

Alternative Simulation Methods for Assessing Aerodynamic Drag in Realistic Crosswind

Andrew D'Hooge, Robert Palin, and Luke Rebbeck
Tesla Motors

Joaquin Gargoloff and Bradley Duncan
Exa Corp.

ABSTRACT

The focus of evaluating yaw characteristics in automotive aerodynamics has been primarily with regards to the effects of crosswind on vehicle handling. However, changes to drag that the vehicle experiences due to prevalent on-road crosswind can also be significant, even at low yaw angles. Using wind tunnel testing, it is possible to quickly determine the static yaw performance of the vehicle by rotating the vehicle on a turntable to different yaw angles during a single wind tunnel run. However, this kind of testing does not account for dynamic crosswind effects or non-uniform crosswind such as with natural on-road turbulence.

Alternatively, numerical simulations using computational fluid dynamics (CFD) can be used to evaluate yaw performance. In this paper, Exa's PowerFLOW is used to examine two alternative methods of simulating aerodynamic performance in the presence of realistic on-road crosswind for the Tesla Model S sedan. These methods provide reduced computational cost compared to simulating a full range of static yaw angles. The first is dynamic vehicle yaw, which rotates the vehicle similar to a wind tunnel turntable during the transient simulation. The second method uses a high onset turbulence level that provides a more realistic on-road wind condition by including non-uniform velocity fluctuations representing natural turbulence from wind. Results are compared to those obtained using the traditional methodology with static yaw sweeps for characterizing performance in crosswind.

CITATION: D'Hooge, A., Palin, R., Rebbeck, L., Gargoloff, J. et al., "Alternative Simulation Methods for Assessing Aerodynamic Drag in Realistic Crosswind," *SAE Int. J. Passeng. Cars - Mech. Syst.* 7(2):2014, doi:10.4271/2014-01-0599.

INTRODUCTION

Automotive aerodynamics ultimately resolves into the six forces and moments that a vehicle moving through an environment experiences. In wind tunnel testing, the typical frame of reference is switched so that the vehicle is stationary and air is moved past it at a velocity to produce forces and moments of the same magnitude. The most obvious complication is in treatment of the ground plane, which on-road has the same speed as the flow around it (normally stationary) and a velocity relative to the vehicle. While stationary floor tunnels have been useful for decades of ground vehicle development, more modern wind tunnels incorporate a variety of means of boundary layer control and ground simulation. Meanwhile, having the ground plane move at a velocity matching the free stream flow in CFD simulations is trivial.

A facet of matching real world conditions that has only recently been investigated is the effect of ambient wind and turbulence conditions. As the effects of aerodynamic forces and moments only become significant to vehicle performance at higher

speeds, much of the focus thus far has been on highway driving conditions. Several of the works in this field have categorized the most typical flow fields into three categories: open terrain, road side obstacle (RSO), and urban/traffic driving. As shown in [1], these three cases represent a broad variety of wind conditions that a vehicle will see in service. However, as shown in [2], it is difficult to completely represent the full spectrum of turbulence scales in the laboratory environment. Passive systems can replicate small scale turbulence, while active systems [3] can replicate large scale turbulence. As with ground plane replication, CFD is able to easily generate any desired turbulence level, usually specified via turbulence intensity and turbulent length scale. The influence of varying these parameters on the aerodynamic performance of a vehicle evaluated in CFD is examined in [4]. Other empirical studies [5, 6] also show that even moderate levels of onset turbulence similar to that seen on-road can cause significant changes to a vehicle's drag coefficient.

