol{u}^{n-q} \tag{2.28}$$

$$\boldsymbol{N}^* = \sum_{q=0}^{J_e-1} \beta_q \boldsymbol{N}(\boldsymbol{u}^{n-q}). \tag{2.29}$$

Equation 2.26 can then be rewritten using these notations:

$$\frac{\gamma_0 \boldsymbol{u}^{n+1} - \boldsymbol{u}^+}{\Delta t} = \boldsymbol{N}^* - \nabla \bar{p}^{n+1} + \nu \boldsymbol{L}(\boldsymbol{u}^{n+1}) \tag{2.30}$$

Karniadakis et al. [25] divided the resolution of equation 2.30 in three sub-steps:

$$\frac{\hat{\boldsymbol{u}} - \boldsymbol{u}^+}{\Delta t} = \boldsymbol{N}^* \tag{2.31}$$

$$\frac{\hat{\hat{\boldsymbol{u}}} - \hat{\boldsymbol{u}}}{\Delta t} = -\nabla \bar{p}^{n+1} \tag{2.32}$$

$$\frac{\gamma_0 \boldsymbol{u}^{n+1} - \hat{\hat{\boldsymbol{u}}}}{\Delta t} = \nu \boldsymbol{L}(\boldsymbol{u}^{n+1}) \tag{2.33}$$

with the following boundary condition:

$$\boldsymbol{u}^{n+1} = \boldsymbol{u}_D \text{ on } \Gamma_D \tag{2.34}$$

$\hat{\boldsymbol{u}}$ and $\hat{\hat{\boldsymbol{u}}}$ are intermediate velocity fields. The velocity field $\hat{\boldsymbol{u}}$ results from a projection of $\boldsymbol{u}^+ + \Delta t \boldsymbol{N}^*$ on the solution space. Two assumptions are now made on the intermediate velocity field $\hat{\hat{\boldsymbol{u}}}$. First, it is assumed to be divergence free, that is $\nabla \cdot \hat{\hat{\boldsymbol{u}}} = 0$ and secondly that is satisfies the impermeability condition, i.e. $\hat{\hat{\boldsymbol{u}}} \cdot \boldsymbol{n} = 0$ on Γ_D. Now applying the divergence

operator on step two gives:

$$\nabla^2 \bar{p}^{n+1} = \nabla \cdot \left(\frac{\hat{\boldsymbol{u}}}{\Delta t} \right) \tag{2.35}$$

which is a Poisson equation ensuring incompressibility at the time step t^{n+1} with the following boundary condition

$$\frac{\partial \bar{p}^{n+1}}{\partial n} = -\frac{\boldsymbol{u}_D \cdot \boldsymbol{n} - \hat{\boldsymbol{u}} \cdot \boldsymbol{n}}{\Delta t} \tag{2.36}$$

However, it was noticed that this boundary condition results in high splitting error. Karniadakis et al. [25] derived a high-order boundary condition with low splitting error:

$$\frac{\partial p}{\partial \boldsymbol{n}} = (\boldsymbol{N}^* - \nu(\nabla \times \nabla \times \boldsymbol{u})^*) \cdot \boldsymbol{n} \tag{2.37}$$

where $(\nabla \times \nabla \times \boldsymbol{u})^*$ is an explicit approximation on the same extrapolation model as $\boldsymbol{N}^*$. The term $\nabla \times \nabla \times \boldsymbol{u}$ is a result of the combination of the formula $\nabla^2 \boldsymbol{u} = \nabla(\nabla \cdot \boldsymbol{u}) - \nabla \times \nabla \times \boldsymbol{u}$ and the incompressibility condition $\nabla \cdot \boldsymbol{u} = 0$. This form is used in order to avoid having to impose an artificial boundary condition on pressure due to the Laplacian operator. Once pressure is computed, the last step is a Helmholtz problem to include viscosity terms and which resolution has been presented in the previous section in one dimension. During this last step, a stabilisation technique called spectral vanishing viscosity is introduced by modifying the viscosity term and helps stabilising spurious oscillations which can affect the stability of the scheme. This technique will now be introduced.

2.4 Stabilisation techniques

2.4.1 Implicit Large Eddy Simulation: Spectral vanishing viscosity (SVV)

Classical turbulence analysis as reported by Sagaut in [26] and Pope [27] show that an estimation of the mesh size to capture the full turbulence spectrum based on the ratio between the largest scales which are typically the size of the studied body and the smallest dissipative scales, the Kolmogorov microscales, scales as $Re^{\frac{9}{4}}$ degrees of freedom where Re is the free stream Reynolds number. Automotive applications involve high Reynolds numbers of order of about 10^6. Therefore capturing the full turbulent spectrum is currently out of reach in terms of computational resources available, at least in an industrial environment where the mesh resolution employed fails at capturing the smallest wavelengths present in the flow. Wavelengths which are not captured are called subgrid-scales (SGS) and can result in high frequency spurious oscillations which are a major source of instability in high-order simulations. Indeed, high-order schemes such as the spectral/hp element method have exponential convergence properties implying low diffusion and fail at attenuating these oscillations. Hence the need to artificially damp the oscillations. The spectral vanishing viscosity (SVV), introduced in high-order methods by Karamanos et al. [28] Kirby et al. in [29], is an artificial diffusion operator added in equation 2.22 for stabilisation purposes, cancelling out the effect of such oscillations by removing the energy at the end of the turbulent cascade. The SVV is not an explicit LES model but it rather *implicitly* models the dissipative action of SGS on lower-frequency resolved scales hence the implicit LES terminology. Still SVV must be tuned not to degrade exponential accuracy properties of high-order schemes. Moura et al. [30] analysed spatial dispersion-diffusion properties of the continuous Galerkin (CG) scheme, which is the same as the one implemented in the incompressible solver of Nektar++, on the linear advection-diffusion equation in one dimension. They used this framework to optimise the

SVV operator in order to match exponential accuracy properties close to the discontinuous Galerkin scheme while fixing instability issues on high-order discretisation.

2.4.1.1 Introduction to the SVV in one dimension

The SVV concept was first presented by Tadmor [31] as a solution to damp oscillations appearing on solutions of non-linear conservation laws exhibiting shocks. Tadmor took the example of the Burger's equation in one dimension to prove the efficiency of this approach. Let us consider the one-dimensional Burger's equation:

$$\frac{\partial u}{\partial t} + \frac{1}{2}\frac{\partial u^2}{\partial x} = 0 \tag{2.38}$$

with appropriate boundary conditions which will not be detailed here. It is well known that this problem admits many weak solutions and that an entropy condition must be specified to select the physically relevant solution:

$$\frac{1}{2}\frac{\partial u^2}{\partial t} + \frac{1}{3}\frac{\partial u^3}{\partial x} \leq 0 \tag{2.39}$$

Moreover, Gibbs oscillations can form around the shock location. The idea of Tadmor is to add an artificial dissipation term which can be tuned ensuring both the relevance of the solution from a physical point of view and the smoothness of the solution:

$$\frac{\partial u}{\partial t} + \frac{1}{2}\frac{\partial u^2}{\partial x} = \epsilon \frac{\partial}{\partial x}\left(Q_\epsilon \frac{\partial u}{\partial x}\right) \tag{2.40}$$

where Q_ϵ is a viscosity kernel. As mentioned before, the interest of this method is to be able to adapt the dissipation to each mode. Therefore for clarity, let us suppose we can decompose the discrete solution basis in N modes and let us write equation 2.40 for the N modes:

$$\frac{\partial u_N}{\partial t} + \frac{1}{2}\frac{\partial}{\partial x}(P_N u^2) = \epsilon \frac{\partial}{\partial x}\left(Q_N \circledast \frac{\partial u_N}{\partial x}\right) \tag{2.41}$$

where P_N is a projection operator onto the modal basis which is a Fourier basis in Tadmor's work, Q_N is a modal viscosity kernel and $\circledast$ is the convolution operator. As mentioned before, Q_N is tuned as a low-pass filter: it damps high frequencies and does not affect lower frequencies. The convolution product of Q_n and u_N in the case of a Fourier decomposition can be written:

$$\epsilon \frac{\partial}{\partial x}\left(Q_N \circledast \frac{\partial u_N}{\partial x}\right) = -\epsilon \sum_{M \leq |k| \leq N} k^2 \hat{Q}_k(t)\hat{u}_k(t)e^{ikx} \tag{2.42}$$

where $\hat{Q}_k(t)$ (resp. $\hat{u}_k$) are the Fourier coefficients of the decomposition of Q_ϵ (resp. u) and M is the cut-offmode from which the SVV becomes active. In the work of Tadmor, the discrete dissipation kernel is defined as

$$\hat{Q}_k = \begin{cases} 1 & \text{if } |k| \leq M \\ 0 & \text{else} \end{cases} \tag{2.43}$$

This technique gave very good result and enabled the authors to remove oscillations of the solution around the location of the shock. It was shown that the resolution of the SVV is improved when the viscosity kernel is C^∞-continuous and several definitions for the SVV kernel have emerged such as the widely used exponential kernel originally proposed by Maday et al. [32]:

$$\hat{Q}_k = \begin{cases} exp\left(-\frac{(k-N)^2}{(k-P_{SVV})^2}\right) & \text{if } k > P_{SVV} \\ \left(\frac{k}{N}\right)^{P_{SVV}} & \text{else} \end{cases} \tag{2.44}$$

where P_{SVV} is the cut-off mode. The technique has then been adapted to spectral/hp methods for LES simulations in the works of Karamanos et al. [28] and Kirby et al. [29].

However, when it comes to applying the CG approach to problems where velocity values and mesh spacing vary throughout the domain, the exponential kernel previously mentioned do not consistently improve stability in the whole domain with increasing polynomial order. The DG kernel has then been developed to overcome this difficulty.

2.4.1.2 The DG kernel

Moura et al. [30] have observed non-smooth dispersion/diffusion features when applying the CG formulation on the one-dimensional advection/diffusion equation in cases where dispersion or diffusion are predominant and that this undesired behaviour persisted even when the exponential kernel was employed. The authors have remarked that this was due to the varying Péclet number, which is the equivalent of a local Reynolds number for the advection/diffusion equation, throughout the domain. The Péclet number Pe depends both on the local advection velocity in the element a, the diffusivity value Γ and the characteristic size of the element h:

$$Pe = \frac{ah}{\Gamma} \tag{2.45}$$

Therefore, depending on the value of the Peclet number, the exponential kernel could be out of its optimal operating point. A constant Péclet number SVV formulation where the scaling factor ϵ in equation 2.40 has been normalized by the local velocity and the element size has been developed ensuring that the SVV kernel operates accordingly to its original design in all the elements of the domain. In a second work, Moura et al. [33] noticed that dissipation bubbles formed in certain frequency ranges in the limit of large Péclet numbers (analogously large Reynolds numbers for the Navier Stokes equations) due to reflections waves and could lead to non-smooth dissipative behaviour in under-resolved simulations. An optimisation has been performed on the dissipation curves of the SVV to counter this effect. The dissipation properties of the Discontinuous Galerkin (DG) scheme have been targeted as it is numerically robust and exhibits low dissipation when the polynomial order is increased.

Yet, the DG dissipation curves were to steep due to DG superconvergence properties and could not be fitted with the CG scheme. The author proposed to match the DG curves at polynomial order $P_N - 2$ with the CG curve at polynomial order P_N in the frequency space in log-log scale. The figure 2.7 shows the DG diffusion curves (in colour) and the CG diffusion curves with SVV (in black) at very high Péclet number ($Pe = 1 \times 10^{10}$). It shows that the optimised diffusion curve without SVV for CG at $P_N = 3$ matches at high frequency the DG curve at $P_N = 1$ and that the CG diffusion curve without SVV at $P_N = 7$ matches the DG diffusion curve at $P_N = 5$. It also shows that the targeted DG diffusion curves completely remove the undesired diffusion bubbles observed, which is not the case for higher polynomial order DG diffusion curves.

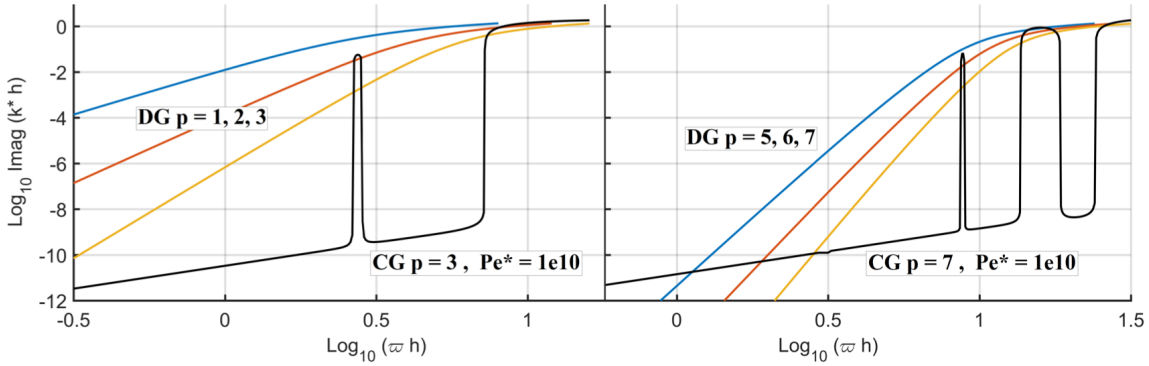


Figure 2.7: Diffusion curves in the frequency space for the CG scheme (in black) at polynomial order $P_N = 3$ and $P_N = 7$ and for the CG scheme (in black) at P_N, P_{N-1} and P_{N-2} (after [33])

The optimisation have been performed with a swarm particle algorithm on the SVV kernel introduced by Moura et al. in [30] to reduce the distance between the CG diffusion curve at polynomial order P_N and DG diffusion curve at polynomial order $P_N - 2$ in the log-log scale in the frequency space for each polynomial order up to $P_N = 8$. Numerical experiments proved excellent agreement between CG with SVV and DG at high polynomial order resulting in superconvergence of the CG dissipation curves with SVV with a slope on $2P_N - 2$.

2.4.2 Dealiasing

Aliasing in spectral methods is caused by inaccurate treatment of high-order non-linear terms, such as the quadratic convective term $\boldsymbol{N}$ in the Navier Stokes equations defined in equation 2.23 and is a common source of instability. In particular, high-order modes which are not consistently integrated by a low-order quadrature rule create spurious energy affecting lower-order modes which is at the source of the instabilities. The super-collocation method presented by Kirby and Karniadakis [34] consists in using a quadrature rule of order Q higher than the polynomial order P representing the solution to avoid aliasing. For instance, a polynomial of order $2P$ requires at least $Q = P + \frac{3}{2}$ quadrature points to be consistently integrated with the Gauss-Lobatto integration rule. In particular, if the solution u of the Navier Stokes equations (equation 2.22) is approximated by a polynomial of order P, then $\boldsymbol{N}(u)$ will be a polynomial of order $2P$ and a quadrature rule of order $Q = P + \frac{3}{2} > P$ is required to consistently integrate this term and maintain the accuracy of the scheme.

There is a second source of aliasing arising from the geometrical mapping. Indeed, the Jacobian resulting from the non-linear isoparametric mapping between the physical space and the standard region is also a polynomial representing the deformation of the elements on the curved boundaries of the studied body which order is usually high for complex industrial geometries with highly curved surfaces (typically between 3 and 6). The order of the Jacobian adds up to the order of the integrand and must be taken into account in the choice of an appropriate order for quadrature rule. This is often referred to as geometrical aliasing. Dealising techniques for high-order spectral/hp in the case of industrial configurations have been discussed by Mengaldo et al. in [35] and will be addressed again later in this work in the case of the model car.

2.4.3 Conclusion

The general spectral/hp element framework, on which Nektar++ is based, has been presented in this chapter. Yet, the application of the methodology to industrial CFD cases with complex three-dimensional geometries at high-Reynolds number is a real challenge. Three aspects have then been developed to address these complex problems. The first one is high-order meshing, taking advantage of high-order polynomial decomposition in the spectral/hp element method to discretise the fluid domain. The second one is the time-marching VCS applied to the incompressible Navier-Stokes equations. The last aspect is stabilisation of the spectral/hp element method. The SVV is designed to damp high-frequency modes appearing at high Reynolds numbers which could potentially deteriorate computation stability. Dealiasing has also been introduced to control the effect of high-frequency non-linear modes, coming from Navier-Stokes equations, on low-frequency modes.

In the next chapter, the methodology is implemented on an industrial case proposed by McLaren Automotive and a workflow is developed with Nektar++ to systematically address complex automotive cases.

Chapter 3

McLaren Automotive geometry: P14

3.1 Literature review

3.1.1 Previous applications of the spectral/hp element method for automotive simulations

Present industrial CFD codes are based on a low-fidelity approach to simulate the mean effects of turbulence by solving the Reynolds Averaged Navier-Stokes (RANS) equations. Although this approach comes at relatively low computational costs, the whole turbulent spectrum modelling approach may lead to misprediction. On the other hand, eddy resolving techniques such as LES are mainly used in combination with wall functions to model the inner layer small turbulent scales in order to keep reasonable computational costs at high Reynolds numbers. Hupertz et al. [36] have compared a wide range of RANS and wall modelled LES (WMLES) simulations on the DrivAer model and shown that WMLES importantly improves wake prediction, especially behind the wheels. Aljure et al. [37] have compared wall resolved LES (WRLES) and WMLES on the DrivAer fastback model. The study has shown that wall modelling effectively allows for coarser mesh to be used resulting in about 70% reduction in

computation time compared to WRLES while preserving accuracy in regions of smooth geometry transition. Yet, the authors point out that the model still fails in complex configurations such as flow separation. When it comes to wall WRLES, Dairay et al. [38] have compared traditional explicit LES approach using the famous Smagorinsky sub-grid scale (SGS) model to the iLES approach using a SVV operator on a three-dimensional Taylor-Green vortex problem. Dairy et al. have shown that the SGS model requires mesh spacings much smaller than the filter size in order to control spurious oscillations coming from small scales affecting the rest of the turbulent kinetic energy spectrum. Therefore, mesh requirements with the LES approach are very close to direct numerical simulations (DNS) mesh requirements, which is unaffordable in an industrial context. On the other hand, the SVV approach allows for more efficient control of SGS with limited impact on the rest of the spectrum, resulting in smoother solutions and more accurate vortical structure prediction. This approach is particularly well suited for high-order methods by tuning the SVV to impact higher-modes more importantly as shown in section 2.4.1. Mesh refinements to track vortical structures on complex industrial geometries using strongly dissipative low-order methods make this type of simulations expensive even on modern supercomputers according to Wang et al. [39]. Low numerical diffusion properties of the spectral/hp method enables it to propagate vortices over longer distances with the same number of degrees of freedom compared to their low-order counterparts which is particularly appealing when it comes to unsteady turbulent vortex dominated flows encountered in automotive applications. Another advantage of high-order methods mentioned by Wang et al. in [39] is their computational efficiency. The more compact discretisation of equations reduces the amount of communications between processes during computations which account for the largest part of computation time. Scale-resolving techniques are increasingly being used in current industrial development processes due to their high demand in computational resources. In addition, low-order techniques tend to be inherently more stable due to the substantial amount of numerical diffusion associated

whereas stabilisation techniques, as those shown in section 2.4, have to be implemented for robust high-order simulations as well as high-order meshing capabilities introduced in section 2.2. A few attempts of high-fidelity simulations of automotive relevant cases with Nektar++ have been carried out and will be reviewed in this section.

First simulations of this type with Nektar++ have been performed on geometries adapted from Formula One. Hoessler [40] simulated the underfloor of a Formula One car at two Reynolds numbers $Re = 100$ and $Re = 250$ based on the typical vortex radius which is 1 cm and compared it to a DES simulation at a Reynolds number of 24,000. The reason for running at two order of magnitude lower Reynolds number values in Nektar++ is explained by the lack of robust stabilisation technique at the time, especially the SVV, preventing from attenuating instabilities caused by unresolved sub-grid scales at high Reynolds number. Results are fairly consistent with DES predictions and Nektar++ was successfully able to convect vortices created by underfloor appendices further downstream than DES even if the vortices are initially less energetic compared to DES, possibly due to lower Reynolds number values employed. The embedded DNS approach introduced by Hoessler which is a local approach consisting in simulating a small portion of the car due to limited computational resources has shown its limits. Indeed in this type of configurations, the upstream flow is unsteady and highly perturbed so its accurate representation as an inlet boundary condition for the domain of interest proves to be a real challenge. Lombard [41] has used more recent developments improving the stability of high Reynolds number simulations with Nektar++ and more particularly the SVV to carry out a fourth polynomial order simulation of a half-body Formula One front section comprising the front wing and the front wheel at a Reynolds number of 200,000 based on the front main wing chord. Although the near wake of the front wing has been qualitatively well predicted, poor agreement has been observed between simulations and experimental results downstream from the wheel which can be partly explained by important approximations in the geometrical representation of the wheel. Buscariolo et

al. [42] and Buscariolo [43] isolated a 50%-scale Formula One front wing called the Imperial Front Wing (IFW) on a rolling belt to model a moving car in a wind tunnel and compared the results with Nektar++ using third, fourth and fifth polynomial orders at a Reynolds number of 220,000 based on the main element chord. Simulations using third order polynomials have marginally resolved vortical structures and separations on the wing elements and good agreement with PIV data and surface flow observations using painting visualisation techniques has been obtained from fourth order polynomial discretisation. The highest level of agreement with the experiment has been reached with fifth order polynomial expansions with the less energetic vortical structures being resolved and better resolution of near-wall flow and separation leading to improved drag predictions.

Automotive bluff bodies are used to mimic the most salient characteristics present in automotive flows while keeping the level of geometrical complexity relatively low and are therefore perfect candidates to assess CFD methodologies before implementing them on industrial cases. An example called the SAE body is presented in Chapter 4. Buscariolo [44], [43] performed CFD simulations of the well-known Ahmed body first introduced by Ahmed et al. [45]. The concept idea of the Ahmed body came from previous observations of Janssen and Hucho [46] who remarked strong variations of road cars drag with changes in the angle of inclination of the rear window. Indeed, Ahmed highlights that most of the drag on road cars is generated by low pressure at the base provoked by the separation of the flow at the rear of the vehicle and that the separation is very sensitive to the slant angle of the back of the body. The variable rear slant angle of the Ahmed body emphasises this phenomenon with different flow regimes associated to different rear window angles and the reader can find more detailed information in [45], [47]. For the complex case of the 25^o slant angle Ahmed body, Buscariolo's investigation [43] shows that even if RANS methodologies are able to approximate forces with 95% agreement with the experiment, flow structures remain poorly predicted. Scale resolving techniques such as LES or hybrid RANS-LES resolve flow

structures more satisfactorily but give less accurate results for the forces compared to the experiment. Simulations of the 25^o configuration are performed with Nektar++ at a Reynolds number $Re = 1.7 \times 10^6$ based on the length of the vehicle with up to sixth order polynomial discretisation for the mesh curvature and for the variables amounting to 21 million degrees of freedom. The sixth order high-order mesh improved the curvature of the front nose resulting in enhanced approximation of the front stagnation. It was found that the finest resolution with a sixth-order polynomial was necessary to accurately capture the complex flow phenomenon at the rear of the vehicle, namely the separation bubble over the back slant and the vortex emerging from the back slant corner. Drag at this polynomial order is underestimated by 13% which is a consequent improvement over the 44% deviation observed in simulations using similar iLES methodology by Minguez et al. [48] and could be partly explained by the missing mounting strut in numerical simulations.

Recently, Mengaldo et al. [49] employed Netktar++ to simulate an industrial car geometry called the Elemental RP1 model [50] with the baseline geometry and with two aerodynamic enhancement devices. Two meshing strategies have been implemented, the first using NekMesh to generate the mesh directly from the CAD definition and the second one using a commercial meshing tool (here StarCCM+) to generate the linear mesh which is then deformed into the the high-order mesh as explained in section 2.2. The second meshing pipeline has turned out to be much more efficient reducing the time to generate a suitable high-order mesh from four months to one week and had no impact on the flow physics compared to the NekMesh embedded approach even though more faulty elements have been turned back into straight-sided elements due to wrong curvature with the second technique. Some simplifications have been operated on the geometry. The wheels have been blanked to avoid having to use intricate modelling for non-axisymmetric parts such as rims. The inlets and outlets of complex diffusers have been replaced by surfaces on which appropriate normal velocities have been imposed and a Helmholtz problem has been solved at the boundaries of the artificially

added surfaces to set the boundary layer thickness. Only half-model has been simulated using the streamwise symmetry plane with a fourth polynomial order approximation at a Reynolds number of one million based on the vehicle length which is half the value of the experimental Reynolds number. Strong vortical structures directly impacting the driver have been identified whereas they are totally absent from the RANS results. The effect of the flow control devices added on the car to counteract the unwanted behaviour and to improve downforce is qualitatively well predicted by the high-order iLES simulation and in agreement with RANS.

3.1.2 Objectives

Previous simulations with Nektar++ have shown that substantial improvement can be achieved in the prediction of automotive flows compared to low-fidelity techniques which are standard in the industry especially when it comes to resolving complex flow features. Recent development in Nektar++, especially the implementation of stabilisation techniques such as the SVV and of a robust high-order meshing pipeline have facilitated the application of spectral/hp methods to high-fidelity simulations of complex industrial geometries. However, limited computational resources have forced authors of previous works to trade-off either in terms of geometry simplification with the use of simplified bluff bodies such as the Ahmed body by Buscariolo et al. [44] and Buscariolo [43] or in terms of lower Reynolds number to reduce mesh refinement requirements to capture small scales and improve stability as in the work of Mengaldo et al. [49].

In this work, we aim at performing simulations of the flow around an automotive industrial configuration adapted from a McLaren Automotive (MLA) production road car using iLES coupled with the spectral/hp element method along with RANS simulations using Star-CCM+ [22], which is the commercial tool used in MLA's numerical development process, at the same Reynolds number which is used as the operating point in the design process of

MLA. The model car is derived from the McLaren 720s shown on figure 3.1 which we will refer to as P14 as it is the internal project reference. The goal of the study is to consolidate and further develop the pipeline partly introduced by Lombard et al. [20] and Mengaldo et al. [49] on an extremely challenging configuration at real operating conditions. The second target is the assessment of the performance of the spectral/hp element method on the P14 model by comparing it with the standard RANS modelling technique. This work particularly focuses on the influence of the flow in the wheel well on the rest of the car. Indeed, the flow around wheel is highly unsteady, turbulent and exhibits large separations making it particularly problematic to correctly resolve for methods based on statistical averaging such as RANS techniques which usually perform poorly on separated flows. Moreover, accurate prediction of the flow downstream from the front wheels is crucial since the engine is placed in a rear position on P14 resulting in engine intakes directly exposed to the flow emerging from the front wheels. For industrial flows at high Reynolds number on complex bodies like P14, the total mesh size is severely restricted by the computational resources and areas which the user is most interested in are prioritised for increase in mesh resolution. For this study, three high-order meshes with different mesh structures have been built. The last goal of the study is to identify the most crucial points to develop to make the new methods suitable in an industrial development process.

First, the model which has been used for the simulations will be presented as well as the workflow which has enabled to perform simulations of the vehicle. Then, the results with Nektar++ will be compared to RANS simulations results to highlight discrepancies in the results between the two modelling approaches. In parallel, the identification of simplifications in the model and domain discretisation will lead to two developments directions to facilitate the deployment of the spectral/hp element method to automotive development process namely the high-order meshing strategy addressed in chapter 4 and accurate wheel modelling investigated in chapter 5.



Figure 3.1: McLaren 720s (or P14) car model from McLaren Automotive from which the geometry is adapted

3.2 Case presentation and workflow

3.2.1 Geometry

The main goal of the study is to prove the feasibility of the airflow simulation around an industrial automotive geometry at high Reynolds number with the spectral/hp element method. On high-performance cars such as P14, flow features emerging from the front of the vehicle dictate the aerodynamics of the whole car especially on crucial parts such as the engine intakes which are placed in rear position or the underfloor which is responsible for an important part of the downforce. Given the complexity of the model and restriction on computational resources, present simulations have been conducted on the front quarter of the car with symmetry and no yaw, as it is the standard in the industry to save computational resources for very complex cases. The influence of the half model simplification has been previously discussed by Kapadia et al. [51] on RANS simulations and DES on the Ahmed body and on an LES of a jet by Schlüter [52]. Authors have concluded that when turbulence levels crossing the plane remain low or when looking at quantities which are averaged over time the simplification is deemed acceptable. Buscariolo [43] conducted simulations of the Ahmed body using both full and half models and deduced that the approximation has little impact on averaged quantities close to the surface of the body which are mostly considered in auto-

motive simulations. The standard convention usually employed for this type of simulations has been used for the axis, the X axis aligned with the streamwise direction from the front to the rear of the vehicle, the Y axis in the spanwise direction positive from the exterior towards the symmetry plane and the Z axis pointing upwards. The origin is placed at the centre of the wheel in the X-Z plane and on the outboard face of the wheel in the Y direction. The geometry is represented on figure 3.2. Aerodynamic optimisation devices characteristic of high-performance vehicles can be seen on the underfloor, such as the guiding vanes or the front diffuser upstream from the wheelhouse.

The rear half part of the car has been removed and replaced by a long tail to smoothly close the body in order to limit induced flow perturbations a sharp end would cause. It was created as a junction surface between the closed curve induced by the trim profile of the front of the car and the symmetry plane and a square surface aligned with the Y-Z plane downstream. The exterior aspect of the front half is very close to the production car even if a few simplifications have been operated. First the radiator in the front duct has been removed. Although it may have an non-negligible impact on the flow feeding the wheel well, it requires particular modelling which is out of scope in this study. The engine intake has been filled with a solid surface at the streamwise location where the car has been trimmed to close the surface.

The wheel well is the region of the car where most of the simplifications have been made. The focus has been laid on the interaction of the wheel rotation and the main body, therefore parts participating in the physical junction between the car and the wheel such as the suspension and the wheel axle creating additional interferences have also been removed resulting in the wheel forming a separate entity. The spokes of the rim have been encapsulated between two surfaces shaped to follow the profile of the spokes to limit the alteration of the geometry. The use of a blanked wheel is a very strong assumption but vastly simplifies the treatment of the wheel rotation which can be imposed using a moving wall technique. Yet,

any crossflow is suppressed what can have an important impact on the flow arrangement around the wheel and thus the whole car. More accurate wheel treatment is necessary for high-fidelity simulation of P14 and is treated in chapter 5.

The junction with the floor has been made through a contact patch which is a standard simplification technique in automotive numerical simulations used in the industry which implications have been discussed in more detail by Diasinos et al. [53]. The wheel is first sunk into the floor for a small distance and the tire is trimmed using the floor as a guide. The wheel is then lifted by the same distance to its original position and a closed surface is created between the floor plane guided by the trim curve. The contact patch can be seen at the bottom of the wheels on figure 3.2. Simulations performed by the authors of [53] on an isolated wheel reproducing Fackrell's pioneering experiments on a isolated wheel [6] showed that the height of the contact patch has a strong influence and must be kept as low as possible. The other constraint comes from meshing. Indeed, the height of the contact patch dictates the maximum cell size and in the limit where this size tends to zero, cells which are generated at the junction between the wheel and the floor are highly skewed which compromises simulation stability. Therefore the contact patch must be kept as small as possible in the limit of meshing capabilities.

Several airflow optimisation devices characteristic of high-performance cars are incorporated in P14. Starting from the front of the car, a horizontal plate called the front splitter slightly extends from the base of the nose in the negative X-direction separating the underfloor from the top of the car. Just above the front splitter at the outboard extremity of the nose is located the first intake of the front duct designed to feed the front radiator - removed in our model. The second intake is located directly below the front light and is thinner and more extended in the spanwise direction. The duct also has two outlets, the first one being visible upstream from the wheelhouse and the second in the wheelarch upstream from the wheel and facing it. The flow in the wheelhouse is normally intended to bring cool air to the brakes to

prevent mechanical deterioration of the brakes assemblies due to the high temperatures at which they operate, especially on high performance cars. On the underfloor, the wheelhouse is preceded by a front diffuser which is shifted towards the symmetry plane in the Y-direction relative to the wheel position. Two fences called guiding vanes are placed between the wheel well and the symmetry plane, which role is to create local negative pressure increasing down-force on the car.

A small surface covers the front door of the car from the underfloor up to the top of the wheelarch following the curvature of the car and narrowing from the bottom to the top. This surface is called the door blade and deflects a part of the flow leaving the wheelhouse towards the small channel between it and the surface of the car. A second channel is visible below the driver's window, starting at the base of the front mirror and extending up to the extension tail. The channel ends at the engine intake and it is blanked by a planar surface at the location where the original model has been trimmed.

An important defeaturing work has been conducted prior to simulations as is the case in standard CFD practice since the final geometry definition which is used for car production contains small geometrical details required for feasibility. These features are removed from the aerodynamic surface as they tremendously add to the complexity of the geometry and therefore represent a challenge for mesh generation and car makers prefer neglecting their impact on the flow field which can be rather limited. Once the surface preparation step is done, the geometry is enclosed within a computational fluid domain and treated as a boundary.

The computational domain is presented in figure 3.3 with all the dimensions based on the wheel diameter D. The model is cut in half at the X-Z symmetry plane and only half of the model is kept in the simulation to reduce computation time which is, as mentioned previously, standard practice in the industry. The section of the computational domain is a

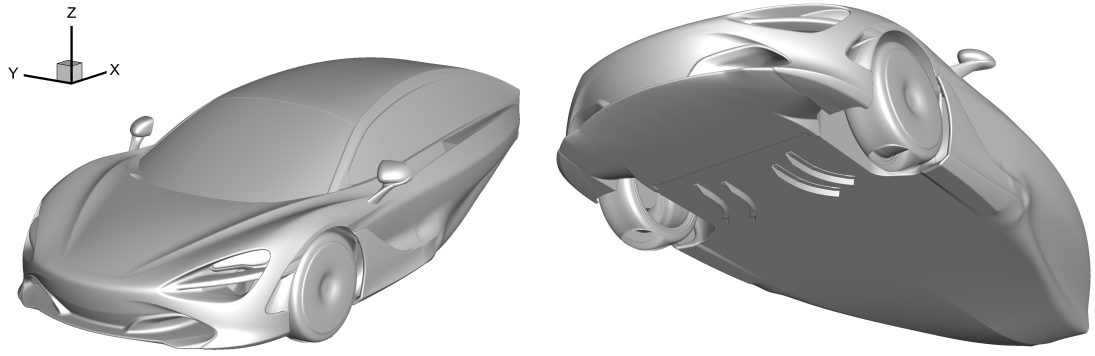


Figure 3.2: P14 model. Simulations have been conducted on half body based on the X-Z symmetry plane

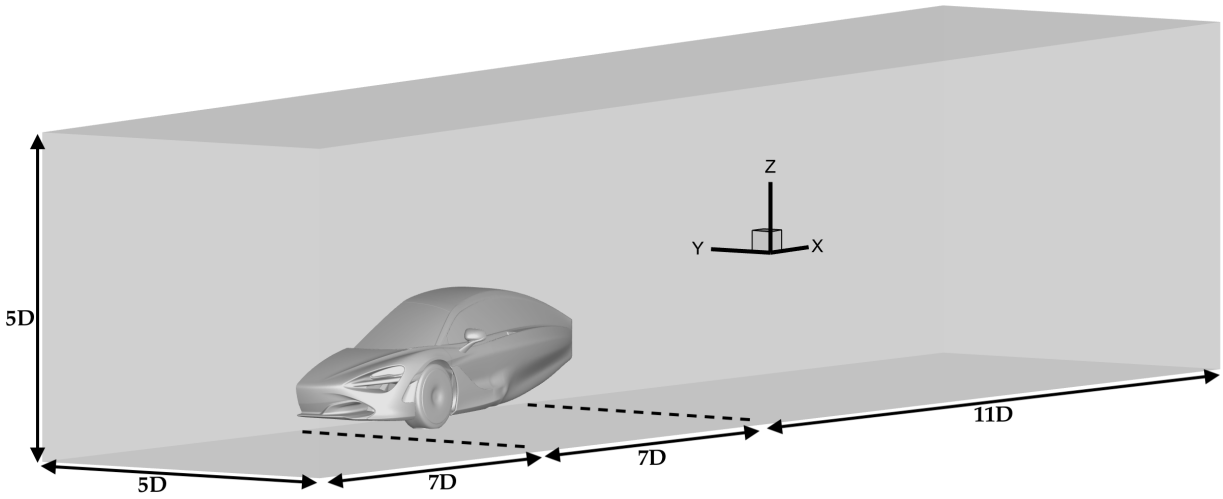


Figure 3.3: Computational domain for P14 simulations. D is the diameter of the wheel

square which edges measure 5D. The total length of the domain is 25D, with 7D upstream distance between the inlet and the car and 11D downstream length between the car and the outlet with the car measuring 7D and the origin of the referential is aligned with the wheel axis.

3.2.2 Mesh generation

The high-order mesh generation process introduced in section 2.2 is now implemented to create the mesh for P14. Starting from the CAD definition of the car in STEP format, a

linear mesh is generated using StarCCM+ where the general layout of the mesh including h-refinements are specified. High-order meshes typically have a lower elements count since the spectral/hp element takes advantage of p-refinement capabilities to increase the number of degrees of freedom within each cell using polynomial expansions. However in practice, P14 entails highly curved and sharp-edged aerodynamic appendices which are intricate to correctly model with a coarse mesh size such as the front splitter or the sharp guiding vanes even after the mesh projection step. Therefore fine mesh refinement is required to model these parts. It restrains p-refinement possibilities due to computational resources limitations. This section introduces the high-order mesh used for spectral/hp element simulation.

3.2.2.1 Linear mesh

The linear base mesh has been generated with the meshing tool embedded within StarCCM+. In order to be valid and converted by NekMesh into a high-order mesh, only tetrahedral and prismatic elements can be used. First, surface refinements are imposed on the car with particular attention to salient edges as well as highly curved features. In particular, the contact patch requires extremely fine mesh resolution due to its small size. Once surface mesh refinement parameters are specified, volumetric refinement boxes are designed to outline the zones of the domain where additional volumetric resolution is desired. By extension of the embedded DNS approach introduced by Hoessler [40] where the precise region of interest is extracted from the geometry and simulated separately, the simulation is performed on a complete geometry and the mesh is set so that the bulk of elements is concentrated in the zones where a more precise insight of the flow physics is sought.

The front of the car on which this study focuses has been favoured compared to the car extension. Mesh resolution has notably been increased around the wheel well where the flow is highly unsteady and perturbed by the rotation of the wheel and due to its strong impact

on the flow downstream and in the underfloor area as it is directly exposed to the wake of the wheel but also to correctly resolve the complex flow dynamics provoked by aerodynamic appendices in this region. Finally, settings for the layer of high-aspect ratio prismatic cells are defined. Splitting of the prismatic layer is operated in the conversion stage with NekMesh so the linear base mesh has to contain only one layer of prismatic cells. Yet, in order to have an accurate evaluation of the total number of cells of the base mesh, a tentative version of the mesh with the final splitting of the layer of prismatic cells into several sub-layers is generated. Three parameters can be tuned to modify the layout of the prismatic layer: the total thickness of the layer, the geometric growth ratio between two adjacent layers and the total number of sub-layers. The total thickness of the prismatic layer is very close to values used by MLA for this car and the sub-layer count has been set so that the global number of elements remains below 8 million, which has been empirically evaluated as the upper limit to allow simulation and post-processing in a reasonable time given computational resources in terms of both CPU and memory limitations.

Wall-resolving simulations require resolution of the flow in the viscous sub-layer. In order to reach this level of resolution near the wall, it may seem appealing to use high values for the geometric growth ratio (GR) between two adjacent sub-layers. Yet, Spalart's study on optimal grid distribution in the boundary layer [54] highlights that using values above 1.5 turns out to be detrimental to the accuracy of the flow resolution in the buffer and logarithmic layers where cells are stretched in the wall-normal direction. Therefore, the value has been fixed to the maximum advised in StarCCM+ best practice guidelines which is 1.5. Spalart also points out that the transition between prisms and elements outside of the layer, which are in our case tetrahedral elements, has to be smooth to ensure good mesh quality metrics what is very demanding. A trade-off has been found by dividing the prismatic layer into 10 sublayers with a geometric growth ratio of 1.5.

Careful attention needs to be paid to the type of elements generated in the meshing process,

especially pyramids which appear when the boundary layer punctually collapses on virtual sharp edges resulting from a insufficient surface mesh refinement of geometric features with a small radius of curvature. The pyramids have been successfully removed by increasing the underlying surface mesh resolution in the close area after highlighting them in the mesh using their identification number which is 14 in StarCCM+. Yet this step requires important amount of handwork and automation is not straightforward.

A baseline RANS simulation has also been performed using MLA standard modelling settings while keeping the same geometry and computational domain to serve as a comparison point for the spectral/hp element simulations. The mesh has been obtained by halving the mesh base size to which all mesh refinements are scaled apart from the prismatic layer settings which are kept identical. The linear base mesh before high-order mesh transformation is shown on figure 3.4. The RANS mesh has not been displayed for readiness since it has the exact same layout as the linear mesh and is only finer away from the car due to the applied scaling. The projections of successive refinement boxes surrounding the front of the car (in green), the underfloor (in yellow) and the wheel region (in red) on the floor can clearly be identified.

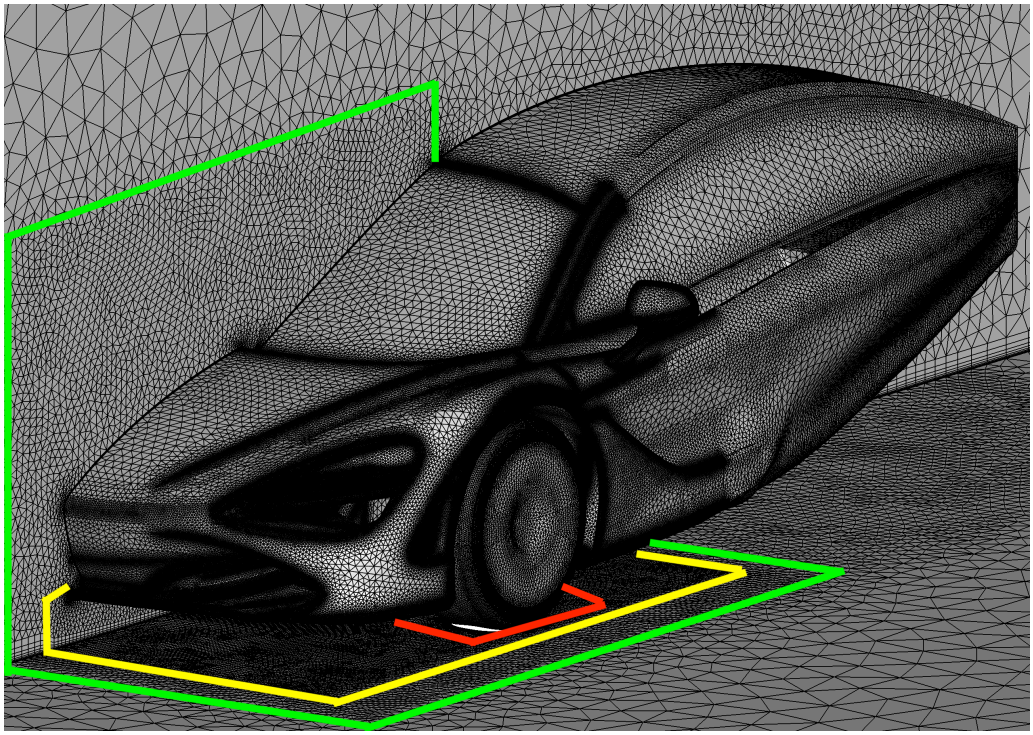


Figure 3.4: General layout of the P14 mesh based on the Nektar++ linear mesh

A streamwise section of the mesh intersecting the wheel is shown on figure 3.5. The mesh refinement boxes indicated in figure 3.4 can clearly be distinguished. A prismatic layer has been generated on the ground. It is thicker than on the car and shrinks when passing underneath. The zoom in the rear wheel region indicates shows better the layers of prismatic cells and how the layer on the floor and the layer on the wheel shrink and join at the contact patch.

3.2.2.2 High-order mesh

A version of the previously introduced mesh without boundary layer splitting is generated to be converted into a high-order mesh. The first step of the conversion process is the projection of boundary elements onto the CAD surface definition, for which the order of the polynomial rule defining the curvature has to be defined. In this case, a fourth order projection rule

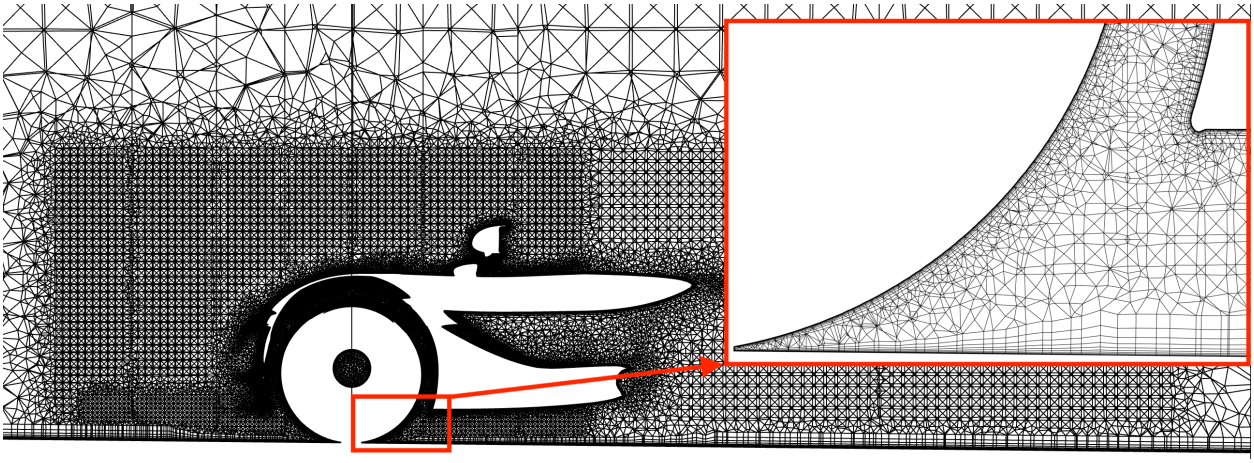


Figure 3.5: Streamwise view of the linear base mesh. The near-wall prismatic layer mesh stucture near the contact patch is highlighted

has been utilised which has resulted in a good approximation of the P14 geometry given the relatively fine surface mesh and prevents from having to use very high-order quadrature rules to address geometrical aliasing which would be computationally prohibitive. The surface mesh obtained is satisfactory in terms of fidelity to the original CAD definition. Invalid elements generated during the projection process are then addressed. The condition derived from best practices is slightly stricter than the negative Jacobian criteria since elements with a Jacobian value smaller than 0.05 are transformed back into their original straight-sided shape. The effect of the projection on the wheelarch is shown on figure 3.6. The surface mesh which looks tessellated on the linear mesh is vastly smoothed by the projection process. Rounded corners with a small radius of curvature that are sharpened on the linear mesh which can have important repercussions on flow physics, causing flow separation which would not exist on the real smooth geometry. A few cells on the zoomed picture appear as straight-sided on the high-order mesh: these are cells which have been turned back to their original linear form by the invalid element correction process. The prismatic layer is finally split into sub-layers using the parameters mentioned in the previous paragraph, resulting in the final high-order mesh. Mesh characteristics are gathered in table 3.1 for more clarity.