Replicating steady side wind conditions, however, is easily accomplished in most wind tunnels. The importance of being able to acquire data at a variety of yaw angles means that a turntable is a requirement of any wind tunnel. Typically, data is gathered for a range of yaw angles specified by the test engineer on each vehicle configuration. It is up to the aerodynamicist to then weigh the importance of yaw performance (the six force and moment components) against other in-service conditions to ultimately make the most effective design decisions. This sort of yaw sensitivity analysis is far more resource-consuming in the digital domain, however. While not difficult to implement in the simulation set-up for most commercial CFD codes, each yaw angle requires its own solution, which can vary from the order of several hours to a day or more for typical modern computing systems. Therefore, developing a complete understanding of a vehicle's yaw performance in a CFD-led program can become very time consuming, to the detriment of exploring further design directions. This paper explores alternative methods of simulating the yaw performance that reduce the computational resource required.

The vehicle used in this study is the Tesla Model S, a four door fastback battery electric vehicle (BEV) shown in Figure 1. The battery is located under the cabin between the wheels, leading to a flat floor. This creates several unique flow features and allows for a functional diffuser aft of the battery. Nineteen inch diameter wheels and tires were used in the simulations, the base wheel package for this vehicle. The wheels were incorporated within a sliding mesh region for true geometry rotation to most accurately capture the aerodynamic properties in this area.



Figure 1. The vehicle used in this study, the Tesla Model S, shown with 19" wheels.

MOTIVATION

One of the primary goals for any road vehicle is to maximize its range for an allotment of energy, which in turn means using the minimum amount of energy for any given operating point. For passenger vehicles, energy consumption is proportional to the overall resistance a vehicle experiences, which varies with speed. The energy use for a vehicle at a steady speed can be broken into four major components: tire rolling resistance, mechanical drivetrain resistance, ancillary power consumption, and aerodynamic drag. At low speed, resistance to maintaining

a speed is largely due to the first two components. At higher speeds, such as during highway travel, aerodynamic drag dominates the energy use of the vehicle. As mentioned previously, it is at highway speeds (typically 60-80 MPH, or 26.8-35.8 m/s) where much of the work regarding on-road turbulence exists as well. Naturally, understanding the on-road conditions are the first step in replicating them both empirically and computationally.

As described in [7], the most statistically significant range of turbulence occurs on the order of several car lengths, specifically described as from two to twenty in [4]. The Tesla Model S used in this study is roughly five meters in length, meaning the turbulence length scale for these eddies will range from ten meters to one hundred meters. At a test speed of 70 MPH (31.3 m/s), it will take more than one second to cover a large percent of flows if the assumption of a Gaussian distribution (from [8]) remains true for this range of turbulence. In the review of cross wind and turbulence studies from [2], this puts the turbulence outside of the fully unsteady region and on the border of being considered quasi-steady ($0.1 < f < 1.0$, Hz). In short, these large scale eddies will often immerse the vehicle for a long enough period for the vehicle's flow field to respond as if it were at a constant yaw angle, the resultant of which is a combination of the vehicle's speed and the velocity vector of the turbulence as shown in [8]. So for both steady wind conditions and a significant portion of on-road turbulence, especially open road driving, being able to characterize the vehicle's yaw performance and understanding its flow field is vital to the greater goal of accounting for in-service energy use.

Prescribing a single, fixed, optimal yaw angle for testing presents two problems. First, vehicle geometries typically have non-linear drag response to yaw angle. In other words, it is not accurate to extrapolate a force coefficient versus yaw angle graph from a small sample of one or two data points. The complicated interactions between major flow structures (wheel wakes, base region, A- and C-pillar vortices, etc.) mean that each vehicle will have unique curve shapes for the yaw performance parameters. Geometry leading into the base region can even create highly non-linear, almost step-like, features in the yaw performance.

The second reason for desiring a more complete understanding of the yaw performance is the known distribution of typical on-road flows. As with the turbulent length scale also described in this section, turbulent intensity (TI) is spread over a wide band. While varied, this band has been characterized in previous literature. Of the three driving scenarios defined in [1] and [7], turbulent intensities between 4% and 10% cover a large majority of the data points, particularly for non-urban driving. The average of this range, 7%, also matches one of the two TIs used in [4].