More information about mesh parameters can be found in appendix section A.1

$N_{elements}$ (million)	$N_{sub-layers}$	GR	Mesh order	N_{DoFs} (million)
8	10	1.5	4	113

Table 3.1: High-order mesh characteristics

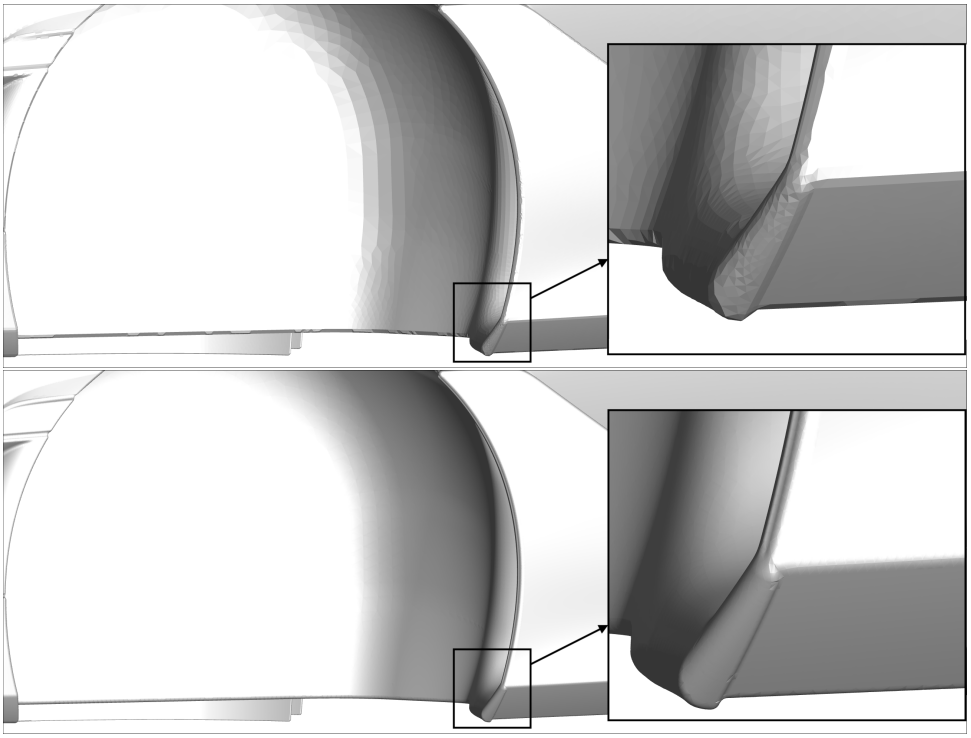


Figure 3.6: Effect of the linear mesh projection upon the CAD surface in the wheelarch region Top: linear mesh. Bottom: high-order mesh

A Y+ map of the car, in particular the underfloor, is shown on figure 3.7. The values are based on the height of the first near-wall layer of prismatic cells, before any polynomial approximation. The maximum y+ value of about 70 is therefore somewhat mitigated by the additional p-type refinement. Although the y+ values for the case remain too high for a wall-resolved simulation at this type of Reynolds numbers which would require significantly more computational resources, this remains acceptable for a proof of concept of the hp/spectral element method on such complex geometry.

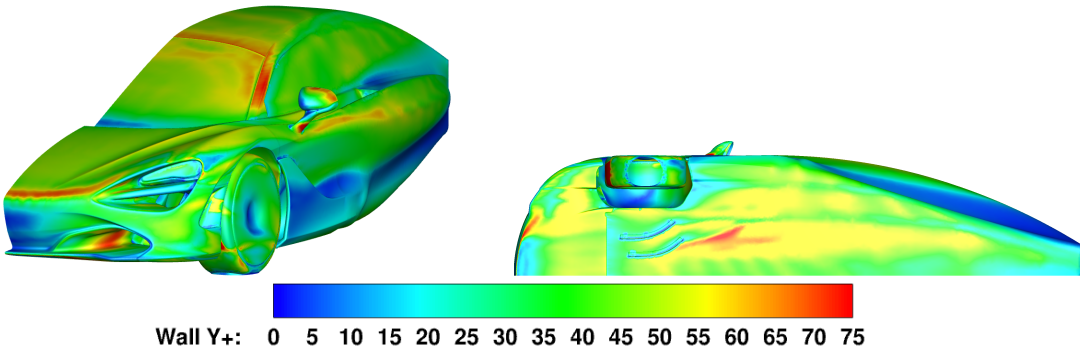


Figure 3.7: Wall y+ distribution based on the fist prismatic cell height for the P14 high-order mesh

3.2.3 Numerical settings

Once the meshing process of the previous section has been completed, numerical parameters need to be set. Boundary conditions are kept the same in both RANS and spectral/hp element simulations to ensure consistent comparisons and are presented simultaneously followed by solver settings which are specific to each methodology.

3.2.3.1 Boundary conditions

Boundary conditions are kept the same both in Nektar++ and StarCCM+. One important specificity of Nektar++ is that calculations are run with normalised quantities based on reference length and velocity values. Although StarCCM+ usually operates with physical quantities, the same normalisation step as in Nektar++ has been performed to facilitate post-processing and direct comparison between the two variants. Velocity values are normalised based on the inlet velocity, therefore a uniform velocity $U_\infty = 1$ is imposed at the inlet in both StarCCM+ and Nektar++. The wheel diameter D is taken as reference length and the numerical model is scaled accordingly so that $D = 1$. The kinematic viscosity ν is then adapted to effectively impose the required Reynolds number $Re_\infty = 2.77 \times 10^6$ as follows

$$Re_\infty = \frac{U_\infty \times D}{\nu} = \frac{1 \times 1}{\nu} \Longrightarrow \nu = \frac{1}{Re_\infty} = 3.61 \times 10^{-7} \tag{3.1}$$

All the walls, excepted the floor are treated as slip walls to limit the influence caused by stagnation over the boundary surfaces on the flow around the car. The floor moves at U_∞ in the positive X-direction by imposing a tangential wall velocity Dirichlet boundary condition. This technique is classical in the automotive industry to mimic the motion of cars which are fixed for wind tunnel testing and where rolling belts are placed over the ground. The tangential velocity boundary condition is also set to U_∞ to match the inflow condition.
Given the simplified blanked wheel geometry, rotation can satisfactorily be modelled by applying a moving wall boundary condition to the wheel and the contact patch. The rotational velocity is defined by $R \times \omega$ where R is the radius of the wheel which is 0.5 after normalisation and ω the angular velocity which is set to 2 in order to recover $U_\infty = 1$ velocity at the junction with the floor. The technique for applying the tangential wall velocity for the wheel is different in StarCCM+ and Nektar++. In the case of StarCCM+, a rotating reference frame centred on the wheel with the required rotation direction and speed is created and the wheel is made part of the reference frame. In Nektar++, the rotational velocity is applied to the wheel by solving a projection problem with the ADR solver embedded within Nektar++, designed to solve a wide range of advection/diffusion equations and in particular projection problems. Further details about the solver can be found in [55]. The procedure, somewhat more intricate than forcing the rotational velocity based on the position of the elements, ensures continuity of the boundary condition that is consistent with the continuous Galerkin formulation utilised by the incompressible Navier-Stokes solver of Nektar++. First, the wheel, the contact patch as well as boundary elements joining the contact patch and the ground are extracted from the mesh. The velocity field associated with rotation $\overrightarrow{U_{rot}}$

is defined at a point $M(x, y, z)$ of the wheel by

$$\overrightarrow{U_{rot}} = \vec{\omega} \times \overrightarrow{CM} = \begin{cases} 0 \\ \omega_y \\ 0 \end{cases} \times \begin{cases} x - x_C \\ y - y_C \\ z - z_C \end{cases} = \begin{cases} \omega_y \times (z - z_C) \\ 0 \\ -\omega_y \times (x - x_C) \end{cases} \tag{3.2}$$

where $C(x_C, y_C, z_C)$ is the centre of the wheel and $\vec{\omega} = (0, \omega_y, 0)$ is the angular velocity vector associated to the rotative motion. The projection problem to solve is $u_i^\delta = U_{rot,i}$ for i=1,2,3 corresponding to each component of $\overrightarrow{U_{rot}}$ with the notation formalism introduced in section 2.1. Resolution is performed with the spectral/hp element method from section 2.1 where continuity is imposed at the interface of adjacent elements, resulting in a continuous field for the velocity boundary condition on the wheel.

The standard outflow modelling used in StarCCM+ is a Neumann boundary condition. In Nektar++, a high-order Neumann pressure boundary condition is imposed at the outlet of the domain. This treatment allows a relatively short distance between the rear-end of the car and the outlet by preventing energy build up due to energetic flow structures impacting the outlet surface as shown by Dong et al. [56].

3.2.3.2 Solver parameters

Given the very different numerical techniques employed in StarCCM+ and Nektar++, solver settings need to be described separately.

3.2.3.3 RANS

The simulation run in StarCCM+ is based on the incompressible RANS equations. They are solved using a segregated solver with second-order upwind discretisation for convective and diffusive terms and a Venkatakrishnan limiter for gradients [57] to improve convergence.

Pressure-velocity coupling is addressed with a SIMPLE algorithm with an under-relaxation factor of 0.7 for the velocity and 0.3 for the pressure in conjunction with a multigrid procedure to accelerate convergence. Turbulence is modelled with the realizable k-ϵ model [58] with a two-layer approach [59] suitable for both low-y^+ ($y^+ \sim 1$) and higher-y^+ ($y^+ > 30$) regions of the mesh which is relevant in our case due to the high difference of mesh resolution between the front of the car which is highly refined and the extension which is coarser. Turbulence intensity at the inlet is set to 1%.

Simulation is initiated with a velocity of U_∞ in the whole domain. Residuals decrease by four orders of magnitude and oscillate with a low amplitude ($\sim 10^{-5}$) that is often observed when steady calculations of wheels with highly unsteady separations are performed [60]. Forces after convergence are averaged over the last 1000 iterations to take into account small fluctuations.

Spectral/hp element method Nektar++ solves the incompressible unsteady Navier-Stokes equations with the added tunable SVV stabilising the numerical scheme. In this case the simulation is unsteady requiring a time-marching scheme and turbulence is resolved rather than modelled with sub-grid scales being treated implicitly by the SVV. Time advancement of the solution is addressed with the second-order velocity correction scheme presented in section 2.3. The explicit part of the scheme forces to fulfil the $CFL < 1$ condition. The timestep is highly conditioned by small-sized cells and has been subsequently adjusted to a value of $\Delta t = 1 \cdot 10^{-5}$ to be as close as possible to this limit to maximise the pace at which the simulation physically advances.

SVV with the DG kernel presented by Moura et al. in [30], [61] used for the P14 simulation is optimised to maintain exponential convergence properties of DG schemes for the CG scheme used in the incompressible Navier-Stokes solver. The technique attenuates sub-grid scales which are not captured by the hp resolution and could deteriorate the stability of the

solution by affecting larger scales which are on the other hand essentially left unaffected. The scaling factor of the SVV is optimised to the lowest value maintaining the stability of the simulation in order to restrain its impact on large scales. The optimal value has empirically been set to be 0.2 which is typical for this type of applications.

The three velocity components are discretised using a third-order approximation and a second-order interpolation is used for pressure forming a Taylor-Hood type expansion ensuring the inf-sup condition is respected and therefore resulting in a well posed problem. The high mesh resolution makes any attempt of increase in polynomial order intractable with the available computational resources. Yet, the polynomial order used in the simulation is deemed sufficient at the proof of concept stage and can be redeemed and easily incremented when larger resources are available. Integration is performed with a fifth-order Gauss-Lobatto rule using five quadrature points in order to correctly integrate high-order expansion terms coming from the convective terms of the Navier-Stokes equations and geometrical transformations. Indeed with the approximation used for velocity components, non-linear terms of the type $(u \cdot \nabla)u$ are of order three. However in the spectral/hp element method, computations are performed on a standard element before being transposed to elements of the physical space by a geometrical transformation and this transformation translates into the inverse of a Jacobian which is in this case a fourth-order polynomial expansion. The subsequent rational function cannot be exactly integrated using numerical polynomial quadrature rules hence the need for a higher-order quadrature rule to counteract the induced geometrical aliasing.

A conjugate-gradient algorithm is used to iteratively solve discretised linear systems with a tolerance of 10^{-11} for velocity components and static condensation with low-energy block preconditioning, described by Sherwin et al. [62] and Grinberg et al. [63] is applied to accelerate convergence. A direct solver is used for pressure due to the linear approximation and because the stiffness of the equation makes low-energy block preconditioning inefficient.

The simulations are initiated from a RANS simulation computed on Star-CCM+ on the linear

mesh with the same boundary conditions. We have recently developed an embedded interpolation routine within Nektar++ to interpolate simulation results coming from third-party software in comma separated values (csv) format into a checkpoint file which can be used to initialise a simulation in Nektar++. The projection of the RANS solution on the linear mesh onto the high-order mesh is used to reduce the transitional time for the flow to develop in the spectral/hp element method which is particularly expensive. After a transitional regime 6 convective time units (CTU) based on the wheel diameter, forces settle to a stable level and results are then time-averaged for a period of $\Delta T = 2$ CTUs. A CTU is the normalised physical time the flow takes to be convected by one wheel diameter based on the normalised inlet velocity, i.e. 1 CTU $= D/U_\infty$. Simulations have been run on CSD3 HPC facility at Cambridge on 2240 CPUs and the wall-time to compute one CTU is 39 hours. The relatively short simulation time is due to the very expensive cost of the simulation. Although this should be run for longer to get more meaningful time-averaging, the simulation has been deemed acceptable for a proof of concept.

3.3 Results and analysis

Simulation results on the adapted P14 geometry described previously both with iLES and RANS modelling are presented and analysed in this section. The emphasis is put on the front of the car as the extension is not representative of the real car. In the first section aerodynamic forces are investigated with particular focus on the underfloor. In fact, drag and lift forces are of primary importance when it comes to car aerodynamics. Aerodynamic drag is the biggest source of resistance to motion past 70km/h and accounts for 80% of total resistance at 100km/h according to Hucho [1] and therefore heavily weighs on fuel consumption. On the other hand, lift impacts vehicle stability and handling characteristics and thus is crucial for high-performance vehicles such as P14. Practically, lift has proved to be more challenging

to numerically predict than drag (see Palaskar and Kumar [64]) especially due to the fact that contributions are spread over larger surfaces compared to drag which is more localised. Moreover, the spatial repartition of lift is also critical since the front/rear lift distribution importantly influences car balance. High-performance car designers have used ground effect since the 70s to increase downforce according to Katz [65] in order to allow higher cornering speeds and concentrating downforce generation on the underfloor. This is the region of the car where the bulk of downforce is created, and several flow structures are generated in this area, on the front diffuser, the wheel or guiding vanes. This is the reason why particularly fine mesh refinement has been applied under the car. The subsequent analysis of the pressure distribution on these parts enables to identify areas from where force discrepancies originate, especially lift given their orientation which is essentially parallel to the ground. A more general overview of the flow structures around the car is finally obtained looking at planes of total pressure contours exhibiting recirculation regions and vortices. The analysis highlights the contrast between the two modelling techniques.

3.3.1 Decomposition of aerodynamic forces

The underfloor of P14 has been decomposed in two main areas: the front floor extending from the front splitter up to the beginning of the wheelhouse and the central floor which joins the front floor to the car extension. The front diffuser has been considered separately from the rest of the front floor to isolate the arising pressure recovery. Forces on the wheel, which is a distinct entity in the P14 model, have also been computed. Comparison between RANS and time-averaged spectral/hp results is shown on figure 3.8. Drag (C_D) and lift (C_L) coefficients are defined as follows

$$C_D = \frac{F_x}{\frac{1}{2}\rho U_\infty^2 S_{ref}}, \tag{3.3}$$

$$C_L = \frac{F_z}{\frac{1}{2}\rho U_\infty^2 S_{ref}} \tag{3.4}$$

where F_x (resp. F_z) is the projection of the integrated forces on the x-axis (resp. z-axis), ρ the density which is set to unity and S_{ref} a reference surface to normalise forces which is also set to unity for consistent comparison. As expected, all the components of the underfloor generate downforce and the central floor is the part where the largest share of downforce is generated, followed by the front floor which generates about half less downforce and the front diffuser which contribution is more than ten times smaller. Downforce predictions are higher in RANS compared to Nektar++ in this case, by 26% on the frontfloor, 15% on the front diffuser and 5% on the central floor. In parallel, drag values are of the same order of magnitude on the front floor and the front diffuser and one order of magnitude lower on the central floor. This can be explained by the fact that the central floor is flat and essentially aligned with the incoming flow reducing the amount of pressure drag acting on it, which is the main source of drag on a bluff body, compared to the diffuser and the front floor which are curved and subject to more important pressure forces exerting in the x-direction, resulting in more important drag values in magnitude. Drag magnitude is greater in RANS than in Nektar++. Yet the trend is not uniform as it is for downforce, the front floor generating negative drag whereas positive drag is obtained on the front diffuser and the central floor. This opposition can allegedly be attributed to the curvature of the front floor, especially near the front splitter which is facing the incoming flow which differs from the diffuser. Drag differences between the two modelling techniques are respectively 22% for the front floor, 18% for the front diffuser and 72% for the central floor. Drag and lift discrepancies between RANS and iLES are consistent on the front floor and diffuser but opposed on the central floor where lift values are relatively close while drag predictions are distant. Still, this has to be put into perspective since the central floor accounts for a negligible part of drag compared to the other surfaces.

Looking at the wheel, the two methods have the same trend with both positive lift and drag which is coherent with literature for shrouded wheels as shown in experimental measurements of Fabijanic [66] of the flow in wheelhouses using a simplified road vehicle model with rolling wheels and of Axon et al. [67] on a wheel enclosed in a wheelhouse. Relative differences are however opposed between the two force components, iLES predicting lower drag and larger lift than RANS. The relative difference is 12% for the drag coefficient and 7% for the lift coefficient which is the overall best agreement among the parts highlighted here.

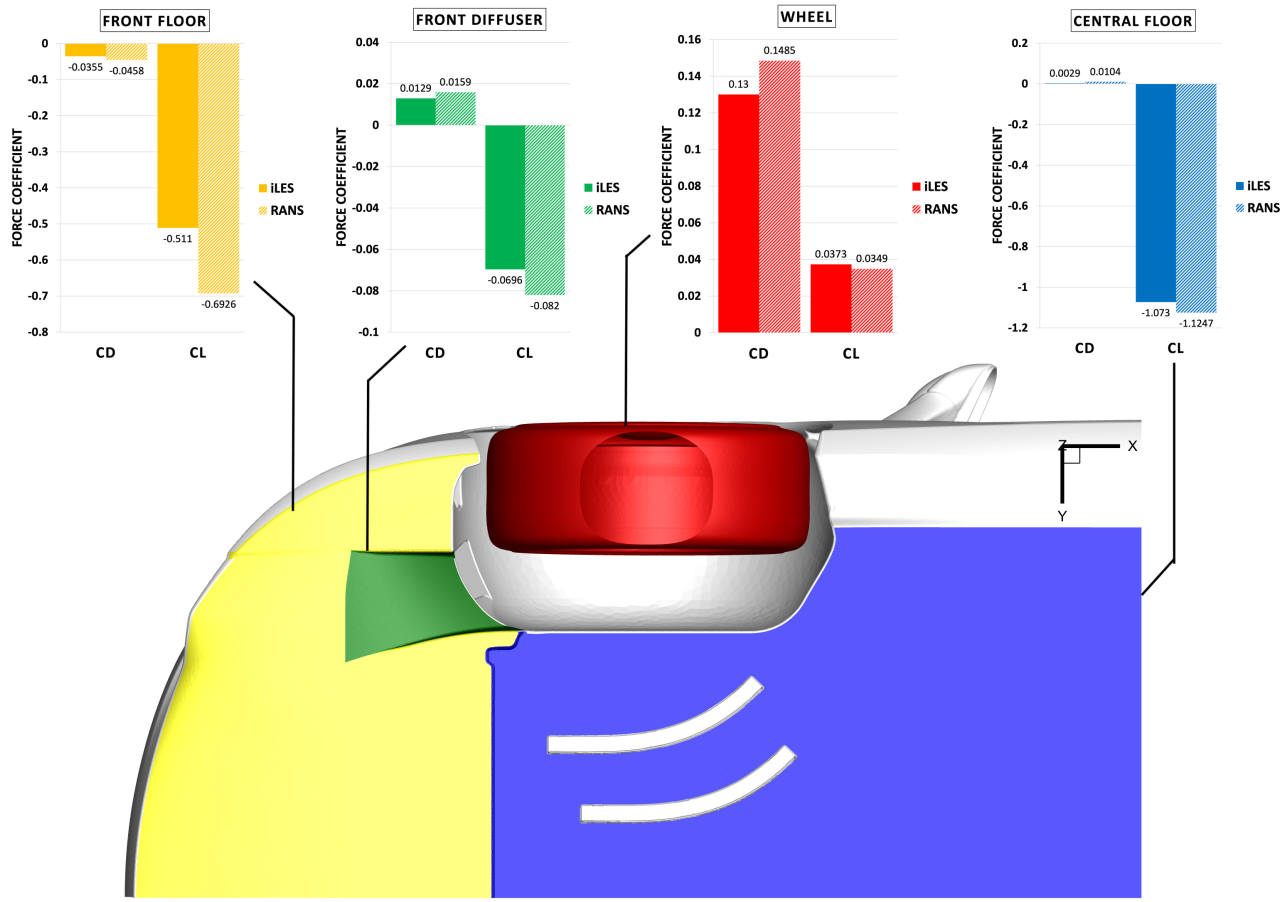


Figure 3.8: Aerodynamic forces distribution on the underfloor of the front part and the wheel of P14. C_L corresponds to the lift coefficient and C_D to the drag coefficient.

The power spectral density of the forces on the wheel is computed for the averaging time

of 2 CTUs and forces are outputted every hundred time steps, giving a sampling frequency $Fs = 833$ Hz. This analysis is only carried out on the iLES simulation which is unsteady whereas it is impossible to operate the same transformation with the RANS steady data. Results are represented in function of the Strouhal number based on the wheel diameter $St_D = \frac{fD}{U_\infty}$ with f the frequency on figure 3.9. Both spectra are regrouped around low frequencies. Two peaks early appear for the drag signal, at $St_D = 0.96$ and its first harmonic at $St_D = 1.9$ which respectively correspond to three and six times the non-dimensional rotation frequency of the wheel $1/\pi$. Both peaks are absent from the lift signal for which the first peak appears later at $St_D = 3.4$. Unmatching peak frequencies suggests that lift and drag generation mechanisms are essentially decoupled on the wheel. For higher frequencies, the signals appear to be more in phase with most peaks matching in terms of frequency. There are very few studies in the literature addressing instantaneous properties of the flow around shrouded wheels at similar Reynolds numbers. Croner et al. [68] evaluated the temporal evolution of lift and drag forces on an isolated blanked wheel using URANS simulations and observed that the two forces are highly correlated with peaks at $St_D = 0.717$. The authors also notice strong correlation between drag/lift and vortical flow structures generation mechanisms. The impact of the shroud on the topology of the flow around the wheel compared to the isolated wheel case could explain the different temporal behaviour in the P14 case. The viscous contribution to forces on P14 is negligible compared to the effect of pressure. The next section focuses on the actual pressure distribution to identify regions causing the differences observed in forces.

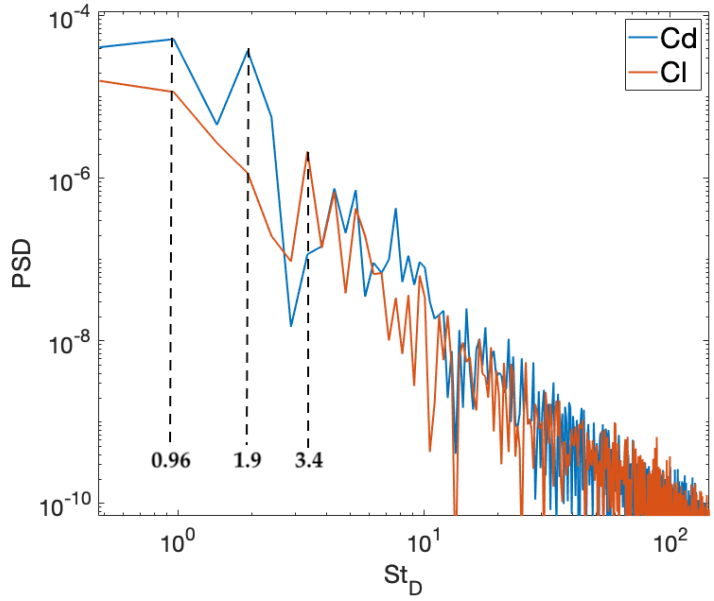


Figure 3.9: Power spectral density of the drag and lift force coefficients on the wheel. The Strouhal number St_D is based on the wheel diameter and defined as $St_D = \frac{fD}{U_\infty}$ where f is the frequency

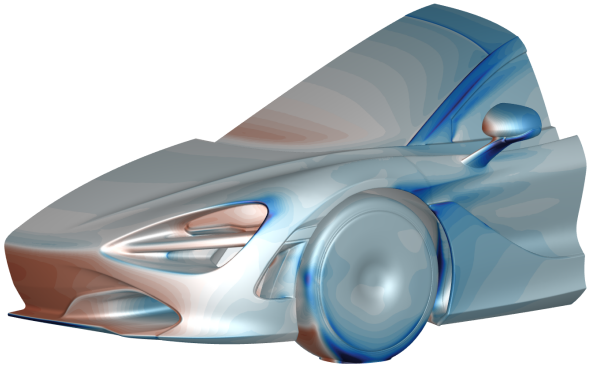
3.3.2 Surface pressure distribution

The pressure coefficient is defined by

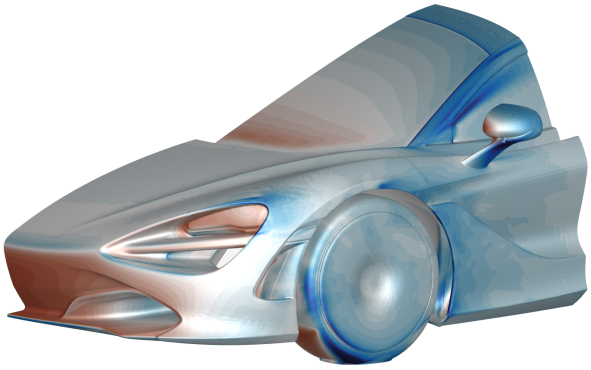
$$Cp = \frac{p - p_\infty}{\frac{1}{2}\rho U_\infty} \tag{3.5}$$

and its distribution, shown for the front of the car on figure 3.10, is directly related to forces as the viscous component is negligible on P14. Results are in overall very similar on the front and the top of the car, except pressure values on the front hood which are slightly lower in the iLES case. There are more noticeable differences on the side of the car yet. The first difference can be spotted at the extreme top edge of the front light, where low pressure values extend further downstream in Nektar++ results than in RANS, where very low pressure values but more spatially restricted are observed. Further downstream, at the

top of the wheel arch, a low pressure area extends on the side of the vehicle up to the base of the side mirror. The low pressure zone starts somewhat earlier in Nektar++ but the magnitude of pressure values is lower than in the RANS case. Pressure values are also lower in RANS at the top of the door blade and the low pressure zone extends further downstream than in the iLES case. There are important variations on the A-pillar joining the top and the side of the vehicle between the wind-screen and the front side-window. A similar low pressure trail is present at the top of the front window although with lower pressure values reported in Nektar++. A thin low pressure streak starts from the base of the A-pillar in RANS and continues up to the end of the wind-screen. Yet in the iLES case, the line is shorter and a second line is visible on the curvature directly upstream from the truncation.



(a) RANS



(b) iLES

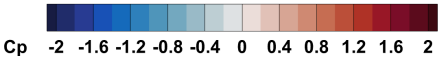
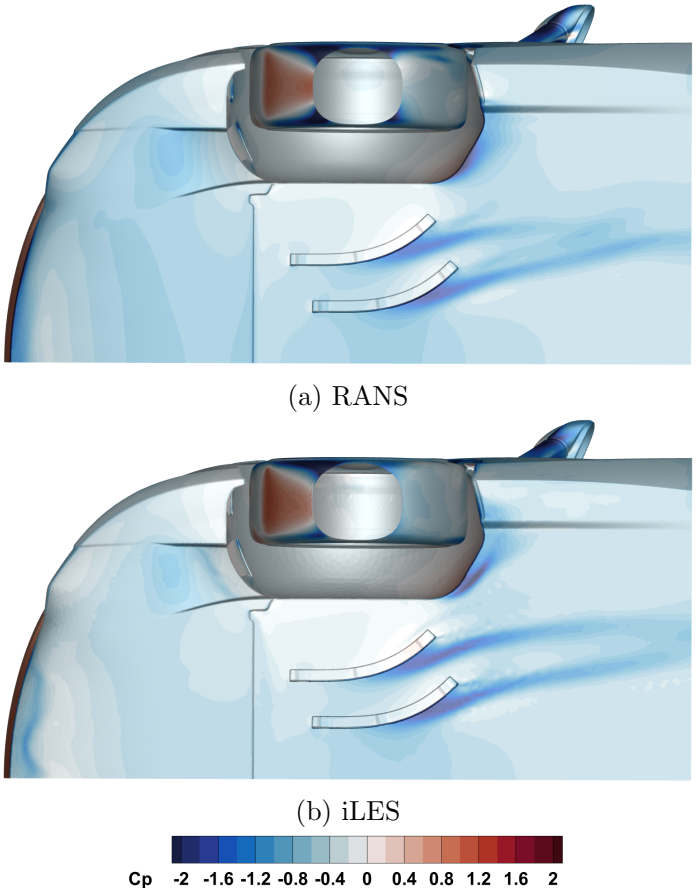


Figure 3.10: Pressure coefficient Cp distribution on the front of P14

As expected, underfloor pressure values are negative hence the computed downforce. The layout of the pressure distributions looks consistent between the two cases, although a few discrepancies can be noticed. Starting from the front splitter, the pressure increase is rapid in the RANS case whereas it is more moderate and delayed in Nektar++ resulting in higher pressure values on the major part of the front floor. The pressure recovery process is monotonic on the front diffuser in RANS with a low pressure peak generating downforce at its start and pressure increasing faster on the outboard side than on the inboard side. The pressure rise with Nektar++ is slower with a less extended low pressure peak at the start. A low pressure band crosses the diffuser interrupting pressure increase before it resumes further downstream up to the base of the diffuser. This causes lower downforce value on the diffuser in Nektar++ compared to RANS.

The most remarkable difference on the central floor is due to the guiding vanes. A low pressure trail forms downstream of each vane in which pressure gradually increases from very low negative values at the vane up to higher values at the junction with the extension. The iLES case exhibits lower starting values. In RANS, the two trails rapidly decay and merge shortly after the middle of the central floor whereas they clearly remain distinct in the iLES case nearly up to the junction with the extension where they finally combine. The flow deviation imposed by the curvature of the vanes results in a higher pressure area between the outboard guiding vane and the wheelarch and lower between the inboard vane and the symmetry plane. In this region, pressures are slightly smaller in the iLES case. Back from the inboard edge of the wheelarch, a low pressure lobe develops on the underfloor and extends hardly longer in the spanwise direction in the iLES case but further downstream up to the guiding vanes trail in RANS ensuing lower pressure values in the last part of the central floor. It is juxtaposed to a high pressure trail downstream from the wheel which despite higher intensity, more rapidly vanishes in RANS compared to the iLES case. The combination of limited differences with overall lower pressure values in RANS on the central

floor result in the light downforce difference between the two cases.



(a) RANS

(b) iLES

Figure 3.11: Pressure coefficient Cp distribution on the underfloor of P14

The comparison between the surface pressure distributions in RANS and iLES enables to map the pressure efforts on the car accounting for the major part of the forces and therefore explain and locate the regions where the downforce gaps originate from. In the next section, the topology of the flow around P14 is analysed to interpret the observed dissimilarities.

3.3.3 Three-dimensional flow structures

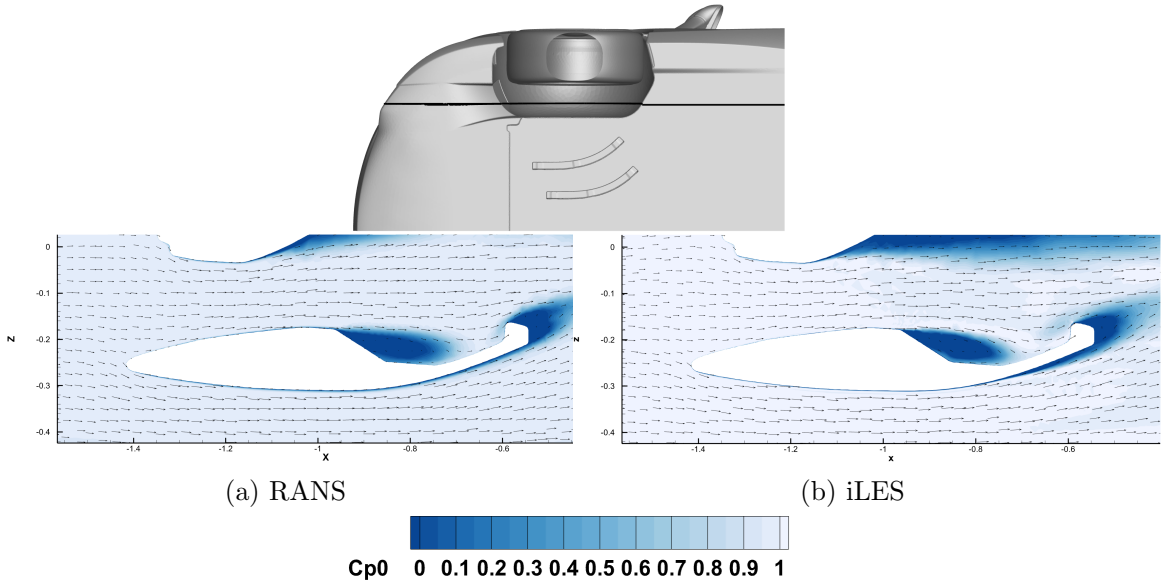
3.3.3.1 Separation on the front diffuser

The first difference which is investigated is the front diffuser, on which a separation line is observed before the end on figure 3.11. The flow in this region has been observed on a Y-normal plane at the spanwise location corresponding to the mid-plane of the diffuser. Total-pressure Cp_0 contours, quantifying energy levels in the flow is defined by

$$Cp_0 = \frac{p + \frac{1}{2}\rho U^2}{\frac{1}{2}\rho U_\infty^2} \tag{3.6}$$

where p is the static pressure and q the dynamic pressure. Cp_0 corresponds to the energy if the flow is virtually brought to rest through an isentropic process i.e. without any energy loss. Cp_0 iso-contours are used by aerodynamicists to highlight flow regions associated with energy deficit and are particularly useful to detect recirculation regions, wakes and vortices. Total-pressure contours are plotted on a plane crossing the front diffuser on figure 3.12 along with in-plane constrained velocity vectors. The front diffuser profile corresponds to a reversed airfoil creating a convergent-divergent airflow section between the car and the road which results in suction below the underfloor due to the ground Venturi effect. A thin dark layer of low-energy flow develops on both sides of the diffuser from the leading edge which corresponds to the boundary layer. Although the thickness smoothly increases up to the trailing edge in the RANS case, the total thickness remains very small in Nektar++ up to $X = -0.75$ where it sharply grows. The flow separates at this location, due to the large adverse gradient imposed by the curvature of the diffuser rear end before the trailing edge in the iLES case whereas it remains attached in the RANS case. The separation causes a pressure drop which curbs the pressure recovery of the diffuser as can be seen on figure 3.11. The absence of experimental measurements prevents from drawing final conclusions at this point, however, the underlying separation process is more deeply explored in the next part of this work, in

the case of the SAE body at the junction between the roof and the backlight (see section 4.4.3)



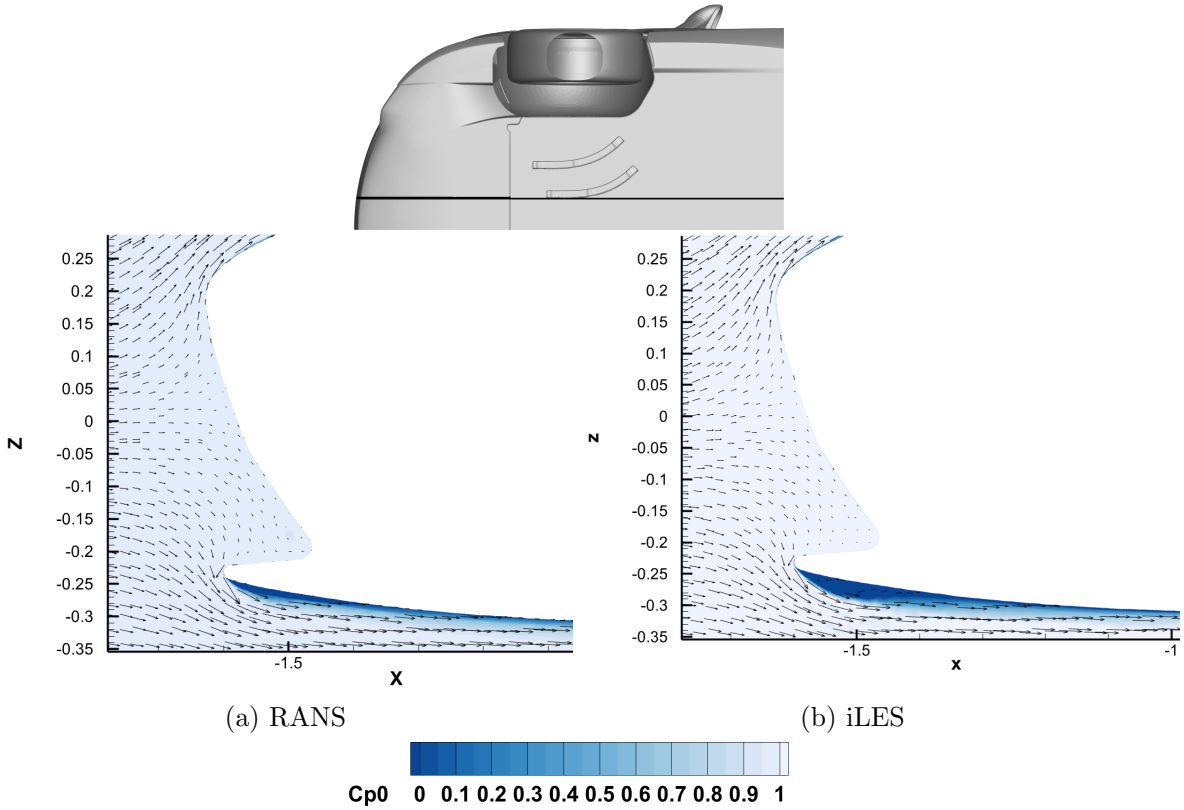
(a) RANS (b) iLES

Figure 3.12: Total pressure coefficient Cp_0 around the diffuser centreline, $X = 0.5$

3.3.3.2 Separation at the front splitter

The delay in pressure rise at the beginning of the front floor downstream from the front splitter in iLES compared to RANS is examined similarly as on the diffuser using total pressure coefficient contours. The total pressure coefficient is represented on a Y-normal plane crossing the vehicle geometry along the black line shown on figure 3.13. In both cases, the dark blue strip of low total pressure coefficient corresponding to the boundary layer first grows rapidly from the leading edge of the front splitter at $X = -1.6$ and then starts reducing in thickness slightly before $X = -1.5$ before finally shrinking to a thin layer at $X = -1.4$ in RANS and later at $X = -1.3$ in the iLES case. The thicker area of low Cp_0 values actually characterises a separation bubble which is larger in the iLES case than in RANS. A closer examination of the surface mesh at this location shows that the high-order mesh projection fails to compensate for the high element size in order to capture the curvature of the front

splitter whereas the reduced size of the elements around the leading edge in the RANS case gives a smoother result. The sharper shape of the leading edge in the iLES simulation is suspected to have a strong impact on the observed separation and therefore the pressure distribution on the front floor.



(a) RANS (b) iLES

Figure 3.13: Total pressure coefficient Cp_0 around the front splitter at $Y = 1.2$

3.3.4 Streamwise evolution of flow structures

The goal of this section is to follow the generation and the evolution of the flow structures generated in the wheelhouse along the vehicle. Not only do they have an impact on the pressure distribution as has been shown on the two previous paragraphs on the front splitter and the diffuser, but also strongly impact the design of the rest of the car since they are convected downstream. The flow is observed in the streamwise normal direction at four

locations from the rear part of the wheel up to the junction with the extension numbered from one to four and displayed on figure 3.14. The progression of the total pressure coefficient Cp_0 as well as in-plane constrained velocity vectors along the car is particularly useful to track down flow structures and RANS and iLES results are compared on figure 3.15 to 3.18.

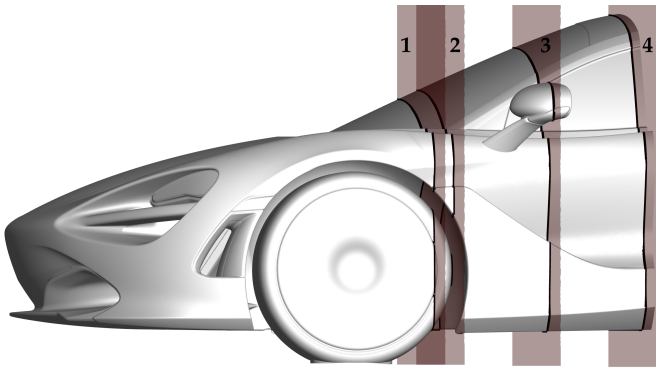


Figure 3.14: Location of the total pressure coefficient X-normal slices along P14

Plane 1 The first plane is located at the rear half of the wheel where the rotation imposes an upwards vertical velocity boundary condition at its surface. At this location, the flow essentially exits the wheelhouse. Spanwise jetting in the negative Y-direction is observed from the wheel arch at the top of the wheel although the direction of the jet is horizontal in the RANS case and rather inclined at 45° towards the top in Nektar++. The different directions can be related to the different application of the moving wall boundary condition on the wheel, where velocity vectors at the surface of the wheel and slightly above are orientated upwards in the iLES case whereas they are slightly slanted towards the negative Y-direction in RANS resulting in an upwards flow pattern in iLES. The jetting discrepancy engenders vast dissimilarities in the outboard wake in dark for $Y < 0$. The separation at the extreme top edge of the wheel arch results in a streamwise travelling vortex which is 45° inclined towards negative Y-values following the jet and therefore largely detached from the surface of P14 in the iLES case whereas it remains close to the surface above of the wheel arch in RANS. The

more important flow entrainment in the RANS case at this particular location can explain the lower negative pressure line observed on the sharp edge above the rear part of the wheel arch in figure 3.10. In the iLES case, a second weaker vortex is formed symmetrically to the first vortex with respect to the main jetting line and is absent from the RANS results. The pair of slanted vortices increases the width of the low energy flow lobe outboard from the top of the wheel arch compared to RANS. The difference in application of the moving wall boundary condition has a second effect outboard of the wheel. The downwash motion of the flow travelling along the side of the car opposes to the upwards motion of the wheel. The flow remains attached to the wheel in RANS whereas it detaches from the surface shortly before the centre of the wheel at $Z = 0.05$ generating a bulkier low total pressure coefficient lobe next to the middle of the wheel in iLES. Deceleration of the attached flow along the whole outboard wheel surface extends the low-Cp_0 zone down to the bottom of the wheel in RANS whereas reacceleration below the separation reduces its vertical extent in the iLES case. Inboard of the wheel, the half bottom part of the flow for $Z < 0$ exits the wheelhouse below the wheel and faces the flow coming from outboard. In iLEs, the inboard flow exiting below the wheel decelerates, resulting in a low total pressure wake, and drives down the low energy detached flow from the outboard side between $Y = 0$ and $Y = 0.1$ where a anti-clockwise vortex is formed due to their relative motion. In the recirculation region created by the two flow components facing each other, an upward motion is provoked by the wheel rotation around its bottom surface. In RANS yet, the flow coming from the outboard side of the wheel has remained attached and is more energetic than in iLES and manages to push the flow coming from the inboard part forming a slanted vertical boundary between $Y = 0$ and $Y = 0.1$ occupying the whole height below the wheel. The flow from the inboard side is forced into a clockwise rotating vortex centred at $Y = 0.2$. The recirculation wake in this case occupies a narrower area shifted towards the inboard part of the wheel whereas it nearly extends over the whole area below the wheel in the iLES case. The two adjacent

vortices formed below the underfloor at $Y = 0.2$ are very similar between the two cases and form at the tip of the guiding vanes due to the inboard/outboard pressure difference. The flow entrained by the vortices near the surface is responsible for the low-pressure trails downstream from the guiding vanes highlighted on figure 3.11.

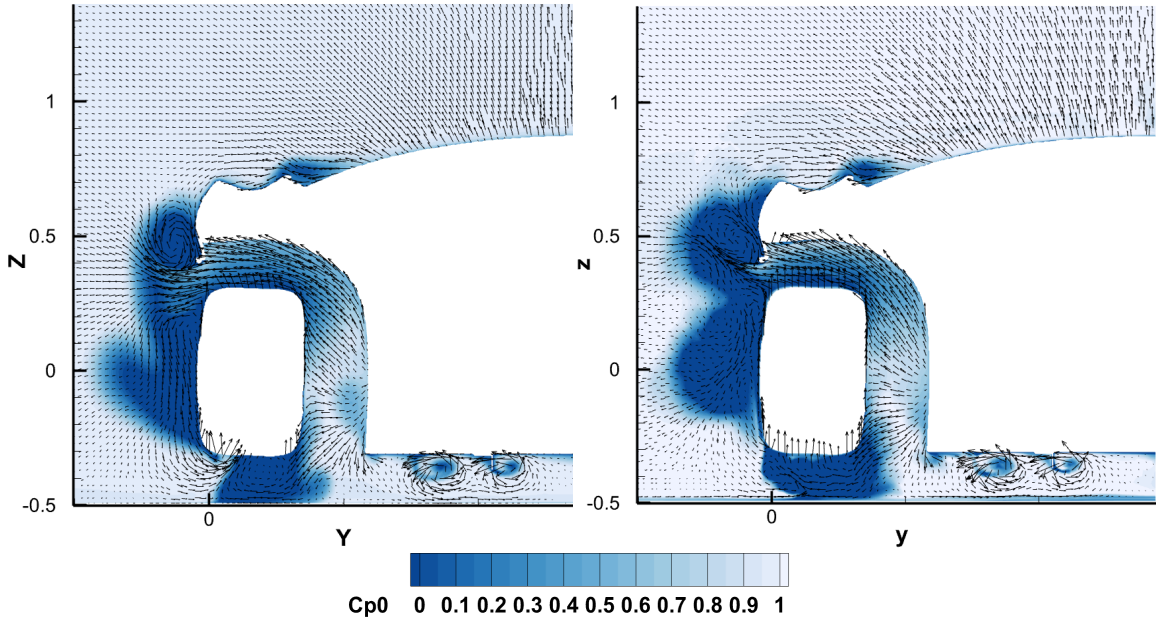


Figure 3.15: Total pressure coefficient and in-plane constrained velocity for plane 1. Left: RANS, Right: iLES

Plane 2 The second plane is located just upstream from the rear end of the wheel. At this location, the velocity of the wheel surface is almost exclusively vertical, the streamwise component being far weaker at this location. The difference of impact of the velocity boundary condition on the surrounding flow is marked by the thin light-coloured layer around the wheel surface visible in the iLES simulation and absent in RANS which corresponds to a layer of energetic flow being entrained by the motion of the wheel. The jetting flow at the top of the wheel arch has gained in intensity compared to the previous plane and is even more inclined towards the positive Z-direction given the larger vertical velocity component on the wheel. Vortices present at the top of the wheel arch have been vertically stretched but the extent

of the low-Cp_0 area is very similar to the previous plane except below the wheel. In RANS, the flow travelling downwards on the outboard side of the wheel is driven upwards below the wheel by both the rotation and the clockwise-spinning vortex already observed on the previous plane resulting in a second significantly larger vortex centred at $(Y, Z) = (0, -0.2)$. In iLES, a similar vortex is observed even if the upwards velocity boundary condition at the wheel, having a wider impact than in RANS , limits downwash on the outboard side of the wheel. The counter-clockwise vortex observed on plane 1 below the wheel between $Y = 0$ and $Y = 0.1$ is still clearly present and a second counter-clockwise vortex starts forming at the same height at $Y = 0.4$ due to the flow being pushed further inboard at the contact patch and the flow exiting the wheelarch below the wheel. This second vortex extends the recirculation region behind the wheel in the positive Y-direction up to $Y = 0.45$ in the iLES simulation whereas it stops at $Y = 0.25$ in RANS. The flow in between these two vortices follows a S-pattern, first being driven towards the inboard vortex and then below close to the ground being directed towards the outboard vortex. The pair of vortices shed from the guiding vanes are very similar between the two cases even if they are slightly more diffused in the RANS case.

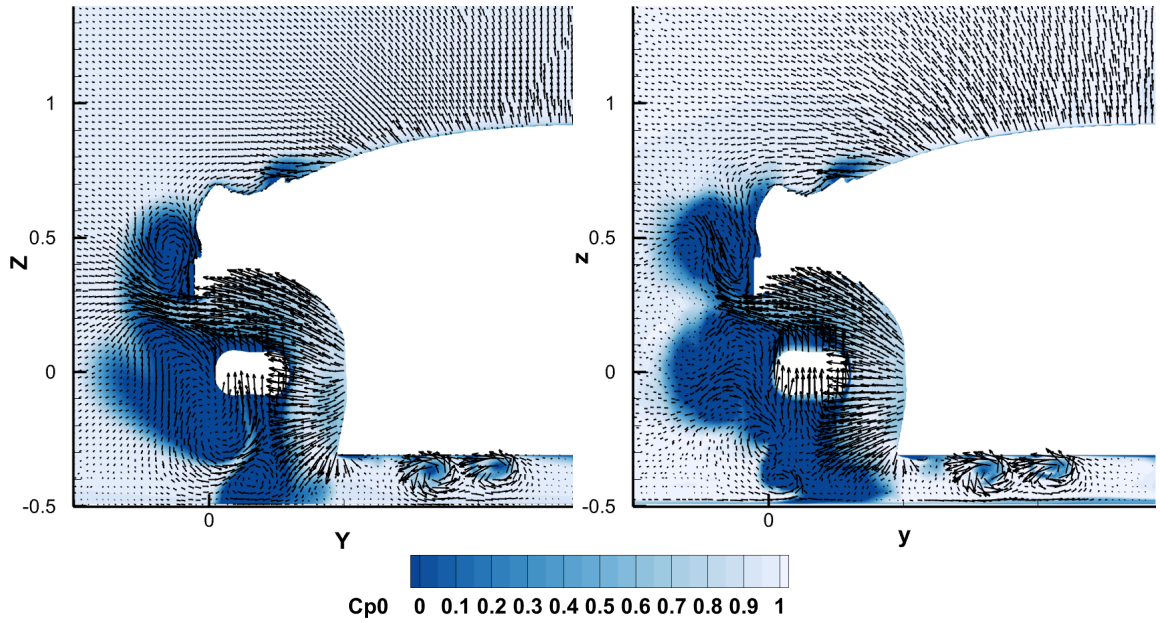


Figure 3.16: Total pressure coefficient and in-plane constrained velocity for plane 2. Left: RANS, Right: iLES

Plane 3 The third plane is located one wheel radius behind the wheel. Outboard from the car, the vortex visible at the top of the wheel arch in previous planes is still visible in both cases, but the centre is slightly higher and further from the surface in iLES at $Z = 0.5$ increasing the width of the wake at the upper part of the side surface. It deviates down the flow coming out of the channel created between the door blade and the door forming a second anti-clockwise vortex below the previously mentioned vortex at $Z = 0.1$ in RANS. The wake arrangement is more complex in iLES. The longer distance between the top vortex and the flow exiting the door blade channel shifts the deviation further away from the surface and a second clockwise rotating vortex forms below the top vortex. A third vortical structure is generated below the blade, at the junction with the underfloor due to the interaction between the middle vortex and the outboard vortex of the wake of the wheel below the car highlighted in planes 1 and 2 which has been convected downstream. The inboard vortex of the wheel wake has more rapidly grown in size relatively to the outboard vortex, fuelled by the flow

downstream from the wheelhouse being direct outboard. The separation line between this pair of vortices is vertical in the iLES case. A pair of vortices is also present in RANS but somewhat more inboard $0.05 < Y < 0.5$. The most outboard vortex corresponds to the large vortex at $Y = 0$ in plane 2 which has been convected below the car and which, contrary to the iLES case, rotates in the anti-clockwise direction. The second vortex observed in RANS is similar to the inboard vortex in iLES and results from the flow being pushed inboard on the contact patch. This vortex is smaller than the outboard counterpart unlike the iLES case but similarly rotates in the clockwise direction. The difference of size between the inboard vortices results in the wake of the wheel being more extended in the positive Y-direction in iLES compared to RANS. The pair of guiding vanes vortices are driven outboard while being convected downstream and become closer to each other in RANS compared to the iLES case which is coherent with the trails observed in figure 3.11.

The flow around the top part of the car for $Z > 0.5$ is very similar between the two cases. The stand-alone part on the top left-hand side is the mirror below which a pair of vortices which has formed upstream on the arm is visible. Looking at the main body, the flow separates at $Y = 0.4$ which is the location of the A-pillar joining the roof and the side window. This separation is recurrently observed for automotive flows and has been emphasised by a low-pressure trail on figure 3.10.

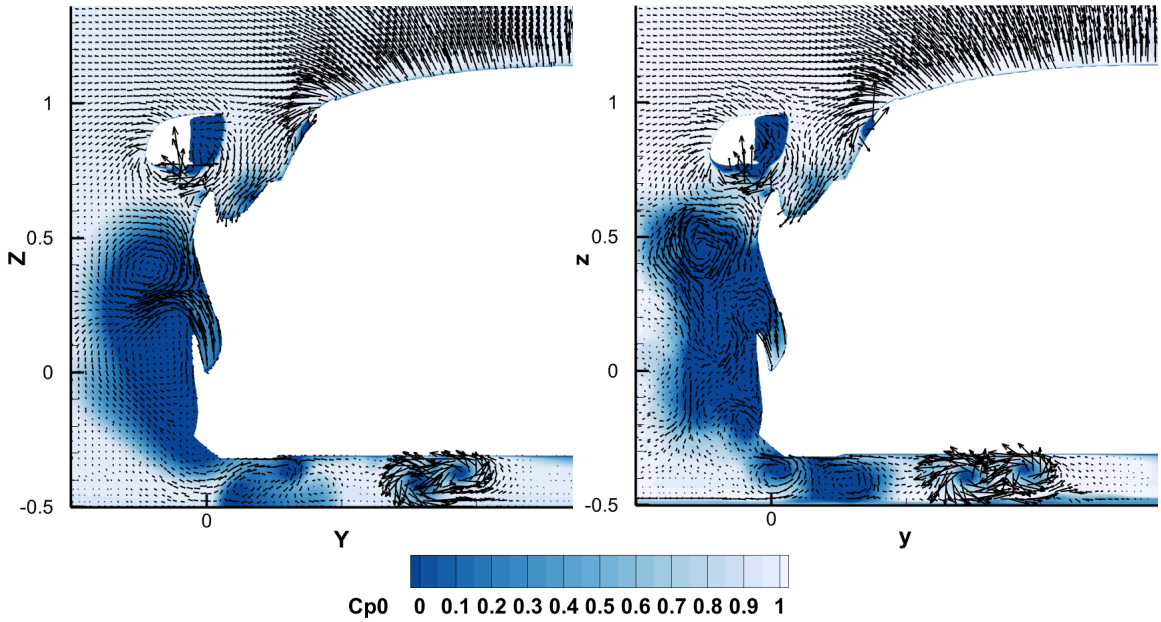


Figure 3.17: Total pressure coefficient and in-plane constrained velocity for plane 3. Left: RANS, Right: iLES

Plane 4 The final plane is situated directly upstream from the junction with the extension. Many of the flow structures observed on plane 3 have been convected and diffused downstream and can still be observed. The vortices outboard of the car have become larger and diffused in both cases. Their position has remained the same as in plane 3 in RANS whereas they have been convected down in iLES. The vortices under the car coming from the wheel have also been diffused, the weaker vortex in each case which is the more inboard one in RANS and the more outboard one have almost vanished compared to the other vortex which has grown slightly wider. The wake of the wheel remains centred on the wheel position in RANS whereas it is more extended inboard in iLES. The most important difference between the two cases is for the guiding vanes vortices which have merged in RANS but can still be distinguished in iLES even if they are very close what is again confirmed by low-pressure trails on figure 3.11 downstream from guiding vanes. The pair of counter-rotating vortices on the top left come from the side mirror. Although their size is comparable, they are more

inclined in RANS than in iLES. Yet more refinement is required for a more accurate study of this region since the mesh is designed to focus on the main body. Slightly higher on the car, the A-pillar vortex is visible in both cases even if it is definitely larger in iLES what explains the lower pressure strip observed on the side window in figure 3.10. This vortex is provoked by the lateral flow separation at the junction between the roof and the side mirror emphasised on plane 3. Experimental data would help assessing the accuracy of the two modelling approaches for this particular area.

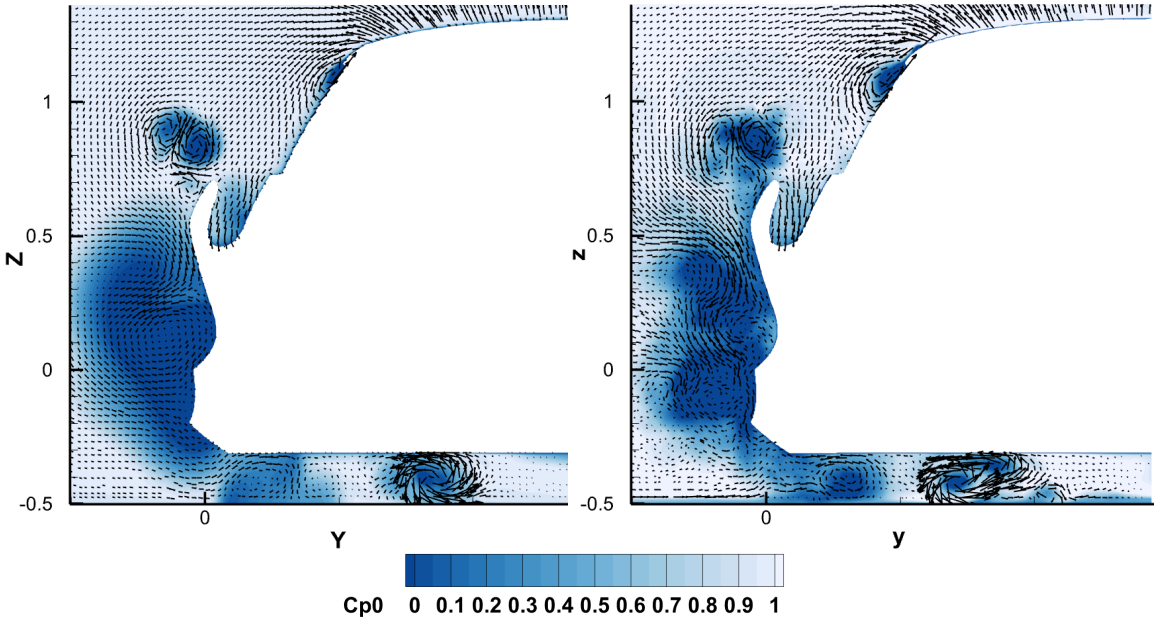


Figure 3.18: Total pressure coefficient and in-plane constrained velocity for plane 4. Left: RANS, Right: iLES

3.3.5 Unsteady flow structures

Iso-surfaces of total pressure $Cp0 = 0.5$ have been captured after 8 flow passes based on the wheel diameter, at $t = 8.3s$. Two different points of view can be seen on figures 3.19 and 3.20. Large separated structures can be seen exiting the wheelhouse and travelling downstream on the side of the car. This chaotic flow arrangement rather indicates disturbed flow jetting

from the wheelhouse which supports the necessity of unsteady simulation techniques in order to obtain an accurate representation of the flow in this region. A periodical pattern seems to be present in the side flow structures from the bottom point of view on figure 3.20 and could be attributed to the wheel rotation as particular excited frequencies close to rotation frequency and first harmonics have previously been identified on forces evolution (see figure 3.9).

Very unsteady flow is also present on the inboard side of the wheel. Highly unsteady wake is generated at the separation already noticed on the front diffuser and propagates in the wheelhouse. The chaotic flow then exits from the rear of the wheelhouse and follows the bottom side edge of the car before merging with the side wake.

Guiding vanes vortices are clearly identifiable confirming the swirling flow is not an averaging effect and underlines their strong impact on the underfloor flow and thus downforce. Vortices are very unsteady and merge in a larger vortex further downstream.

Separated flow structures are also shed from the front lip, at the location of the separation pointed out on the time-averaged results and travel downstream along the underfloor. The separation bubble can reasonably be considered to be unsteady and turns out to perturb the flow further downstream on the underfloor.

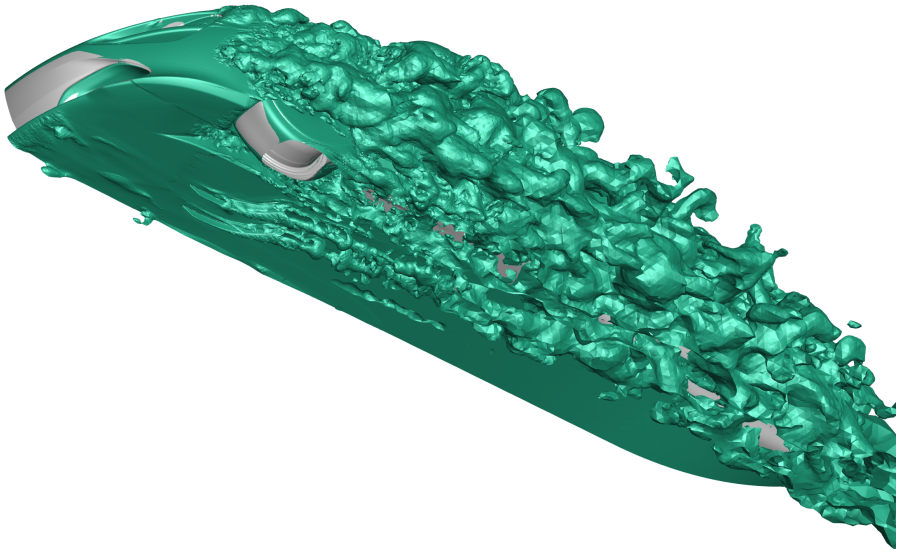


Figure 3.19: Iso-surface of total pressure coefficient $Cp0 = 0.5$ in the iLES case at $t = 8.3s$ (side bottom view)

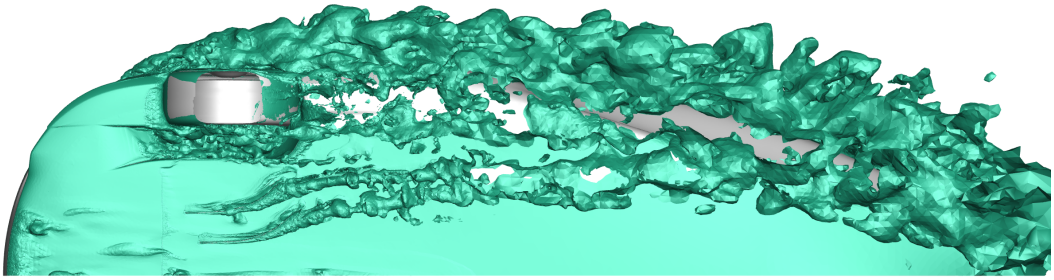


Figure 3.20: Iso-surface of total pressure coefficient $Cp0 = 0.5$ in the iLES case at $t = 8.3s$ (bottom view)

3.4 Conclusion

This first study on the P14 case has shown that the spectral/hp element method can be successfully applied to compute the flow around complex industrial geometries at high Reynolds numbers. A robust and repeatable workflow has also been developed, from pre-processing with high-order mesh generation, combining the external meshing tool from StarCCM+ and NekMesh, and interpolation of a low-order solution upon the high-order mesh to generate

a good starting point for the high-order simulation to actually running the simulation on a large number of CPUs. Most of the pipeline is ready to be automated and the last part requiring human intervention is the removal of pyramidal elements from the layer of prismatic elements in the linear mesh. Despite our recent implementation of the conversion of linear pyramids into high-order pyramidal elements, the splitting into several layers results in highly skewed elements which impose very restrictive constraints on the time step of the semi-explicit time-marching scheme.

The comparison of the RANS and iLES methodologies has shown noteworthy differences. Although the RANS simulation can be easily run on a small HPC machine in a few hours, the time to compute a sufficiently time-averaged solution from the unsteady solution with the scale-resolving technique to conduct a fair comparison is far larger. Yet this is expected since the resolution required to resolve small turbulent scales is much higher than the resolution required for turbulence-modelling techniques. The embedded DNS approach, provided adapted boundary conditions are employed at the inlet and outlet of the domain, or further simplifications of the model can however be used to isolate a portion of interest of the car, such as the wheelhouse. The discrepancies observed in the results on the front splitter and diffuser are due to separation prediction which is an area where RANS wall-modelling techniques usually perform poorly and in this case underpredict separation compared to the iLES simulation. The second area is the unsteady flow around the wheel. The rotation of the wheel is arguably more correctly represented on the surface of the wheel in the iLES simulation and has a tangible impact on the surrounding flow on the side of the car and on the underfloor downstream from the wheel. The flawed application of the boundary condition in RANS can be due both to wall modelling where the boundary condition is set further away from the wall in the logarithmic layer and not directly on the surface of the wheel. Yet, further testing with other CFD packages using the same turbulence modelling technique would validate whether this feature is inherent to StarCCM+ or to the modelling approach. Finally, slower decay

of the guiding vanes vortices in the iLES case tends to indicate that numerical diffusion is more limited in Nektar++, even if the polynomial order of the expansions used in this work is relatively low.

This work has emphasised the main shortcomings of the methods preventing from the integration in an industrial development process. The complexity of the geometry imposed high levels of h-refinement and subsequently p-refinement has been kept relatively low to keep the simulation tractable. However, when both type of refinements are possible, for instance when a larger HPC facility is utilised, balancing both type of refinements is still an open problem. Chapter 4 compares two hp-refinement approaches on an simplified automotive bluff body to derive guidelines. In terms of modelling, the blanking of the wheel is the most restrictive simplification and has allegedly a tangible impact on the flow field. When employing a real wheel geometry with an enclosed rim, particular treatment must be applied to address the rotation of non-axisymmetric parts. A particular immersed boundary method based technology called smoothed profile method has been developed for threedimensional simulations of complex geometries and assessed on the case of an isolated rotating wheel in chapter 5.

Chapter 4

Influence of hp refinements: SAE body

4.1 Literature review

4.1.1 Automotive bluff bodies

The previous chapter has shown the possibility to perform simulations of automotive industrial configurations with the spectral/hp element method. However, due to the high complexity of the P14 geometry, very high level of h-refinement had to be imposed to correctly mesh sharp edges. Therefore, given the size of the mesh, p-refinements attempts are out of reach in terms of computational resources available for the study. Yet, p-refinement is an essential component of the spectral/hp element method when it comes to scale resolving simulations due to exponential convergence properties as well as the compact formulation reducing the amount of inter-process communication accounting for the major part of computation efforts. Therefore, when sufficient computational resources are available, it is desirable to combine h- and p-type refinements. Alternative automotive bluff bodies are candidates to carry out more fundamental study of the two types of refinements. Indeed the relative sim-

plicity of the geometry compared to industrial standards such as P14 loosen the constraint on h-refinement resulting in lower element count enabling the possibility to increase p-refinement despite restricted resources. Le Good and Garry [69] have summarised the different reference models used in automotive aerodynamics development and emphasised their importance for fundamental understanding of car aerodynamics but also for the progress in experimental and validation of computational methods. There are a few well-established automotive bluff body models in the literature, the most famous of them being the Ahmed body introduced by Ahmed [45] and the Windsor body which has been internally used by Rover for many years before being introduced in the literature by Howell et al. [70]. Automotive bluff bodies are designed to model and isolate complex flow features encountered on production cars such as flow transition, separation and reattachment or vortices but also to evaluate the impact of aerodynamic optimisation devices with minimal geometry details. For instance, Hanfeng et al. [71] equipped the Ahmed body with deflectors to quantify their effect on drag reduction and Buscariolo [43] studied the impact of diffusers on the wake and aerodynamic forces by implementing diffusers to the Ahmed body. Fundamental studies have also been conducted on the Windsor body, such as Varney [72] who added wheels and evaluated their impact on the flow field in comparison to the original body. Another advantage of automotive bluff bodies is that the geometry is often open-access and easily replicable. The automotive industry is highly competitive and carmakers are reluctant to share sensible geometries as they represent a strategic advantage. On the other hand, production cars already on the market are very difficult to model numerically or experimentally making any attempt of aerodynamic study extremely challenging. Moreover, car models designed by the various carmakers are very different in terms of design which makes even direct comparison of cars of the same type issued by distinct manufacturers delicate. Automotive bluff bodies provide generic geometries allowing one to compare different modelling techniques, be they experimental or numerical and extensive experimental and numerical database can be found in the literature

for the most famous cases such as the Ahmed and Windsor bodies.

The SAE notchback body has been selected for this work. It is a generic bluff body designed to replicate the main flow features around automotive geometries which was originally used for wind tunnel interference investigations by the SAE "Open Jet Interference Committee" [12] and in different wind tunnels for comparison purposes. The geometry has been taken as a reference case for the first Automotive CFD Prediction Workshop in Oxford in December 2019 based on an extensive experimental campaign performed by Wood et al. [73] on a 20%-scale model and spectral/hp element method simulations of this configuration have been performed with Nektar++ by Buscariolo, Hambli, Slaughter and Sherwin [74] and presented at the workshop.

4.1.2 Notchback flow topology

Understanding of the wake structure of notchback has proved extremely challenging compared to other typical car shapes due to the very complex flow physics at the rear of the vehicle resulting from the combination of the slanted rear window followed by the bootdeck. Studies to gain knowledge of the flow topology were started with Carr [75] who conducted force measurements and flow visualisation on the MIRA notchback geometry enabling him to establish the first topology map of the flow in the notch region. Carr's observations led him to propose a flow structure similar to a backward facing step with the flow separating at the top of the rear window and reattaching just upstream from the bootdeck trailing edge. The flow in the separation bubble is entrained by the flow coming from the sides of the vehicle and deflected towards the symmetry plane, resulting in a transverse vortex. A different flow structure has been later proposed by Nouzawa et al. [76]. Using experimental flow visualisations and numerical simulations, the authors have observed two main flow features in the wake. When the flow on the sides of the vehicles reaches the notch, it is deflected towards the low pressure wake and separates and rolls up over the side edges of the bootlid

resulting in a pair of trailing vortices. An arch vortex forms over the rear window and is caused by the flow separating from the sharp trailing pillars and from and the flow separating at the end of the roof. The arch vortex drives the flow down separating at the trailing edge of the roof, which reinforces the strength of the pair of trailing vortices as it passes between them. The wake structure is summarised on figure 4.1.

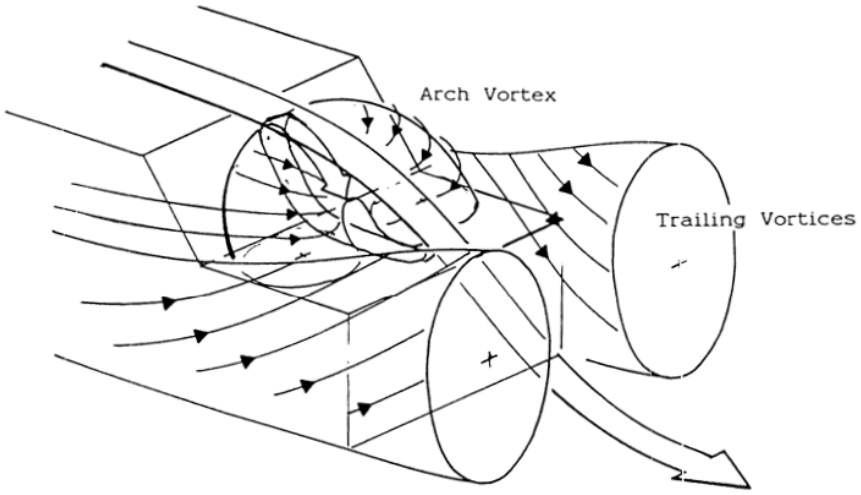


Figure 4.1: Notchback flow topology proposed by Nouzawa et al. [76]

For slant angles below 25^o, Nouzawa et al. state that the separated flow reattaches before the trailing edge of the bootdeck which is confirmed by Sims-Williams et al. [77] who experimentally measured and visualised the flow and performed numerical simulation on simplified notchback geometries. Beyond this critical slant angle values, both Nouzawa et al. [76] and Sims-Williams et al. [77] notice that the flow does not reattach anymore after separating from the roof with a sudden fall of drag which they infer is due to the important rearrangement of the wake structure. Both studies have been conducted with sharp-edged geometries that causes non-negligible differences with more realistic shapes with rounded edges. Gilhome et al. [78] experimentally investigated the near-wake structure on production notchback cars. Although the general flow structure is very close to the previous observation of Nouzawa and

Sims-Williams, the arch vortex is reversed which can be explained, according to Wood [73], by the progressive separation mechanism induced by smooth edges in Gilhome's study. Indeed, less aggressive pressure gradients on smooth geometries delay transverse vortex formation on the backlight and the bootdeck resulting in an inverted arch vortex. Wood's argument is backed by observations of Fuller and Passmore [79] who observed slower and weaker vortex roll-up on the trailing pillars of fastback models with rounded edges when comparing to a sharp-edged model.

The SAE body enables researchers to focus on the notch physics without interferences arising from flow structures generated by other parts of the car and provides fundamental understanding of the flow in this region. The geometry is presented in the next section.

4.1.3 The notchback SAE body

The SAE notchback body, which can be seen on figure 4.2, is a generic bluff body designed to build fundamental understanding of complex notchback flow. The original model is 4.2m long, 1.6m wide and 1.2m high. It has a flat underfloor on which four cylindrical pins are mounted to keep ground clearance when the model is tested in wind tunnel. The front slant is inclined by 30^o and the model is equipped with a 6° angle rear diffuser which can be seen on the rear underfloor view on the second picture of figure 4.2 balancing the pitching moment created by the A-pillar. The rear slant is inclined by 20^o relative to the horizontal which is typical of vehicles produced in the 2000s according to Wood [73] since shape optimisation aiming at reducing drag on commercial cars has led to increasingly smaller rear window inclination angles. Except for the base, all the edges are rounded which is more consistent with real car geometries. However, a slightly sharper shape can be noted for the edges joining the side surfaces and top/bottom surfaces of the vehicles. As highlighted by Fuller [80] and Wood [73], sharper longitudinal shapes are used on production cars to fix the location of separation preventing unsteady oscillations and therefore improve the stability of the vehicles. Several

studies have been carried out on this geometry using experimental and numerical techniques with some variations from the original design and are presented in the next section.

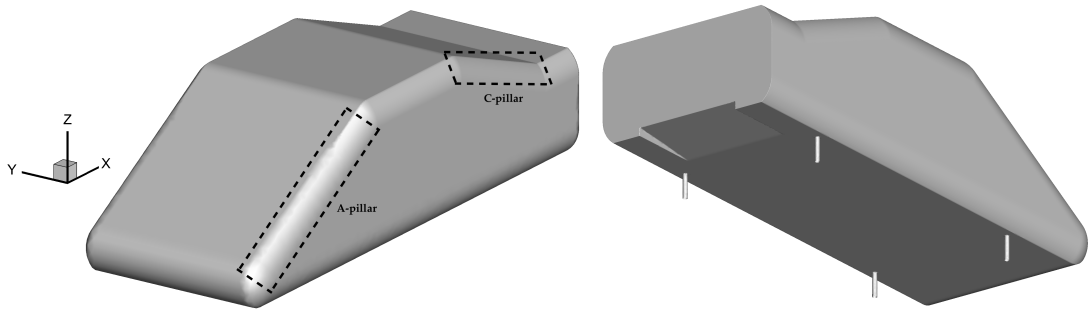


Figure 4.2: SAE body geometry

4.1.4 Previous studies on the SAE notchback body

4.1.4.1 Experimental results

Early work focusing using the SAE notchback is Cogotti's [81] who built a full size model to conduct a parametric study on the underfloor flow and measured forces and centreline pressure values. The geometry deviates from the original concept as the Cogotti equipped it with wheels and performed tests in the Pininfarina wind tunnel with the Ground Effect Simulation System to assess the effect of a moving floor on the aerodynamics. He observed in particular a reduction of the drag when the ground clearance was reduced. Ishima et al. [82] extended the study by performing 2D and 3D-PIV measurements on a 25% SAE body model equipped with wheels and combined it with an interpolation technique to recover data on noisy points and flow regions where this type of visualisation techniques usually fail such as in the vicinity of the body in order to reconstruct streamlines. Ishima et al. [82] remarked the flow remained attached to the backlight and only separated at the trailing edge of the bootlid. The streamlines wrapped the vehicle, first being pushed outwards of the centreplane on the A-pillar before being deflected inwards on the C-pillar, probably due to

the low pressure wake. A large 2D vortex structure is observed at the streamwise middleplane in the wake when looking at averaged velocity vectors. Yet, Ishima et al. [82] argue that this picture can be misleading since the smaller eddies responsible for large velocity perturbations and associated to higher frequency motion which are visible in the instantaneous snapshots are missed during the averaging process due to the high time period between two velocity measurements.

An extensive study has been conducted by Wood et al. [73] and Wood [83] on a 20%-scale model of the SAE body in wind tunnel. Measurements have been carried out in a one-fourth scale wind tunnel at a velocity of 40 m/s resulting in a Reynolds number of 2.3 million based on the vehicle length. Although the 20^o rear slant model reference model has been used, the back of the body has been made interchangeable to study the wake structure with different C-pillar inclinations and the yaw angle has been swept. Forces as well as pressure distribution over the centreline and the base of the vehicle have been measured and PIV planes at several locations have also been captured. Confirming Ishima's observations [82], the pressure build up at the centreline of the backlight indicates mainly attached flow for the 20^o slant case but it is not the case for higher angles where separation occurs at the top of the backlight and its extent grows with the inclination up to 30^o and remains constant above. Larger separation for higher angles results in lower pressure values over the notch and therefore increased drag and rear lift. The impingement of the flow coming from the C-pillar results in a pressure peak at the beginning of the bootlid and the pressure then decreases with the flow accelerating on the deck up to the trailing edge. Looking at the transverse pressure coefficient at the base of the backlight, pressure drop has been observed at the extreme outboard location coming from vortices generated on the C-pillars and entraining flow near the surface. For higher angles, the drop is either of lesser magnitude or absent which can be explained by trailing pillar vortices which fail to follow the slant and remaining further from the surface. PIV observations allowed the authors to create a map of the topology of the flow on the notch

which is summarised on figure 4.3. A pair of vortices generated at the outboard edge at the top of the backlight and fuelled by the flow separating on the side edges of the slant in red deflects the flow from the sides of the vehicles towards the central region of the slant and the backlight. These two vortices have already been detected by Nouzawa [76]. The flow over the rear remains mainly attached (in green) except for a small region where it detaches (in blue) when encountering the transverse flow over the bootdeck generated by the relative strength difference of the trailing vortices in red. Wood [83] also reports that the flow is weakly unsteady and that the variation of the location of the separation is responsible for the unsteadiness.

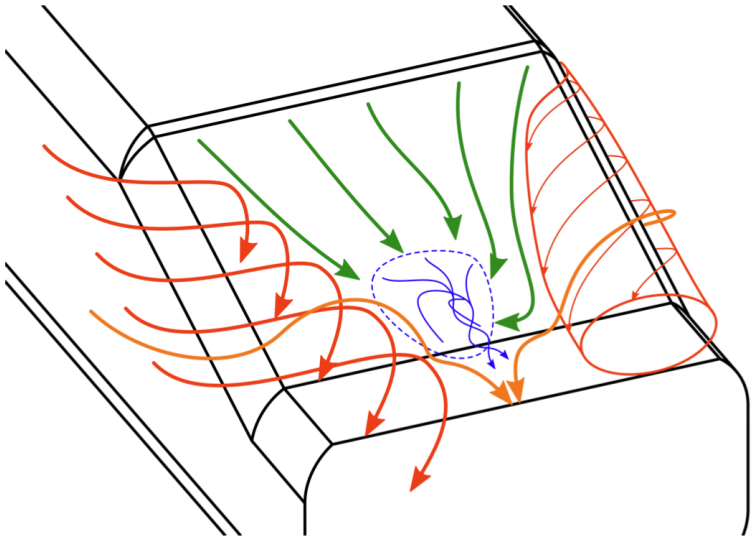


Figure 4.3: Notchback flow topology observed by Wood [83]

4.1.4.2 Numerical simulations

Ishima et al. [82] compared the experimental results with RANS simulations of the same 25% scale model and compared streamwise velocity profiles at the centreplane for different longitudinal locations. Although CFD results are in good agreements with the experimental results up to the end of the front slant, the results start to diverge from the roof and agree-

ment becomes worse on the backlight especially close to the car surface and in the wake. Detachment is overpredicted, with lower streamwise velocity values computed at the vicinity of the wall which is due to both wall-models and the lack of resolution for the PIV close to the surface. The length of the wake is overpredicted and slightly shifted upwards compared to the PIV. The authors also point out that small scale motions which are visible on the instantaneous snapshots vanish in the simulation and are replaced by two longitudinal vortices due to the steady approach just like when PIV data is time averaged and conclude that turbulence modelling as well as the steady approach must be reconsidered.

A previous RANS simulation was performed on a full scale model by Basara et al. [84] with the $k-\epsilon$ turbulence model and a Reynolds stress model (RSM) and the results compared to industrial experimental data. The pair of longitudinal vortices observed by Ishima et al. [82] due to stationary modelling is present and shifted upwards which is due to the rear diffuser inducing upwash according to Basara et al. [84] even if weaker upwards offsetting is noticed for the RSM. The two models predict the same pressure distribution over the vehicle on the left-hand side of figure 4.4 which agrees well with measurements except on the notch of the car starting from $x = 2.75m$. The pressure recovery forecast over the slant is not as accurate, especially for the RSM predicting sensibly lower pressure values. The explanation can be found on the right hang side pictures of figure 4.4 representing the velocity field over the slant. The simulation results show a separation of the flow at the centerplane directly after the trailing edge on the roof for the RSM before reattaching shortly before impinging the bootdeck whereas the boundary layer smoothly grows with $k-\epsilon$ and more pressure is built up in the process. Very different flow fields can also be seen near the C-pillar where the trailing vortex generated at the edge is visible since the top of the slant for RSM but is delayed with the $k-\epsilon$ model where it starts being visible at the middle of the slant and remains weaker. The shape of the pressure recovery also differ. The impingement of the weakly detached flow in the case of the $k-\epsilon$ model severely impinges the bootlid resulting in a pressure peak. Yet

for the separated flow obtained with the RSM, the impingement is weaker and the peak is not as sharp. On the underbody, both turbulence models underestimate the downforce generated by the diffuser which may indicate separation, reducing the efficiency of the diffuser. Finally, the authors state that the wake qualitatively correlates better with the experiment when using the RSM compared to the $k - \epsilon$ model in which the impact of the underbody flow is predominant.

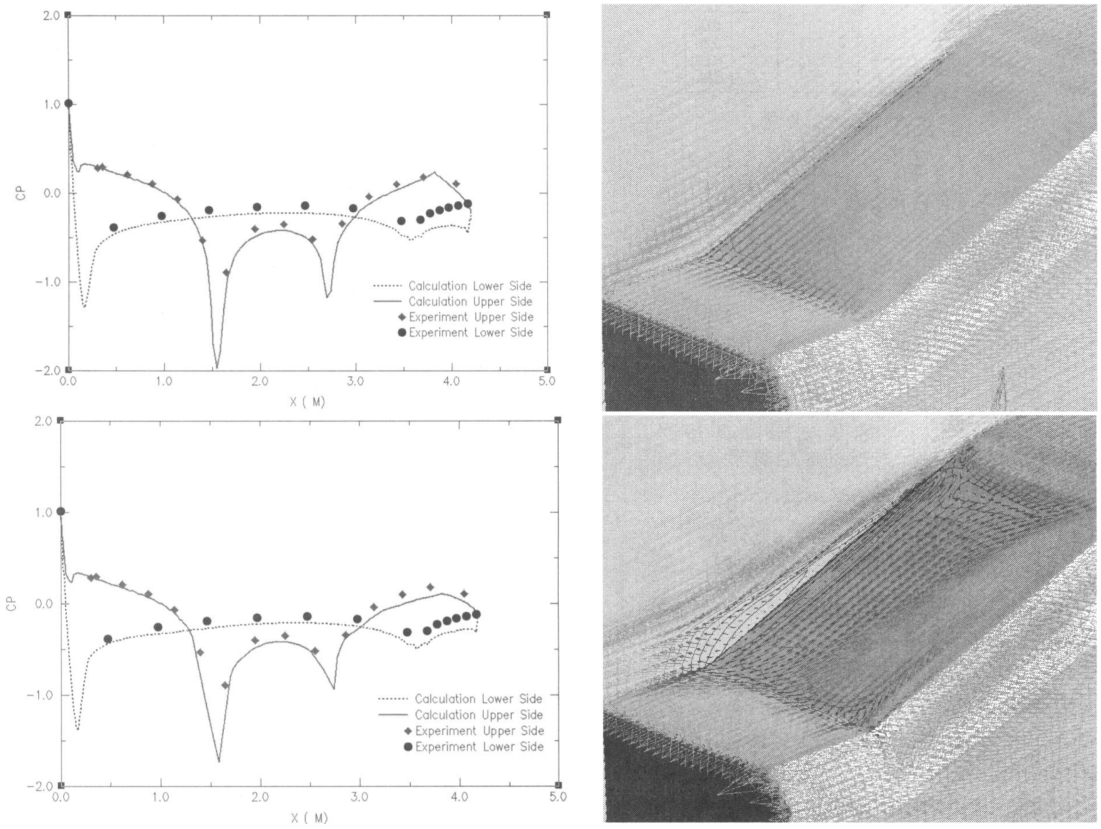


Figure 4.4: Basara's et al. [84] simulations results on the SAE body with the $k-\epsilon$ turbulence model on the top and the RSM on the bottom. **Left**: Pressure coefficient over the centreline. **Right**: Velocity field over the slant

More recently, Nader et al. [85] performed a RANS and DES simulation of the SAE body at a Reynolds number of 2.5×10^5 and compared their results to the experimental data of Wood et al. [73] and Wood [83] which are carried out at a Reynolds number of

2.3×10^6 meaning one order of magnitude higher. The pressure distribution over the centreline agrees satisfactorily with the experiment up to the notch. The flow then detaches from the middle of the slant interrupting the pressure recovery and reattaches approximately at the middle of the bootdeck generating a pressure peak higher than the measurements. The base pressure predictions with both turbulence modelling techniques do not provide accurate results although the DES turns out to output a more reliable trend.

4.1.5 Objectives

The literature review has shown that the flow topology over the notch is particularly complex and that the SAE body successfully conserves the flow arrangement while reducing the interferences coming from other regions on production cars. Previous simulations based on Reynolds averaging predict the flow satisfactorily on the front of the car but results on the notch region turn out to be less accurate especially due to separation physics and vortex generated on the rear slant edges which represent a real challenge to resolve correctly. This is particularly because of the use of rounded edges on which separation is more challenging to predict compared to sharp edges which fix the separation. Buscariolo [43] has shown that high-fidelity simulations performed with the spectral/hp element method have more successfully predicted wake structures on the Ahmed body than RANS simulations.

In the previous chapter, the size of the linear mesh was dictated by the extremely complex geometry of P14 imposing very high h-refinements requirements to model sharp and highly curved edges correctly. Computational resources restrictions forced the polynomial order to remained caped.

In this chapter we apply the spectral/hp element method to the SAE body and compare the results with the experimental data of Wood [73] on the 20%-scale SAE body model at a Reynolds number of 2.3 million. The first goal is to compare the results yielded by the high-order simulations to previous numerical simulations in the literature. A first set of

simulations, which has been presented at the first Automotive CFD Prediction Workshop in Oxford, utilises a committee grid designed to keep a standard reference for participant on which second, third and fourth polynomial orders for velocity have been run to increase resolution. The results are presented and the second objective is to assess whether p-refinement effectively improves the precision of predictions on a fixed predesigned grid. In a second study, two hp-refinement approaches are compared, one using a very fine linear base mesh coupled to a second polynomial order for velocity and the second one using a coarser mesh but a fourth order polynomial expansion velocity, both having similar number of degrees of freedom. The last target is to compare the two strategies and extract guidance for hp-refinement settings for complex simulations.

4.2 Simulation settings

4.2.1 Computational domain

For this study, the 20^o rear slant SAE body has been utilised. Like the P14 case, all dimensions are expressed relatively to a reference length which is in this case the length of the SAE body L. The main geometrical parameters of the body are represented on figure 4.5. The A pillar is slanted with a 30^o angle. As shown on figure 4.2, a rear diffuser inclined by 6^o is centred relatively to the symmetry plane. The model is connected to the floor with fours cylindrical pins mounted on the underfloor measuring $40mm$ which can be seen on figure 4.2. All the simulations presented in this part are run on the full body.

The outer domain used for this study is adapted from Loughbourough one-fourth scale wind tunnel. The dimensions of the working section are $2.31L \times 1.57L$ and the total length is $7L$ with the body centred, meaning the upstream and downstream lengths are both $3L$ long. The relatively short outlet distance compared to the workshop model requires the use of the high-order pressure outlet boundary condition developed by Dong et al. [56] preventing the

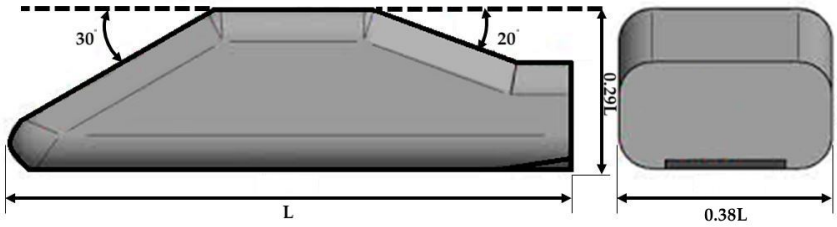


Figure 4.5: Dimensions of the 20° rear slant angle SAE body

kinetic energy build up due to energetic flow structures impacting the outflow boundary and thus allowing shorter outlet lengths. The dimensions of the domain are summarised on figure 4.5.

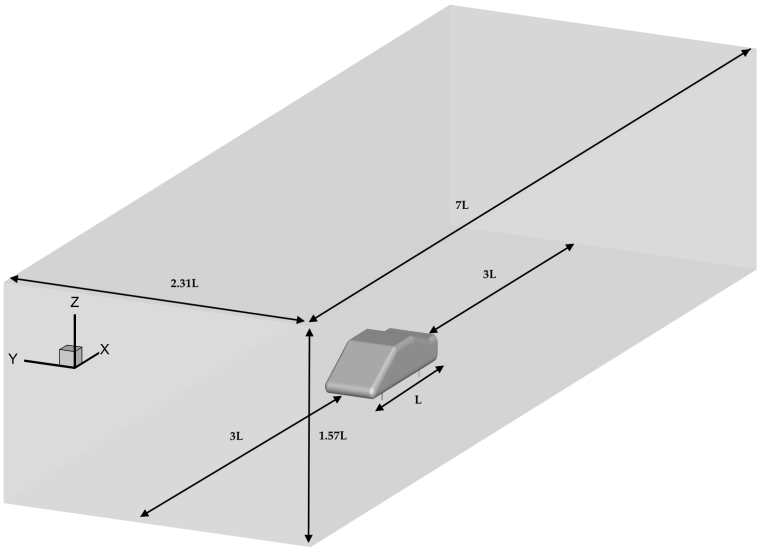


Figure 4.6: Computational domain for the SAE body based on the Loughborough wind tunnel

4.3 Mesh

In this section, two linear meshes are presented. The first one is the committee grid mesh used for the first Automotive CFD Prediction Workshop adapted to scale-resolving simulations

on which the p-refinement study has been carried out. The first results indicate a lack of resolution in the notch region encouraging a mesh rearrangement transferring more mesh elements at the rear of the vehicle. This second mesh has served for the hp-refinement strategies comparison. Both meshes are now presented.

4.3.0.1 Adapted committee mesh

As an effort to make results of different participants consistently comparable, the organisers of the workshop have designed two cartesian grids which needed to be used by the participants, one of which is designed for RANS simulations with four million elements and the other finer mesh for DES with 30 million elements. The meshing strategy here consists in generating an unstructured linear base mesh suitable for NekMesh, the meshing tool embedded within Nektar++ introduced in section 2.2, i.e. composed of tetrahedral elements and prismatic element on the surface of the body, while keeping the same mesh structure in terms of relative h-refinement and refinement zones and counting the same number of degrees of freedom as the DES mesh once polynomial expansions are applied. The original committee mesh is a trimmer mesh which is not compatible with Nektar++. Polynomial orders for velocity components expansions range from two to four, higher polynomial orders being too expensive to run with available computational resources. The linear mesh has then been designed to match the 30 million of degrees of freedom of the DES mesh using fourth order expansion for velocity components estimated using the elemental number of degrees of freedom for tetrahedrons

$$N_{DOF_{tet}} = \frac{P_N(P_N+1)(P_N+2)}{6} \tag{4.1}$$

and for prisms

$$N_{DOF_{prism}} = \frac{P_N^2(P_N+2)}{2} \tag{4.2}$$

where P is the number of modes of the polynomial expansion. It was adjusted by applying a scaling factor to the base size all mesh refinement parameters refer to while the prism layer total thickness is kept the same. The high-order mesh is obtained by specifying a sixth order projection rule for the elements in contact with the surface of the body and splitting the layer of prismatic elements in five sublayers on the floor and eight elements on the surface of the SAE body. The final element count for the high-order mesh is of 1.2M elements. This mesh will be referred to as CM.

4.3.0.2 Custom meshes

The two custom meshes are used for comparison of hp-refinement strategies and are an iteration over the mesh designed in the previous section. As will be discussed in section 4.4, the results of simulations with the adapted committee grid point towards a lack of resolution in the region of the notch, therefore a mesh rearrangement has been operated to alleviate this issue. Yet, instead of comparing results based on the same mesh with increasing polynomial order resulting in different number of degrees of freedom between successive cases, we fix the number of degrees of freedom and compare results obtained with two different p and h-refinement levels. Unlike low-order methods typically used in the industry where only mesh size can be changed, both h and p parameters can be modified. The custom meshes aim at comparing two methodologies: using a fine base linear mesh and keeping the polynomial order low, or designing a coarser base mesh with higher polynomial expansion order.
The meshes are inspired from the committee grid with an increased resolution in the notch region. The fine-h mesh, noted HFP3, has been created first and counts 7.8 million elements and 10 prismatic sublayers on the surface of the car and this case has been intended to be run with a second order polynomial approximation for velocity. Then, using the formula expressing the number of degrees of freedom for base elements in equations 4.1 and 4.2, an enlargement factor has been calculated and applied to the mesh to keep a similar number

of degrees of freedom with a fourth order expansion instead of a second order expansion. The number of prism sub-layers has been reduced to 6 to keep the product of the number of sub-layers and interpolation points for the polynomial expansion constant. The product is 10×3 for the fine mesh and 6×5 for the coarse mesh which is called HCP5. The final coarse mesh counts 1.7 million element, equivalent to 35 million degrees of freedom using the fourth order expansion rule. Although the number of elements are not exactly equal for HFP3 and HCP5, due to the difficulty to control exactly the number of elements with a single scaling factor while keeping the structure of the boundary layer equivalent, they are consistent enough to carry out a meaningful comparison. Both meshes have projected with a second order rule for the curvature of the edges due to the relative simplicity of the geometry which has given satisfactory approximation of the SAE body and has reduced geometric aliasing. In particular, this has allowed lower order integration quadrature rules to be used and has therefore reduced computation time. Main parameters for the three meshes are summarised in table 4.1 and centreplane longitudinal slices of each high-order mesh are shown on figure 4.7. More information about linear mesh is given in appendix section A.2.

Mesh name	$N_{elements}$ (million)	$N_{sub-layers}$	Mesh order	Maximum N_{DoFs} (million)
CM	1.1	5-8	6	30
HFP3	7.8	10	2	32
HCP5	1.7	6	2	35

Table 4.1: High-order mesh summary for the SAE body

The general layout of the CM mesh is different since a more recent version StarCCM+ v15 has been used to generate the mesh, whereas HFP3 and HCP5 have been generated with StarCCM+ v13. In general, the older version outputs better quality unstructured tetrahedral volume meshes compared to newer versions, this is why it has been utilised for the hp-refinement study.