While it would be nearly impossible to justify optimizing for a very precise yaw angle given the variety of conditions an internationally homologated vehicle could see in service, further data suggests that the majority of yaw seen falls within

the band of zero to six degrees. Meteorological data, typically gathered at a height of ten meters, can be interpolated to a height more representative for ground vehicles using a power law function. Combining this wind speed probability with a vehicle speed of 70 MPH and an arbitrary direction gives the distribution shown in Figure 2. With this information, the yaw range of interest chosen was from zero to six degrees, which covers more than 88% of the probability distribution. Further, the chosen TI of 7% and free stream vehicle speed of 70 MPH creates an effective local yaw angle of four degrees, which is within the same band.




Figure 2. Yaw angle probability distribution for a ground vehicle travelling at 70 MPH.

As mentioned in the introduction, gathering the required information to evaluate yaw performance via a series of fixed yaw angle simulations is not any more difficult than a non-yawed simulation. However, the time required to solve such a case is increased over a straight ahead case where the assumption of a symmetric flow field can be used to cut the domain in half. While the inherent lateral wake oscillations will be lost, previous study of the vehicle used in this simulation [9] shows that the symmetric boundary condition at the vehicle centerline plane produces no significant change to the time-averaged results apart from side force. Furthermore, the electric drivetrain of Model S means that the floor of the vehicle is actually symmetric with the exception of a few flow independent features, making the symmetric simulation an obvious choice for the standard design simulation.

METHODOLOGY

The CFD software used in this study is Exa's PowerFLOW, a Lattice Boltzmann Method solver. Its inherently transient solution offers the ability to vary boundary conditions and move geometry during the simulation. This feature has been previously used to simulate wheel geometry and its effects on aerodynamic power consumption [10]. These two capabilities allow for two different methods of changing the flow direction relative to the vehicle. The first is rotating the vehicle according to a time-varying function, similar to a turntable in a wind tunnel. Alternatively, the vehicle can be kept stationary while the boundary conditions, namely the magnitude direction of onset flow, are varied. With this option, the side boundaries of the domain are changed to be periodic rather than frictionless walls. This latter approach was chosen for two reasons.

First, it allows the yawing motion to approximate a steadily increasing cross wind approaching a vehicle of constant speed, more closely matching the on-road scenario than if the car was rotated in a constant speed flow. In that case, and with a wind tunnel turntable, the velocity component in the direction of travel changes with the yaw angle to give a constant resultant speed. However, in reality, the vehicle speed is typically constant and ambient wind augments the apparent wind, a state which the second set-up represents. Specifically, V_x is kept constant while V_y is used to control the yaw angle. To compensate for the time required for flow to reach the vehicle from the inlet, a probe in the same longitudinal location as the furthest forward point vehicle but far enough away laterally and vertically to not be influenced by the local flow field was used during post-processing to judge flow angularity at the vehicle. In principle the same sorts of practices could have been used in a simulation where the vehicle was rotated, but the added complication could be an added source of uncertainty and error. Figure 3 details the location of this probe within the computational domain.




Figure 3. Overview of computational domain and location of far field probe.

The second reason for choosing to control the onset flow via the inlet was for proper development of the ground boundary layer in the lateral direction. Again, the set-up should replicate on-road conditions as nearly as possible. Going back to the discussion of reference frames shows that the wind has a velocity with respect to both the ground and the vehicle. The relative difference between the wind and ground will result in a boundary layer with an increasing momentum deficit approaching the ground. For a case with the vehicle rotated on a simulated turntable, the oncoming flow will always be in the same direction and magnitude as the ground meaning a boundary layer cannot form. Again as in the first rationale, it would have been possible to vary the ground speed as a function of the flow approaching the vehicle during the simulation but at the cost of increased simulation complexity. The existence of this sheared yaw profile in ground proximity is noted in [8], however insufficient empirical data exists to fully describe it across a range of driving conditions for correlation in this study. [8] also mentions that this effect is an important part of yacht design, where lower travel speed combines with a much taller vehicle to create a larger gradient of yaw over the height of the yacht as compared to automotive cases.