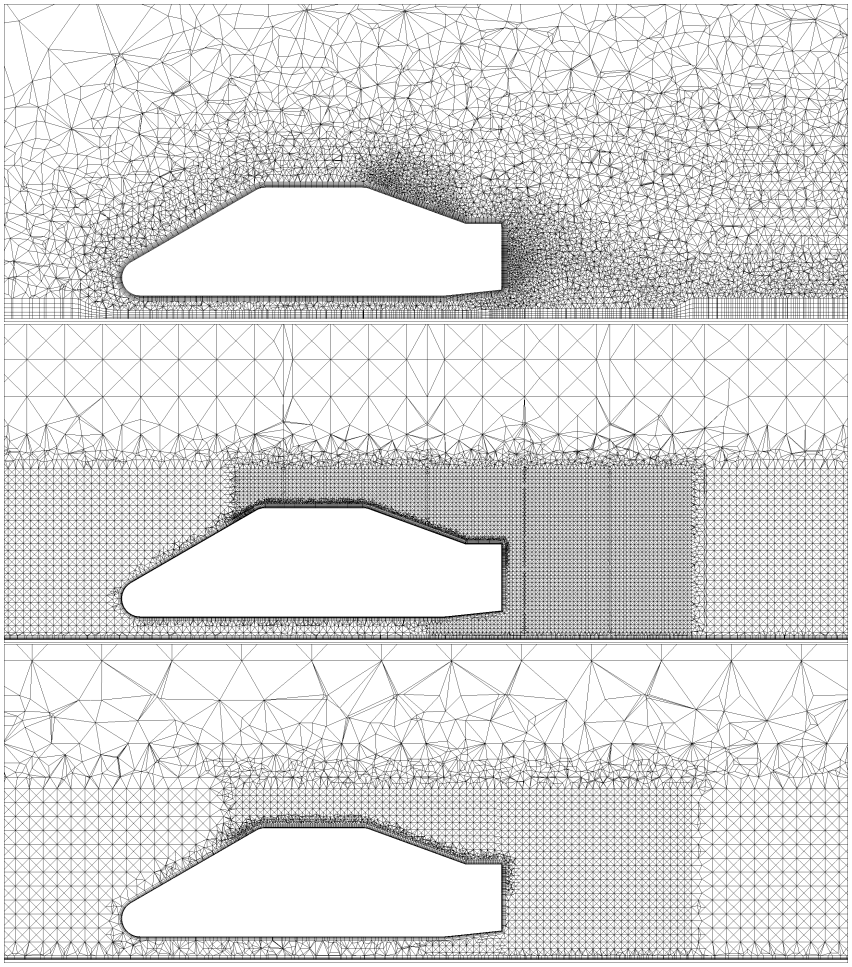


Figure 4.7: Mesh on the centreplane. **Top**: CM, **Middle**: HFP3, **Bottom**: HCP5

As previously explained, figure 4.7 confirms that the HFP3 and HCP5 meshes have the same structure, the latter being an enlargement of the coarse mesh. This is particularly visible when comparing surface meshes on figure 4.8.

A zoom in the junction region between the roof and the backlite is shown on figure 4.9 to compare near-wall mesh. Prismatic layers have the same total thickness in both cases but the number of layers is higher in the HFP3 case with 10 sub-layers compared to the HCP5 case where there are 6 layers, resulting in a smaller first sub-layer height in the HFP3 case. The goal is to compensate for lower near-wall h-resolution with higher p-refinement in the

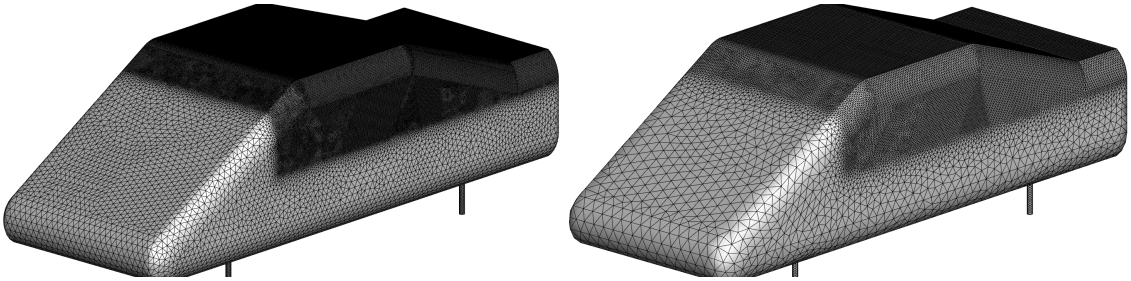


Figure 4.8: Surface mesh of the SAE notchback body. **Left**: HFP3, **Right**: HCP5

HCP5 case.

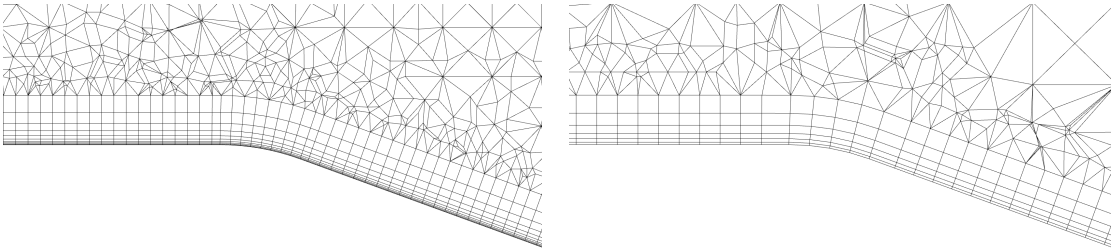


Figure 4.9: Structure of the prismatic elements layer at the roof/backlite junction. **Left**: HFP3, **Right**: HCP5

The corresponding y+ maps are presented for HFP3 and HCP5 case on figure 4.10. It is important to note that values are based on the first sub-layer height of the mesh, that is only h-refinement is taken into account and p-refinement is omitted which automatically results much lower y+ values in the HFP3 case. The comparison shows that y+ values are approximately five times higher in the HCP5 case than in the HFP3 case. The distribution is quite similar in the two cases, with high y+ values at the edges and corners and lower y+ values on the backlite. The low y+ area on the backlite occupies a larger portion of the backlite in the HFP3 case compared to the HCP5 case. High y+ value stripes at the edges are more extended in the HCP5 case than in the HFP3 case. Two low y+ spots are present at each side of the backlite/notch junction in the HFP3 but are not visible in HCP5.

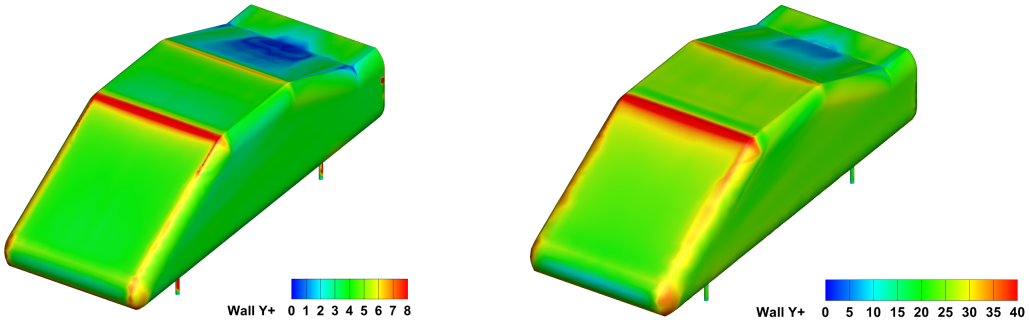


Figure 4.10: Y+ distribution based on the first layer height for the SAE notchback body. **Left**: HFP3, **Right**: HCP5

4.3.1 Numerical parameters

Wood's experiments [73], [83] are conducted with a stationary floor, therefore the ground is modelled as a non-slip wall. The rest of the walls of the domain are set as slip walls. The walls of the SAE body as well as the pins are set as non-slip walls. A logarithmic velocity profile is defined and tuned at the inlet since the boundary layer thickness with a uniform velocity inlet underestimates the thickness of the boundary layer at the location of the vehicle, which is estimated between 40mm and 60mm et al. [73]. The free-stream velocity is set to unity and the kinematic viscosity is set to match the Reynolds number of the experiment which is 2.3 million, based on the length L of the model. The high-order outflow boundary condition defined by Dong et al.[56], permitting energetic vortical structures to leave the outflow boundary without generating additional energy when using the spectral/hp element method. The same set of boundary conditions have been used for all the simulations presented in this part.

Spectral/hp simulations have been run with different polynomial orders depending on the mesh. The Taylor-Hood type expansions is still adopted, with the order of the polynomial expansion for the pressure being one order lower than for velocity. For the p-refinement study with the CM mesh, simulations have been run with second, third and fourth order polynomial approximations for velocity component. When it comes down to general hp-refinement

strategy comparison, the simulation on HFP3 has been run with a second order polynomial approximation for velocity and fourth order for HFP5. The reason for the naming convention of the two meshes is that second order polynomial approximation gives a third order precision scheme in a L^2 norm and a fourth order polynomial approximation corresponds to a fifth order error. Fifth to seventh order quadrature rules have been used for CM with depending on the polynomial order and a fourth order and a sixth order quadrature rule have been used for HFP3 and HCP5 respectively. The difference between the cases is due to the different mesh and approximation orders causing PDE and related mesh aliasing if insufficiently high-order quadrature rules are used causing instabilities due to energy build up as explained by Mengaldo et al. [35]. Temporal discretisation is computed with the second order implicit-explicit velocity correction scheme presented in section 2.3 and time steps have been tuned to keep the CFL condition below and close to unity, due to the explicit part of the scheme. The corresponding time steps are $\Delta t_{CM} = 2.5 \times 10^{-5}$ for CM, $\Delta t_{HFP3} = 1.0 \times 10^{-4}$ for HFP3 and $\Delta t_{HCP5} = 5.0 \times 10^{-5}$ for HCP5 which is typical for this type of applications. The scaling parameter of the SVV, based on the DG kernel introduced by Moura et al. [33] is manually set to its minimum possible value ensuring stable simulations and is of 0.2 for CM and 0.35 for both HFP3 and HCP5.

Simulations have been run for 10 CTUs $= 10\frac{L}{U_\infty}$ based on the length of the vehicle L to allow the flow to develop and then averaged for 2 CTUs. They have been run using Cambridge Tier2 high-performance computing facility CSD3 on 1280 Intel skylake cores. The wall time to compute one CTU is 26.6 hours for CM with fourth order polynomial expansion, 5.2 hours for HFP3 and 16.3 hours for HCP5. It is remarkable that a large time step can be used for HFP3, explaining the gain in computation time comparatively to the two other cases. The very large computation time required for CM can be explained by the use of a sub-stepping scheme detailed by Sherwin [86] and Xiu et al. [87] to stabilise the simulation at the cost of a higher computation time. Physical averaging time, even if it is consistent between the

cases presented in this chapter, is relatively short due to the high computational cost and the number of cases run. Yet, time-averaged results turn out to be reasonably smooth to allow to draw conclusions and compare mesh refinement strategies.
In the rest of this part, we adopt the following convention for p-refinement reference. A case run with a polynomial expansion order N is referred to as P(N+1). For instance, P5 will be used to designate a simulation run with a fourth order polynomial approximation.

4.4 Results and discussion

In this section, the results obtained with the SAE body on each mesh are presented with a particular focus on the notch region. First, the aerodynamic forces exerting on the body are compared and differences are then explained using surface pressure distributions as this counts for the major part of aerodynamic forces for this type of body. Then, differences of flow structures predictions are highlighted and investigated using PIV measurements of Wood [83]. In the rest of this part, we adopt the following convention for p-refinement referencing: a case run with a polynomial expansion order N is referred to as P(N+1). For instance, P5 will be used to designate a simulation run with a fourth order polynomial approximation.

4.4.1 Forces

As a bluff body, forces on the SAE body are almost exclusively originating from the pressure distribution, viscous forces having a negligible impact. Forces in the experiment by Wood [83] are measured with a balance connected to the four pins mounted to the underfloor of the SAE body. Drag and lift coefficients, based on the frontal surface area of the vehicle and freestream velocity upstream from the model, are compared to experimental measurements. Drag is generally underestimated in simulations compared to the experiment except for the CMP4 case. However, the results switch from overshooting to underprediction when p is

increased to P5 on CM. Results with HFP3 and HCP5 are almost identical and below the prediction with CMP5. When it comes to lift values, once again the CMP4 case stands out since it predicts positive lift like experimental measurements whereas all the other cases predict downforce. Yet, downforce predictions for HFP3 are well below HFP5 which is in this case very close to CMP5. It is worth emphasising here that all the other thirty simulations presented at the first Automotive CFD Prediction Workshop with different type of methodologies predict downforce in this case and that doubts about lift measurements have been voiced by the authors. Pressure distribution over the SAE body will now be investigated to understand differences in force predictions.

Case	C_D	C_L
Experiment	0.213	0.055
CMP4	0.247	0.036
CMP5	0.185	-0.080
HFP3	0.176	-0.003
HCP5	0.177	-0.084

Table 4.2: Force coefficients on the SAE body for the different cases

4.4.2 Pressure distribution

Previous simulations on the SAE body present in the literature report a rather good agreement with the experiment up to the notch where large discrepancies are observed. Wood [83] equipped the model with pressure probes on the centreline in a longitudinal arrangement on the upper surface of the body. Measured pressure values have then been corrected using Mercker's [88] correction formula for automotive wind tunnel testing taking into account the wake generated by the separation of the flow at the rear of the car. Experimental and numerical results are compared on figure 4.11.

The values of the pressure at the front stagnation point is 1 for the CMP5 case, but

less than one for both HFP3 and HCP5 and a value above one can be noted for the CMP4 case. Pressure then rapidly decreases downstream from this point and even quicker on the underbody due to the higher deviation imposed to the flow by the flat underfloor compared to the 30^o slant on the top on which the recompression is well predicted in all the simulations. Under the body, pressure rapidly reaches a local minimum slightly after getting around the rounded nose at $x = -0.4$ before the pressure increase imposed by the transition to the flat underfloor where pressure stabilises. The magnitude of the local minimum is greater for CMP5 and HCP5 compared to the two other cases. During the recompression phase at the beginning of the underfloor at the junction with the rounded nose, a negative pressure bump is observed on all the cases except CMP5. The pressure drop is very localised for HFP5, more extended for HFP3 and extending over an even larger distance for CMP4. On the upper surface, pressure drops to a global minimum at the edge between the A-pillar and the roof, at $x = -0.12$. The magnitude is the same for HFP3 and HCP5, higher for CMP5 and lower for CMP4 however the absence of pressure probe at this particular location prevents from having an experimental reference. Pressure increase over the first part of the roof is predicted differently depending on the cases. HFP3 and HCP5 show again the same behaviour with a second local pressure peak while pressure keeps increasing for simulations performed on the CM. The pressure variation on the roof from $x = -0.07$ up to the the junction with the C-pillar is in excellent agreement with the experiment for custom mesh simulations and CMP5 even if a small inflexion for pressure recovery occurs around $x = -0.06$ and computed pressure values are slightly higher than experimental values on the last part of the roof. Yet, CMP4 forecasts a completely different pressure evolution over the roof with a slight pressure drop after $x = -0.06$ followed by a pressure increase and stabilisation up to the middle of the rear slant at $x = 0.25$.

The notch region is the location where the most important changes are observed between

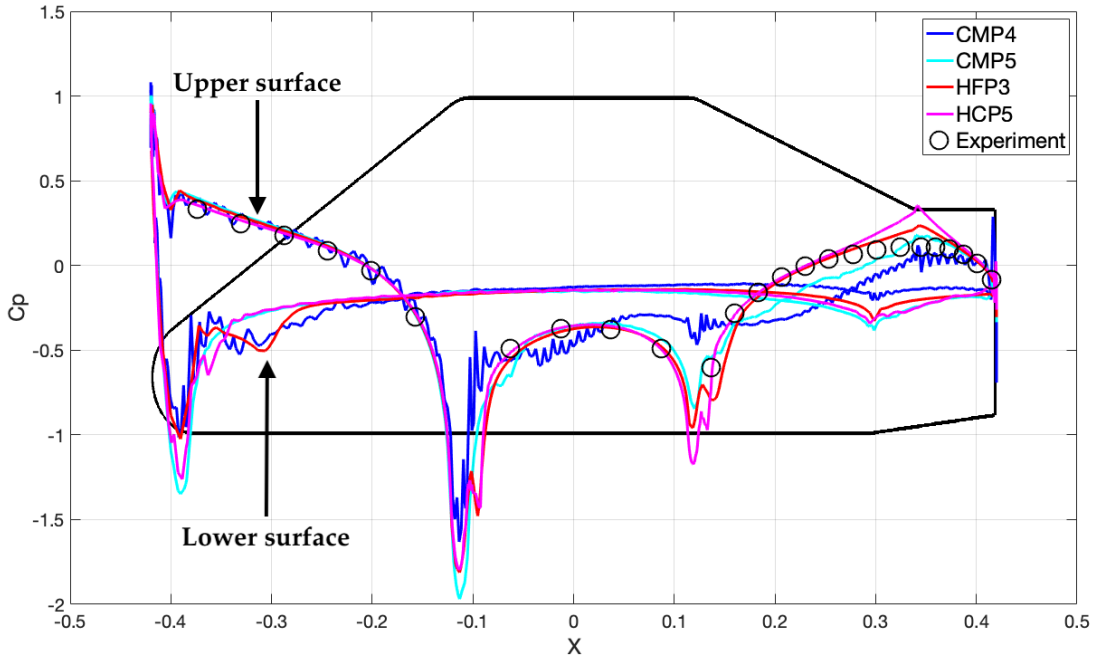


Figure 4.11: Pressure coefficient distribution on the centreline compared to experiments by Wood [83]

the four cases. An abrupt pressure drop happens at the rear edge of the roof, which magnitude increases from CMP5 to HFP3 to HCP5. The first negative pressure peak is followed by a second lower-magnitude peak. The lower the extension of the pair of negative peaks, the higher the longitudinal extent. Pressure then increases on the C-pillar. Pressure recovery is best predicted by HCP5 and slightly less by HFP3. However, pressure recovery is missed by both cases run on the committee grid, even if predictions with P5 are more accurate than P4. Resulting lower pressure values over the rear slant are responsible for higher drag values for CMP4 and CMP5 cases and also the positive lift in the CMP4 case given the severe pressure underprediction. From $x = 0.27$, computed pressure field for the HFP3 and HCP5 cases start diverging from the experimental values, overestimating them. Pressure then peaks at the same location, $x = 0.35$, in experimental measurements and all simulations. The peak is fairly well predicted by CMP5 but underpredicted by CMP4 and overpredicted by HFP3

and even more by HCP5. The overshooting of the pressure peak magnitude, corresponding to the impingement on the bootlid, induces larger downforce values on custom meshes compared to other cases. The following pressure decrease over the rest of the bootlid is fairly well predicted by all the models up to the trailing edge even if the level of agreement with experimental results is higher with CMP5 due to the better peak prediction.
Under the car, pressure decreases after the flat underfloor directly upstream of the rear diffuser and reaches a minimum approximately between $x = 0.28$ and $x = 0.29$ in all simulations. The suction is weaker in CMP4, and has approximately the same strength but occurs slightly earlier for HCP5 than HFP3, CMP5 reaching the lowest pressure values among simulations. Sudden variations in pressure are present at the extreme rear of the body, where flow detaches. The lack of experimental measurements at this location prevent from drawing final conclusions on the accuracy of simulations on this part of the vehicle but simulations and experimental results of [84] qualitatively agree on the presence of a negative pressure peak at this location of the body.

In summary of the first results, CMP5, HFP3 and HCP5 show rather consistent predictions whereas very important discrepancies are observed for the CMP4 case. However, discrepancies of the results on custom meshes are somewhat reduced when performing a p-refinement from P4 to P5. Unsurprisingly, larger differences appear in the notch region for pressure recovery over the rear slant and impingement on the bootlid resulting in a pressure peak. Although the results with HFP3 and HCP5 are rather close, localised differences remain on the front nose, the diffuser and pressure recovery over the rear slant which seems to be slightly better predicted by HFP5. Still, the magnitude of the impingement peak is in better agreement with the experiment for the HFP3 case. A particular focus is made on this zone of the flow using PIV results of Wood [83] to give a more precise insight of flow physics and interpret pressure distribution results.

4.4.3 Flow visualisation in the notch region

Wood [83] produced a series of 2D-PIV planes in the notch region of the SAE body providing a better qualitative understanding of the flow mechanics in this intricate area to explain differences in pressure distributions observed on the centreplane. Longitudinal in-plane velocity magnitude field over the notch region and in the wake are shown for all the cases and compared to PIV visualisations of Wood on figure 4.12. Looking at the experimental data, despite poor resolution at the vicinity of the wall, the flow seems to remain attached on the rounded corner between the floor and the rear slant and starts separating slightly afterwards at $x = 0.13$ on the C-pillar. In simulations, the flow strongly detaches from the corner for both CM cases hence the lower pressure peak values compared to custom meshes where the separation is weaker and the deviation imposed by the rounded edge results in a deeper decrease in pressure coefficient observed for the negative peak at $x = 0.12$. For the committee grid, the separation is far more important for the P4 case than P5 case, hence the very mild negative pressure peak in the P4 case. The separation is less important for HCP5 than HFP3 where a separation bubble accompanied by low-velocity values can be noticed between $x = 0.13$ and $x = 0.17$. The recirculation bubble also exists for HCP5 but due to its tiny size, it can barely be spotted. In general, separation in simulations run on the committee mesh is more important than on custom grids which translates in more moderate (for CMP5) to quasi-absent (for CMP4) negative pressure peaks at the trailing edge of the roof.

Planes with higher resolution but focusing only on the rear slant are shown on figure 4.13. In-plane velocity magnitude contours indicate that a low velocity area develops slightly upstream from the base of the C-pillar in the experimental visualisations. The weak separations provokes an inflexion and thus slowing pressure recovery starting from $x = 0.18$ up to the impingement on the bootdeck at $x = 0.35$. Simulations results are quite differ-

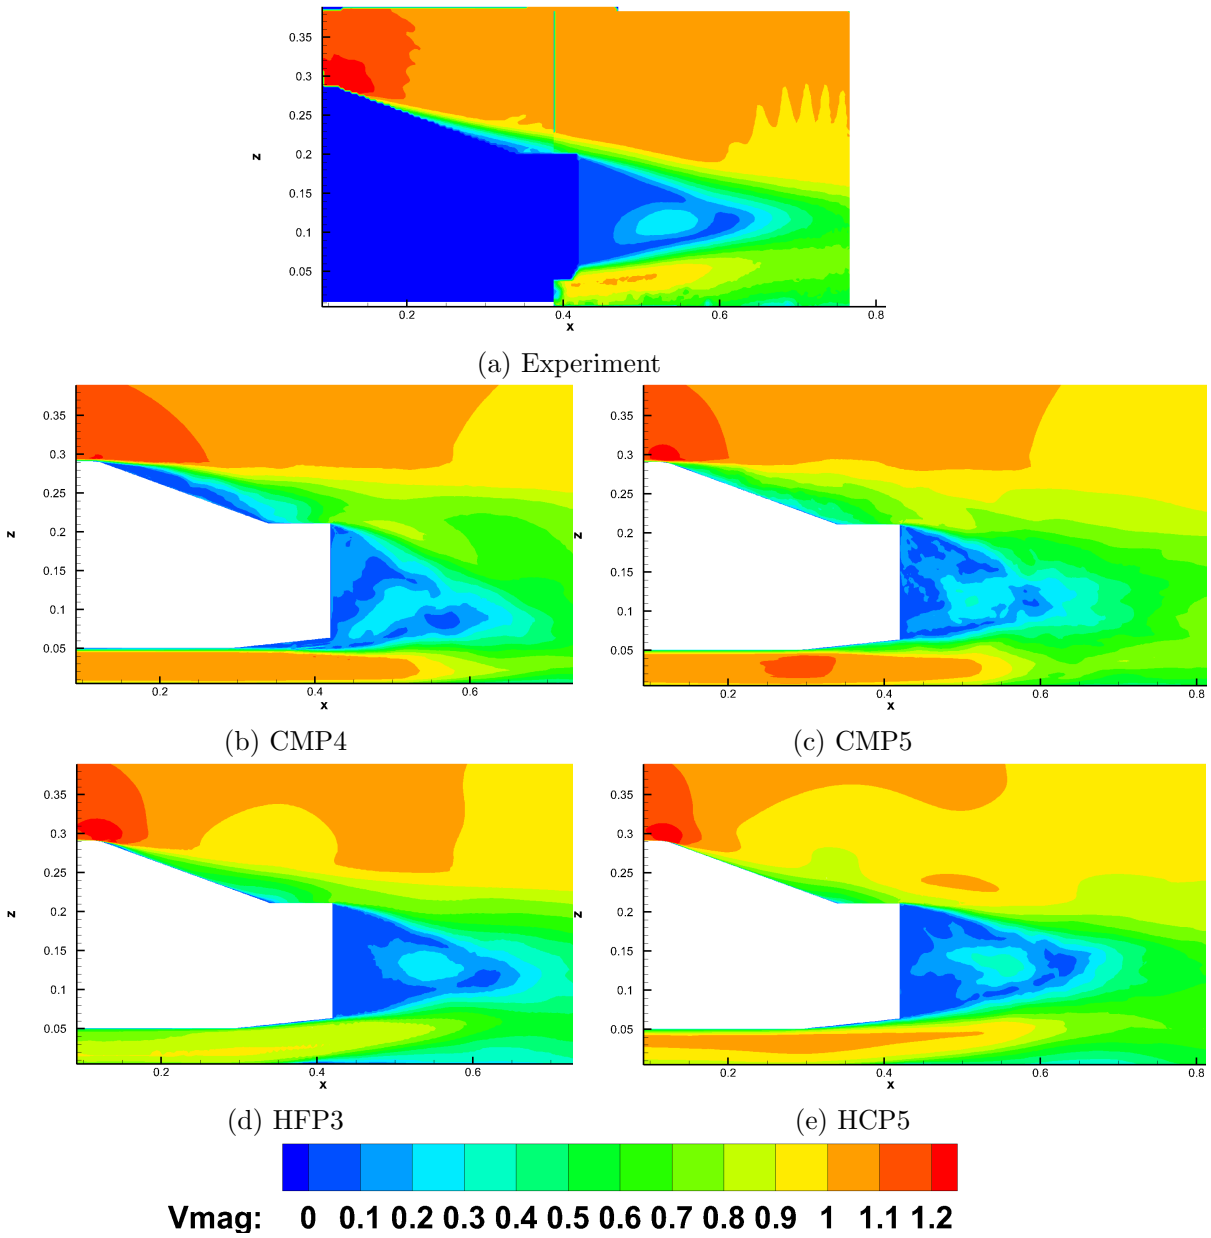


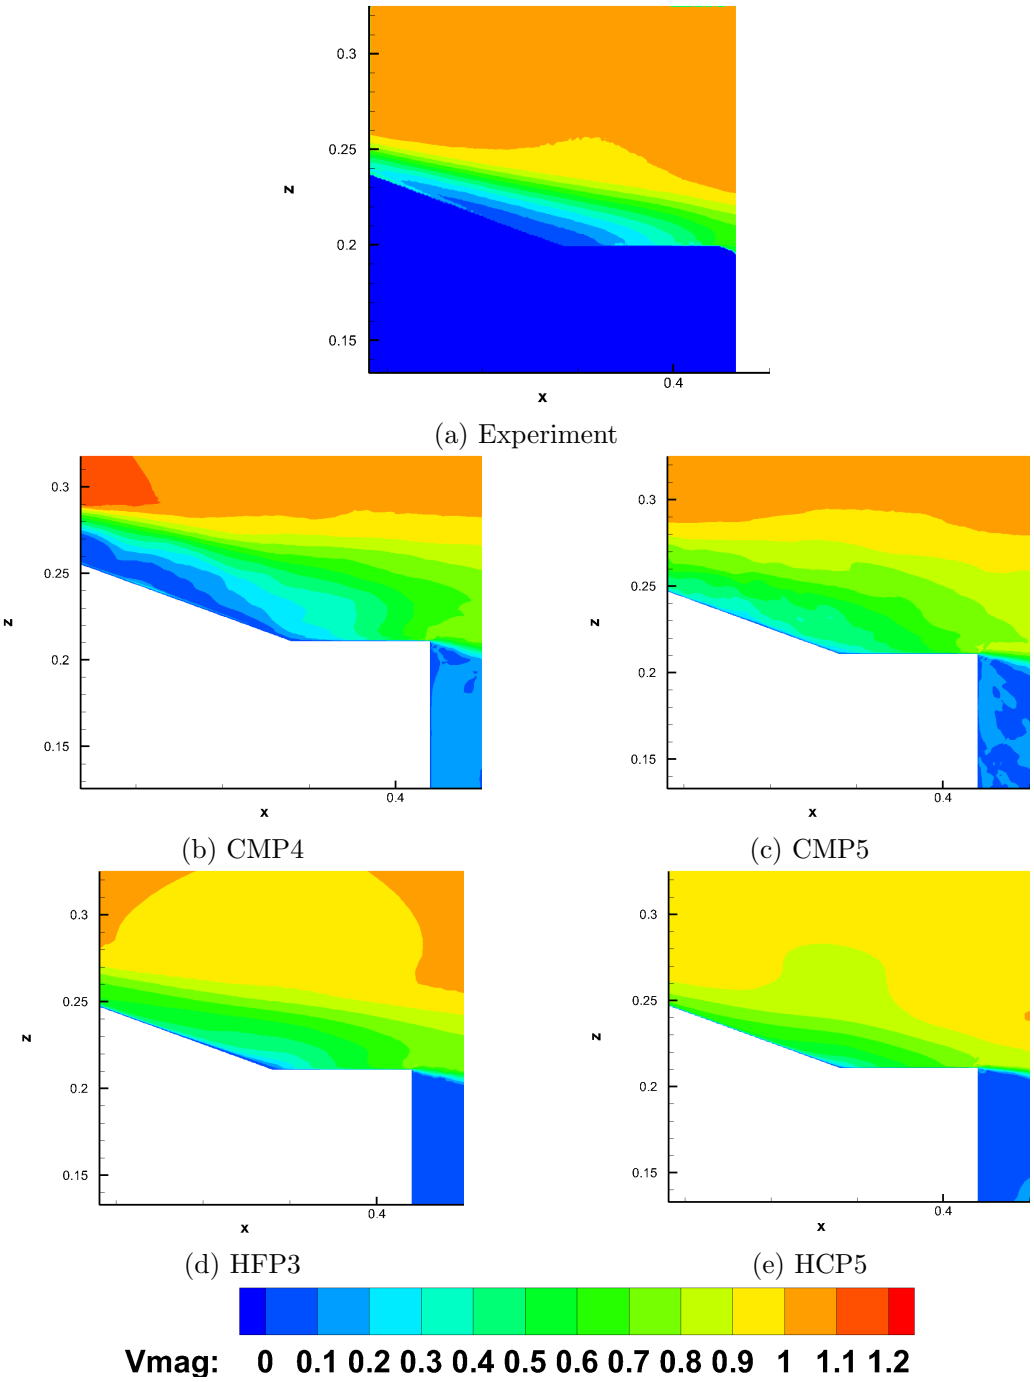
Figure 4.12: In-plane velocity magnitude on the centreplane in the notch region

ent between committee and custom meshes. Flow is essentially detached on the whole rear slant for CMP4 and CMP5 with low velocity layer extending from the top of the slant to its base. The separation is more important for the P4 case with a thicker low-velocity layer compared to other cases. When it comes to custom meshes, the very low-velocity magnitude

area is considerably smaller. Figure 4.12 shows that after flow reattachment at the top of the rear slant, the flow remains attached to the rear slant and that boundary layer only slightly thickens compared to experimental data. The very low-velocity area extent at the base of the slant with velocities below 0.2 is tremendously shrunk for custom meshes especially for HCP5, resulting in strong impingement and the large pressure peak observed at the base of the C-pillar on figure 4.11 which magnitude is logically higher for HCP5.

Coming back to figure 4.12 and looking at the wake, the length is underpredicted on CM cases and well predicted on custom meshes, especially for HFP3. The height of the wake is exaggerated in CMP4 compared to the experiment and the other simulations which is mainly due to the underbody flow on the rear diffuser that is severely detached in the CMP4 case increasing the size of the recirculation zone from the bottom. The detached flow over the diffuser can be inferred form the pressure contour over the centreline where the negative pressure peak at the start of the diffuser is weaker than in other cases. It can also be remarked that the averaged results are smoother on custom meshes than on committee grids creating more uniform contour plots.

Velocity profiles directly upstream from the junction between the roof and the C-pillar are plotted on figure 4.14 and help to explain the separation. Low resolution of the PIV near to the surface of the body prevent from having accurate values for velocity near the wall. However, slightly above the wall, streamwise velocity is underpredicted in all simulations. From $Z = 0.31$ and above, velocity values in CMP5 and HFP3 are close to experimental values while CMP4 and HCP5 are equal and smaller. Simulations all match after $Z = 0.37$ while remaining slightly below experimental data. Looking closer to the wall, two outliers can be spotted blow $Z = 0.3$ due to the evoked poor resolution issue close to the wall. Comparing simulations shows that near-wall velocity vastly varies between cases. CMP4 is the case with lower velocity values followed by CMP5. Both cases have the same shape for the



(a) Experiment

(b) CMP4

(c) CMP5

(d) HFP3

(e) HCP5

Figure 4.13: Zoom on the in-plane velocity magnitude on the centreplane over the rear slant

velocity profile, with a velocity peak at $Z = 0.297$. Increase in polynomial order from P4 to P5 results in an increase in near wall velocity by $0.05 \times U/U_\infty$ at the peak location and higher difference below the peak.

On the other hand, the shape of the velocity profiles of HFP3 and HCP5 between $Z = 0.29$ and $Z = 0.32$ are sensibly different which contrasts with results of p-refinement preserving the shape of the velocity profile and highlighting the important impact of h-refinement on the velocity distribution in the boundary layer. The HCP5 velocity profile peaks at $Z = 0.293$ with a normalised velocity of 1.27 and the HFP3 reaches a local maximum of 1.21 at this height. Closer to the wall, the difference increases to $0.2 \times U/U_\infty$ between the two cases. Pursuing away from the wall, streamwise velocity then decreases in the HCP5 case while it keeps increasing in the HFP3 case which passes over at $Z = 0.307$ where its starts decreasing. The near-wall shape of the velocity profiles for CMP4, CMP5 and HCP5 look more in agreement with the experiment than the HFP3 case and the tendency is reversed further from the wall. The velocity profiles, especially the near-wall region directly correlate to separation mechanism. Indeed, higher velocity gradients at the vicinity of the wall enables the boundary layer to withstand adverse pressure gradients caused by the transition to the C-pillar and therefore remain attached over a longer distance. Therefore it is consistent to see that separation is stronger and happens earlier in the case of the committee grid than custom meshes due to the lower near-wall velocities observed. However, an improvement of the boundary layer velocity profile is observed when increasing p-refinement from P4 to P5 leading to a significant improvement of the prediction of the separation and pressure distribution in the notch region. From a hp-refinement perspective, better forecast of the pressure distribution directly after the roof trailing edge and the first half of the rear slant for the HCP5 are due to higher near-wall velocity values obtained in this case shrinking separation which is closer to experimental observations. Yet, further from the wall, above $Z = 0.31$, the rapid decrease in streamwise velocity values in the HCP5 lead to better agreement of the HFP3 case with

experimental measurements, which is visually confirmed in velocity contour plots of figures 4.12 and 4.13 where below one velocity magnitude contours extend higher in the HCP5 case than in the HFP3 case and in the experiment.

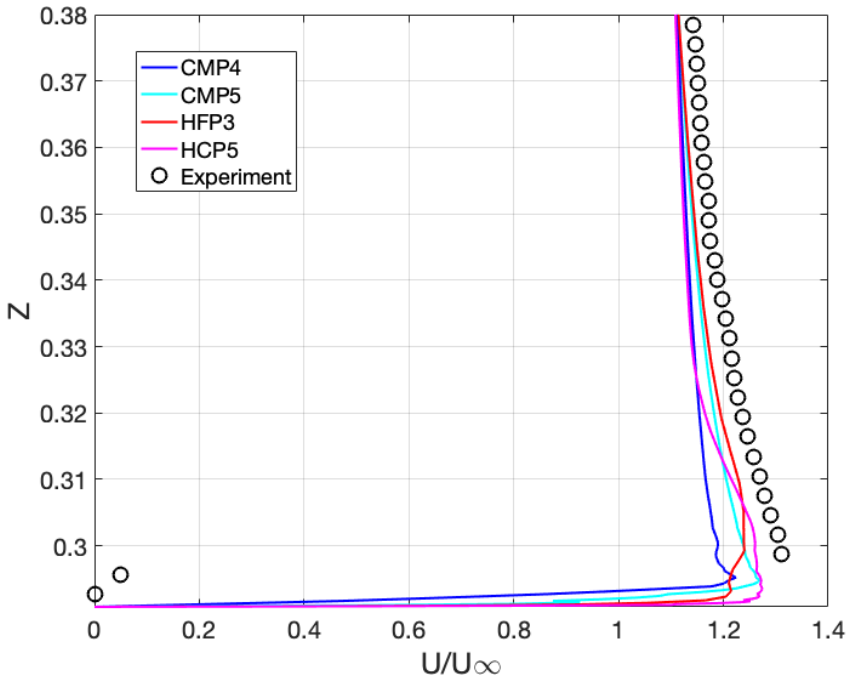


Figure 4.14: Normalised streamwise velocity profiles upstream from the roof trailing edge at $x = 0.11$ compared to experiments by Wood [83]

The velocity field on the centerplane has been investigated to explain the differences in pressure distribution, especially the separation and impingement in the notch region. However, the flow around automotive bluff bodies is highly three-dimensional. The object of the next section is to give a tridimensional view of the flow field and compare the resolution of flow structures in the four cases.

4.4.4 Three-dimensional flow structures

The previous sections focused on overall quantities and pressure and velocity fields on the centreplane. Automotive flows are characterised by small aspect ratio bodies with flow detachment generating three-dimensional flow structures, especially vortices and wakes. The prediction of these flow structures is of foremost importance as they dictate the aerodynamic behaviour of the car as previously shown with P14.

Total pressure coefficient Cp_0 quantifies the energy in the flow and is defined as follow

$$Cp_0 = \frac{p + 0.5\rho U_{mag}^2}{q} = \text{constant} \tag{4.3}$$

where p is the static pressure, ρ is the air density, U_{mag} the velocity magnitude and q the free-stream dynamic pressure. Cp_0 corresponds to the energy if the flow is virtually brought to rest through an isentropic process i.e. without any energy loss. Cp_0 iso-surfaces are used by aerodynamicists to highlight flow regions associated with energy deficit and are particularly useful to detect recirculation regions, wakes and vortices. Iso-surfaces of total pressure coefficient of 0.75, represented on figure 4.16, qualitatively highlight the flow topology around the SAE body. As the pressure field is not measured in the experiment, total pressure coefficient has been replaced by X-normal velocity magnitude PIV planes on the rear slant and the bootdeck to identify flow structures.

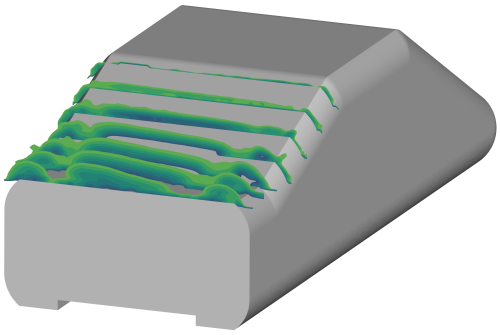
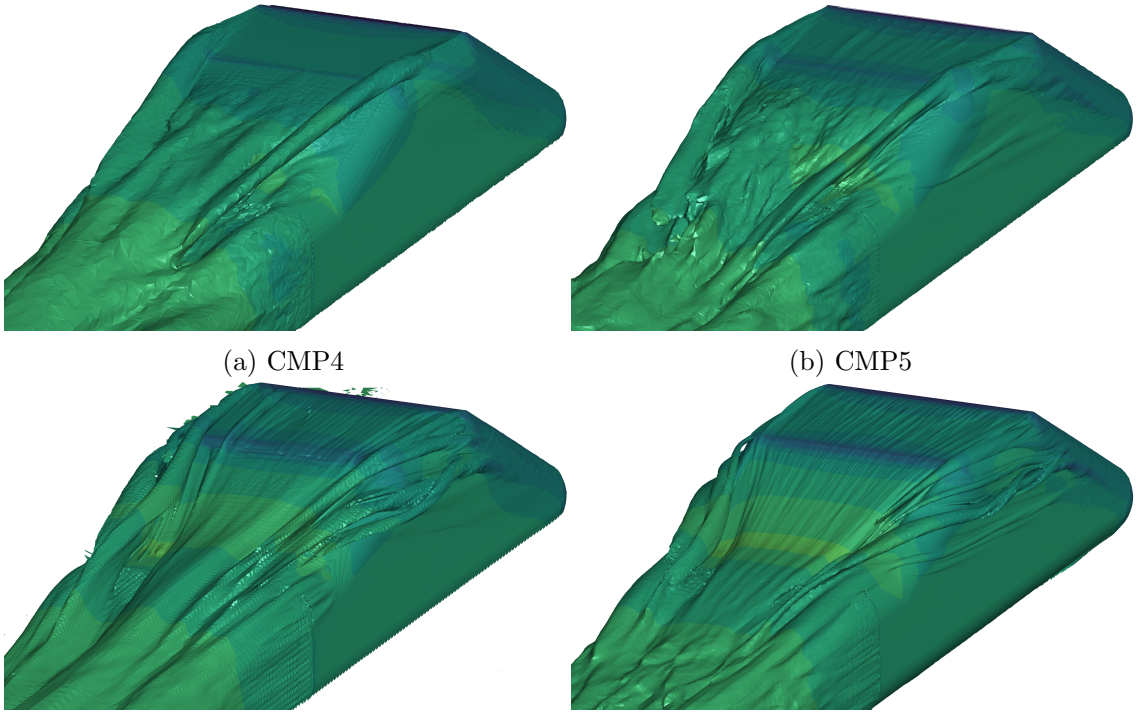


Figure 4.15: X-normal planes of velocity magnitude over the notch of the SAE body extracted from experimental results of Wood [83]



(a) CMP4 (b) CMP5

(c) HFP3 (d) HCP5

Figure 4.16: Time averaged iso-surface of total pressure coefficient $Cp_0 = 0.75$ coloured by pressure coefficient showing the flow topology around the SAE body

In the experimental data, the boundary layer thickens while the flow travels down the rear slant. Thickening is uniform on the first half of the C-pillar and progressively splits in three parts further down while two low-velocity lobes grow on the side edges causing the mid-

dle section to become narrower. At the trailing edge of the bootdeck, the three regions are clearly distinguishable with the middle part becoming asymmetric in the spanwise direction, the right part extending higher than the left part of the central lobe. Iso-contours of total pressure show large discrepancies in the flow arrangement around the vehicle for numerical simulations. Starting from the front, simulations on the committee grid predict only one vortex forming along the A-pillar and further developing on the top side of the vehicle up to the rear slant for CMP5 whereas no vortical structure emerges from the A-pillar in the CMP4 case. Multiple vortical structures are shed from the A-pillar in simulations on custom meshes. They are very similar between the two cases with three main vortices which can be identified. The top vortex is similar to the one observed on CMP5 and originates from the corner of the junction between the front slant and the roof. The vortex travels up to the notch and merges with the rear wake in the HFP3 case but cannot be tracked after the middle of the rear edge in HCP5. A second vortex forms below on the A-pillar. It merges with the precedent vortex at the top of the rear slant in HFP3 whereas it remains distinct in HCP5 and can also be tracked up to the middle of the rear slant. An intricate set of vortices forms at the middle of the front slant, finally merging in one large vortex when they reach the rear slant before merging with the wake. In the HFP3 case, the vortex has two superposed branches, the bottom one immersing in the wake at the side edge of the base of the rear slant. A last vortex develops only in the HCP5 case on the side of the vehicle, downstream from the A-pillar and plunges into the wake at the base of the C-pillar.

Focusing on the rear, the important separation highlighted on the centreplane for CMP4 actually extends over the whole backlight. For the CMP5 case, separation is more important on the left-hand side of the wake where a large recirculation bubble is visible. When it comes to custom meshes, separation is definitely more limited even if more present in the HFP3 case compared to the HCP5 case. The separated region is centred, becoming narrower from the top of the rear slant to the trailing edge of the vehicle. Although the separated region on

the rear slant is symmetrical in HFP5, an asymmetry with the right-hand side being taller can be seen in HFP3.

A focus on the wake at the base is made in figure 4.17 where X-normal plane of normalised streamwise velocity above the trailing edge of the bootdeck are represented. Experimental results have already been described on figure 4.16 with a three-part flow arrangement, an asymmetric central zone surrounded by two lobes of vortical flow. Large separations in the case of committee meshes result in a very different shape for the central part, with the zone of normalised streamwise velocity u/U_∞, extending up to $Z = 0.28$ whereas it does not go beyond $Z = 0.24$ in the experiment. Both cases are slightly asymmetric, with the flow converging and exiting the bootdeck at $Y = -0.035$ in the CMP4 case and $Y = -0.045$ in the CMP5 case which is highly asymmetrical in the Y direction. The asymmetrical flow on the bootlid has been pointed out by Wood [83] with the flow leaving the vehicle with an offset from the centreplane at $Y = -0.03$. In both cases, small lobes with weakly swirling flow are present at the side edges of the bootdeck, which are quasi-symmetrical in the CMP4 case whereas the left-hand side lobe can barely be spotted in the CMP5 case. Finally, a large vortical structure in the CMP5 case is centred at $Y = -0.03$ corresponding to the bubble observed on figure 4.16. Simulations on custom meshes conserve the tree-part layout. Still, vortices on the side edges of the bootlid are slightly off the sirface in the Z direction compared to simulations on committee grids hence a better agreement with the experiment on this point. The separation in the central zone is smaller and weaker in CMP5 compared to HFP3 even if it is slightly wider which is coherent with the smaller separation previously remarked on the rear slant. Yet, in both cases the flow is very symmetrical and leaves the bootdeck at the centrelplane location.

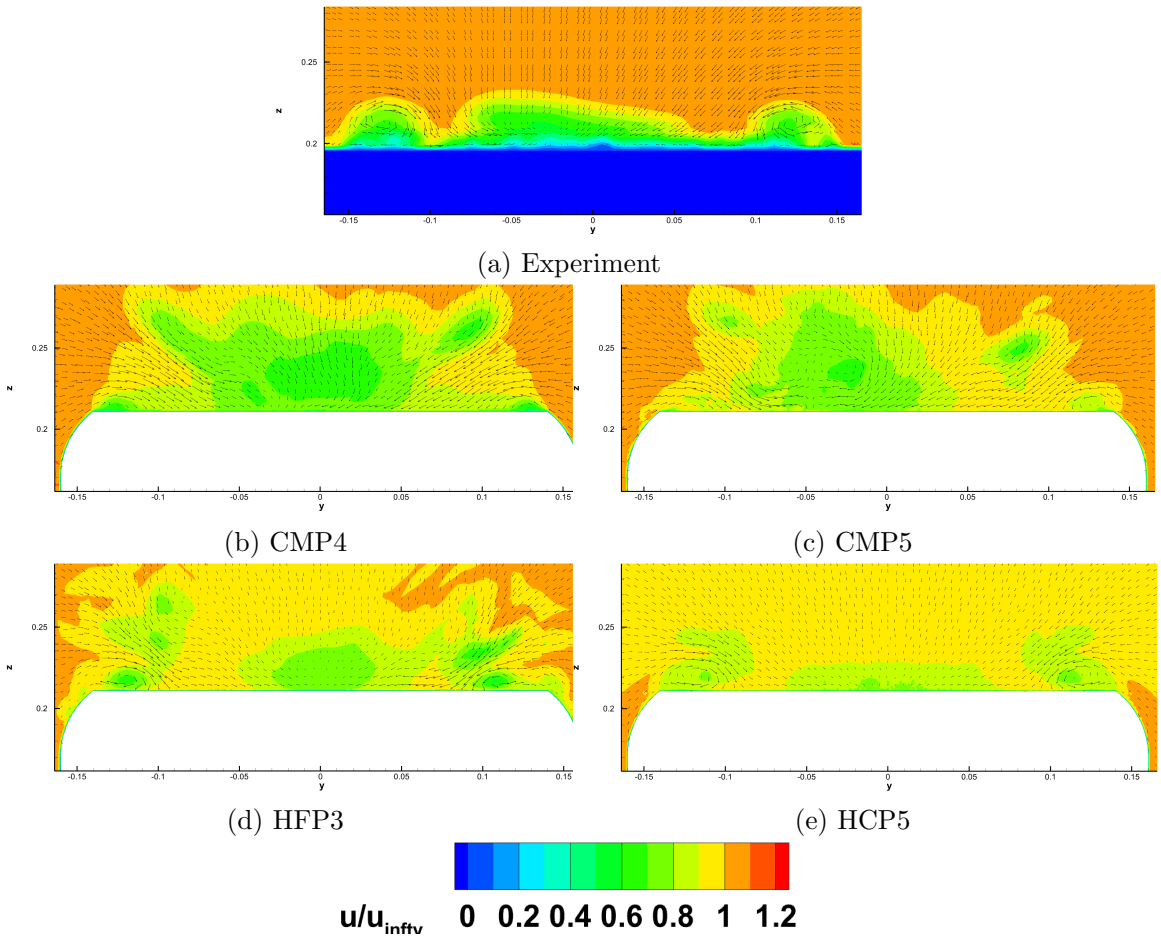
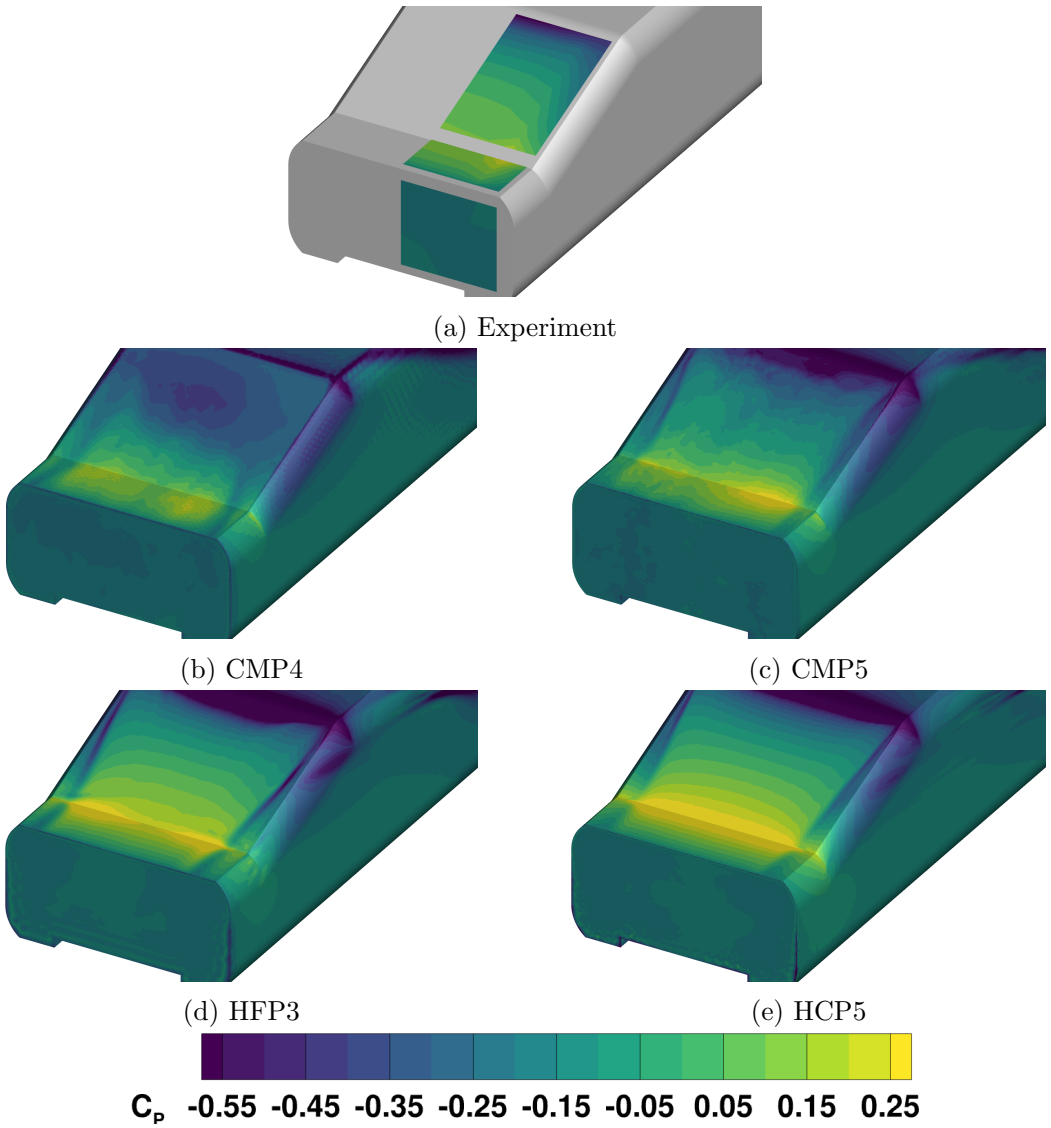


Figure 4.17: X-normal planes of streamwise normalised velocity (using the inlet freestream velocity) at $X = 0.42$ corresponding to the trailing edge of the SAE body. Experimental results are extracted are extracted from Wood [83]

A more global view of the pressure distribution on the centreline introduced on figure 4.11 is now represented on the whole notch region on figure 4.18 and quantifies the effect of the flow structures described previously on the pressure distribution. In the experimental data, the pressure recovery over the rear slant is clearly visible. Howerver, it is not uniform over the whole spanwise section, the pressure build up suddenly falling near the side edge which is due, according to Wood [83], to the lateral flow entrainment caused by the vortices rolling up over the sharp side edges of the C-pillar after separation which are indicated by

the side lobes on figure 4.15. When approaching the base of the backlight, a pressure peak is present between the centreline and the side edge. Wood claims that is corresponds to the downwash induced by the vortex coupled to its downward travelling direction inducing strong impingement on the surface at this location. Pressure on the bootdeck is marked by a line of high-pressure coefficient aligned with the pressure peak previously observed which the author assumes is due to the stronger impingement generated by the thickening and increase in intensity of the vortices rolling up over the side edges. Between the side edge and the peak, the increasing extent of low-pressure area is hypothetically attributed to a raise in angular momentum after the impingement the junction with the backlight.



(a) Experiment

(b) CMP4

(c) CMP5

(d) HFP3

(e) HCP5

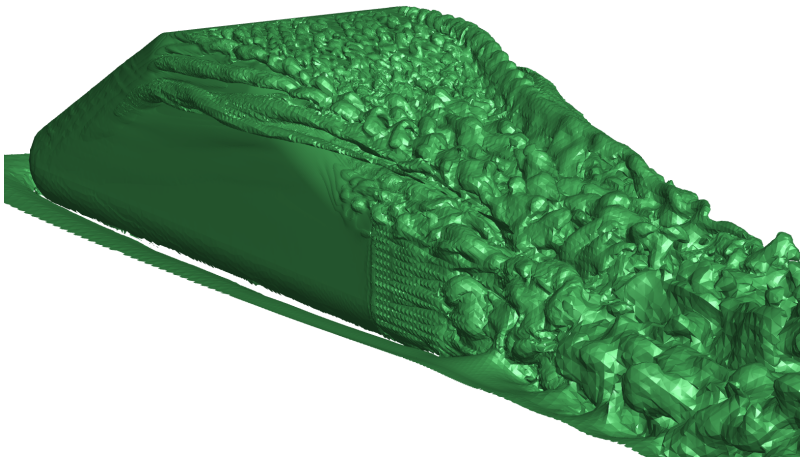
Figure 4.18: Pressure coefficient distribution on the notch of the SAE body. Experimental results from Wood [83] are acquired for only half body

Pressure recovery over the backlight happens in all simulation cases but important differences exist. In CMP4, pressure build up only starts from the middle of the rear slant due to the large separation on the first part resulting in a large low-pressure region occupying the first half. The right and left-hand sides of the backlight are not symmetrical, with the delay in pressure recovery observed due to vortex roll up on the side edge appearing on the

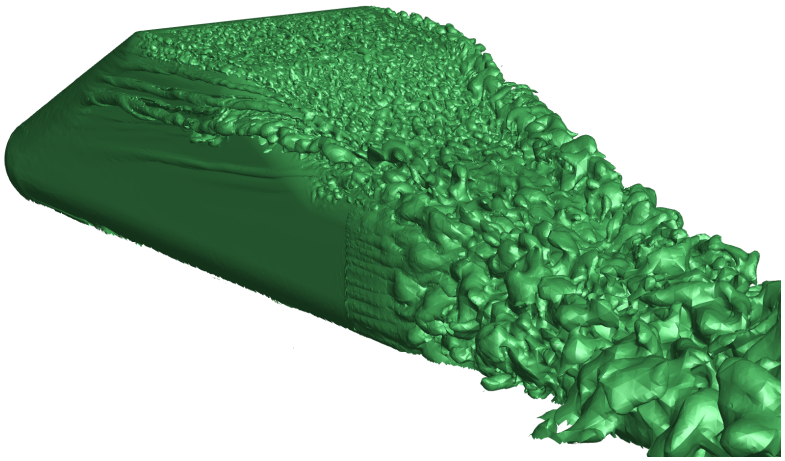
left part from the junction with the roof and on the right part from the middle of the rear slant but unlike experimental data, pressure then rapidly rises up to the edges. The pressure peak observed at the rear slant/bootlid junction is shifted towards the bootdeck and more centred compared to the experiment. When increasing polynomial order to P5, the subsequent reduction of the separation at the trailing edge of the roof enables pressure recovery to happen notably earlier and the positive pressure peak is translated further upstream at the junction between the rear slant and the bootdeck and towards the side edges of the vehicle which is closer to the experiment. However, the peak on the right hand side of the vehicle is considerably higher than in P4 which is caused by the left/right unbalance of the separation on the notch already observed on figure 4.15. Pressure drop associated with the side vortices is marked by a low-pressure line following the vortices and separating on one side the pressure peak due to impingement and on the other side a low pressure area. The lines are also present in CMP4 but can only be seen from the middle of the backlight due to the larger separation and propagate in the opposite direction, outwards. On custom meshes, pressure recovery is similar between the two cases on most of the rear slant and is more uniform compared to CMP4 and CMP5 because of the weaker flow detachment off the surface compared to these cases. There are however some differences between HFP3 and HCP5. The longer separation in HFP3 already noticed at the centreplane on figure 4.12 at the trailing edge of the roof causes the low pressure area at the central part of the top of the backlight to extend further downstream than for HCP5. Secondly, the high pressure at the backlight/bootlid junction is not restricted to two peaks but extends on the whole span what is in accordance with the symmetrical flow at the trailing edge on figure 4.17 yet singular among the cases presented. The airflow remains attached to the surface on the major span of the backlight and strongly impacts the bootdeck, generating the wide high-pressure distribution. Comparatively, the two lobes in the HFP3 case are wider than in the experiment and their position more central. Similarly to the CMP5 case however, a very low-pressure trail caused by the flow deviated

laterally by the trailing pillar vortices force the central zone of pressure rise to narrow and delay the pressure recovery on the zone directly next to it. On the outboard side of the trail, there is a zone of very limited pressure fluctuations which is called a "dead flow zone" by Wood even if a small increase in pressure can be remarked just before the bootlid. The trails in both HFP3 and HCP5 are better defined and extend further than in CMP5 and follow the same inwards path up to the trailing edge of the SAE body. Negative pressures induced by the side vortices are lower in the HFP3 case than in HCP5 where they are more dissipated. The trails are not present in experimental data, which can also be due to the lack of resolution of the pressure sensors grid over the surface given the very localised aspect of the phenomenon.

A comparison of instantaneous iso-surface of total pressure coefficient $Cp_0 = 0.75$ at $t = 16s$ for HFP3 and HCP5 cases can be seen on figure 4.19. The general flow layout looks similar between the two cases. Vortices generated at the side edges of the front slant and travelling downstream on the sides of the car and merging with the rear wake are present in both cases indicating swirling flow at these locations is not present only after time-averaging. Yet, these vortices look more perturbed and better resolved in the HCP5 case. This can be put in parallel with the flow on the roof of the car where separated flow starts from the font slant/roof junction and extends up to the backlight where more massive separations appears up to the rear wake. Near-wall flow is more finely resolved in this location as well confirming that even if h-refinement is lower in the vicinity of the wall in the HCP5 case, the additional p-refinement compensates the lack of geometrical resolution.



(a) HFP3



(b) HCP5

Figure 4.19: Instantaneous iso-surface of total pressure coefficient $Cp_0 = 0.75$ at $t = 16s$

4.4.5 Discussion

Simulation results on the SAE body have highlighted important differences depending on h- and p-refinement parameters, all the other parameters being fixed. In particular, the notch region and especially separation has turned out to consequently vary between the four cases presented. Indeed, the rounded junction between the roof and the backlight engenders a considerably more intricate detachment physics compared to cases with sharp edges forcing separation. Considering first the p-type refinement study carried out on committee meshes,

the increase in polynomial order from P4 to P5 vastly improved the accuracy of predictions with an important reduction of separation on the rear edge of the roof and on the rear diffuser enhancing wake forecast. P-refinement therefore appears as a powerful tool to increase the prediction accuracy without having to redesign the underlying mesh which is an arduous process. However, the polynomial expansion being applied to each element, raising the polynomial order greatly increments the total number of degrees of freedom, in this particular case doubling the total number of degrees of freedom from CMP4 to CMP5. Still, the comparison between the results obtained on committee grids and custom meshes shows that p-type refinement is more efficient when combined to an adapted h-type refinement. In fact, even if the number of degrees of freedom for CMP5 is in the same range as HFP3 and HCP5, the separation on the backlight is largely overestimated compared to these two cases. This indicates that an optimal meshing procedure for the spectral/hp element differs from standard techniques and that standard mesh best practices cannot directly be transposed into high-order meshing. The limiting criteria for scale resolving simulations of three-dimensional cases at high Reynolds numbers is the total number of degree of freedom due to computational limitations, leaving the choice of h and p-refinement to the user. Results have shown with the same number of degrees of freedom and mesh layout, different settings for h and p lead to distinct outcomes. Separation prediction requires a particularly fine resolution at the wall to predict the interaction between the boundary layer and the pressure gradient and it turns out that the HCP5 case better predicts the phenomenon. This is mainly due to the near-wall velocity values upstream from the roof/backlight junction which are higher compared to the HFP3 case. Indeed, as mentioned previously the flatter velocity profile enables the boundary layer to better resist adverse pressure gradient and therefore it is more likely to stay attached in the HFP5 case, what is observed in the results. Higher velocities at the vicinity of the wall in HCP5 indicate that smaller turbulent scales are resolved in this case. For under-resolved simulations as it is the case here, sub-grid scales are filtered by the grid and damped

by the SVV in the iLES approach. Moura et al. [30], [33] have analysed the eigenmodes of the spectral/hp method solutions of one-dimension linear advection-diffusion equations and shown that the SVV strongly dissipates high-frequency modes, corresponding to small scales, compared to lower order modes leaving larger scales relatively unchanged. Therefore, for higher p-order simulations as HCP5, large and intermediate scales carrying an important share of turbulent kinetic energy are less dissipated by the SVV comparatively to lower order simulations such as HFP3 explaining the deficit in near-wall velocity and subsequently larger separation. Therefore, p-refinement turns out to be an efficient way to capture the turbulent spectrum in the boundary layer. H and p-type refinements have a complementary role when it comes to off-body flow structures. Increasing polynomial order from CMP4 to CMP5, flow structures appear to be more finely resolved as can be seen with the A-pillar vortex propagating up to the wake. The impact of higher-order polynomial approximation is also present on simulations on custom meshes, where the resolutions of vortical structures is clearly better in the HCP5 than in HFP3 case where complex arrangements of vortices on the A-pillar, being close one to another, interacting and finally merging can be better distinguished in the HCP5 case. Yet, it turns out that vortices can be tracked over a longer distance in the HFP3 case, up to the wake, possibly meaning that energy of vortical structures diffuses at a slower rate in this case. These observations indicate that h- and p-type refinements have to be used in conjunction. Therefore, hp-refinement should be regarded as a two step process with first the linear mesh resolution which needs to be set in order to capture flow structures of interest. P-refinement can then be tuned to improve their resolution within the bounds imposed by restrictions on computational resources.

4.5 Conclusion

Simulations on the 20^o rear-slant SAE body have been performed with Nektar++ using different h- and p-type refinement parameters. Although simulations rather agree on the major front part of the vehicle, very important differences appear in the notch region especially due to the prediction of the separation on the rounded edge joining the roof and the rear slant. It has been shown that on a given base mesh, incrementing polynomial approximation order effectively increases resolution and has a direct positive impact on the accuracy of the predictions, shrinking the size of the separation. Yet, the global p-refinement is accompanied by a very heavy cost in terms of the size of the case, more than doubling the number of degrees of freedom when going from third to fourth order polynomial expansions. The direct transposition of meshes used with traditional industrial techniques into a base mesh for the high-order simulations proves to be inefficient and the mesh has to be completely redesigned taking into account both h and p parameters. The study has also show that for a given numbers of degrees of freedom, different settings for h- and p-refinement lead to different results. P-type refinement has proved to be particularly effective in the near-wall region as it preserves large and intermediate scales carrying large shares of the turbulent kinetic energy from numerical diffusion, extending the range of resolved turbulent scales compared to lower order polynomial expansions, what is crucial when it comes to the accurate resolution of complex phenomena such as the interaction between a boundary layer and adverse pressure gradient as it is the case in the study. Mesh resolution is yet important to capture flow structures as a certain level of h-refinement needed to be imposed to track vortices over a long distance, p-refinement then improving the resolution of the flow structures.

First general guidance for high-order mesh generation have been laid out, highlighting the interaction between h- and p-type refinement. It would be particularly interesting to perform simulations with increased p-type refinement to improve the resolution of the separation at the trailing edge of the roof and evaluate how high polynomial expansion needs to be incre-

mented for the kinetic energy in the boundary layer to be sufficiently high to overcome the adverse pressure gradient. Variable p-refinement would be a clear asset in this task, enabling a finer polynomial resolution localised in the boundary layer without vastly impacting the total computational cost.

Chapter 5

Rotating wheel modelling: Smoothed Profile Method

5.1 Literature review

In the P14 study in chapter 3, wheel modelling has been identified as a crucial point for the validity of industrial automotive simulations. Tyre grooves are often removed and the junction with the floor is modelled with a contact patch making the moving wall modelling suitable for the tyre. Despite the simplicity of this approach, it is rather limited and does usually fail when it comes to non-axisymmetric parts such as rims as it does not account for the added momentum due to the motion and the varying angular position of the spokes in unsteady simulations. In order to focus on the effect of wheel rotation and eliminate perturbations arising for instance from the front of the car, the case of an isolated wheel on a moving ground has been studied.

The literature review first focuses on describing flow features around isolated wheels in previous studies, both experimental and numerical. Emphasis is then laid on the various techniques used in CFD simulations to model rotating non-axisymetric parts. Finally, the Smoothed

Profile Method (SPM) which has been developed within Nektar++ for more accurate modelling is introduced and applied on an industrial wheel geometry.

5.1.1 Isolated wheel aerodynamics

Wheels have been overlooked in early development of car aerodynamics. It is only in the late sixties and beginning seventies that researchers have started focusing on their impact on the overall car aerodynamics (see Morelli [89], Stapleford and Carr [90]). Research efforts have then rapidly focused on the aerodynamics of isolated wheels to have a better grasp of the complex flow features induced by this complex case. Studies on isolated wheels will be reviewed in this section, mainly focusing on the importance of operating conditions, wheel rotation and the contact with the floor. Although early studies relied exclusively on wind tunnel experiments, rapid development of CFD in the nineties led to growing use of numerical tools in automotive aerodynamic development to support or complete experimental observations.

5.1.1.1 Experimental works

The reference study on wheels in isolation has been carried out at Imperial College by Fackrell in 1974 [6] who measured pressure and aerodynamic forces on isolated wheels rotating on a moving belt in a wind tunnel. Fackrell highlighted the impact of appropriate experimental representation of the wheel, e.g. rotation and contact with the ground, on the accuracy of the results. Two important flow features were reported on figure 5.1: the separation point on the top of the wheel at Region A moved forward compared to non-rotating wheels and a pair of vortices originating from the contact patch at Region B. There is a rapid change of direction of the flow around region A where the flow in the thin layer entrained by the rotation of the wheel is moving opposite to the incoming flow progressing along the wheel,

resulting in strong shear stress. The boundary layer usually forming directly on the surface of non-moving bodies is slightly shifted from the wheel in this case. Once the momentum of the incoming flow fails to overcome the adverse pressure gradient, the flow separates from the shear layer between the fluid regions with opposed directions resulting in the free separation point in region A. The latter phenomenon in region B is referred to as jetting. According to the author, jetting is due to the boundary layer on the bottom of the wheel converging with the flow entrained by the moving ground and creating a pumping phenomenon which pushes the flow towards the bottom side edges of the wheel in two vortices convected by the moving floor. Jetting is at the origin of the pressure peak of $Cp = 2$ at 90^o. This feature does not appear in the MIRA case run by Stapleford et al. [90] for which the wheel was put in a hole in a steady floor to allow rotation showing that this is necessary to have a rotating wheel in contact with a moving floor to observe jetting.

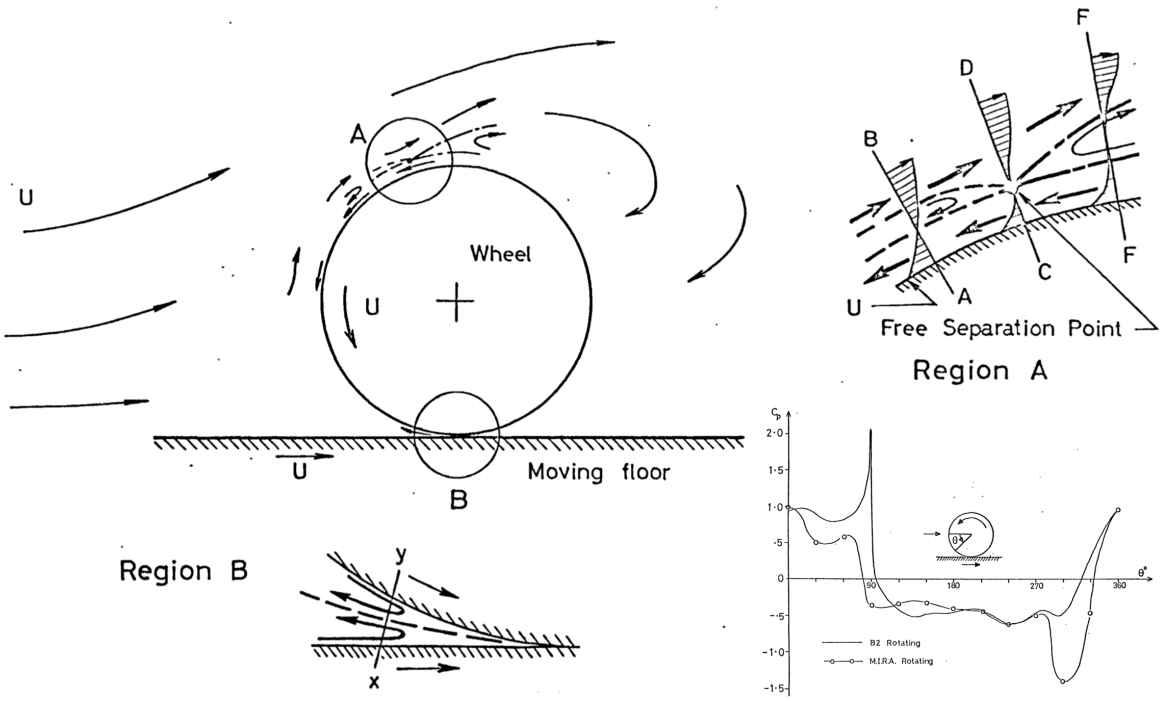
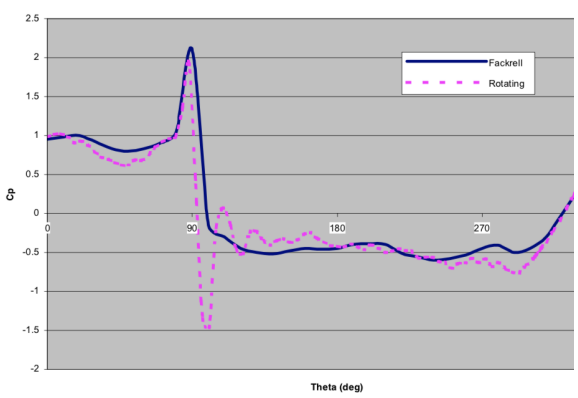


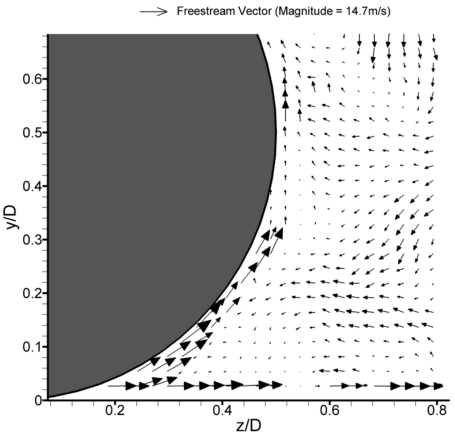
Figure 5.1: Flow features on an isolated rotating wheel as predicted by Fackrell (from Fackrell [6])

The importance of the sealing of the contact region was further emphasised by Cogotti [91] who measured the pressure on a steady floor near the contact region for a rotating wheel for different ground clearances. When the size of the gap was reduced, the associated Venturi effect between the wheel and the floor led to negative pressure values on the bottom surface of the wheel. The pressure upstream from the junction became positive only when the gap was sealed. Moreover negative pressure under the wheel resulted in negative lift contrary to what is expected for a rotating wheel in contact with the ground which should generate positive lift. When it comes to drag, it was observed that a rotating wheel produces less drag than a steady wheel. Results were confirmed by Mears and Williams [92] who measured the pressure distribution around the centreline of a racing wheel. The more recent experimental apparatus enabled them to measure a negative pressure peak directly downstream from the contact zone which was forecasted by Fackrell even if it was not observed in the measurements (see figure 5.2a). Angles are measured using the same convention as Fackrell on figure 5.1, that is anticlockwise with the origin at the apex of the wheel facing the flow. A region of relatively constant pressure coefficient can be observed between 100^o and 280^o, corresponding to separated flow. The decrease of pressure coefficient indicates that flow remains attached on the front of the wheel up to 280^o. In a second study, Mears et al. [93] observed using Particle Image Velocimetry (PIV) that the origin of the negative pressure peak downstream from the contact patch is a reverse jetting. Instead of converging like on the front of the contact patch, the boundary layer is split between the rotating wheel and the moving floor as can clearly be seen on figure 5.2b resulting in very low pressure on the bottom of wheel after the contact patch.

The wake structure of isolated wheels has been extensively studied. Fackrell [6] used Kiel probes to plot iso-contour of 90% total pressure and observed a taller wake in the case of rotating wheels which is due to earlier separation. Yet, the wake was wider near the floor in the case of the steady wheel because of a horseshoe vortex around the contact region which



(a) Pressure coefficient around the wheel. In blue: Fackrell's measurements [6]. Pink: racing wheel



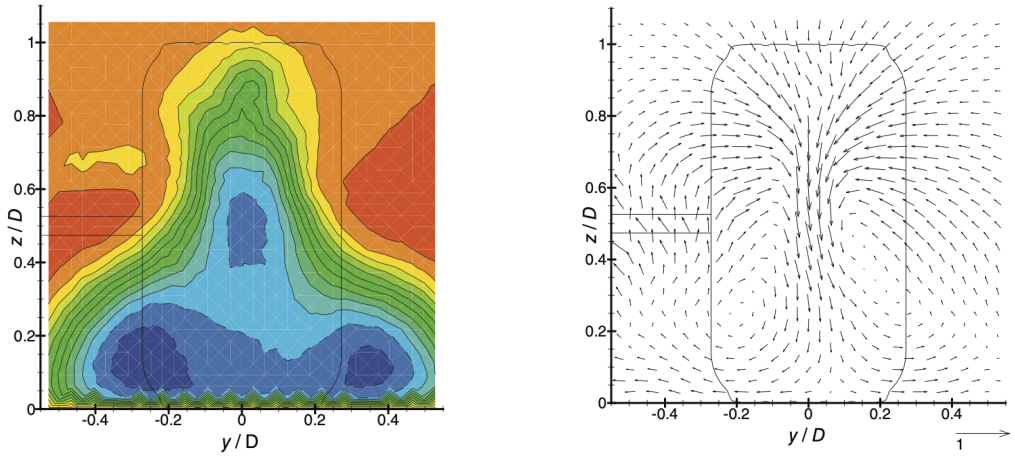
(b) PIV at the back of the wheel

Figure 5.2: Measurements around an isolated racing wheel by Mears [92],[93]

disappears when the wheel is rotating. Mercker [94] and Cogotti [91] postulated a similar wake structure with three pairs of contra-rotating vortices, one on each side of the wheel. The pair of bottom vortices has been experimentally observed whereas the top and middle pairs of vortices are only theoretical models. Mears used total pressure measurements to identify two low-total pressure lobes downstream from the contact line associated to two vortices, one on each side of the wheel. The size of the two regions, the strength of vortices and the magnitude of depression were larger in the case of a static wheel resulting in higher drag. A taller wake was also observed for the rotating wheel which is coherent with the separation point being shifted towards the front of the wheel when it is rotating.

For most studies cited above, the wheel is blanked but more recent studies include through-hub flow. Knowles et al. [95] have tested a 40%-scaled rigid champ car wheel rotating on a moving belt both without camber and with 4^o camber. The flow structures were very similar with and without camber but the drag was higher for the cambered wheel by 10.5%. Only the two jetting vortices and the top right vortex were observed in the LDA measurements compared to the six vortices predicted by Mercker. The authors assume the middle and

top left vortices are absent due to the sting and the flow through the hub which strongly influences the flow structures around the wheel by introducing asymmetry in the flow and causing part of the vortices analytically predicted to collapse. Knowles et al. [96] have conducted a second test campaign on a 50%-scale Formula One wheel using LDA in a wind tunnel and observed an inverted T-shape wake profile associated with a strong downwash region downstream from the wheel centreline with two recirculation regions near the ground, one on each side of the wheel as shown on figure 5.3a. Two vortices were also observed, one on each side of the wheel below the wheel axle and entraining flow therefore reducing the width of the wake from the top to the middle of the wheel (see figure 5.3b).



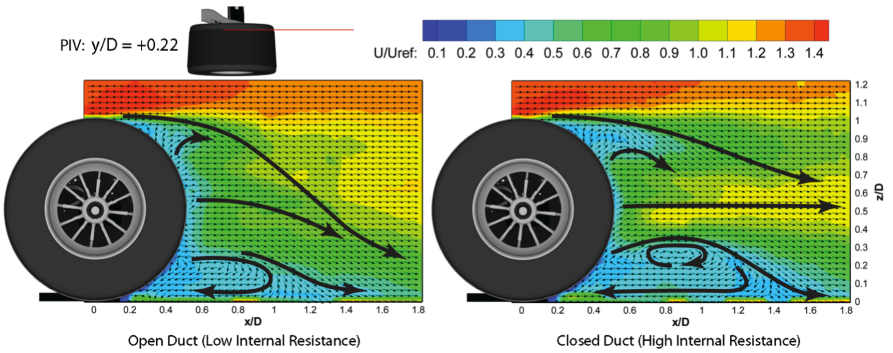
(a) Streamwise velocity

(b) In-plane velocity

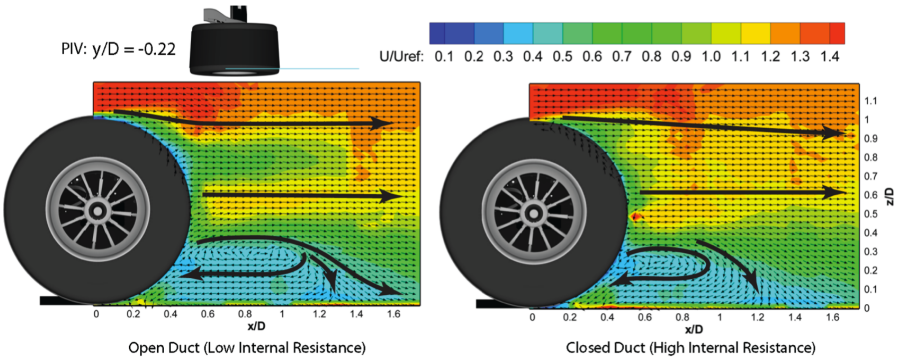
Figure 5.3: LDA observations in the wake of a 50% Formula One wheel by Knowles [96]

Sprot et al. [97], [98] have performed PIV and total pressure measurements on a 50% scaled Formula One rigid wheel cambered by 3.5^o with variable hub opening to control the flow through the wheel hub. Longitudinal (e.g. Y-normal) PIV planes on the inboard and outboard sides of the wheel represented on figure 5.4 show that higher blockage due to the closed duct caused a longer wake on the inboard side of the wheel (figure 5.4a). The effect of jetting flow exiting from the through-hub duct can clearly be seen on the open duct case at the height of the wheel axle on the outboard-side plane where it creates a region of high velocity

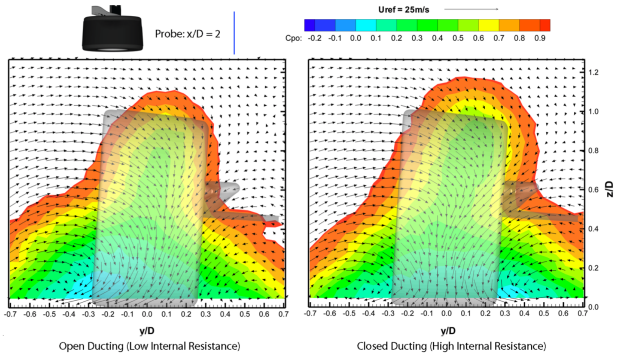
flow inside the wake starting from $x = 0.6$ on figure 5.4b. The middle plane however remains unaffected. The wake was measured two wheel diameters downstream from the wheel axle in the X-normal direction and is represented on figure 5.4c. The global inverted T-structure of the wake is coherent with the observations of Knowles on figure 5.3a. Comparing the closed and open duct cases shows that the size of the wake is reduced on the inboard side of the wheel and increased on the outboard side when the duct is open due to cross-flow jetting from the wheel hub. Indeed, part of the flow on the inboard side of the wheel is vacuumed by the duct bending the flow on the inboard side of the wheel towards the middle plane and this effect is emphasised by the low-total pressure wake of the inlet of the duct. Comparatively, the flow exiting from the duct on the outboard side of the wheel deviates the flow on the outboard side further away from the wheel which expands the size of the wake on the outboard side. Although downwash above the ground at the middle plane looks nearly vertical in the closed duct case, it is inclined towards the outboard side in the open duct case. The authors also measured a reduction of drag by 2.9% when the duct is open compared to the closed duct case.



(a) Inboard plane (y/D = 0.22)



(b) Outboard plane (y/D = -0.22)



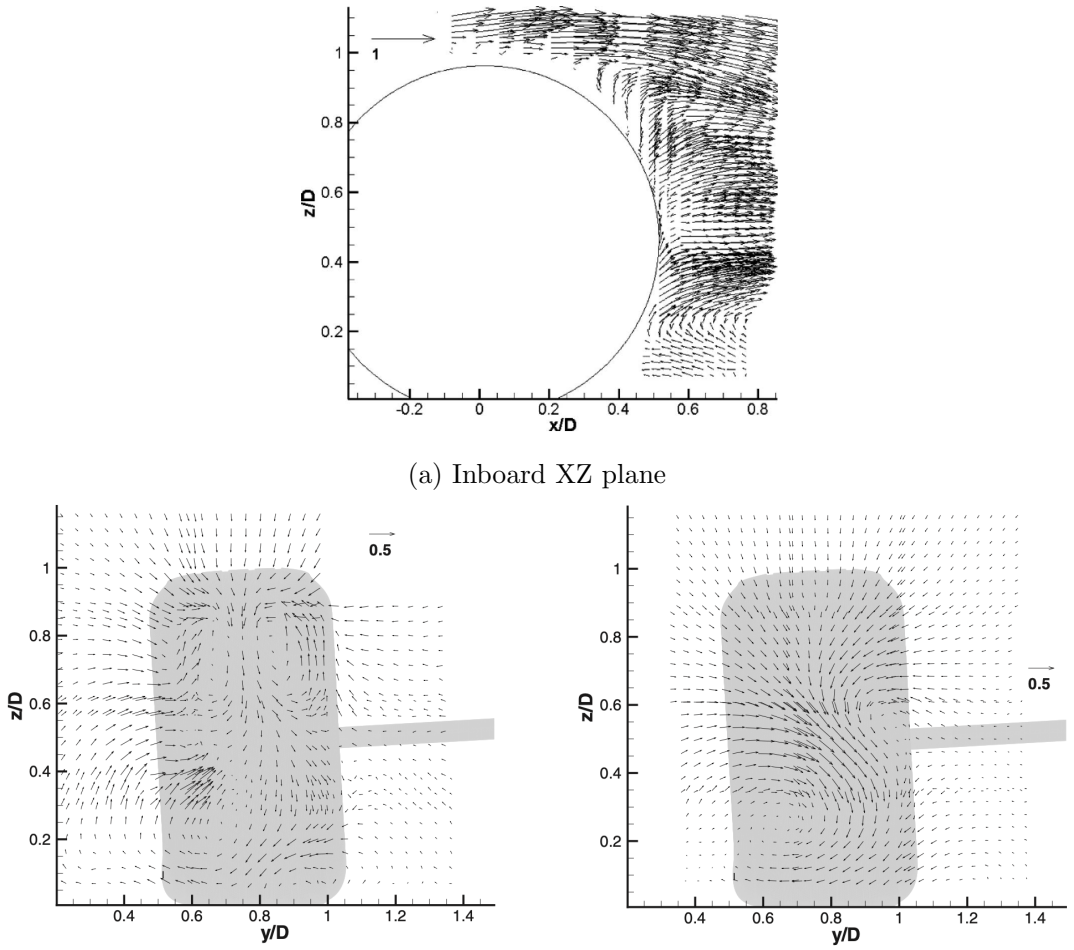
(c) Wake total pressure measurements (y/D = 2)

Figure 5.4: Wake measurements by Sprot [98] on an isolated rigid Formula One wheel. D is the diameter of the wheel

Issakhanian et al. [99] have carried out PIV measurements on a 60%-scale deformable Formula One isolated wheel with a camber of 2.5 degrees in wind tunnel. It was equipped with a spoked rim and a complete brake assembly with a cooling duct. The longitudinal wake

structure behind the inboard side of the wheel shown on figure 5.5a is in accordance with Sprot's findings with a wake about the height of the wheel rapidly decreasing in height along the X-axis and recirculating flow above the floor for $z/D < 0.4$ extending slightly higher than in Sprot's experiment. X-normal in-plane velocity vectors are plotted 0.57D and 1.14D downstream from the wheel axle taken as the origin on figure 5.5. Two pairs of counter-rotating vortices are identified just behind the wheel on the plane $x/D = 0.57$ on figure 5.5b, one on the top of the wheel and one on the bottom. The top pair of vortices comes from the flow detaching from the shoulders of the wheel and is symmetrical with respect to the centreline. The bottom pair is far from symmetrical since the inboard vortex, fed by the wake of the inlet of the through-hub duct and downwash, is bigger and located higher, nearly at the wheel axle position, than the outboard vortex. Indeed, the open through-hub duct draws the flow from the inboard side of the wheel and exits on the outboard side so that penetration of the outboard flow into the wake is limited. Global downwash at the mid-section in interaction with the system of vortices is once again observed. Although it is rather vertical at the top part of the wheel due to the two balanced top vortices, it is inclined towards the inboard side above the ground. Further downstream, at $x/D = 1.14$ on figure 5.5c only two vortices remain, one on the outboard side at the height of the wheel axle and a larger one on the inboard side closer to the floor. A third vortex seems to remain on the top of the ground on the outboard side but the limited vertical extent of the PIV plane near the ground does not allow complete observation. The authors suppose that the inboard lower vortex has absorbed the upper one, and that on the outboard side, the upper vortex has kept convecting downstream whereas the lower vortex has been dissipated. The S-shape of downwash intensifies compared to the $x/D = 0.57$ plane due to the highly unbalanced system of vortices and their relative position and strength. This arrangement of vortices is somewhat different from Sprot's study [98] where the inboard vortex vanishes by $x/D = 2$ and the outboard vortex grows larger and keeps convecting towards the floor. This can be

due to the different positions of the sting in the two experiments which is placed on the inboard side of the wheel in Sprot's experiment and results in lower strength of the inboard vortex.



(a) Inboard XZ plane

(b) YZ plane at $x = 0.57D$

(c) YZ plane at $x = 1.14D$

Figure 5.5: Wake structure observed with PIV by Issakhian et al. [99] on an isolated Formula One wheel. D is the diameter of the wheel

Wäschle et al. [100] previously observed wake structures behind an isolated 1:3-scaled rotating Formula One wheel with brake disks on a moving belt using LDA and also highlighted a strong downwash line downstream from the centreline of the wheel but only the bottom pair of symmetric vortices were spotted.

Croner et al. [68] studied the flow around a blanked full scale wheel rotating on a moving belt using PIV and hot-wire anemometry. Strong downwash was again observed behind the wheel centreline. Six vortical structures could be observed in the wake, two on the top due to boundary layer separation, two vortices on the ground coming from a horseshoe vortical structure which is formed downstream from the contact region and two vortices originating from the wheelbase at the height of the wheel axle and rapidly merging with the top or bottom pair of vortices. Spectral analysis of the velocity in the streamwise-normal plane one wheel diameter downstream from the wheel showed two large peaks appearing at the frequency corresponding to the rotation of the wheel and its first harmonic and another peak in between. The latest peak corresponds to vortex shedding from the contact region and dictates temporal evolution of lift and drag. Schnepf et al. [101], [102] compared the wake structure of a smooth and a treaded full-scale isolated rotating wheel on a moving belt with PIV and total pressure probes. The wheel was blanked on the outboard side and cambered by 0.25^o and a vertical load of 5500N was applied to mimic real conditions in a vehicle. The velocity field indicates strong downwash in the middle section separating the wake defined by low total pressure values in two parts. On the ground, the separation lobe coming for the outboard side is significantly bigger than its counterpart on the inboard side of the wheel. However no vortex coming from the top of the shoulders of the wheel was detected.

5.1.1.2 Numerical simulations

Axon et al. [9], [67] carried out the first numerical study of a simplified isolated wheel consisting in a low aspect ratio cylinder with slanted edges in contact with the ground using RANS and compared their results against Fackrell's [6] experiment. They observed the positive pressure peak in front of the contact patch and the forces were in good agreement with measurements. However separation was predicted 30^o later than in the experiment due to a incorrect pressure gradient in the simulation resulting in a misprediction of the wake.

In 2006, McManus and Zhang [60] performed unsteady RANS (URANS) simulations of Fackrell's wheel model both stationary and in rotation on a rolling belt. They deduced a flow topology in agreement and completing the observations of Fackrell et al. [6] presented on figure 5.6. The authors remarked that the separation on the back of the stationary wheel is only partial and that the flow remains attached at the middle section and is energised by the flow coming from both sides of the wheel. They also noticed, in agreement with Fackrell [6] that separation happened at the top of the wheel in the rotating wheel case and that the wake is approximately the same size as the wheel. An arch vortex is formed at the top of the wheel due to the separation which is very similar to the arch vortex forming on the backlight of notchback vehicles mentioned in chapter 4. Although Fackrell supposed the larger jetting vortices near the ground in the stationary case is due to the presence of a horseshoe vortex forming at the front of the contact patch, neither Fackrell or McManus and Zhang observed it. McManus et al. [60] proposed a different explanation. The higher velocity near the ground in the case of the rotating wheel due to the moving belt reduces lateral deflection compared to the stationary wheel and floor case. They also suggested different forming processes for the jetting vortices in the two cases. The jetting vortices are formed when the attached flow at the rear of the wheel reaches the floor and rolls up in the stationary case whereas they are due to the recirculating flow in the separated area behind the wheel which is convected downstream by the freestream in the rotating case. Computed surface pressure values and forces are in good agreement with the experiment for the stationary case and reasonably agree in the rotating case apart from some inconsistencies near the contact patch which is due according to the author to faulty measurements in this region. The claim is backed by the negative pressure peak behind the contact patch which is not observed in Fackrell's data although generally present in the literature.

Wäschle et al. [100] used their LDA visualisations to asses predictions of a Lattice-Boltzmann code (EXA PowerFLOW V3.4p2a) and finite volume commercial code (STAR-

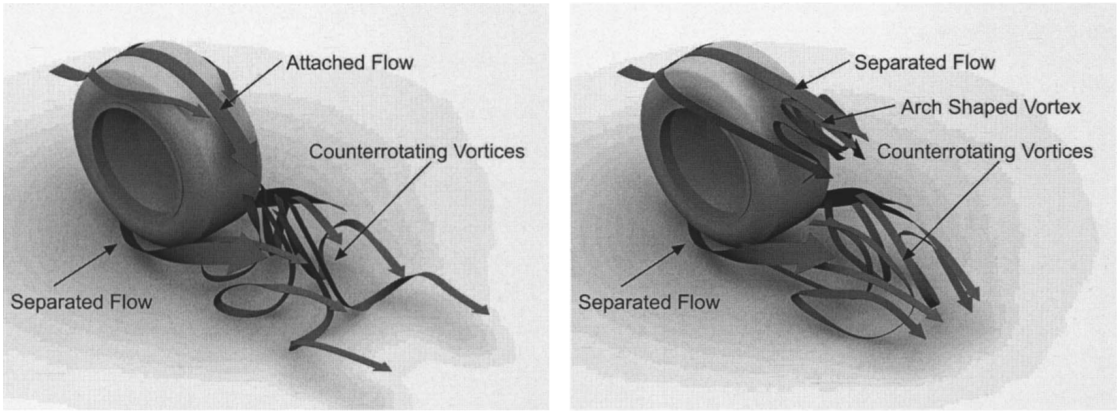


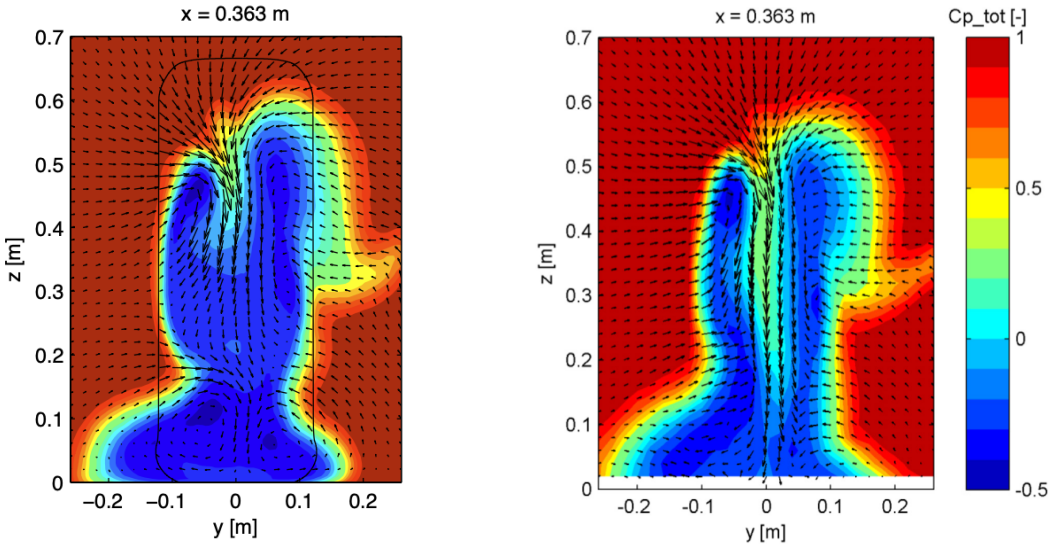
Figure 5.6: Flow topology around Fackrell's wheel predicted by URANS simulations by McManus et al. [60]. Left: stationary wheel. Right: Rotating wheel

CD V3.15) with two turbulence closure models ($k-\epsilon$ and two-layer treatment). The moving reference frame technique (MRF), which is defined in section 5.1.2, was used to model the rotative entrainment of the flow between the spokes of the rim whereas a moving wall boundary condition was imposed on the spokes in the Lattice-Boltzmann code. Although downwash is well predicted by PowerFLOW and in a lower magnitude in STAR-CD with the two-layer treatment, the recirculation region does not have the right shape in PowerFLOW and is too wide in STAR-CD.

Sprot [98] performed a RANS simulation of the rigid wheel measured in wind tunnel already presented and also noticed later separation on the wheel leading to poor wake predictions and a discrepancy of about 10% in forces compared to the experiment. A second simulation was carried using EXA PowerFLOW on a deformable Formula One wheel. Even if the numerical results are closer to experimental measurements in terms of flow structures, computed total pressure is lower in the bottom pair of vortices and higher downstream from the inboard upper shoulder of the wheel, near the inboard upper vortex. The author also remarked higher strength of the vortices in CFD. Although forces predictions were in agreement with the experiment, higher discrepancy was observed when the though-hub duct was fully open due to a lower amount of through-hub flow in numerical simulations compared to

experiments.

Schnepf et al. [101] and Schnepf [102] compared their experimental results on the rotating blanked wheel with EXA PowerFLOW CFD simulations. Results on figure 5.7 show a relatively good agreement on the envelope of the wake even if downwash is weaker in the CFD results. On the other hand, vortical structures especially just above the floor are stronger in simulations.



(a) CFD results

(b) Experimental results

Figure 5.7: Comparison of the total pressure coefficient 0.55D downstream from an isolated rotating wheel in CFD and wind tunnel by Schnepf et al. [101]

An extensive numerical study of the Formula One wheel studied by Issakhanian et al. [99] has been carried out by Axerio-Cilies and Iaccarino [103] with a stationary wheel with and without fairings to control through-hub flow. RANS and URANS simulations have been performed with several turbulence models but poorly converged due to the unsteadiness of the wake whereas the LES simulation, which has only been performed on the faired wheel, lead to more satisfactory results. The asymmetry of the bottom pair of vortices has once again been observed when the through-hub duct is open in RANS simulations even if it was minimal compared to PIV results. Moreover unsteady wake structures such as the horseshoe

vortex forming in front of the contact patch could only be observed with LES. The authors insist on the importance of accurately predicting downwash fuelling the vortical motion to correctly predict the location and strength of the bottom vortices. Axerio-Cilies et al. [104] have studied the same geometry with URANS modelling the wheel in rotation by applying a rotating wall boundary condition to the rotating parts (tyre, spokes and rotor) and a MRF to the corresponding cell zones to model the rotating volumes of fluid between the spokes and the rotor. The MRF resulted in a reduced jetting on the outboard side of the wheel due to the boundary condition between the spokes, reducing the size of the wake. The positions of the cores of the vortices 1.14D downstream from the wheel predicted by the different turbulence models are marked on figure 5.8. The locations are poorly predicted and there are important variations between the models, even if the unsteady $k-\omega$ SST gives better predictions than the other models.

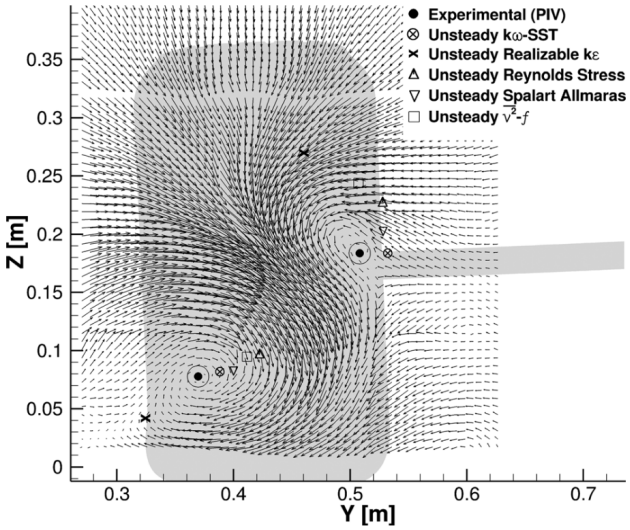


Figure 5.8: Comparison of the locations of the cores of the vortices at x = 1.14D in CFD and PIV. The circles around the PIV location represent the measurement uncertainty

5.1.2 Rotating wheel modelling techniques

5.1.2.1 Theory background

While experimental observations tend to indicate that wheels must be rotated and in contact with a moving floor, by using for instance a moving belt, to recreate realistic operating condition in wind tunnel as shown by Stapleford et al. [90] or Cogotti [91], different choices for wheel rotation modelling techniques have been used in numerical studies. Three main approaches can be identified:

Moving wall approach (MW) The first and simplest one is the moving wall approach (MW) already used in the P14 study and consists in applying a Dirichlet rotational velocity boundary condition V_s directly at the solid surface of the wheel of type

$$V_s = \omega \times r \tag{5.1}$$

where ω is the angular velocity at which the wheel rotates and r the radial position of a point on the surface of the wheel. Although it correctly models the tangential velocity on the surface of the wheel and can therefore successfully be applied to blanked wheels, the main shortcoming of this method is that the normal component on internal surfaces of the spokes, which are not rotating, is set to zero which is inaccurate.

Moving reference frame (MRF) A second approach called the moving reference frame technique (MRF) has been developed to correct this inconsistency. The approach is also called frozen rotor approach in turbomachinery applications. The rotating zone of interest, for instance the rim, is encapsulated within a subdomain to which a moving referential following the rotative motion of the wheel is attached. Navier-Stokes equations are solved in the whole domain and the inertial and Coriolis acceleration terms are treated as source terms

in the subdomain associated to the moving reference frame. The following mathematical development is adapted from [105]. Indeed, by denoting with a subscript r quantities in the moving reference frame, the material derivative of the velocity $\mathbf{U}$ in the absolute reference frame can be expressed with the velocity $\mathbf{U_r}$ in the relative reference frame moving with a fixed angular velocity vector $\mathbf{\Omega}$ as

$$\frac{D\mathbf{U}}{Dt} = \frac{\partial \mathbf{U}}{\partial t} + \mathbf{U} \cdot \nabla \mathbf{U} = \frac{D\mathbf{U_r}}{Dt} + \mathbf{\Omega} \times \mathbf{\Omega} \times \mathbf{r} + 2\mathbf{\Omega} \times \mathbf{U_r} \tag{5.2}$$

The term $\mathbf{\Omega} \times \mathbf{\Omega} \times \mathbf{r}$ is the centrifugal force and $2\mathbf{\Omega} \times \mathbf{U_r}$ the Coriolis force. The steady momentum can be rewritten in the moving reference frame

$$\mathbf{U_r} \cdot \nabla \mathbf{U_r} = -\frac{1}{\rho}\nabla p + \nabla^2 \mathbf{U_r} - (\mathbf{\Omega} \times \mathbf{\Omega} \times \mathbf{r} + 2\mathbf{\Omega} \times \mathbf{U_r}) \tag{5.3}$$

Using the velocity transfer identity between the two referentials

$$\mathbf{U} = \mathbf{U_r} + \mathbf{\Omega} \times \mathbf{R} \tag{5.4}$$

where $\mathbf{R}$ is the position vector from the absolute to the rotating reference frame, the momentum equation can be rearranged in conservative form using the absolute velocity

$$\nabla \cdot (\mathbf{U_r} \otimes \mathbf{U}) = -\frac{1}{\rho}\nabla p + \nabla^2 \mathbf{U} - \Omega \times \mathbf{U} \tag{5.5}$$

The source term $\Omega \times \mathbf{U}$ is then only applied in the moving reference frame and fluxes at the interface between the two regions are corrected using the source term to connect the rotating reference frame to the rest of the domain. The technique is quite simple and well suited for stationary approaches such as RANS while it compensates for the zero normal velocity observed on parts facing the rotation direction in MW and alleviates the lack of momentum in the wheel hub. It is therefore extensively used in industry. Yet, coupling between the

rotating and steady regions highly depends on the definition of the rotating zone boundaries and the method is not appropriate for unsteady approaches since the position of the wheel is fixed.

Sliding mesh (SM) With the increase in computational power and the rise of high-fidelity CFD simulations, unsteady representation of wheel rotation has become necessary. The sliding mesh method (SM) is different from the MW and MRF techniques since it is inherently unsteady and physically rotates the wheel geometry during the simulation. In this approach, the mesh is built in blocks and rotating parts are enclosed in a specific mesh block and interfaces with surrounding steady mesh blocks in the rest of the mesh are created. Although the approach is more correct than the two previously mentioned techniques, it is also far more computationally demanding and coupling with unsteady simulations methods makes it very expensive. It also requires important implementation efforts in particular to transfer quantities between the different mesh blocks since mesh interfaces may be non-conformal. Similarly to the MRF approach, the definition of the mesh block can have an impact on the results.

The three techniques introduced in the previous paragraph have been applied to rotating wheel simulations in the literature, and their results are presented in the next paragraph.

5.1.2.2 Application to rotating wheel simulation

Numerical experiments have been conducted to assess and compare the impact of different wheel modelling strategies on the flow around wheels. Bienz et al. [106] compared the three methods on a brake disk cooling case for a Formula One car and concluded that when only steady data is needed, MRF represents a good alternative to SM. Moreover, due to abutting parts such as brake disks and pads or the contact patch and the road, the SM implementation proves to be challenging. Kandasamy et al. [107] performed Lattice-Boltzmann simulations

of a BMW commercial car with interchangeable front and rear parts using PowerFLOW 4.3. They used three methodologies: MW, SM and a combination of MRF for the volume of fluid around rim spokes and MW for the tyre and compared the forces to experimental measurements. MW and MRF turned out to be less accurate for drag prediction with a maximum deviation of 5% compared to SM which is of 2%. When it comes to lift on the front and rear axles, the error is the largest for MW which differs from experimental values by up to 27 drag counts, double than MRF and SM approaches. Total pressure iso-contours and static pressure contour planes in the left-front wheel region indicate that the through-hub flow close to the wheel in the MW case is highly impacted by the fixed spokes geometry compared to the SM case where the wheel actually rotates and therefore results in a more accurate representation of the flow in this region. Yet further from the wheel region, the authors observe very little differences between the two cases. They also mention that the modelling choice impacts the pressure distribution on the wheel and the wheelarch. MRF predict front and rear axle lift as precisely as SM since the pressure distributions on the wheelarch are very similar. Still, the pressure distribution on the wheel is predicted with less accuracy in the MRF case hence higher error in drag prediction.

Similarly, Lewis et al. [108] performed Delayed Detached Eddy Simulations (DDES) of a full Aston Martin production vehicle. Tyres are modelled using MW and the spokes are encapsulated between two surfaces and MRF and SM are distinctly applied in the isolated region and compared. Although total drag values are very close between the two methodologies, lift is far lower in the SM case by 83% compared to MRF. The drag distribution is higher on the wheels with MRF but it is balanced by the higher drag generated on the rest of the vehicle surface with SM. Important differences in the pressure distribution on the wheels, especially near the contact patch, suggest a change in wheel wakes. Examination of the flow through the wheelhub using the cross-velocity component on the outboard side of the left-front wheel on figure 5.9 confirms Bienz' observations that it is strongly marked by the rim geometry.

The higher flowrate jetting through the rim results in a wider wake downstream from the front wheel in the MRF case. When it comes to computation time, the authors point out that SM, although more accurate, increasing computation time by more than 45% which is attributed to mesh rotation and information transfer at the interface between the two mesh blocks with non-conformal mesh elements.

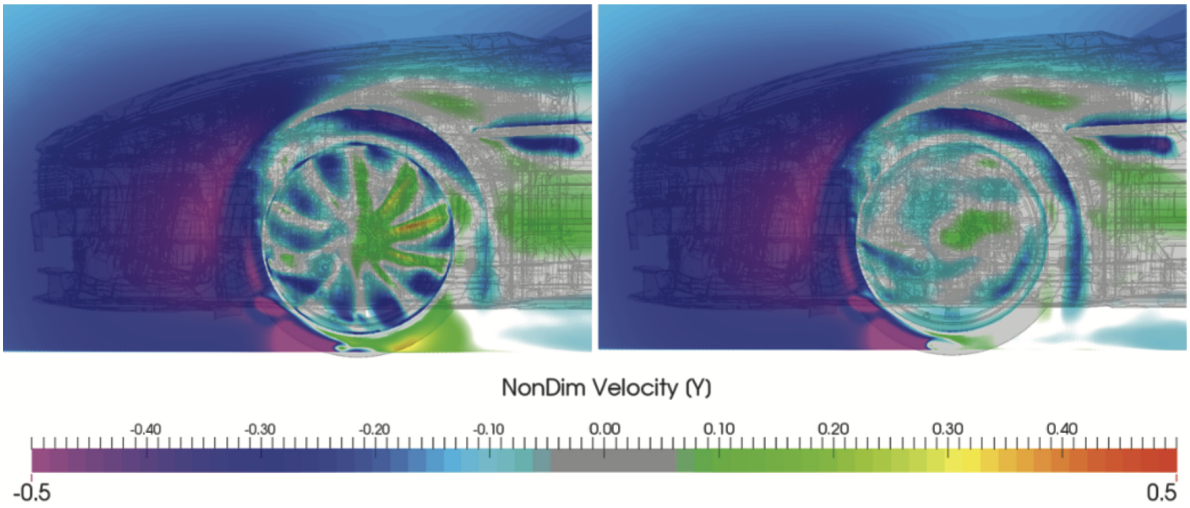


Figure 5.9: Cross-velocity component through the rim of the left-front wheel of a production car using two rim rotation modelling techniques. Left: MRF. Right: SM. Study by Lewis et al. [108]

Yu et al. [109] compared the three modelling techniques on a generic fastback vehicle and the DrivAer notchback geometry, which is an open-source realistic automotive geometry developed by the Technische Universität München (TUM) and introduced by Heft et al. [110]. The MRF and SM approaches are applied to rim spokes which are encapsulated between two planar surfaces and coupled to MW on the tyre. In agreement with the work of Lewis et al. [108], drag values for the car are very similar between the three methodologies and differences are below 4%. Yet, more severe discrepancies are noted for lift values, about 10% for the sedan fastback car. SM predicts downforce for the DrivAer model whereas positive lift is computed for the two other methodologies. In particular, the author remark significantly higher flow velocities in the wheelhub for the MRF technique due to artificially

added momentum, resulting in higher velocities downstream from the wheels and reduced rear wheel wake and rear wake compared to the two other methods. Plotting the evolution of the lift coefficient for a wheel during 5 wheel rotations reveals large amplitude oscillations, of about 50 lift counts, which have the same period as the wheel rotation. Results are reinforced by additional steady simulations with MW and MRF using reversed wheel angular positions in which an increase in lift by 50 counts is witnessed whereas the effect on drag is more limited. Yu et al. confirm that SM has to be used for unsteady simulations, especially when lift and flow topology are investigated. Although the authors do not dissociate the additional cost coming from the unsteady simulation when SM is used and the cost of the SM itself, they notice that computation time for SM is at least ten times higher than for the two other methods what makes it unsuitable for industrial applications.

Several studies focus rather on the variants for tyre modelling. Hobeika et al. [111] compared the combination of MW and MRF to SM on a Volvo commercial car. The wheels have been blanked to remove different spokes treatments, the vehicle lifted by 10mm above the ground to avoid tyre deformation at the ground which would make the SM approach impossible to implement and both slick and treaded tyres have been used in the simulations. Results show that when the reference frames used for the MRF encapsulated the whole wheel, large deviations in aerodynamic forces from SM are observed. The difference is importantly reduced when the MRF zone is limited to the air volume within the grooves. A change in forces has also been observed for the slick tyre between MW, which should be physically correct, and SM techniques raising questions about the influence of the extent of the SM zone. In a later study in 2018, Hobeika et al. [112] implemented combinations of the three modelling techniques to perform improved DDES (IDDES) of the rim and the tyre of a real wheel geometry both isolated away from any surface and in the Volvo production car already used in the precedent study. Forces comparison shows that SM is required to correctly predict the aerodynamic efforts on the wheel and that is is mainly due to the impact of the entrainment of the air in

the volume between spokes which is underestimated when MW is applied. The difference can mainly be explained by the velocity boundary condition on the surface of the spokes arising from the choice of a modelling technique. The MW approach actually correctly represents tangential velocities on the wheel but fails at correctly imposing the normal component on the internal surfaces of spokes as can be seen on figure 5.10. The authors also confirm that MRF or SM needs to be applied to the grooves in order to correctly predict forces on the treaded tyre. Therefore, MRF in the grooves (called MRFg by the authors) is a suitable approach when the wheel is in contact with the floor, a condition where SM is inapplicable due to the contact patch. Yet, MW is deemed sufficient to satisfactorily model the slick tyre. The authors also report an increase of 20% in computation time if only the rim is rotated and 50% if the whole wheel is rotated with the SM approach compared to MW.

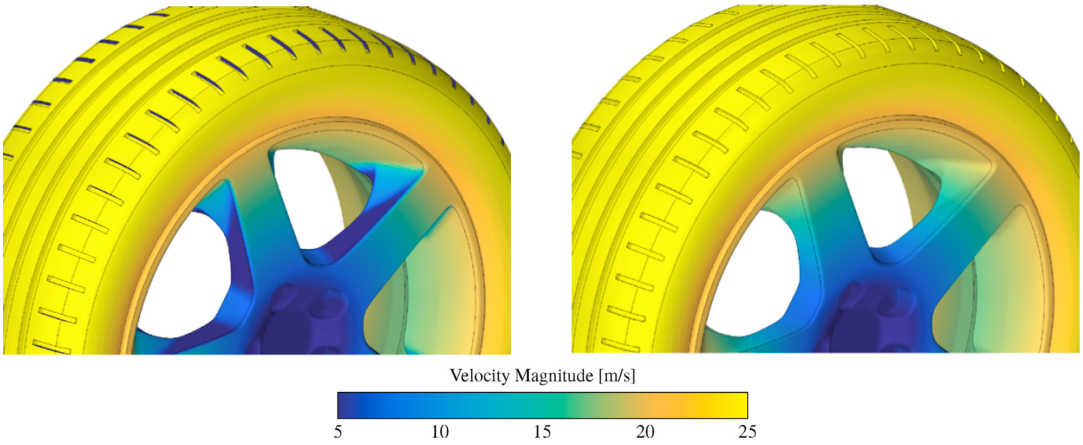


Figure 5.10: Comparison of the velocity boundary conditions for the MW (on the left) and SM (on the right) modelling techniques after Hobeika et al. [112]

Fu et al. [113] used the isolated blanked wheel model and experimental results of Schnepf et al. [101] and Schnepf [102] as a baseline for correlation with URANS simulations using MW, MRF and SM to model tyre rotation. Reference simulations with MW show a similar near wake arrangement at the top of the wheel with two lobes separated by a downwash strip but the bottom wake is significantly wider than in the experiment, resulting in higher computed

drag. Differences are attributed to different strut geometries as well as uncertainties in the definition of the shape of the contact patch. A second set of simulations is conducted with an open wheel using MW for tyre modelling and successively MW, MRF and SM for the rim. Both MW and MRF output larger drag values than SM, what put into perspective with the results on Schnepf's model [101], suggests that SM is more accurate than the two other methodologies. The largest discrepancy in terms of drag is observed for the MRF case. The author claims that unrealistic pressure gradients within the MRF zone are the cause. Indeed, looking at equation 5.5, the source term can be taken as an additional pressure gradient ∇p_{MRF} defined as

$$\nabla p_{MRF} = -\rho\Omega \times \mathbf{U} \tag{5.6}$$

which is zero only when velocity vectors are aligned with the rotation vector, ie. crossing the rim plane. This of course is not the case everywhere hence the rise of artificial pressure gradients.

Haag et al. [114] compared experimental data to simulations on several isolated wheel geometries in isolation and mounted on the DrivAer model [110] using the MW and SM. They first studied the impact on the extend of the SM zone on force predictions on an isolated blanked wheel, taking the MW results as baseline and concluded that is has to be as restrained as possible, only enclosing the spokes to provide the most accurate results. The accuracy of the MW approach has been assessed against PIV measurements 0.65 wheel diameters downstream from the wheel on the same blanked wheel put in rotation on a moving ground. Although the hub vortex and the pair of ground vortices have been successfully captured, the hub vortex and bottom inboard vortex location are poorly predicted, being shifted upwards in simulations compared to experimental observations. The downwash direction is also straighter than in the experiment especially, at the top of the wheel, where it follows a highly curved S-shape and drag prediction is 34 drag counts higher than in the experiment. The combination of MW on the tyre surface and SM in the spokes area developed in the first

part has then been applied to an industrial open wheel geometry. Only the bottom pair of vortices is present in this case both in the experiment and simulation. The upper part of the wake is in very good agreement but the bottom part is vastly wider than in the experiment and more symmetrical in simulations. Indeed, downwash in the bottom part of the wake is vertical in the simulation whereas it is inclined towards the outboard side of the wheel in the experiment. Drag difference is in the same range as for the blanked wheel, that is 37 counts higher than in the experiment. Authors also point out, using results on a rotating cylinder, that the use of SM instead of MW introduces an increase in computation time by 1.5 to 1.8. Misar et al. [115] replicated Fackrell's wheel model [6] and performed RANS, URANS and IDDES simulations using MW, MRF and SM. Best agreement for the forces is obtained for the IDDES simulation with the combination of MW on tyre surfaces and MRF on the rim, which is about 1% for drag and 7% for lift which is more intricate to predict and mainly due to the different representation of the ground contact. Indeed, in Fackrell's experiment, the wheel is rigid and the contact zone is thus limited to a line whereas a contact patch is used in simulations. Centerline pressure distribution correlates well with the experiment except for two areas: at the contact patch which is represented differently, explaining the larger difference observed for lift values, and at the rear-upper side and the top of the wheel where pressure is underestimated in the simulation. This corresponds to a detached region and it can be assumed that it is not correctly predicted by the IDDES turbulence model. Focusing on the wake region emphasises the fact that unsteady simulations with either MRF or SM for the rim, which is blanked in Fackrell's model, generally perform better in this region, especially at the bottom part of the wake. For the IDDES simulation retained, separation at the top of the wheel is minimised which can directly be related to the lower surface pressure values obtained in this region.

The review of previous studies focusing on wheel aerodynamics has proved that the flow

topology is complex, associated with separation and recirculation commonly observed on bluff bodies, and reinforced by the particular operating conditions which are rotation and proximity with the ground. Both conditions must be fulfilled, most frequently using a rolling belt in wind tunnel, in order to obtain meaningful output. Deviations from these conditions, such as stationary or lifted wheels engender important modifications of the flow field and inconsistencies in the pressure distribution and aerodynamic forces. Although the authors do not fully agree on flow features around an isolated wheel, the bottom pair of jetting vortices formed at the contact with the ground, flow separation at the top of the wheel and downwash in the centreplane of the wake are widely accepted. Notable variations in terms of strength of the vortices and downwash inclination are observed, mainly due to differences in the employed wheel models. Yet, measurement remains a real challenge due to the rotating motion of wheels and the proximity with the ground, and this is further reinforced in the case of shrouded wheels due to spatial constraints in the wheelhouse.

Numerical simulation therefore not only appears as a way to reduce computational costs of industrial design processes in this case, but also as an alternative to acquire data which are currently very difficult to collect experimentally. Yet, proper wheel modelling is also crucial for relevant CFD simulations. Numerical studies present in the literature indicate that turbulence resolving techniques usually perform better on this type of configurations than steady methods massively used in the industry because of highly separated recirculating flow coupled to strong adverse pressure gradients. Wheel rotation modelling is essential for accurate simulations, especially when it comes to real open wheels with non-axisymmetric internal parts such as spoked rims. Imposing a tangential velocity Dirichlet boundary condition on the solid surfaces is sufficient for the tyre and the moving floor but leads to inconsistencies when it comes to imposing the correct velocity on the internal faces of spokes normal to the plane of rotation which is set to zero. This issue is corrected in steady simulation by the MRF technique where an additional rotating momentum source term is added in the spokes

region. Yet, unphysical pressure gradients arise from the source term introduction and the position of the rim remained fixed which impacts cross flow through the rim and the wake structure. The SM approach is the closest to reality as the wheel is physically rotated with the mesh and results match the best with experiment among the three methodologies presented. Still, the method comes at higher computational cost compared to MW and MRF and implementation is intricate due to the information transfer between the rotating and steady mesh which often results in non-conforming interfaces and generates numerical errors and therefore flow misprediction. This issue can potentially be even more problematic when high-order meshing is employed in the case where curvature between two adjacent elements does not match. Hence the need for a different wheel rotation modelling technique when it comes to high-fidelity simulation of rotating wheels with the spectral/hp technology.

Immersed boundary methods (IBM) is a different class of approaches where the solid geometry is implicitly represented using an additional localised body force in the momentum equation. They are particularly convenient when it comes to moving bodies since the geometry can be physically moved without any mesh transformation. The smoothed profile method (SPM) belongs to this family and has properties which are well adapted to the spectral/hp element method. It is introduced in the next section.

The literature review has shown that contact patch treatment, tyre deformation and longitudinal grooves and lateral tread patterns are also important when it comes to rotating wheel modelling. Yet, these parameters have not been investigated in this study and are outside of the scope of this work.

5.1.3 Smoothed profile method

5.1.3.1 Introduction

The SPM has been presented by Yamamoto et al. [116], Nakayama and Yamamoto [117] and Yamamoto et al. [118] to address a hydrodynamic interaction problem between particles immersed in a fluid. It has been proposed as a replacement of the traditional Eulerian-Lagrangian treatment to alleviate computational costs associated with the simulation of a large number of interacting particles which need to be independently tracked. Indeed, boundary conditions for the Navier-Stokes equations are imposed for each particle which is a very computationally intensive process. The SPM overcomes this issue by imposing a forcing term to model fluid/particle interactions directly in the Navier-Stokes equation and thus tremendously reduces the whole set of boundary conditions to just those on the boundaries of the domain. The correct boundary condition is enforced through the incompressibility equation on all the particles. Solid bodies are removed from the domain and integrated in the fluid domain so that the need for a body-conforming mesh adapted to the shape of the particles is avoided. In this method, the interface between solid particles and the fluid domain is represented by a smooth function which prevents from having to readapt the mesh to align it with particles boundaries in the case when the boundary thickness is reduced to zero. The second advantage is that it is well suited for the CG approach of the incompressible Navier-Stokes solver of Nektar++. Indeed, zero-thickness interfaces imply jumps in the boundary condition that clashes with the continuous spectral/hp element formulation. In this method, particles are represented by a concentration function Φ which smoothly varies from one in the solid domain to zero in the fluid domain away from the particle and which rate of change can be tuned via a scaling parameter ξ. A schematic representation of the evolution of the concentration function is shown on figure 5.11 highlighting that the thickness of the transition region between the fluid and solid domains can be controlled using the scaling factor ξ.

The value of ξ is crucial in the method and takes into account mesh typical size in the region of the immersed body and the Reynolds number. This point is discussed later in section 5.2.3.

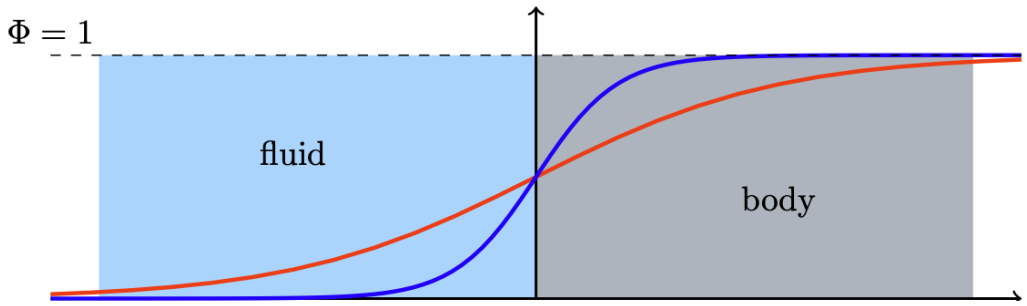


Figure 5.11: Evolution of the concentration function in one dimension with different values of the scaling factor ξ. High value (in red) and a low value (in blue). Extracted from [55]

5.1.3.2 Theoretical background

The SPM has been adapted to the high-order spectral/hp element framework by Luo et al. [119] and is introduced in this section. As seen previously, the concentration function Φ smoothly transitions from one in the solid body to zero in the fluid domain. The function proposed by Yamamoto et al. [116] and definitely adopted for the method later by Nakayama and Yamamoto [117] is the following

$$\Phi(\mathbf{x}, t) = \frac{1}{2}\left[\tanh\left(\frac{-d(\mathbf{x}, t)}{\xi}\right) + 1\right] \tag{5.7}$$

where d is the signed distance between the solid and the element at the spatial location $\mathbf{x}$ at time t meaning it is negative when the element is within the solid body and positive in the surrounding fluid away from the solid body. The scaling factor ξ controls the extent of the smooth transition zone between the solid and the fluid. A lower value of ξ reduces the thickness of the transition region. The chosen concentration function is C^{∞} and therefore

particularly well adapted for the CG formulation used in the incompressible Navier-Stokes solver from Nektar++. The generation of the concentration function in the case of complex geometries in three dimensions is tackled in section 5.2.

Once the solid and fluid regions have been identified by the concentration function, the numerical SPM scheme can be applied. It is implemented as an extension of the velocity correction scheme, presented in section 2.3, to take into account the body force $\mathbf{f}_s$ modelling the interaction between the solid body and surrounding fluid by enforcing the no-slip and no-penetration conditions at the solid boundary through addition of the forcing term on the righ-hand side of the momentum equation

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}_s \tag{5.8}$$

along with the incompressibility condition which remains unchanged

$$\nabla \cdot \mathbf{u} = 0. \tag{5.9}$$

The SPM starts similarly as the VCS for the three first steps.

Step 1: Explicit integration of the non-linear terms Starting from the solution at time $t^n = n\Delta t$, that is u^n, p^n, Φ^n, the non-linear advection terms are integrated using a stiffly stable integration scheme

$$\frac{\tilde{\boldsymbol{u}} - \sum_{q=0}^{J-1} \alpha_q \boldsymbol{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J-1} \beta_q [-\boldsymbol{u} \cdot \nabla \boldsymbol{u}]^{n-q} \tag{5.10}$$

where α_q and β_q are the coefficients of the stiffly-stable scheme. It is worth noting here that Φ is allowed to change during the simulation hence the temporal discretisation for the concentration function.

Step 2: Poisson equation for pressure The intermediate pressure p^{*n+1} is then obtained by solving a Poisson equation using the intermediate velocity $\tilde{\mathbf{u}}$

$$\nabla^2 p^{*n+1} = \nabla \cdot \left(\frac{\tilde{\boldsymbol{u}}}{\Delta t} \right) \tag{5.11}$$

with the following boundary condition on all velocity Dirichlet boundaries

$$\frac{\partial p^{*n+1}}{\partial \boldsymbol{n}} = -\sum_{q=0}^{J-1} \beta_q [\boldsymbol{u} \cdot \nabla \boldsymbol{u} + \nu(\nabla \times \nabla \times \boldsymbol{u})]^{n-q} \cdot \boldsymbol{n} \tag{5.12}$$

where $\boldsymbol{n}$ is the normal unit vector pointing outside the boundary.

Step 3: Integration of viscous terms The intermediate pressure field computed at the previous step is used to update velocity. The new velocity field $\hat{\boldsymbol{u}}$ is defined by

$$\frac{\hat{\boldsymbol{u}} - \tilde{\boldsymbol{u}}}{\Delta t} = -\nabla p^{*n+1}. \tag{5.13}$$

Viscous terms are then implicitly introduced to obtain the corrected velocity $\boldsymbol{u}^*$

$$\left(-\nabla^2 + \frac{\gamma_0}{\nu \Delta t} \right) \boldsymbol{u}^* = \frac{\hat{\boldsymbol{u}}}{\nu \Delta t}. \tag{5.14}$$

The SVV can also be added at this stage by addition of the SVV operator ν_{SVV}, already introduced in section 2.4.1, to the molecular viscosity term ν. Up to this point, the three first steps are sensibly the same as the VCS. New terms arising from the application of the SPM are now going to be introduced.

Step 4: Solve the SPM pressure The body force $\mathbf{f}_s$ representing the fluid/solid interaction is an inertial force proportional to the difference between the velocity of the solid boundary $\boldsymbol{u_p}$ and the velocity predicted by the VCS $\boldsymbol{u}^*$. If the solid body is moving with

the velocity $\mathbf{V}_s$, the immersed body velocity field can be written

$$\boldsymbol{u_p} = \Phi \mathbf{V}_s \tag{5.15}$$

and the body force

$$\mathbf{f}_s^{n+1} = \frac{\gamma_0 \Phi^{n+1}(\boldsymbol{u_p}^{n+1} - \boldsymbol{u}^*)}{\Delta t}. \tag{5.16}$$

Far away from the immersed body, the concentration function Φ is equal to zero and the body force is also zero. The force $\mathbf{f}_s$ is then expressed in the form of a pressure gradient $\nabla p_p = \mathbf{f}_s$. The extra pressure field p_p arising from the impermeability of the solid body is determined by solving the following Poisson equation

$$\nabla^2 p_p^{n+1} = \nabla \cdot \left[\frac{\gamma_0 \Phi^{n+1}(\boldsymbol{u_p}^{n+1} - \boldsymbol{u}^*)}{\Delta t} \right] \tag{5.17}$$

with the following boundary condition on Dirichlet velocity boundaries

$$\frac{\partial p_p^{n+1}}{\partial \boldsymbol{n}} = \frac{\gamma_0 \Phi^{n+1}(\boldsymbol{u_p}^{n+1} - \boldsymbol{u}^*)}{\Delta t} \cdot \boldsymbol{n}. \tag{5.18}$$

The total pressure field taking into account both hydrodynamic pressure and fluid/solid interaction is then

$$p^{n+1} = p^{*n+1} + p_p^{n+1}. \tag{5.19}$$

Step 5: Velocity update The velocity $\boldsymbol{u}^*$ is finally updated using the additional pressure field p_p induced by the solid body interaction with the fluid

$$\frac{\gamma_0(\boldsymbol{u}^{n+1} - \boldsymbol{u}^*)}{\Delta t} = \frac{\gamma_0 \Phi^{n+1}(\boldsymbol{u}_p^{n+1} - \boldsymbol{u}^*)}{\Delta t} - \nabla p_p^{n+1}. \tag{5.20}$$

Setting Φ to one shows that the velocity may be slightly different from u_p inside the solid. This effect can be avoided by scaling the additional pressure gradient ∇p_p^{n+1} by $1 - \Phi^{n+1}$. Equation 5.20 then becomes

$$\frac{\gamma_0(\boldsymbol{u}^{n+1} - \boldsymbol{u}^*)}{\Delta t} = \frac{\gamma_0 \Phi^{n+1}(\boldsymbol{u}_p^{n+1} - \boldsymbol{u}^*)}{\Delta t} - (1 - \Phi^{n+1})\nabla p_p^{n+1} \tag{5.21}$$

which cancels the contribution of the pressure gradient on the velocity field within the solid body. Yet, the incompressibility condition is not any more ensured with the last modification. Summing up all the steps described above results in the following total equation

$$\frac{\boldsymbol{u}^{n+1} - \sum_{q=0}^{J-1} \alpha_q \boldsymbol{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J-1} \beta_q [-\boldsymbol{u} \cdot \nabla \boldsymbol{u}]^{n-q} - \nabla(p^* + p_p) + \nu \nabla^2 \boldsymbol{u}^{n+1}. \tag{5.22}$$

No-slip and no-penetration conditions have been verified by Luo et al. [119]. The velocity field in the SPM is a linear combination of the surrounding fluid velocity $\boldsymbol{u}_f$ and the solid velocity $\boldsymbol{u}_p$ as follows

$$\boldsymbol{u} = \Phi \boldsymbol{u}_p + (1 - \Phi)\boldsymbol{u}_f. \tag{5.23}$$

The no-slip condition can be derived by applying the curl operator to equation 5.23

$$\nabla \times \boldsymbol{u} = \nabla\Phi \times (\boldsymbol{u}_p - \boldsymbol{u}_f) + \Phi \nabla \times \boldsymbol{u}_p + (1 - \Phi)\nabla \times \boldsymbol{u}_f \tag{5.24}$$

Given that $\nabla\Phi$ is normal to the body surface, no-slip condition can be expressed as

$$\nabla\Phi \times (\boldsymbol{u}_p - \boldsymbol{u}_f) = 0. \tag{5.25}$$

Rewriting equation 5.24 taking into account this condition, one obtains

$$\nabla \times \boldsymbol{u} = \Phi \nabla \times \boldsymbol{u}_p + (1 - \Phi)\nabla \times \boldsymbol{u}_f \tag{5.26}$$

which has been numerically validated by the authors on an example. On the other hand, the no-penetration condition can be derived by taking the divergence of equation 5.23

$$\nabla \cdot \boldsymbol{u} = \nabla\Phi \cdot (\boldsymbol{u}_p - \boldsymbol{u}_f) \tag{5.27}$$

since the surrounding flow field is supposed incompressible. Recalling that $\nabla\Phi$ is normal to the solid surface, the no-penetration condition is fulfilled if $\nabla \cdot \boldsymbol{u} = 0$ which is the incompressibility condition of the total flow field. In other words, incompressibility of the velocity field $\boldsymbol{u}$ automatically ensures that the no-penetration condition is validated.

5.1.3.3 Previous applications

The SPM has first been applied on canonical problems for validation purposes. Nakayama et al. [117] computed the drag coefficient of a two-dimensional circular cylinder at low Reynolds number, $Re = 1$ and found agreement within 5% compared to theory imposing a scaling parameter ξ of the same size as the mesh characteristic length. They then extended it to successfully simulate a two-dimensional sedimentation process using cylindrical particles at a Reynolds number of 0.1.

Luo et al. [119] have experimented the SPM on more complex particle shapes and have obtained very good agreement with DNS results at Reynolds number values up to 200. They have also pointed out that a mesh resolution comparable to the scaling ξ has to be imposed to correctly resolve the flow field close to immersed particles. A non-monotonic behaviour of the numerical error of the SPM scheme has been highlighted, on the one hand due to the more accurate modelling of the additional body force $\mathbf{f}_s$ which depends on Δt in equation 5.16 and on the other hand the thickening of the Stokes layer near particles scaling as $\sqrt{\nu\Delta t}$ and becoming severely under-resolved when Δt becomes too small. The optimal balance happens when the *effective smooth interface* l_e which represents the thickness over which the

SPM boundary thickness is deemed to have a strong impact balances the Stokes layer size. The optimal numerical error can be therefore found for a particular value of $\frac{\sqrt{\nu\Delta t}}{\xi}$ and varies depending on the considered case and temporal advancement scheme used.

Wang et al. [120] implemented the SPM in Nektar++ to perform a simulation of a simplified representation of an oscillating riser with buoyancy modules consisting in a dual-step cylinder at a Reynolds number of 1,000 based on the small cylinder diameter. Results obtained are very close to experimental measurements and simulations have enabled the authors to obtain a good characterisation of the turbulent wake and vortex shedding dynamics. The method has first been assessed on a stationary sphere and a cylinder, both stationary and oscillating perpendicularly to the incoming flow in two and three dimensions at Reynolds number values up to 10,000. Good agreement has been observed with DNS simulations. In particular, a correlation for the choice of ξ depending on the Reynolds number value only and not on the time step as proposed by Luo et al. [119] has been empirically deduced and validated. It relates the scaling parameter ξ to the minimum momentum thickness on a sphere or a cylinder δ_1 estimated by Schlichting and Gersten [121]

$$\xi = 0.8 \times \frac{0.664}{\sqrt{\pi Re}} \tag{5.28}$$

in three-dimensional cases. Once ξ has been calculated, the mesh resolution can be deduced. The authors claim that one quadrature point is required in the solid/fluid transition region to ensure simulation stability. It is worth mentioning that the empirical estimations have been numerically verified for Reynolds number values below 10,000.

In a second study, Luo et al. [122] coupled the SPM to URANS, LES and iLES in order to perform high-fidelity simulation of a waterjet modelling the interaction between rotor and stator blades at higher Reynolds number. The methodology has first been validated on a NACA0012 airfoil at different angles of attack at Reynolds number values ranging from 4,000

to 20,000 and has correlated very well with DNS and experimental results. In the waterjet simulation, the rotation of stator and rotor blades has been modelled with the SPM at a Reynolds number of more than 230,000 with URANS simulation and nearly 58,000 with LES and iLES. The simulation successfully allowed to capture turbulent structures in the water jet with a very simple mesh compared to the traditional body-conforming mesh approach. Yet, the authors warn that the SPM comes at high computational costs when applied to industrial configurations at high Reynolds number and that further work is required to validate the approach.

Cassinelli et al. [123] have very recently implemented the SPM within Nektar++ to perform 2.5D simulations, which consist in two-dimensional simulations extended in the third direction using Fourier expansions, of a low pressure turbine cascade at Reynolds numbers up to 297,000 based on the the chord of the blades. Yet, the equivalent Reynolds number for the upstream cylinders modelled with the SPM is lower by an order of magnitude. The SPM is used to generate periodic wakes upstream from turbine blades using cylinders translating normally to the incoming flow direction in order to replicate the rotor/stator interaction. The interface thickness ξ has been fixed and the diameter of cylinders varied and best agreement with DNS results has been obtained when the diameter with the SPM is 60% smaller than for the original cylinder. SPM simulations results are in very good agreement with experimental data, within 1% for the wake velocity profile peak and less than 5% difference for the pressure loss prediction.

Mohaghegh and Udaykumar [124] compared the SPM to a sharp interface method (SIM) using a second-order accurate finite volume method on a sphere, both stationary and translating, at low to moderate Reynolds number values up to a few hundreds and proposed the following correlation for ξ in three-dimensional cases

$$\xi = 6\Delta x(0.20 + 1.7Re^{-0.4})(10CFL)^{(0.65+0.1/Re)}Re^{-0.11}. \tag{5.29}$$

The form looks however less applicable than the one proposed by Wang et al. [120] since it also depends on the time-step. Moreover, the authors noticed a lower influence of the time step at higher Reynolds number values. Numerical results indicate that although SPM, which requires to solve two Poisson equations to update pressure, is in average 33% more expensive in terms of computations, the SIM requires higher mesh resolution to reach the same level of accuracy. Yet, SIM generally provides better results on fine grids. The authors also emphasise that the clear distinction of the fluid and solid domains in SIM produces spurious pressure oscillations, due to the transformation of solid mesh nodes into fluid mesh nodes when the immersed body is moving. This behaviour can be severely constraining in the low numerical dissipation framework of the spectral/hp element method. A study at higher Reynolds number values up to 10,000 has been conducted by Mohaghegh and Udaykumar [125] on a cylinder and spherical and elliptical particles at rest and moving. It has been pointed out that although SPM performs well on coarse grids even at low Reynolds numbers compared to SIM, finer grids are required at higher Reynolds number values to capture flow features near the surface of the solid body which are smoothed by the SPM. Yet, forces, flow separation and wake topology are well predicted by the method.

5.1.4 Objectives

The literature review has shown that accurate modelling of wheel rotation is essential to precisely capture the wake structure past wheels. MW turns out to be a suitable solution to model slick tyres but it fails at representing non-rotationally invariant parts such as rim spokes. Although MRF alleviates this shortcoming in the case of stationary simulations, it introduces unrealistic pressure gradients due to the added source term and omits actual wheel rotation. SM is the most accurate technique but implementation is arduous since it requires complex mesh transformation implying non-conforming interfaces, which is prob-

lematic when high-order curved meshes are employed. Moreover, parallelisation represents a challenge since computational domains enclosed in the rotating area constantly change neighbouring domains in the stationary zone during the simulation, which makes communication between processors complex. The SPM constitutes a potential alternative to rotate rims realistically without any mesh transformation since solid bodies are implicitly represented by a body force in the fluid domain. The main goal of the present study is to evaluate whether the SPM can accurately model rim rotation in high-fidelity simulations. The smooth interface between solid and fluid areas is particularly well adapted to the spectral/hp element framework as pressure oscillations that would arise from a sharp boundary are avoided and the very smooth transition function preserves exponential convergence properties. Still, the method has principally been used to simulate particles with simple shapes immersed in fluids at low to moderate Reynolds number values. The SPM numerical scheme has already been implemented within Nektar++ and used in 2.5D turbomachinery applications at relatively high Reynolds numbers but limited to a simplified cylindrical stator geometry. The first step of the study is to extend the concentration field generation process to arbitrary 3D shapes. Simulations with Nektar++ of industrial geometries are performed on unstructured meshes. Therefore, analytically defining the evolution of the concentration field for complex immersed body shapes which cannot be defined with a closed equation form is impossible, hence the need of a motion reconstruction process which is built in a second step. Finally, the resulting extended SPM strategy is applied on an isolated wheel case to evaluate its performance and impact on the prediction of the flow, especially wake structures and compared to other standard wheel rotation modelling techniques.

5.2 Concentration field generation

The concentration field is an essential part of the SPM as it defines the immersed solid body. The most widely used function is defined in equation 5.7. In the case of simple shapes which can be defined analytically in a closed form such as cylinders or sphere, the generation of the concentration field is straightforward once the scaling factor ξ is defined. Yet, when it comes to more complex shapes which are defined with CAD files, the concentration function must be built numerically. Considering a point of the mesh, three inputs are required to compute the value of the concentration function at this location: the scaling factor ξ, the information whether the point is located inside or outside the solid body and the distance to the boundary surface. The first input ξ is directly specified by the user whereas the signed distance has to be calculated. Meshes employed for wall resolving simulations of industrial geometries usually count several million of degrees of freedom. Therefore, one important requirement of the concentration function generation algorithm is the capacity to be run in parallel on many CPU cores.

5.2.1 Distance from a mesh point to the immersed body surface

The first step is the computation of the distance between the closed surface of the immersed body and an arbitrary mesh point. The format used for the definition of the immersed body is an STL file which consists in a tessellation of the boundary surface where each triangle is defined by its three vertices and the normal vector pointing outwards. Starting from a mesh point the problem then comes down to computing and selecting the smallest distance between the mesh point and triangles from the STL definition. The set of triangles to consider can be reduced to a subset of neighbouring triangles only to save computation time. The tessellation is therefore parsed using an octree data structure type in order to efficiently

isolate triangles located in the direct neighbourhood of the mesh point. Once the subset of triangles has been identified, the distance to each of the triangles of the subset is computed. Möller and Trumbore [126] developed an algorithm to check whether a ray traced from the mesh point aligned with the direction of the normal vector of a triangle hits the triangle without having to store the parametric equation of the triangle plane, which results in a lower memory demand. The algorithm also computes the distance to the triangle. There are actually two possible cases. If the ray hits the triangle, then the shortest distance is given by the orthogonal projection on the concerned triangle. Else, the shortest distance is subsequently to one of the edges or vertices of the triangle. This algorithm has previously been implemented in the concentration function generation method of Nektar++.

Once the shortest distance from the mesh point to the STL surface is calculated, the relative position of the point respectively to the closed surface needs to be determined to assign the correct sign to the distance d in equation 5.7 determining whether the concentration is closer to zero or one. The algorithm developed by Möller and Trumbore and implemented previously to this work in Nektar++ also determines from which side of the triangle the intersection occurs using the fact that normals are pointing outwards from the body. Yet, in practice, some miscalculations are observed resulting in points considered inside the solid body although they are located outside. This is illustrated on the left-hand side picture in figure 5.12 representing an isosurface of $\Phi = 0.5$ computed for a simplified rim geometry. Although the shape can be clearly distinguished and is well approximated, several faulty mesh elements outside of the rim with a positive Φ value are present.

Two reasons can be suggested to explain the inconsistencies. The first one can be found in O'Sullivan's master's thesis we have supervised [127] and is related to the STL tessellation. Indeed, triangles are added in the octree depending on the position of their centroid. When the tessellation is not uniform and highly skewed triangles with a needle shape are present, the centroids may be far from the vertices and they may therefore be absent of the list of

neighbouring triangles although this is the closest triangle to the mesh point. The closest triangle from the incomplete list may then give wrong information about the position of the mesh point relatively to the immersed body. The situation is likely to happen on corners where two adjacent triangles can have completely different normal vectors. This is coherent with the fact that many of the flawed mesh points are in the same plane as the side surface of the rim and therefore in the same plane as corners. The user thus needs to be particularly cautious about the quality of the immersed body tessellation.

The second reason has a numerical origin and has been highlighted by Woop et al. [128]. The way an edge shared by two triangles is parsed by the ray intersection algorithm can be different depending on the considered triangle. The test condition for each edge of a triangle in Möller's and Trumbore's algorithm is different depending on the order of the edge in the execution of the algorithm. Indeed, the last edge test is different that the two firsts due to the use of barycentric coordinates which makes the algorithm cost effective. Yet, this also implies that the same edge can be treated in two different ways depending on the triangle at stake. The different treatments can lead to different results due to the limits in terms of numerical precision of the computations. This is exacerbated when large and small values are mixed, for instance when the coordinates of a small triangle are large when it is far from the origin of the domain. Faulty mesh points result in spikes in the concentration field which have been noticed to deteriorate simulation stability. In order to resolve inconsistencies, a second test has been imposed to check the sign of the distance.

5.2.2 Second condition to determine the location of the mesh point relative to the immersed body

A second condition has been imposed in order to alleviate issues explained in the previous section. Starting from a mesh point, a ray is traced between the mesh point and the STL surface and the number of times the ray intersects the surface is counted. If the number of

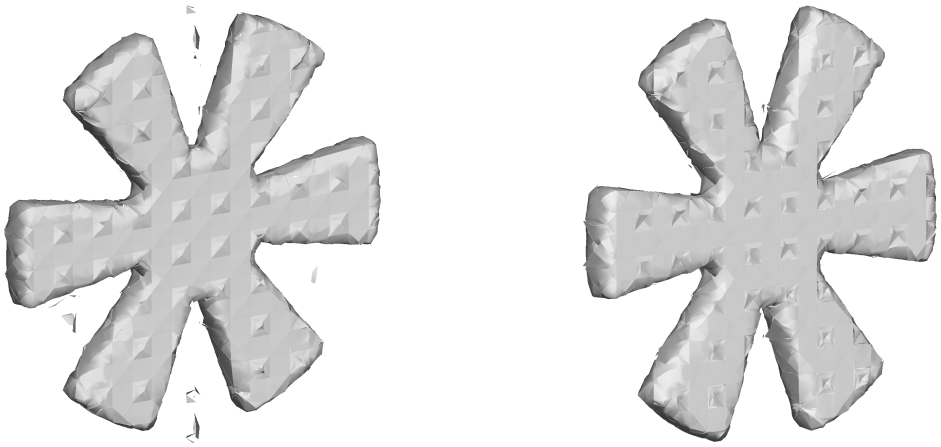


Figure 5.12: Isosurface of concentration function $\Phi = 0.5$ on a simplified rim geometry. Left: before ray hit count correction. Right: after ray hit count correction

intersections is odd, then the point is inside the closed surface and a negative sign is applied to the distance. Else, the point is considered outside and the distance is positive. The ray hit count algorithm has been implemented in the concentration field generation process of Nektar++ and has proved to greatly improve the fidelity of the concentration field as shown on the right-hand side image of figure 5.12. The algorithm does actually loop over the whole tessellation and errors arising from the restriction to neighbouring triangles are removed. The very few small inaccuracies are due to the computational precision problem already mentioned in the previous section, but they have far less impact and simulations are kept stable. Still, the algorithm is computationally less efficient since the whole tessellation needs to be tested but it can be alleviated by the fact that it can be run in parallel.

5.2.3 Mesh resolution in the transition interface

Concentration fields for arbitrary shapes can now be generated effectively but the necessary mesh resolution to accurately capture the transition remains to be defined. The literature review has shown that authors point out a relationship between the Reynolds number and

the interface thickness ξ. Yet, they do not agree on the form of this dependency, whether it depends on the time step and the range of validity in terms of Reynolds number. For moderate Reynolds number values however, proposed correlation indicate that the higher the Reynolds number, the smaller the optimal interface thickness. The large majority of correlations have been established on cylinder and sphere cases. The interface thickness is particularly important in the rim case since it controls the area of the opening between the spokes which impacts the amount of through flow and thus the wake of the wheel with the risk of artificially obstructing the wheel when ξ is too large. A clear example is shown on figure 5.13 where cross-velocity vectors are plotted for a ring case using the SPM at two different Reynolds numbers $Re = 225$ and $Re = 100,000$ with the same $\xi = 0.2D$ parameter, where D is the external diameter of the ring. It can be seen that for both cases, the zero to very low velocity area extends largely beyond the surface of the ring which is represented only for clarity but is absent from the domain in the actual simulation and represented as an immersed body. When it comes to the stream through the hole, a contraction is observed in the high Reynolds number case resulting in a smaller stream surface at the middle plane of the ring compared to the low Reynolds number case. This numerical experiment indicates that when the Reynolds number is increased, the immersed boundary has an effect at a broader distance from the surface, what is coherent with correlations proposed in the literature suggesting to use sharper interfaces when Reynolds number grows.

A sharper solid/fluid interface is accompanied by a rise in mesh resolution necessary to model the boundary accurately. Indeed, at least one interpolation point is required at the interface. Extrapolating the estimation of Wang et al. [120] at a Reynolds number about 100,000 which is the order of magnitude at which most of the isolated wheels simulations reviewed are run results in a scaling factor $\xi = 0.95$mm compared to the non-dimensional wheel diameter which is 1m. In first attempts, the whole rim has been modelled as an immersed body but required h-refinement is prohibitive since the whole rim needs to be

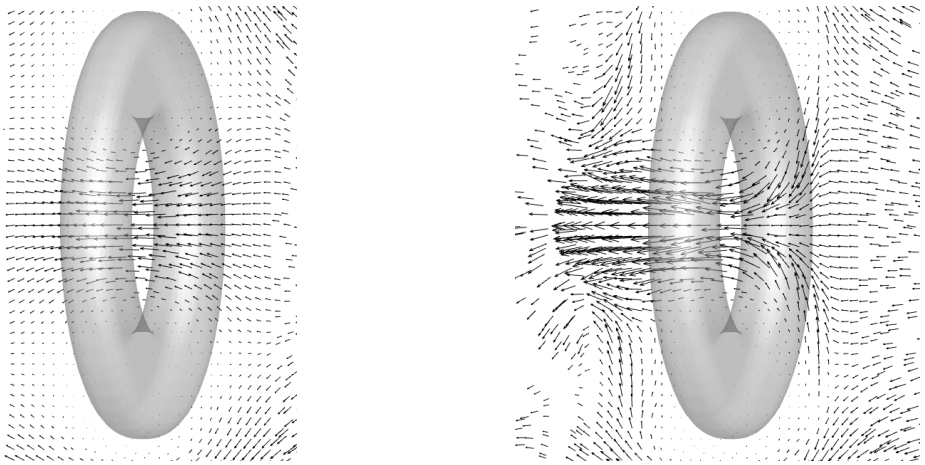


Figure 5.13: Cross-flow velocity vectors through a ring using the SPM (the solid surface is added only for clarity purposes). Left: $Re = 225$. Right: $Re = 100,000$. The interface thickness parameter ξ is kept the same in both simulations. A contraction of the flow going through the ring can be observed in the high Reynolds number case.

encapsulated in a cylindrical refinement box to anticipate physical rotation. Yet, the only non-rotation invariant parts are the spokes. Therefore, the rim shroud can be treated similarly as tyres, that is with the MW wall approach and the SPM limited to the spokes which vastly reduces the extent of the refinement box and consequently the total number of elements. An example of preliminary and simplified rim geometries are represented on figure 5.14. In the second design, the shroud is removed from the assembly reducing the extent of the fine resolution zone required for the concentration field generation. The SPM region should in general be kept as limited as possible to reduce the artificial mesh count increment coming from the concentration field generation.

5.3 Concentration field rotation

5.3.1 Strategy

In the previous section, a routine to create accurate concentration fields of complex geometries starting from a STL file has been implemented. In this second step, the motion of the

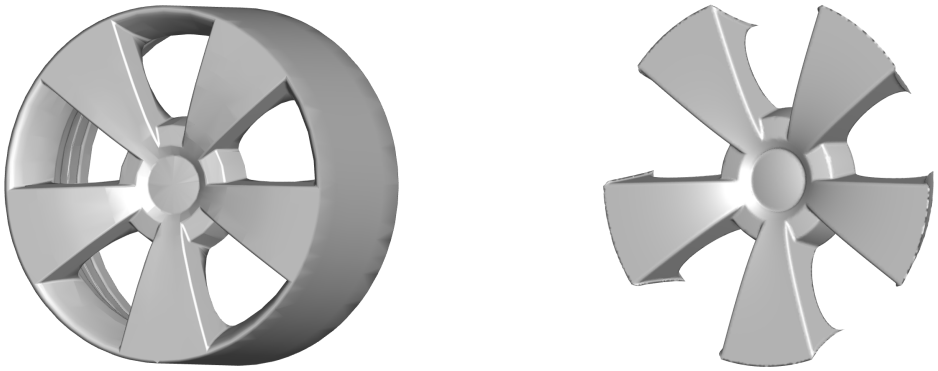


Figure 5.14: DrivAer wheel extracted from the geometry designed by Heft [110]. Left: complete rim. Right: isolated spokes.

concentration field is imposed. In the case of simple geometries with a motion which both can be defined analytically in closed form, the temporal evolution of the concentration field can relatively simply be predicted. For instance, in the work of Cassinelli et al. [123], the position of the cylinders representing the stator move at constant velocity in the vertical direction, thus the specification of their position during the simulation is rather straightforward. When it comes to complex shapes, the concentration field cannot be defined analytically and therefore, no simple equation can be derived for the motion defining the concentration field at each time step. This work focuses on rim rotation which is a periodic motion. The idea which is proposed to mimic the actual rotation is to generate the concentration field representing the rim at predefined regular angular positions and to interpolate between the samples at each time step to reconstruct the actual motion. In addition to the spatial resolution requirement to correctly approximate the smooth transition interface, a temporal resolution estimation, that is the number of required samples to smoothly reconstruct the motion of the rim is necessary.

This evaluation is conducted on the wheel of the DrivAer geometry [110] which is generic, open-source and representative enough to evaluate the impact of the rotation modelling techniques. It is represented on figure 5.14. The wheel geometry is split in two parts, the tyre and the rim. Literature shows that MW is well adapted to model rotating tyres, especially

when they are slick as is the case here. The rim is however problematic and is therefore extracted from the wheel geometry to be integrated as an immersed body. As discussed in the previous section, only the spokes of the rim are extracted and modelled with the SPM in order to limit element count since SPM requires rather fine mesh resolution to capture the fluid/solid interface.

5.3.2 Interpolation

The DrivAer rim is composed of five spokes resulting in an angular periodicity of 72^o. Several concentration fields are first generated for different wheel angular positions between 0^o and 72^o. In each element, the concentration field Φ created from the STL file is projected upon the polynomial basis and the interpolation between the different samples is operated on the resulting expansion coefficients.

Discrete Fourier interpolation appears as a relevant candidate to reconstruct the rim periodic motion due to its convergence properties. It has been implemented in the simulation process to compute the position of the spokes at each time step. Let us assume a set of N concentration field samples is created a regular angular intervals $\Delta\alpha$ resulting in a set of samples $\Phi = (\Phi_1, .., \Phi_N)$. In a preliminary step, a discrete Fourier transform of the expansion coefficients in each element for the samples in the set Φ is performed engendering a discrete basis of Fourier coefficients set $\widehat{\Phi} = (\widehat{\Phi_1}, .., \widehat{\Phi_N})$. Finally, at each time step of the simulation, the concentration field corresponding to the position of the rim is calculated using the Fourier set before flow field calculations and a real time update of the angular position of the rim is subsequently obtained. The result is shown on a vertical plane crossing the spokes using 24 samples, that is one sample every 3^o on the left-hand side of figure 5.15. The concentration field gives a good representation of the shape of the spokes on this plane. However, streaks can be seen in the radial direction caused by oscillations of the concentration field. The concentration field has been plotted using curvilinear coordinates along the black dashed circle

of figure 5.15 on one spoke to observe this phenomenon more precisely and results are shown on figure 5.16. The black line represents the exact STL definition, that is one within the solid and zero outside. Fourier interpolation has been carried out with three different sets of regularly spaced samples: 6 samples (one sample every 12^o), 12 samples (one sample every 6^o) and 24 samples (one sample every 3^o). The temporal resolution with 6 samples is far from the STL definition, below one in a large part of the solid domain and about 15% above in the rest of the solid domain and it does not recover the zero value in the fluid domain. Increasing the number of samples to 12 improves the quality of the approximation but large differences with the STL remain and large spatial oscillations of the concentration field appear. The interpolation using 24 samples is the most accurate but higher frequency modes are present in the reconstructed concentration field with relatively large amplitude of about 5 to 10% of the ideal STL concentration field. These oscillations are characteristic of Gibbs oscillations and appear when the function to interpolate is sharp, as is the case for the fluid/solid transition. Increasing the resolution further will limit the amplitude of oscillations in the central part of the fluid and solid domains but spikes around the transitions will grow. The oscillations and spikes directly impact the pressure field through equation 5.17 and provoke large instabilities in the simulation.

Linear interpolation naturally prevents from overshooting and oscillation. Given the relatively fine resolution of the samples, it has been assessed on the DrivAer rim case. For each time step, using the rotational speed of the rim, the two closest bounding samples are selected and a linear interpolation is operated in between. The result on the same plane as the Fourier interpolation is shown on the right-hand side of figure 5.15. The concentration field turns out to be a good approximation of the real rim shape and oscillation streaks have been removed. The result is plotted along the same circle on figure 5.16. Oscillations effectively disappear when linear interpolation is employed. It also turns out to be closer to

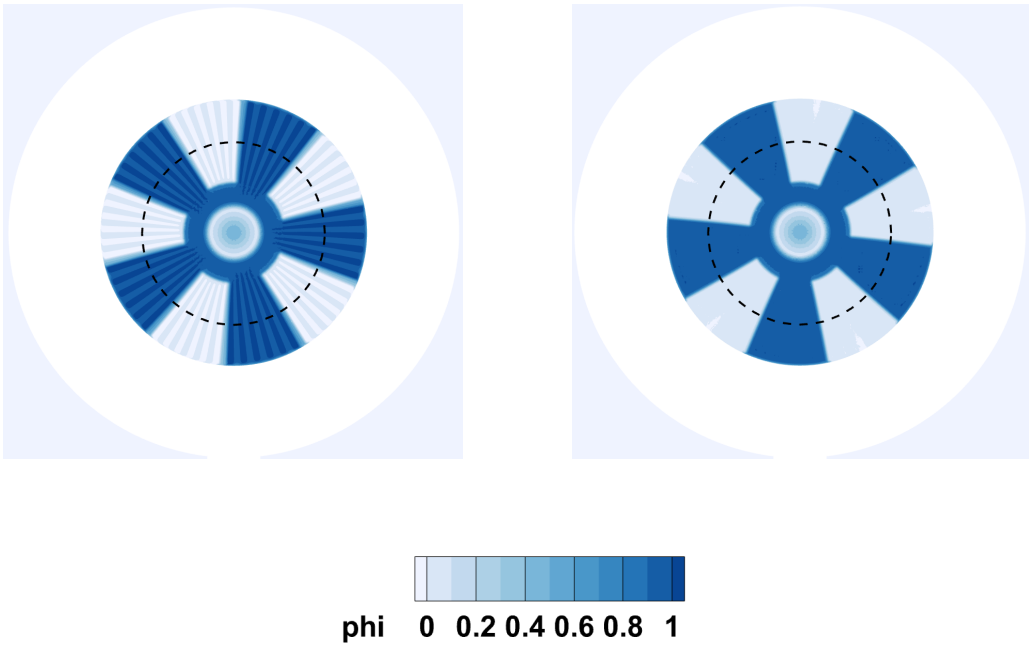


Figure 5.15: Concentration field contour on a vertical plane through the spokes using 24 samples. Left: with Fourier interpolation. Right: with linear interpolation.

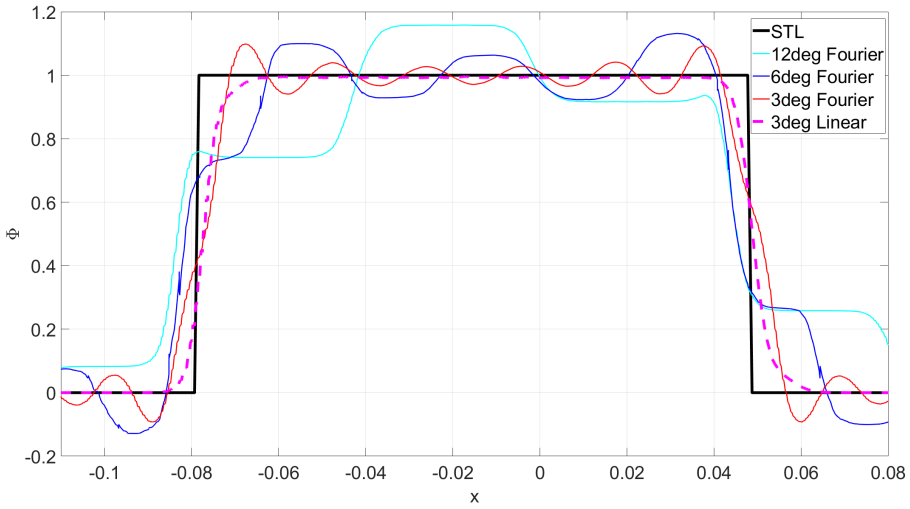


Figure 5.16: One-dimensional representation of the concentration field evolution on one spoke along the back dashed circle on figure 5.15 for different number of samples and interpolation methods.

the actual STL surface on the major part of the spokes and prevents from overshooting. It is therefore used in the following SPM simulations.

5.4 Isolated wheel simulation

In the previous section, a procedure to generate an evolving concentration field to model rotating rims with the SPM has been developed. The methodology is now applied to an isolated wheel geometry. It is intended to remove interferences which would arise from other parts of the car if a simulation of a full vehicle would have been performed. The SPM is also compared to the MW and the MRF methods which are the most widely spread techniques in the industry and a qualitative comparison against references present in the literature, both experimental and numerical to assess possible benefits.

5.4.1 Geometry and mesh

The geometry selected for this work is the DrivAer wheel [110], as mentioned in the previous section, which presents the benefits of being both open-source and close to industrial standards. The wheel is split into two distinct entities, the wheel and the rim spokes as can be seen on figure 5.17, what is useful for different treatments of the two parts. Indeed, although the slick tyre is addressed as a solid wall on which MW is applied, both the MW or the SPM may be used on the rim spokes. Therefore, two different solid geometries are employed in simulations. When rim spokes are modelled with MW, the solid geometry consists in the complete wheel whereas when the SPM is applied, only the tyre is physically present as a solid in the computational domain, rim spokes being represented implicitly.

The computational domain with the complete wheel is represented on figure 5.18. The dimension of the section, which is 3.65D x 3D where D is the wheel diameter, is inspired from previous works in the literature, particularly Fackrell's [6] which has been extensively used

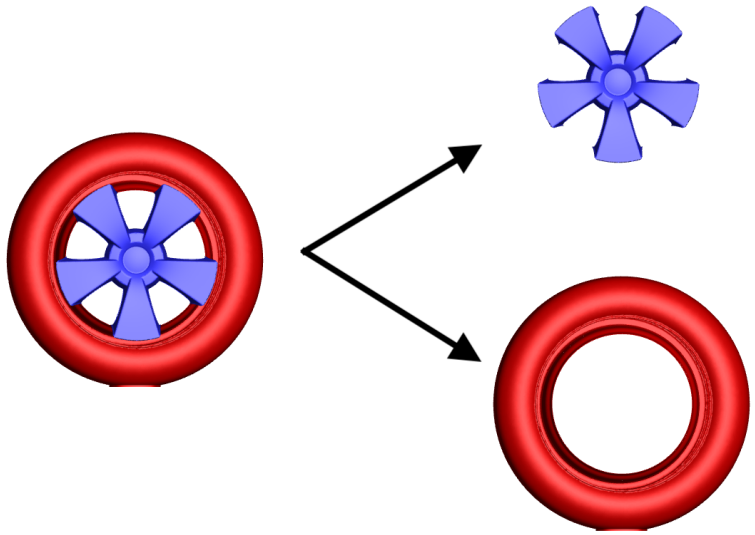


Figure 5.17: Decomposition of the DrivAer wheel geometry into two sub-parts: the tyre (in red) and the rim spokes (in blue).

in more recent numerical investigations. The length has however somewhat been reduced to 3D measured from the extreme front and rear of the wheel to the inlet and outlet surfaces respectively, especially behind the wheel taking advantage of the high-order outflow boundary condition developed by Dong et al. [56] allowing for shorter outlet distances, and since the focus is laid on the near wake. The wheel is centred in the tunnel, both in the streamwise and spanwise directions and the junction with the floor is made through a contact patch. In the complete wheel case, used in conjunction with MW or MRF, the rim counts five spoke equivalent to a 72^o angular periodicity. It is taken out of the geometry when the SPM is employed, as described on figure 5.17 and only the tyre is kept as a solid physical boundary.

Two meshes have been generated to take into account the two cases and a plane normal to the spanwise direction crossing the spokes is shown on figure 5.19. Three levels of refinements can be distinguished and are kept the same in both cases for consistency purposes. The far field is kept to a coarse level. A first level of refinement can be seen around the wheel and extends slightly further at the back of the wheel to increase the resolution in the near wake. The finest level of refinement can be observed in the rim region and is delimited

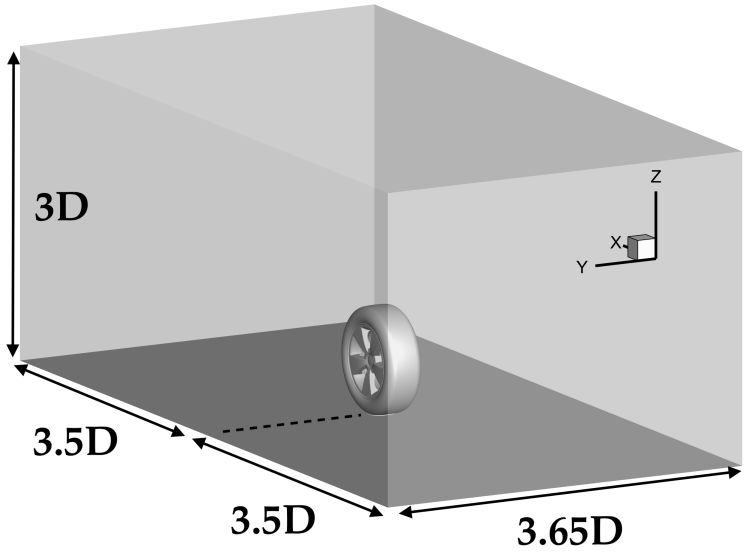


Figure 5.18: Computational domain used for the DrivAer isolated wheel simulation.

by a circular cylindrical box encapsulating the spokes. This refinement area is designed for the SPM. Indeed, correlations at low to moderate Reynolds number values emphasise the need for steeper transition when Reynolds number is increased what subsequently raises the recommended mesh resolution to capture the fluid/solid interface. As no established relationship exists between the Reynolds number and the interface thickness, the choice is empirically guided by the cross-flow through the interspace between the spokes. Indeed, if ξ is too large, the gap between the spokes is virtually filled by the immersed boundary and through-flow disappears. ξ has been set to 0.035 and the size of the elements in the region of the spokes adapted to have an interpolation point in the transition region. H-type refinement has been privileged in the study with second order p-refinement to alleviate differences arising from high-order discretisation since both high-order iLES and low-order finite volume RANS techniques are leveraged. A second order mesh curvature rule is used to generate high-order meshes and the prismatic layer has been split into 6 sub-layers with a geometrical growth rate factor of 1.5 between two successive layers. The near wall mesh can be observed more precisely on figure 5.20. The final mesh count is of 4.6 million elements for the complete

wheel case and of 5.1 million elements for the tyre case which is logically higher due to the incorporation of the spokes as part of the fluid domain. More information about the mesh can be seen in appendix section A.3.

Values of $y+$ for both linear meshes before p-refinement are below 25. The high values are due to high strain rate at the top of the wheel where the rotative boundary condition is opposed to the incoming flow. Yet, the impact of relatively high $y+$ values at this location do not have a very strong influence on the conclusion of this study, since the focus lays on the impact of rim rotation modelling techniques.

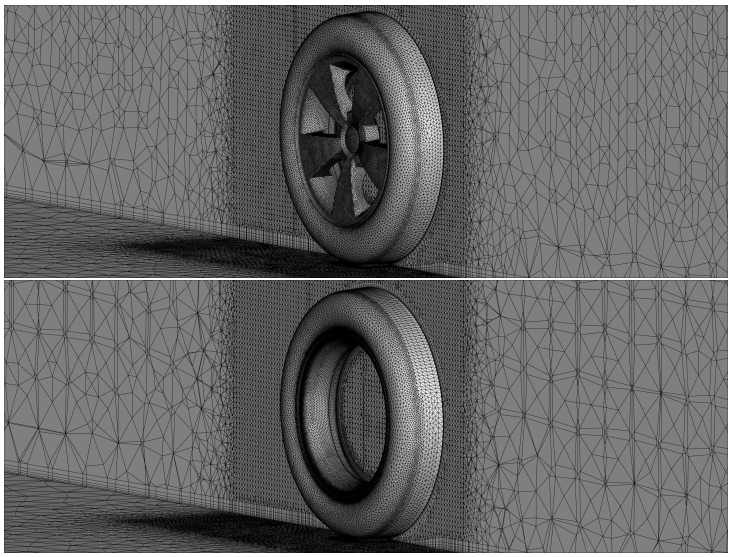


Figure 5.19: Mesh for the two modelling approaches for the isolated wheel. Top: rim physically present in the domain (MW and MRF). Bottom: rim removed from the wheel assembly (SPM case).

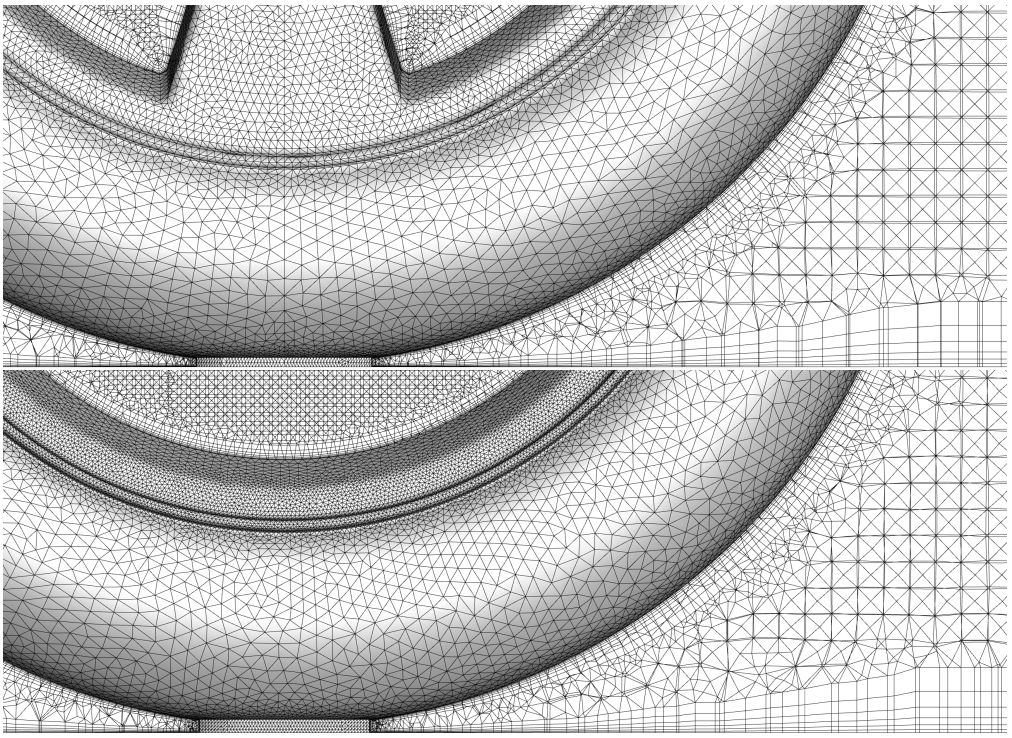


Figure 5.20: Near wall mesh structure. Top: MW and MRF. Bottom: SPM.

5.4.2 Numerical settings

The boundary conditions applied in the isolated wheel case are very similar to those used in the P14 case. Side walls and the ceiling are set as slip walls to reduce their influence and the floor moves at $U_\infty = 1$ in the positive X-direction, matching the velocity at the inlet. The high-order outflow boundary condition [56] is imposed at the outlet. MW is used for the tyre and set to match the floor velocity at the junction which is made through a contact patch. Three different boundary condition are compared for the rim: MW, MRF and SPM. For the last method, the rotative motion is reconstructed using 24 samples in the 72^o angular period of the DrivAer wheel and the period is set to match the angular velocity of the tyre.

Simulations are also performed with StarCCM+ which is the standard at McLaren automotive coupled with the MRF and used as a baseline for assessment of results obtained with Nektar++. Spectral/hp element method simulations are performed with a second order polynomial expansion for velocity components and linear approximation for pressure for two reasons. It first limits differences with low-order finite volume RANS simulations but also facilitates the imposition of a suitable mesh resolution to capture the transition interface, alleviating the coupling with polynomial order of the expansions.

The Reynolds number used for this study is $Re_\infty = 500{,}000$ based on the inflow velocity and the wheel diameter which is in the same range as previous simulations of isolated wheels presented in the literature. For spectral/hp element method based unsteady simulations, the non-dimensional time step is set to $\Delta t^* = \Delta t \times \frac{U_\infty}{D} = 2 \times 10^{-5}$ to keep the Courant number around one what is necessary to ensure stability of the partially explicit time-integration scheme of the incompressible Navier-Stokes solver of Nektar++. The SVV scaling parameter is also set to the lowest value possible guaranteeing the stability of simulations, which has been found to be 0.25 in this case.

RANS simulations are run until residuals decrease by four order of magnitude which is the minimum that could be reached. This phenomenon is well documented for this type of simu-

lations where the steady solution is actually unstable as pointed out by McManus and Zhang [60]. Simulations with Nektar++ are first run for four CTUs, where one CTU is defined by $\frac{U_\infty}{D}$, that is the time for the flow to pass the wheel four times in the streamwise direction and is required to pass the transitory regimes once initiated from uniform fields. Results are then averaged for 8 CTUs, corresponding to more than two complete wheel rotations. Simulations have been run on the CSD3 Tier2 HPC facility based at Cambridge using 2240 Intel Cascade Lake cores. When it comes to computation time, wall time to compute one CTU with MW is 4 hours and doubles when the SPM is employed, requiring 8 hours to compute one CTU. The large difference is principally explained by the second Poisson equation coming from the body force modelling the immersed body (see equation 5.17) which proves to be highly resource intensive since the time spent on the Poisson equation solver is multiplied by ten between the two cases. This can be allegedly attributed to the relatively large gradients caused by the steep transition profile Φ in equation 5.17.

First, cross-flow distribution through the spokes of the wheel at its initial position with the standard solid body and the SPM are compared to estimate the effect of the artificially augmented thickness of the spokes. Then, the close wake structure is investigated with the MW, MRF and SPM to evaluate the effect of each of these techniques on the flow topology downstream from the wheel and assess the ability of the SPM to effectively model wheel rotation. Finally, forces evolution is first investigated to exhibit unsteady behaviour. Cases are designated with two-part naming, the first part referring to the type of simulation (RANS/iLES) and the second to the wheel rotation modelling technique applied (MW/MRF/SPM). The SPM is subdivided in two methods: the steady SPM (SPMs) where the concentration field is fixed and the moving wall velocity $\omega \times r$ applied to the immersed body, and the unsteady SPM (SPMu) which is similar to SPMs but the concentration field is rotated.

5.4.3 Results

5.4.3.1 Cross-flow through the spokes

The SPM represents the boundaries of solid bodies with a smooth transition function Φ introduced in equation 5.7. The function virtually thickens the solid body since non-zero Φ-values extend beyond the actual immersed body, in the fluid domain. This deviation from the geometry can, in the limit when the thickness parameter ξ is too large, obstruct the space between the spokes and block any eventual through-hub flow. In this study, ξ has been set to 0.035 to allow the resolution of the solid/fluid transition with a mesh of a reasonable size. A first computation has been performed with the SPM on the original sample of the wheel fixed to its intial position at 0^o to compare it with the classical MW technique with Nektar++. Cross-flow contours are shown on the $Y = 0.11$ plane passing through the spokes on figure 5.21. Negative cross-velocity v values correspond to flow driven from the outboard side of the wheel towards the wheel hub through the spokes and positive values to flow exiting the wheel hub through the spokes. The through-flow distribution is very similar between the two cases with inflow at the left-hand side rear part of the rim which is especially strong at the top left and outflow from the right-hand side front part of the wheel. It is interesting noticing that a significant cross-flow component exists without any yaw-angle resulting from the front of the car on a shrouded wheel or through-hub flow forcing duct. A thin layer of zero-velocity flow surrounds the spokes in the iLES-SPMs case which is due to the smooth fluid/solid interface resulting in small deviations from the MW case. Weak inflow is present on the right-hand side of the rim next to the centre with the SPM whereas only low positive velocity values can be observed in the MW case. Slightly higher outflow can be noticed on the bottom of the left-hand part of the rim with the SPM than with MW. Still, velocity distribution remains globally consistent between the two cases indicating correct modelling of the through-hub flow with the SPM although a more satisfactory reproduction can probably be obtained with

a steeper concentration function at the cost of higher computational resources demand to resolve the interface.

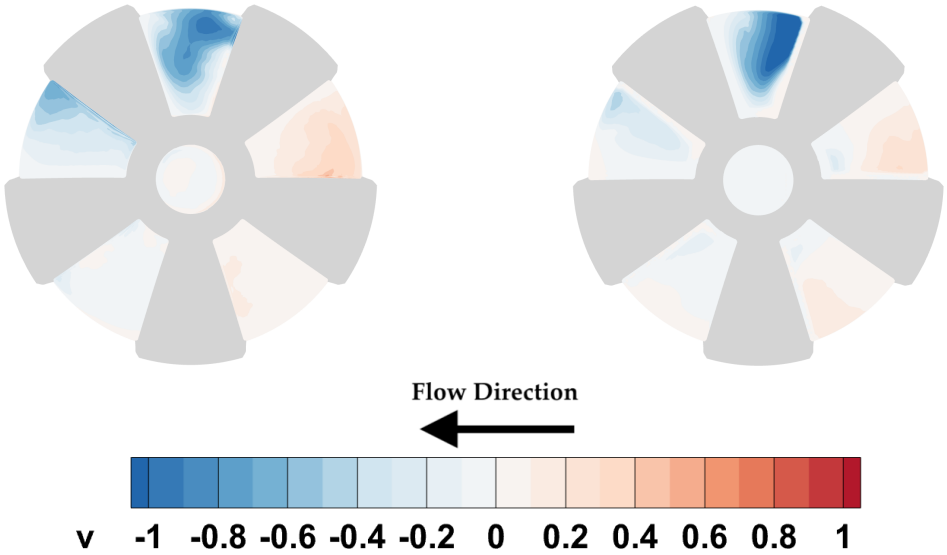
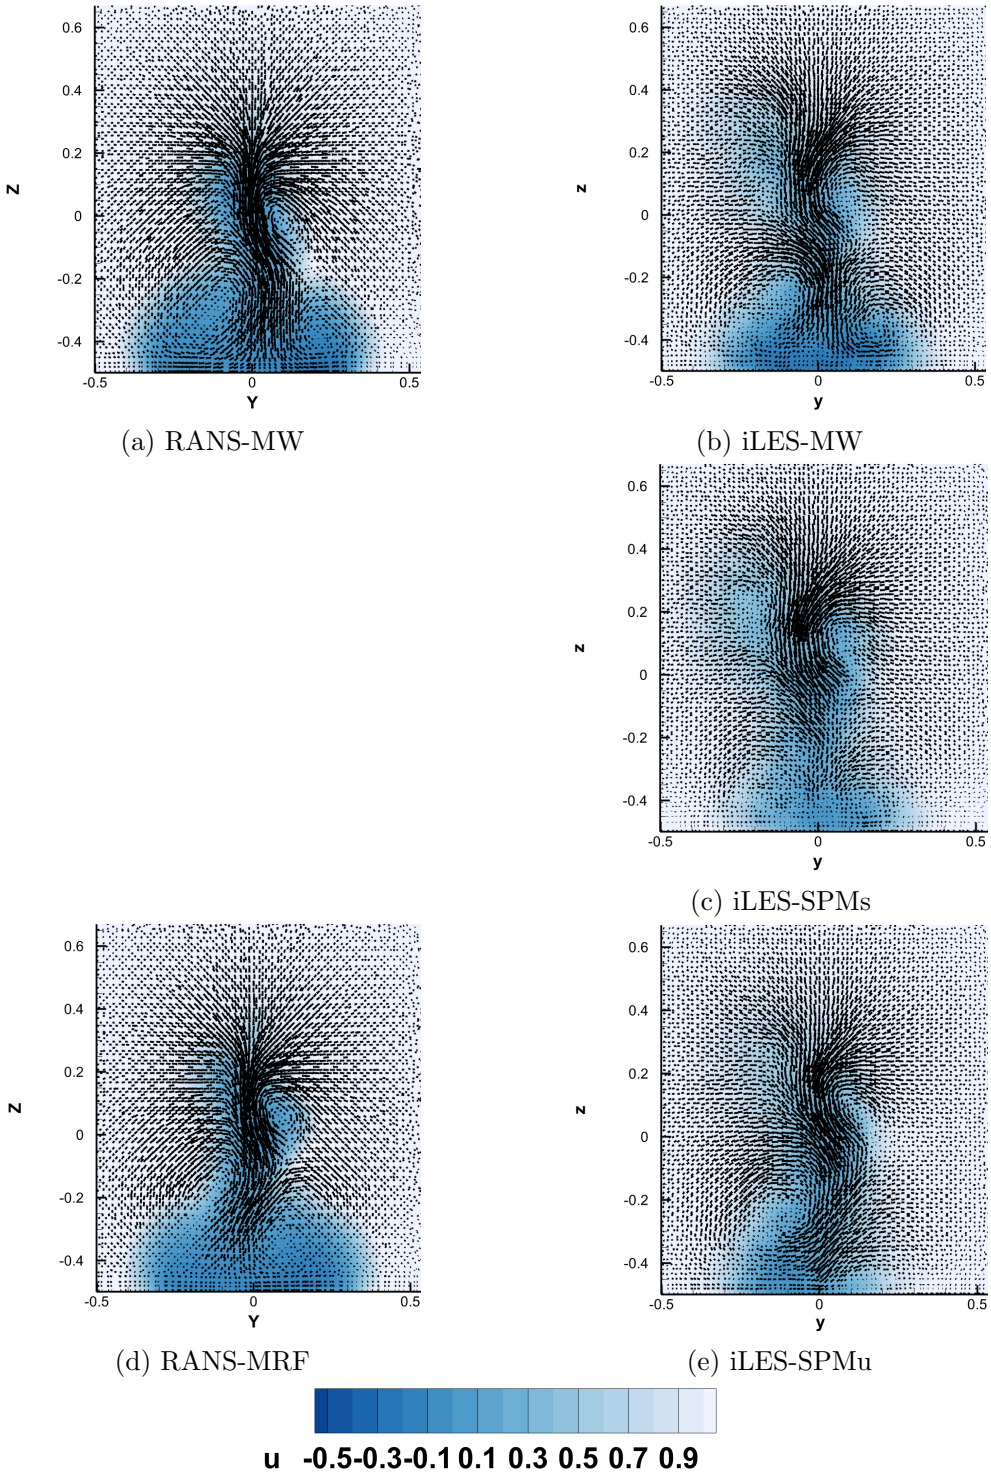


Figure 5.21: Cross-velocity v through the spokes of the DrivAer wheel on the $Y = 0.11$ plane. Left: iLES-MW. Right: iLES-SPMs.

5.4.3.2 Near wake

The previous section has shown that cross-flow distribution through the spokes with the SPM at the same fixed position is consistent with the usual MW solid boundary case. The concentration field is now rotated to mimic the spinning wheel realistic condition and assess the effect of different modelling techniques. The focus is laid on flow topology of the near-wake since it is a better representative than far-wake when the wheel is enclosed in the wheel well but also confines refinement to a region close to the wheel, alleviating computational costs. Streamwise velocity contours and in-plane velocity vectors 0.25D downstream from the wheel are represented on figure 5.22 for both RANS and iLES with the MW, MRF and SPM.



(a) RANS-MW

(b) iLES-MW

(c) iLES-SPMs

(d) RANS-MRF

(e) iLES-SPMu

Figure 5.22: Streamwise velocity and in-plane velocity vectors in the near wake of the DrivAer wheel at $X = 0.75$ with different modelling techniques.

Focusing first on the results with MW modelling on the two first pictures, the wake can be separated in two parts: the upper part for $Z > -0.2$ and the bottom part for $Z < -0.2$. The wake is also divided vertically by a downwash stream which is bent in both RANS and iLES cases even though the downwash stream is more curved in the iLES case. The wake is taller in the iLES case with both MW and SPM techniques, and is principally due to different separation point predictions by the two different turbulence treatments. Three vortical structures can be identified. The first one is located in the upper side of the wake at $Z = 0$ on the outboard side of the wheel where the spokes are located and is generated at the central part of the rim and forces the downwash stream to bend at this vertical location. The vortex is stronger in the RANS case and centred on $Z = 0$ whereas it is located slightly above in the iLES case but also larger, inducing a larger deviation of the downwash stream at is location. When it comes to the bottom part, downwash is more vertical in both cases separating two jetting vortices generated at the contact patch. The outboard side vortex is smaller in size and closer to the ground than its counterpart in both cases. The recirculation region is wider in the RANS case due to a stronger separation at the contact patch than in the iLES case where the flow is re-energised by the tyre motion due to different boundary condition implementations as has already been highlighted in the P14 case.

The iLES-SPMs is a replication of the MW case with the concentration field fixed at the same position as the solid rim in the MW case. Flow structures are consistent between the two cases with the outboard vortex in the upper part and the pair of bottom vortices in the lower part. The recirculation region looks similar on the upper part but thinner in the lower part, confirming that the upper part of the wake is mainly influenced by the separation on the tyre. The difference observed in the lower part originates from the inflow/outflow balance in the wheelhub. The higher blockage due to the SPM at the location of the spokes has an impact on the inboard side of the wheel. This is illustrated on figure 5.23 where the cross-flow velocity is represented along with in-plane velocity vectors on the $X = -0.2$ plane for both

the iLES-MW and iLES-SPMs cases. At this streamwise location, the deficit of outflow in the space between the spokes due to artificially increased thickness in the SPMs case results on an additional outflow component at $Z = 0$ on the inboard side. The deviation combined to the flow pushed out of the wheel generates a vortex on the outboard side at $Z = -0.1$ which is then convected downstream creating a strong downwash component on the bottom part of the rim which is already visible on the $X = -0.2$ plane. Downwash reenergises the flow and weakens the separation on the bottom part of the tyre and the contact patch, leading to a smaller recirculation region on the outboard part of the wheel. On the outboard side, thicker spokes reduce the amount of jetting outflow deviating the incoming flow on the bottom part of the wheel and therefore help the flow to remain attached to the wheel, shrinking separation on this side as well.

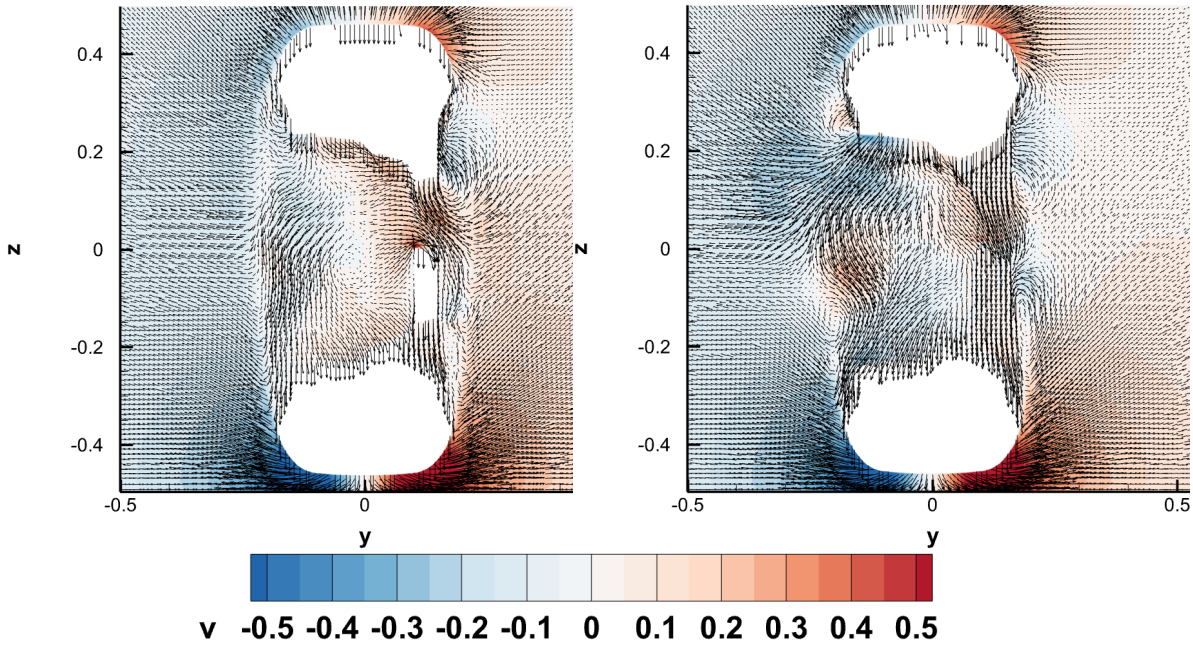


Figure 5.23: Cross velocity and in-plane velocity vector on the plane $X = -0.2$. Left: iLES-MW. Right: iLES-SPMs.

The two last cases are the RANS-MRF, which is used by MLA to model rotating wheels, and the iLES-SPMu case where the concentration field is physically rotating. No significant

changes can be noticed on the upper part of the wake indicating once again that this part of the wake is mainly driven by the flow on the tyre. Important changes can however be observed on the lower part of the wake. The vertical downwash component on the centreline leading to a rather symmetrical position of the jetting vortices on each side of the wheel with the MW is tilted in both MRF and SPM cases from the outboard side towards the inboard side of the wheel, with a larger slant in the SPMu case. In fact, in the SPM case, the outboard vortex is translated closer to the ground, reducing the size of the recirculation region on the right-hand side of the bottom part of the wake, whereas the inboard vortex is somewhat lifted causing the low velocity lobe to grow, which results in imbalances in the bottom part of the wake which extends further towards the inboard side of the wheel. In the MRF case, the bottom left-hand-side vortex also moves up compared to the MW case but the outboard vortex is also lifted and moved towards the centreplane what causes the corresponding recirculation region to grow.

The cross-velocity on the outboard side of the wheel next to the spokes for the iLES-SPMu is shown on figure 5.24. Compared to the iLES-SPMs case on the right-hand side of figure 5.21, velocities are more uniform with the inflow and outflow spikes which have disappeared. The inflow/outflow regions are however coherent between the two cases with the outwards jetting at the front and bottom of the rim and inflow at the top and rear part of the wheel. The physical rotation of the wheel in this case creates artificial blockage limiting cross-flow at peak regions due to the spokes regularly passing at this locations.

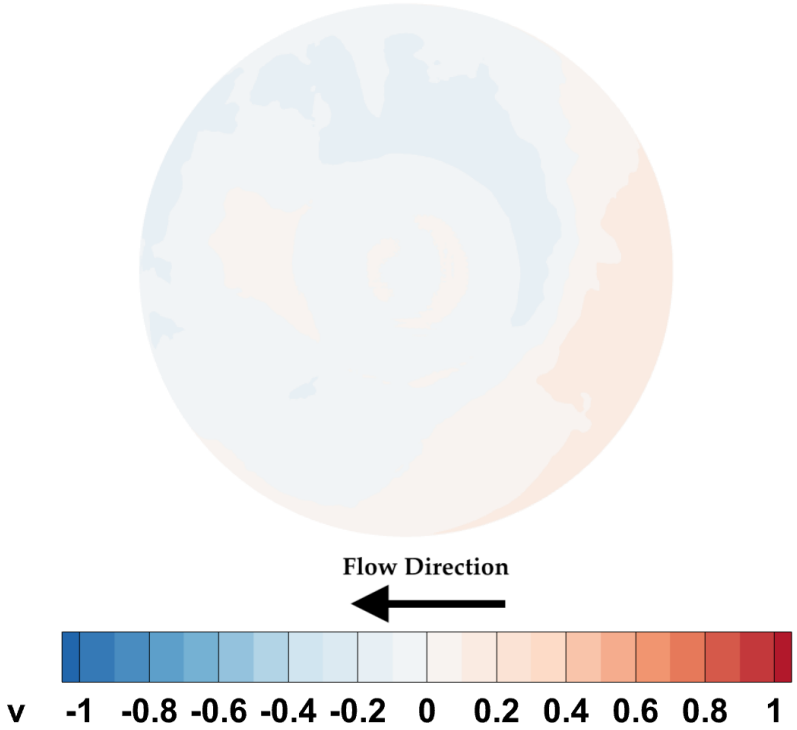


Figure 5.24: Time-averaged cross-velocity v in the spokes region of DrivAer wheel ($Y = 0.11$) in the iLES-SPMu case.

The cross-flow modification around the spokes modifies the inflow/outflow balance between the inboard and outboard sides which tangibly impacts the wake. This is illustrated at the location $X = 0.2$ on figure 5.25 where through-hub flow is compared in the iLES-SPMs and iLES-SPMu cases. The flow on the top of the tyre is very similar in both cases. When it comes to the wheel hub, substantial inflow can be observed at the top in the iLES-SPMs case corresponding to the peak already pointed out on figure 5.21. Inflow on the other side of the hub is then reduced as stronger positive cross-velocity which extends over the whole height of the hub is present in the iLES-SPMu. In particular, high v values at the bottom of the rim draws flow from the bottom part of the tyre towards the wheel hub resulting in a recirculation region on inboard side of the tyre. On the other hand, in the iLES-SPMs case, the flow is directed from the bottom part of the wheel hub towards the bottom part

of the tyre due to weaker inflow and recirculation is significantly shrunk at the bottom of the tyre. When it comes to the bottom outboard part of the rim, inflow can be noticed in the iLES-SPMs case whereas it is absent in the iLES-SPMu case once again due to regular passing of the spokes at this location. Deviation of the flow in the iLES-SPMs case causes small deceleration at the bottom part of the tyre whereas it remains strongly attached in the iLES-SPMu case. The differences between the inboard and outboard part of the wheel explain the wake topologies variations between the two cases emphasised on figure 5.22 where the wake is rather balanced between the inboard and outboard sides in the iLES-SPMs case and larger on the inboard side in the iLES-SPMu case with the strong cross-component of downwash in the bottom part of the wake from the outboard towards the inboard side of the wheel.

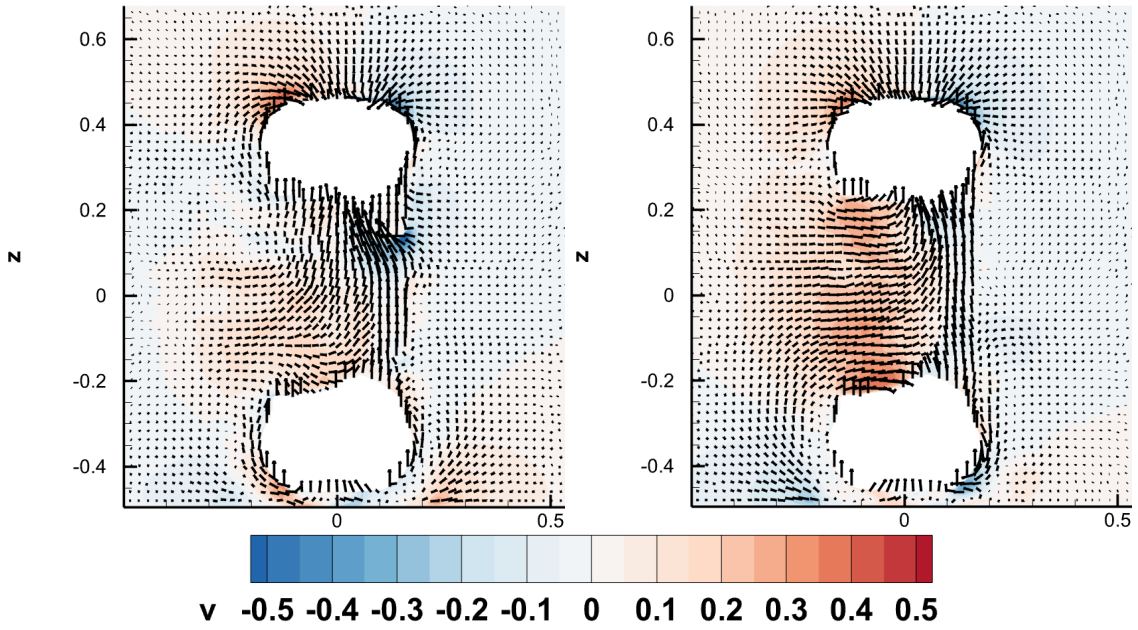


Figure 5.25: Cross velocity and in-plane velocity vector on the plane $X = 0.2$. Left: iLES-SPMs. Right: iLES-SPMu.

The evolutions of wake topology from the MW cases to the RANS-MRF or iLES-SPM are coherent and the obtained wake structure is in accordance with observations of Axerio-

Cilies at al. [104] on an isolated stationary Formula One wheel and later [103] on the same rotating wheel. In particular, the great similarity between the MW and stationary cases of the mentioned study emphasises the lack of momentum in the spokes region when only wall treatment is implemented.

5.4.3.3 Evolution of forces

A description of the time-averaged wake arrangement with the different wheel rotation modelling techniques has been provided in the previous section highlighting their respective impact. In this section, the time-evolution of the forces on the tyre is compared for a steady modelling technique, iLES-SPMs and its unsteady counterpart, the iLES-SPMu. The iLES-SPMs case has been chosen for the comparison rather than the iLES-MW case since the geometry of the solid tyre is the same as in the iLES-SPMu case due to the decomposition, allowing a fair comparison. The power spectral density for the drag and lift coefficients are shown in figure 5.26. In the SPMs case, the bulk of the energy is concentrated in low frequencies and both drag and lift forces turn out to be highly correlated with most of the peak frequencies matching for the two cases. The three first peaks are located at $St_D = 0.5$, the harmonic $St_D = 1.0$ and the third one located at $St_D = 2.0$ for C_D accompanied by a large drop in energy for C_L. Force coefficients turn out to be far less correlated in the SPMu case and a large share of the energy is still contained in higher frequency modes. The first peak for C_L appears for $St_D = 0.5$ before other peaks present at higher frequencies. C_D peaks are shifted by $0.25St_D$ compared to the steady modelling case. The important dissimilarities in power spectra suggest different lift and drag generation process when the more realistic unsteady rim rotation modelling is employed.

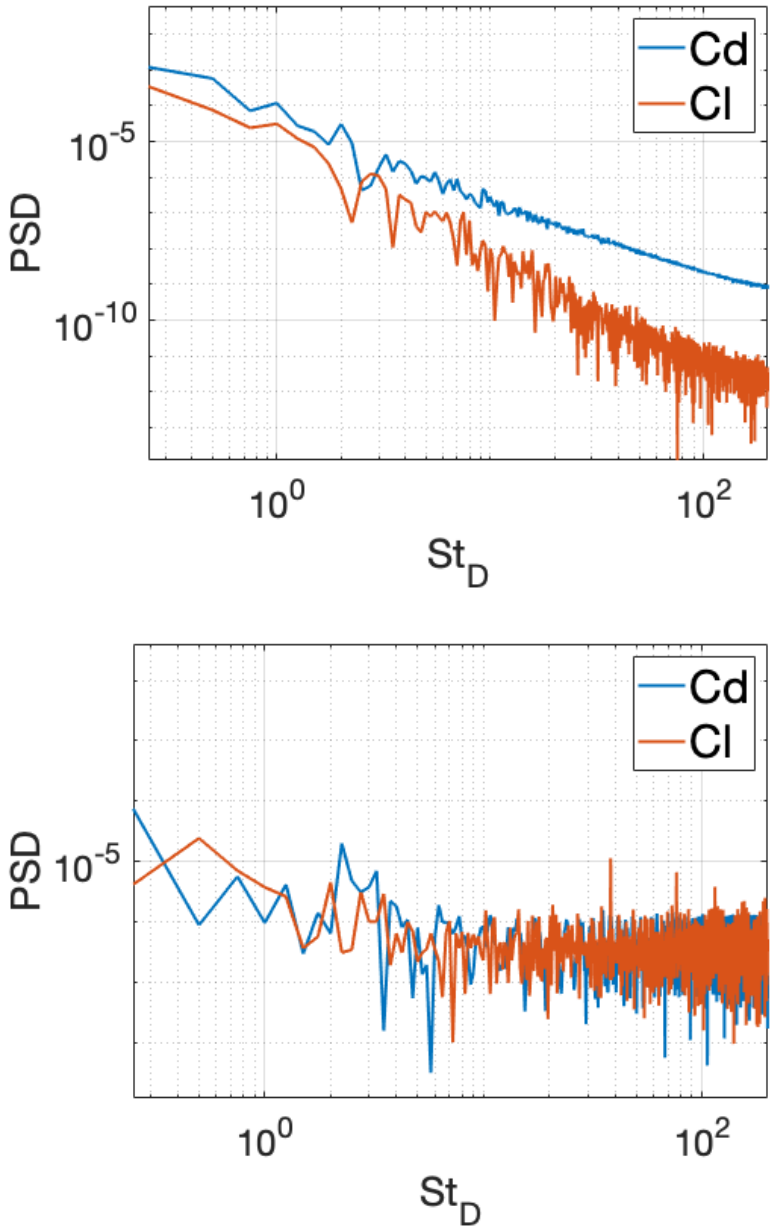
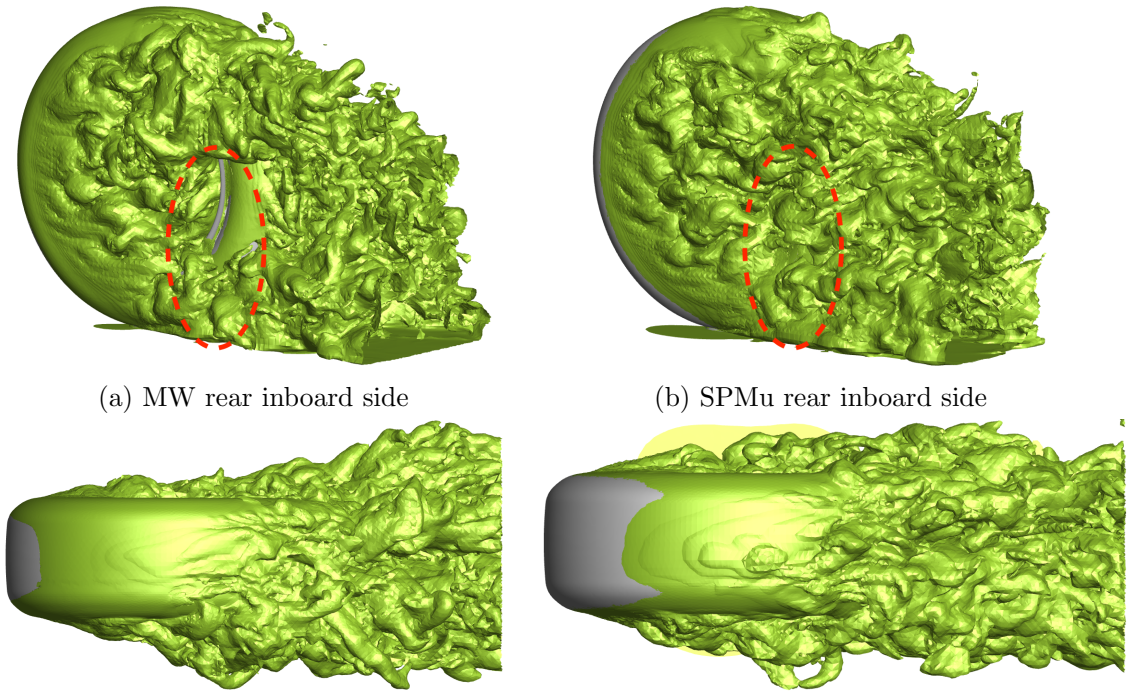


Figure 5.26: Power spectral density of the drag and lift force coefficients on the tyre. The Strouhal number St_D based on the wheel diameter is defined as $St_D = \frac{fD}{U_\infty}$. Top: iLES-SPMs. Bottom: iLES-SPMu.

5.4.3.4 Unsteady wake structure

A three-dimensional comparison of the unsteady near wake of the MW and SPMu cases after 9 CTUs (i.e. $t^* = 9$) using iso-surfaces of total pressure $Cp0 = 0.35$ is shown on figure 5.27. Similarities can be observed between the structures in the MW and SPMu cases, with a separation at the top of the tyre where flow structures are shed. The flow structures travel downstream towards the floor and merge with the wake coming from the rear side flanks and the wheel hub. A first difference is present downstream from the wheel hub, in the region circled in red, where separated flow structures recover the wheel in the SPMu case whereas there is no shedding at this location in the MW case. This is due to the higher inboard bottom vortex in the SPMu case compared to the MW case highlighted on figure 5.22. The larger separation between top and bottom vortex in the MW case makes the jetting vortices generated at the contact patch more distinguishable compared to the SPM-u case where they are quickly driven into the wake.

The top view shows the strongest differences in wake width between the two cases. The flow separation at the beginning of the wheel hub is more contained in the MW case and rapidly expands along the wheel ending with a very wide wake. On the other hand, the separation is larger at the beginning of the hub in the SPMu case and slightly grows larger along the hub due to the varying cross-flow component giving a more uniform time-averaged through-hub flow. The wake does not increase in width downstream from the wheel compared to the MW case.



(a) MW rear inboard side

(b) SPMu rear inboard side

(c) MW top

(d) SPMu top

Figure 5.27: Total pressure iso-surface $Cp_0 = 0.3$ for the isolated wheel in the MW (left) and SPMu (right) cases.

5.5 Conclusion

Several wheel rotation modelling techniques have been explored and compared. Steady techniques such as the MW where the velocity is imposed on the rim surface or the MRF adding momentum to the flow between the spokes are well adapted to steady simulation technologies such as RANS. Yet, the added momentum in the MRF technique used to recover the motion between the spokes generates artificial unphysical pressure gradients and the fixed position of the spokes concentrates the cross-flow in parts of the rim resulting in high-velocity localised peaks.

Physical motion of the spokes is actually required to correctly represent the flow in this highly unsteady situation and is possible only leveraging time-dependent modelling technologies. The SM is the most widely employed solution for this type of cases but the im-

plementation within the spectral/hp element framework proves to be challenging. The SPM which is an immersed boundary technique enabled rotating the wheels without complex mesh transformation. A routine has successfully been implemented within Nektar++ to generate concentration fields for arbitrary complex shapes based on a hit-ray algorithm with a satisfactory level of accuracy. The rotative motion is then reconstructed using an interpolation between pre-generated regular samples. Simulation on an isolated wheel close to industrial standards has shown that the cross-flow distribution through the rim is far more uniform in this case than when only MW is applied due to obstruction of the regularly passing spokes. The SPM effectively corrects the through-hub flow balance and the vortical structure arrangement in the wake from a rather symmetrical layout to tilted jetting vortices due to the slanted downwash showing better qualitative agreement with experimental results on similar isolated wheels compared to the other modelling strategies investigated in this work. This demonstrates that SPM is an effective approach to accurately model rotating wheels and predict the complex flow arrangement caused by the highly unsteady flow.

Yet, there is no systematic method to set the interface thickness parameter ξ since the relationship with the Reynolds number is still uncertain. It turns out that for high Reynolds number values, very small interface thickness values are needed to closely represent the solid surface, which requires very fine mesh resolution to capture. An increase in computational cost has also been noted due to the second Poisson equation for the body force of the immersed body. Pre-processing time, in particular time to generate concentration fields for complex shapes on very fine meshes is also significantly increased compared to classical modelling techniques. Finally, the SPM is a promising technique to accurately represent rotating wheels and to better understand complex flow they generate, thanks to accurate unsteady modelling and the ease of implementation. In fact, once the mesh is generated, there is no need to go through the whole pre-processing step to simulate different rim designs. Rim geometries can be simply swapped by generating the corresponding concentration field, what

is beneficial when it comes to automation and can importantly reduce preparation time when many designs are to be tested. The SPM can also straightforwardly be extended to other rotating components, such as brake disks or tread patterns making it a very convenient tool for wheel studies. Yet, optimisations must be carried out to alleviate associated computational costs, especially on the Poisson solver efficiency.

Chapter 6

Conclusion and further work

This work focused on transferring the spectral/hp element technology to the automotive industry leveraging Nektar++. The transition from academic cases to industrial workflows represents a real challenge and requires important development efforts in terms of reliability and modelling capabilities of numerical tools due to complex geometries and real-world operating conditions. In this view, a workflow has first been developed on the McLaren Automotive P14 adapted geometry to systematically address industrial automotive geometries with the spectral/hp element method. The process shed the light on important differences with classical low-order numerical tools, in particular high-order meshing, turbulence resolution and boundary condition settings. This first study of such complex geometry at realistic Reynolds number with Nektar++ helped to identify two crucial points which need particular attention to enhance capability of the code to tackle such cases with a high level of accuracy and fidelity. P-refinement, introduced by high-order methods, adds up to the traditional h-refinement and raises the question of the combination of these two refinement types. Two different hp-refinement strategies have been assessed on the SAE body with similar numbers of degrees of freedom and highlighted complementarity of the two approaches. The second point which has been considered is wheel rotation modelling which has a very strong impact

on flow prediction around wheels. Turbulent scale-resolving techniques are intrinsically unsteady, therefore real-time rotation of the wheel is physically more correct in this case. The SPM has been further developed to be used in the case of moving complex three-dimensional shapes and employed to perform a simulation of a real-time rotating isolated wheel. The technique proved to be both well adapted to the spectral/hp element framework and capable of qualitatively predicting the correct wake arrangement downstream from the isolated wheel. The two improvements constitute we believe, an important step towards the integration of the numerical tool in automotive development processes.

6.1 Application to the P14 geometry

The simulation of the McLaren Automotive adapted P14 model with the spectral/hp element method can be considered a success in terms of feasibility, showing that simulations of even extremely complex geometries at industrial Reynolds number values can be performed with Nektar++. Comparison of the iLES results to the RANS approach shows notable differences, in particular the imposition of the moving wall boundary condition on the wheel looking more realistic in Nektar++ simulations as well as vortical structures generated by the guiding vanes which can be tracked over a longer distance in Nektar++, possibly due to the lower numerical diffusion of the high-order method.

High-order meshing has been identified as the most critical step of the workflow, sensibly deviating from the usual low-order meshing process. Although geometrical h-refinement is also present, it has to take into account the high-order p-type refinement resulting in lower total element count. Yet for complex geometries, despite the use of high-order elements, high levels of h-refinement need to be kept in very detailed and highly curved regions of the geometry to represent it correctly reducing p-refinement possibilities on the HPC facility. Computation time is also considerably higher than the RANS simulation due to the finer

resolution necessary to solve turbulence scales.

6.2 hp-refinement strategy

The implementation of two different hp-refinement combinations with a fixed number of degrees of freedom on the SAE notchback body has given important insight on how to set the two types of refinements. Higher p-refinement proved more efficient at solving turbulence near solid boundaries than increased h-refinement. Off-body flow features such as vortices however required a minimal level of h-refinement to be effectively captured. Increasing p-refinement upon the already refined mesh improves resolution of these flow features once they have been captured. Yet, the fact that separation over the notch did not match with experimental observations tends to indicate that the turbulent boundary layer is still not fully resolved in our setup. The phenomenon has already been observed on the front splitter and the front diffuser of P14. Near-wall velocity profiles suggest the employed resolutions may still fall short of capturing smallest energetic vortices helping the boundary layer to withstand adverse pressure gradients.

6.3 Rotating wheel modelling

The proposed SPM strategy constitutes an alternative to the SM approach for real-time wheel modelling without associated complex mesh transformations. The enhanced parallelised concentration field generation process, with additional conformity checks, coupled to temporal interpolation allowed us to reconstruct a smooth rotative motion of the rim. A stable simulation of an industrial isolated wheel at high Reynolds number has been effectively performed with the SPM. The modified through-hub flow fluxes have had a strong impact on the wake structure, changing the symmetrical position of the jetting vortices in the MW case to a tilted position with the SPM, which is more in line with MRF results and experimen-

tal observations. The more realistic boundary condition and tangible improvement of flow prediction is to mitigate with the substantially increased computation time. It is doubled compared to MW principally due to the second Poisson pressure equation for the immersed body force and to the important mesh refinement level needed to capture steep concentration field transitions to accurately represent complex shapes at high Reynolds number.

6.4 Further work

We believe that the efforts to facilitate the implementation of Nektar++ in automotive development processes need to be pursued further. Automation is a key point when it comes to industrial workflows, enabling engineers to assess designs with minimum human intervention. In particular, the only step requiring manual intervention in the developed workflow is the treatment of pyramid elements in the base mesh. Although we have recently participated in developments to convert linear pyramids into high-order elements, the boundary layer splitting process still needs to be readapted for this type of elements. Indeed very skewed sub-elements impose restrictive conditions on the maximum time step of the semi-explicit scheme.

Simulations on complex geometries helped realising that, despite their integration in computationally efficient schemes such as spectral/hp element methods, scale resolving techniques remain expensive on HPC structures of reasonable size typical in industrial environments. High-fidelity CFD should, for the moment and near future, be used as a precision tool to provide a better understanding of the flow in particular regions where scale modelling techniques usually fail, such as wheel houses. One important point to consider however is how to define the extent of these regions and the adapted inflow and outflow boundary conditions on truncated geometries.

Although the SPM rather successfully mimics rotating wheels, optimisations can be operated,

in particular for the generation of concentration field samples during the pre-processing step. More efficient ray-triangle intersection algorithms are actively being developed for graphical applications (see for instance the latest work of Pichler et al. [129]) and could help saving time.

The next step for the SPM would be the implementation in a shrouded wheel case to evaluate differences in wake structures predictions compared to steady-state techniques.

Appendices

Appendix A

Mesh metrics

A.1 P14

	Nektar++	RANS
Surface mesh		
Maximum size (mm)	75	112.5
Minimum size (mm)	5	7.5
Volume mesh		
Maximum size (mm)	300	450
Minimum size (mm)	30	45
Prism layer		
Total thickness (mm)	100	100
Number of layers	10	10
Stretching	1.5	1.5

Table A.1: Linear mesh metrics for P14

A.2 SAE body

	HFP3	HCP5
Surface mesh		
Maximum size (mm)	10	17
Minimum size (mm)	2.5	4.25
Volume mesh		
Maximum size (mm)	100	170
Minimum size (mm)	10	17
Prism layer		
Total thickness (mm)	10	10
Number of layers	10	6
Stretching	1.5	1.5

Table A.2: Linear mesh metrics for the SAE body

A.3 DrivAer wheel

	RW/MRF/SPM
Surface mesh	
Maximum size (mm)	20
Minimum size (mm)	2
Volume mesh	
Maximum size (mm)	200
Minimum size (mm)	7
Prism layer	
Total thickness (mm)	20
Number of layers	6
Stretching	1.5

Table A.3: Linear mesh metrics for the isolated DrivAer wheel

Bibliography

[1] W. Hucho, *Aerodynamics of Road Vehicles*. Elsevier, 1987.

[2] M. Nakagawa, S. Kallweit, F. Michaux, and T. Hojo, "Typical velocity fields and vortical structures around a formula one car, based on experimental investigations using particle image velocimetry," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 9, pp. 754–771, apr 2016.

[3] P. Elofsson and M. Bannister, "Drag reduction mechanisms due to moving ground and wheel rotation in passenger cars," in *SAE Technical Paper Series*, SAE International, mar 2002.

[4] A. Schwarczkopf, T. Regert, and T. Lajos, "Investigation of simple possibilities for reduction of drag due to the wheels of road vehicles," in *European Automotive Simulation Conference (EASC 2009), Munich, Germany*, 2009.

[5] B. Agathangelou and M. Gascoyne, "Aerodynamic design considerations of a formula 1 racing car," in *SAE Technical Paper Series*, SAE International, feb 1998.

[6] J. E. Fackrell, *The aerodynamics of an isolated wheel rotating in contact with the ground*. PhD thesis, Imperial College London, 1974.

[7] A. Wäschle, "The influence of rotating wheels on vehicle aerodynamics - numerical and experimental investigations," in *SAE Technical Paper Series*, SAE International, apr 2007.

[8] M. Aultman, R. Auza-Gutierrez, K. Disotell, and L. Duan, "Effects of wheel rotation on long-period wake dynamics of the DrivAer fastback model," *Fluids*, vol. 7, p. 19, dec 2021.

[9] L. Axon, K. Garry, and J. Howell, "An evaluation of CFD for modelling the flow around stationary and rotating isolated wheels," in *SAE Technical Paper Series*, SAE International, feb 1998.

[10] S. Krajnović, "What can LES do in vehicle aerodynamics?," in *The Aerodynamics of Heavy Vehicles III*, pp. 311–326, Springer International Publishing, aug 2015.

[11] C. Cantwell, D. Moxey, A. Comerford, A. Bolis, G. Rocco, G. Mengaldo, D. D. Grazia, S. Yakovlev, J.-E. Lombard, D. Ekelschot, B. Jordi, H. Xu, Y. Mohamied, C. Eskilsson, B. Nelson, P. Vos, C. Biotto, R. Kirby, and S. Sherwin, "Nektar++: An open-source spectral/hp element framework," *Computer Physics Communications*, vol. 192, pp. 205–219, jul 2015.

[12] "Aerodynamic testing of road vehicles–open throat wind tunnel adjustment," Tech. Rep. SAE Report J2071, SAE International, 1994.

[13] G. Karniadakis and S. Sherwin, *Spectral/hp Element Methods for Computational Fluid Dynamics*. Oxford University Press, jun 2005.

[14] S. Sherwin and G. Karniadakis, "A triangular spectral element method: applications to the incompressible navier-stokes equations," *Computer Methods in Applied Mechanics and Engineering*, vol. 123, pp. 189–229, jun 1995.

[15] W. J. Gordon and C. A. Hall, "Transfinite element methods: Blending-function interpolation over arbitrary curved element domains," *Numerische Mathematik*, vol. 21, pp. 109–129, apr 1973.

[16] P. M. Knupp, "Algebraic mesh quality metrics for unstructured initial meshes," *Finite Elements in Analysis and Design*, vol. 39, pp. 217–241, jan 2003.

[17] S. J. Sherwin and J. Peiró, "Mesh generation in curvilinear domains using high-order elements," *International Journal for Numerical Methods in Engineering*, vol. 53, pp. 207–223, oct 2001.

[18] Turner, *High-Order Mesh Generation For CFD Solvers*. PhD thesis, Imperial College London, 2017.

[19] D. Moxey, M. Green, S. Sherwin, and J. Peiró, "An isoparametric approach to high-order curvilinear boundary-layer meshing," *Computer Methods in Applied Mechanics and Engineering*, vol. 283, pp. 636–650, jan 2015.

[20] J.-E. W. Lombard, D. Moxey, S. J. Sherwin, J. F. A. Hoessler, S. Dhandapani, and M. J. Taylor, "Implicit large-eddy simulation of a wingtip vortex," *AIAA Journal*, vol. 54, pp. 506–518, feb 2016.

[21] M. Turner, D. Moxey, J. Peiró, M. Gammon, C. R. Pollard, and H. Bucklow, "A framework for the generation of high-order curvilinear hybrid meshes for CFD simulations," *Procedia Engineering*, vol. 203, pp. 206–218, 2017.

[22] Siemens Digital Industries Software, "Simcenter STAR-CCM+, version v13.02," Siemens 2018.

[23] A. Johnen, J.-F. Remacle, and C. Geuzaine, "Geometrical validity of high-order triangular finite elements," *Engineering with Computers*, vol. 30, pp. 375–382, dec 2012.

[24] A. Gargallo-Peiró, X. Roca, J. Peraire, and J. Sarrate, "Defining quality measures for validation and generation of high-order tetrahedral meshes," in *Proceedings of the 22nd International Meshing Roundtable*, pp. 109–126, Springer International Publishing, 2014.

[25] G. E. Karniadakis, M. Israeli, and S. A. Orszag, "High-order splitting methods for the incompressible navier-stokes equations," *Journal of Computational Physics*, vol. 97, pp. 414–443, dec 1991.

[26] P. Sagaut, *Large Eddy Simulation for Incompressible Flows*. Springer-Verlag, 2006.

[27] S. B. Pope, *Turbulent Flows*. Cambridge University Press, aug 2000.

[28] G.-S. Karamanos and G. Karniadakis, "A spectral vanishing viscosity method for large-eddy simulations," *Journal of Computational Physics*, vol. 163, pp. 22–50, sep 2000.

[29] R. M. Kirby and S. J. Sherwin, "Stabilisation of spectral/hp element methods through spectral vanishing viscosity: Application to fluid mechanics modelling," *Computer Methods in Applied Mechanics and Engineering*, vol. 195, pp. 3128–3144, apr 2006.

[30] R. Moura, S. Sherwin, and J. Peiró, "Eigensolution analysis of spectral/hp continuous galerkin approximations to advection–diffusion problems: Insights into spectral vanishing viscosity," *Journal of Computational Physics*, vol. 307, pp. 401–422, feb 2016.

[31] E. Tadmor, "Convergence of spectral methods for nonlinear conservation laws," *SIAM Journal on Numerical Analysis*, vol. 26, pp. 30–44, feb 1989.

[32] Y. Maday, S. M. O. Kaber, and E. Tadmor, "Legendre pseudospectral viscosity method for nonlinear conservation laws," *SIAM Journal on Numerical Analysis*, vol. 30, pp. 321–342, apr 1993.

[33] R. Moura, G. Mengaldo, J. Peiró, and S. Sherwin, "On the eddy-resolving capability of high-order discontinuous galerkin approaches to implicit LES / under-resolved DNS of euler turbulence," *Journal of Computational Physics*, vol. 330, pp. 615–623, feb 2017.

[34] R. M. Kirby and G. E. Karniadakis, "De-aliasing on non-uniform grids: algorithms and applications," *Journal of Computational Physics*, vol. 191, pp. 249–264, oct 2003.

[35] G. Mengaldo, D. D. Grazia, D. Moxey, P. Vincent, and S. Sherwin, "Dealiasing techniques for high-order spectral element methods on regular and irregular grids," *Journal of Computational Physics*, vol. 299, pp. 56–81, oct 2015.

[36] B. Hupertz, N. Lewington, C. Mockett, N. Ashton, and L. Duan, "Towards a standardized assessment of automotive aerodynamic CFD prediction capability - AutoCFD 2: Ford DrivAer test case summary," in *SAE Technical Paper Series*, SAE International, mar 2022.

[37] D. Aljure, J. Calafell, A. Baez, and A. Oliva, "Flow over a realistic car model: Wall modeled large eddy simulations assessment and unsteady effects," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 174, pp. 225–240, mar 2018.

[38] T. Dairay, E. Lamballais, S. Laizet, and J. C. Vassilicos, "Numerical dissipation vs. subgrid-scale modelling for large eddy simulation," *Journal of Computational Physics*, vol. 337, pp. 252–274, may 2017.

[39] Z. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. Huynh, N. Kroll, G. May, P.-O. Persson, B. van Leer, and M. Visbal, "High-order CFD methods: current status and perspective," *International Journal for Numerical Methods in Fluids*, vol. 72, pp. 811–845, jan 2013.

[40] F. A. J. Hoessler, *Direct numerical simulations and stability analysis of vortex dominated flows around complex geometries*. PhD thesis, Imperial College London, 2011.

[41] J.-E. W. Lombard, *A high-fidelity spectral/hp elementLES study of Formula 1 front-wing and exposed wheel aerodynamics*. PhD thesis, Imperial College London, 2017.

[42] F. F. Buscariolo, J. Hoessler, D. Moxey, A. Jassim, K. Gouder, J. Basley, Y. Murai, G. R. S. Assi, and S. J. Sherwin, "Spectral/hp element simulation of flow past a formula one front wing: validation against experiments," Sept. 2019.

[43] F. Fabian Buscariolo, *Spectral/hp large eddy simulation of vortex-dominated automotive flows around bluff bodies with diffuser and complex front wing geometries*. PhD thesis, Imperial College London, 2021.

[44] F. F. Buscariolo, S. J. Sherwin, G. R. S. Assi, and J. R. Meneghini, "Spectral/hp iLES-SVV simulation methodology study on an ahmed body squared back," in *SAE Technical Paper Series*, SAE International, sep 2018.

[45] S. Ahmed, G. Ramm, and G. Faltin, "Some salient features of the time-averaged ground vehicle wake," in *SAE Technical Paper Series*, SAE International, feb 1984.

[46] J. L. J. and H. W. H., "Aerodynamische formoptimierung der type vw-golf and vw-scirocco.," in *Kolloquium über Industrie-aerodynamik, Aachen, Part 310.4271/840300*, pp. 46–49.

[47] T. Morel, "Aerodynamic drag of bluff body shapes characteristic of hatch-back cars," in *SAE Technical Paper Series*, SAE International, feb 1978.

[48] M. Minguez, R. Pasquetti, and E. Serre, "High-order large-eddy simulation of flow over the "ahmed body" car model," *Physics of Fluids*, vol. 20, p. 095101, sep 2008.

[49] G. Mengaldo, D. Moxey, M. Turner, R. C. Moura, A. Jassim, M. Taylor, J. Peiro, and S. J. Sherwin, "Industry-relevant implicit large-eddy simulation of a high-performance road car via spectral/hp element methods," Sept. 2020.

[50] "Elemental cars10.1063/1.2952595." http://elementalcars.co.uk/the-rp1/, 2018.

[51] S. Kapadia, S. Roy, and K. Wurtzler, "Detached eddy simulation over a reference ahmed car model," in *41st Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, jan 2003.

[52] J. U. Schlüter, "Influence of axisymmetric assumptions on large eddy simulations of a confined jet and a swirl flow," *International Journal of Computational Fluid Dynamics*, vol. 18, pp. 235–245, apr 2004.

[53] S. Diasinos, T. J. Barber, and G. Doig, "The effects of simplifications on isolated wheel aerodynamics," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 146, pp. 90–101, nov 2015.

[54] P. Spalart, "Trends in turbulence treatments," in *Fluids 2000 Conference and Exhibit*, American Institute of Aeronautics and Astronautics, jun 2000.

[55] *Nektar++: Spectral/hp Element Framework User Guide*. Department of Aeronautics, Imperial College London, UK and Scientific Computing and Imaging Institute, University of Utah, USA.

[56] S. Dong, G. Karniadakis, and C. Chryssostomidis, "A robust and accurate outflow boundary condition for incompressible flow simulations on severely-truncated unbounded domains," *Journal of Computational Physics*, vol. 261, pp. 83–105, mar 2014.

[57] V. Venkatakrishnan, "Convergence to steady state solutions of the euler equations on unstructured grids with limiters," *Journal of Computational Physics*, vol. 118, pp. 120–130, apr 1995.

[58] T.-H. Shih, W. W. Liou, A. Shabbir, Z. Yang, and J. Zhu, "A new k-epsilon eddy viscosity model for high reynolds number turbulent flows," *Computers & Fluids*, vol. 24, pp. 227–238, mar 1995.

[59] W. Rodi, "Experience with two-layer models combining the k-epsilon model with a one-equation model near the wall," in *29th Aerospace Sciences Meeting*, American Institute of Aeronautics and Astronautics, jan 1991.

[60] J. McManus and X. Zhang, "A computational study of the flow around an isolated wheel in contact with the ground," *Journal of Fluids Engineering*, vol. 128, no. 3, p. 520, 2006.

[61] R. C. Moura, M. Aman, J. Peiró, and S. J. Sherwin, "Spatial eigenanalysis of spectral/hp continuous galerkin schemes and their stabilisation via DG-mimicking spectral vanishing viscosity for high reynolds number flows," *Journal of Computational Physics*, vol. 406, p. 109112, apr 2020.

[62] S. J. Sherwin and M. Casarin, "Low-energy basis preconditioning for elliptic substructured solvers based on unstructured spectral/hp element discretization," *Journal of Computational Physics*, vol. 171, pp. 394–417, jul 2001.

[63] L. Grinberg, D. Pekurovsky, S. Sherwin, and G. Karniadakis, "Parallel performance of the coarse space linear vertex solver and low energy basis preconditioner for spectral/hp elements," *Parallel Computing*, vol. 35, pp. 284–304, may 2009.

[64] P. M. Palaskar, V. Kumar, and R. Vaidya, "Methodology development to accurately predict aerodynamic drag and lift for passenger vehicles using CFD.," in *SAE Technical Paper Series*, SAE International, apr 2016.

[65] J. Katz, "Aerodynamics of race cars," *Annual Review of Fluid Mechanics*, vol. 38, pp. 27–63, jan 2006.

[66] J. Fabijanic, "An experimental investigation of wheel-well flows," in *SAE Technical Paper Series*, SAE International, feb 1996.

[67] L. Axon, K. Garry, and J. Howell, "The influence of ground condition on the flow around a wheel located within a wheelhouse cavity," in *SAE Technical Paper Series*, SAE International, mar 1999.

[68] E. Croner, H. Bézard, C. Sicot, and G. Mothay, "Aerodynamic characterization of the wake of an isolated rolling wheel," *International Journal of Heat and Fluid Flow*, vol. 43, pp. 233–243, oct 2013.

[69] G. M. L. Good and K. P. Garry, "On the use of reference models in automotive aerodynamics," in *SAE Technical Paper Series*, SAE International, mar 2004.

[70] J. Howell, A. Sheppard, and A. Blakemore, "Aerodynamic drag reduction for a simple bluff body using base bleed," in *SAE Technical Paper Series*, SAE International, mar 2003.

[71] W. Hanfeng, Z. Yu, Z. Chao, and H. Xuhui, "Aerodynamic drag reduction of an ahmed body based on deflectors," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 148, pp. 34–44, jan 2016.

[72] M. Varney, "Base drag reduction for squareback road vehicles," 2020.

[73] D. Wood, M. A. Passmore, and A.-K. Perry, "Experimental data for the validation of numerical methods - SAE reference notchback model," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 7, pp. 145–154, apr 2014.

[74] F. F Buscariolo, W. Hambli, J. Slaughter, and S. Sherwin, "Using a spectral/hp element method for high-order implicit-les of bluff automotive geometries," 2020.

[75] G. W. Carr, "Influence of rear body shape on the aerodynamic characteristics of saloon cars," Tech. Rep. 2, Motor Industry Research Association, 1968.

[76] T. Nouzawa, S. Haruna, K. Hiasa, T. Nakamura, and H. Sato, "Analysis of wake pattern for reducing aerodynamic drag of notchback model," in *SAE Technical Paper Series*, SAE International, feb 1990.

[77] D. Sims-Williams, D. Marwood, and A. Sprot, "Links between notchback geometry, aerodynamic drag, flow asymmetry and unsteady wake structure," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 4, pp. 156–165, apr 2011.

[78] B. R. Gilhome, J. W. Saunders, and J. Sheridan, "Time averaged and unsteady near-wake analysis of cars," in *SAE Technical Paper Series*, SAE International, mar 2001.

[79] J. Fuller and M. A. Passmore, "The importance of rear pillar geometry on fastback wake structures," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 125, pp. 111–120, feb 2014.

[80] J. B. Fuller, *The unsteady aerodynamics of static and oscillating simple automotive bodies*. PhD thesis, Loughborough University, 2012.

[81] A. Cogotti, "A parametric study on the ground effect of a simplified car model," in *SAE Technical Paper Series*, SAE International, feb 1998.

[82] T. Ishima, Y. Takahashi, H. Okado, Y. Baba, and T. Obokata, "3d-PIV measurement and visualization of streamlines around a standard SAE vehicle model," in *SAE Technical Paper Series*, SAE International, apr 2011.

[83] D. Wood, *The effect of rear geometry changes on the notchback flow field*. PhD thesis, Loughborough University, 2015.

[84] B. Basara, V. Przulj, and P. Tibaut, "On the calculation of external aerodynamics: Industrial benchmarks," in *SAE Technical Paper Series*, SAE International, mar 2001.

[85] A. Nader, A. Islam, and B. Thornber, "A comparative aerodynamic investigation of a 20° SAE notchback model," *Applied Mechanics and Materials*, vol. 846, pp. 79–84, jul 2016.

[86] S. Sherwin, "A substepping navier-stokes splitting scheme for spectral/hp element discretisations," in *Parallel Computational Fluid Dynamics 2002*, pp. 43–52, Elsevier, 2003.

[87] D. Xiu, S. J. Sherwin, S. Dong, and G. E. Karniadakis, "Strong and auxiliary forms of the semi-lagrangian method for incompressible flows," *Journal of Scientific Computing*, vol. 25, pp. 323–346, oct 2005.

[88] E. Mercker, "A blockage correction for automotive testing in a wind tunnel with closed test section," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 22, pp. 149–167, jun 1986.

[89] A. Morelli, "Aerodynamic effects of an automobile wheel," tech. rep., MIRA Translation number 47/69 of A.T.A, vol. 22, 6, pp. 81-288, 1969.

[90] W. Stapleford and G. Carr, "Aerodynamic characteristics of exposed rotating wheels," tech. rep., Motor Industry Research Association, 1970.

[91] A. Cogotti, *Aerodynamic Characteristics of Car Wheels*, ch. International Journal of Vehicle Design, Technological Advances In Vehicle Design Series, Special Publication SP3, Impact of Aerodynamics on Vehicle Design,, pp. 173 – 196. 1983.

[92] A. P. Mears, R. G. Dominy, and D. B. Sims-Williams, "The air flow about an exposed racing wheel," in *SAE Technical Paper Series*, SAE International, dec 2002.

[93] A. P. Mears, S. C. Crossland, and R. G. Dominy, "An investigation into the flow-field about an exposed racing wheel," in *SAE Technical Paper Series*, SAE International, mar 2004.

[94] E. Mercker, N. Breuer, H. Berneburg, and H. J. Emmelmann, "On the aerodynamic interference due to the rolling wheels of passenger cars," in *SAE Technical Paper Series*, SAE International, feb 1991.

[95] R. Knowles, A. Saddington, and K. Knowles, "On the near wake of rotating, 40%-scale champ car wheels," in *SAE Technical Paper Series*, SAE International, dec 2002.

[96] R. D. Knowles, A. J. Saddington, and K. Knowles, "On the near wake of a formula one front wheel," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 227, pp. 1491–1502, sep 2013.

[97] A. J. Sprot, J. B. Minto, D. B. Sims-Williams, and R. G. Dominy, "Aerodynamic investigation on the effect of varying through-hub flow on a formula one front wheel assembly," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 4, pp. 929–944, apr 2011.

[98] J. Sprot Adam, *Open-Wheel Aerodynamics: Effects of Tyre Deformation and Internal Flow*. PhD thesis, Durham University, http://etheses.dur.ac.uk/7292/, 2013.

[99] E. Issakhanian, C. J. Elkins, K. P. Lo, and J. K. Eaton, "An experimental study of the flow around a formula one racing car tire," *Journal of Fluids Engineering*, vol. 132, jul 2010.

[100] A. Wäschle, S. Cyr, T. Kuthada, and J. Wiedemann, "Flow around an isolated wheel - experimental and numerical comparison of two CFD codes," in *SAE Technical Paper Series*, SAE International, mar 2004.

[101] B. Schnepf, T. Schütz, and T. Indinger, "Further investigations on the flow around a rotating, isolated wheel with detailed tread pattern," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 8, pp. 261–274, apr 2015.

[102] B. H. Schnepf, *Untersuchung von Einflussfaktoren auf die Umströmung eines Pkw-Rades in Simulation und Experiment.* PhD thesis, TECHNISCHE UNIVERSITÄT MÜNCHEN, Fakultät für Maschinenwesen, 2016.

[103] J. Axerio-Cilies and G. Iaccarino, "An aerodynamic investigation of an isolated rotating formula 1 wheel assembly," *Journal of Fluids Engineering*, vol. 134, nov 2012.

[104] J. Axerio-Cilies, E. Issakhanian, J. Jimenez, and G. Iaccarino, "An aerodynamic investigation of an isolated stationary formula 1 wheel assembly," *Journal of Fluids Engineering*, vol. 134, feb 2012.

[105] "Moving reference frame description." $https://openfoamwiki.net/index.php/See_the_MRF_developme$ Modified: 13-06-2009.

[106] T. Larsson, C. Bienz, T. Sato, and B. Ullbrand, "In front of the grid - cfd at sauber petronas f1 leading the aerodynamic development," 06 2003.

[107] S. Kandasamy, B. Duncan, H. Gau, F. Maroy, A. Belanger, N. Gruen, and S. Schäufele, "Aerodynamic performance assessment of BMW validation models using computational fluid dynamics," in *SAE Technical Paper Series*, SAE International, apr 2012.

[108] R. Lewis, M. Cross, and D. Ludlow, "The influence of rotating wheels on the external aerodynamic performance of a vehicle," in *The International Vehicle Aerodynamics Conference*, pp. 161–173, Elsevier, 2014.

[109] X. Yu, Q. Jia, D. Bao, and Z. Yang, "A comparative study of different wheel rotating simulation methods in automotive aerodynamics," in *SAE Technical Paper Series*, SAE International, apr 2018.

[110] A. I. Heft, T. Indinger, and N. A. Adams, "Experimental and numerical investigation of the DrivAer model," in *Volume 1: Symposia, Parts A and B*, American Society of Mechanical Engineers, jul 2012.

[111] T. Hobeika, L. Löfdahl, and S. Sebben, "Study of different tyre simulation methods and effects on passenger car aerodynamics," in *The International Vehicle Aerodynamics Conference*, pp. 187–195, Elsevier, 2014.

[112] T. Hobeika and S. Sebben, "CFD investigation on wheel rotation modelling," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 174, pp. 241–251, mar 2018.

[113] G. Fu and A. Untaroiu, "Investigation of tire rotating modeling techniques using computational fluid dynamics," *Journal of Fluids Engineering*, may 2021.

[114] L. Haag, M. Kiewat, T. Indinger, and T. Blacha, "Numerical and experimental investigations of rotating wheel aerodynamics on the DrivAer model with engine bay flow," in *Volume 1B, Symposia: Fluid Measurement and Instrumentation; Fluid Dynamics of Wind Energywable and Sustainable Energy Conversion; Energy and Process Engineering; Microfluidics and Nanofluidics; Development and Applications in Computational Fluid Dynamics; DNS/LES and Hybrid RANS/LES Methods*, American Society of Mechanical Engineers, jul 2017.

[115] A. S. Misar, M. Uddin, A. Robinson, and C. Fu, "Numerical analysis of flow around an isolated rotating wheel using a sliding mesh technique," in *SAE Technical Paper Series*, SAE International, apr 2020.

[116] R. Yamamoto, Y. Nakayama, and K. Kim, "A smooth interface method for simulating liquid crystal colloid dispersions," 2003.

[117] Y. Nakayama and R. Yamamoto, "Simulation method to resolve hydrodynamic interactions in colloidal dispersions," *Physical Review E*, vol. 71, p. 036707, mar 2005.

[118] R. Yamamoto, Y. Nakayama, and K. Kim, "A method to resolve hydrodynamic interactions in colloidal dispersions," *Computer Physics Communications*, vol. 169, pp. 301–304, jul 2005.

[119] X. Luo, M. R. Maxey, and G. E. Karniadakis, "Smoothed profile method for particulate flows: Error analysis and simulations," *Journal of Computational Physics*, vol. 228, pp. 1750–1769, mar 2009.

[120] Z. Wang, M. S. Triantafyllou, Y. Constantinides, and G. E. Karniadakis, "A spectral-element/fourier smoothed profile method for large-eddy simulations of complex VIV problems," *Computers & Fluids*, vol. 172, pp. 84–96, aug 2018.

[121] K. G. Hermann Schlichting (Deceased), *Boundary-Layer Theory.* Springer-Verlag GmbH, Oct. 2016.

[122] X. Luo, B. P. Epps, C. Chryssostomidis, and G. E. Karniadakis, "Comparison of turbulence models for simulating flow in waterjets," in *Proceedings of the eleventh inter-national conference on fast sea transportation. Honolulu, Hawaii, USA978-3-662-52919-5*, (Honolulu, USA), 2011.

[123] A. Cassinelli, A. M. Gabín, F. Montomoli, P. Adami, R. V. Díaz, and S. J. Sherwin, "Reynolds sensitivity of the wake passing effect on a LPT cascade using spectral/hp element methods," *International Journal of Turbomachinery, Propulsion and Power*, vol. 7, p. 8, feb 2022.

[124] F. Mohaghegh and H. Udaykumar, "Comparison of sharp and smoothed interface methods for simulation of particulate flows i: Fluid structure interaction for moderate reynolds numbers," *Computers & Fluids*, vol. 140, pp. 39–58, nov 2016.

[125] F. Mohaghegh and H. Udaykumar, "Comparison of sharp and smoothed interface methods for simulation of particulate flows II: Inertial and added mass effects," *Computers & Fluids*, vol. 143, pp. 103–119, jan 2017.

[126] T. Möller and B. Trumbore, "Fast, minimum storage ray-triangle intersection," *Journal of Graphics Tools*, vol. 2, pp. 21–28, jan 1997.

[127] A. O'Sullivan, "Cfd simulation of a rotating isolatedcar wheel with smoothed profile method in nektar++," Master's thesis, Imperial College London and EPFL, 2021.

[128] S. Woop, C. Benthin, and I. Wald, "Watertight ray/triangle intersection," *Journal of Computer Graphics Techniques (JCGT)*, vol. 2, pp. 65–82, June 2013.

[129] T. A. Pichler, A. Ferko, M. Ferko, P. Kán, and H. Kaufmann, "Precomputed fast rejection ray-triangle intersection," *Graphics and Visual Computing*, vol. 6, p. 200047, jun 2022.