

Using Vehicle Simulation to Investigate Controllability

Michael Ellims and Helen Elizabeth Monkhouse
Protean Electric Ltd.

Damian Harty and Teena Gade
Coventry Univ.

ABSTRACT

All functional safety standards have some definition of “risk” and the automotive standard ISO 26262 is no exception. Risk is related to the exposure, the severity of the outcome, and in the case of ISO 26262, the controllability in relation to a specific vehicle hazard or hazards associated with the behavior of the vehicle or part of the vehicle. Thus hazards are central to understanding the risk associated with systems.

When considering traditional power train systems, based on internal combustion engines or centralized electric motors, hazards are most usually limited to unintended acceleration and deceleration. The situation is complicated somewhat with the introduction of electronically controlled differentials, which can induce limited amounts of induced yaw, as can ABS and ESC. In a similar manner, replacing the centralized driveline system with in-wheel electric motors brings with it a similar set of issues.

In this paper we describe the work undertaken to qualitatively identify the hazards associated with in-wheel motors and to quantify the vehicle level effects that could be expected. With this being done to ensure that, when realized as an engineering object, the level of controllability and hence the residual risk to a vehicle fitted with in-wheel motors remains within tolerable bounds.

CITATION: Ellims, M., Monkhouse, H., Harty, D. and Gade, T., "Using Vehicle Simulation to Investigate Controllability," *SAE Int. J. Alt. Power.* 2(1):2013, doi:10.4271/2013-01-0180.

INTRODUCTION

This paper is intended to document the activities, tools and techniques that have been used to *quantitatively* assess the risk associated with an in-wheel electric motor in line with ISO 26262 Part 3[19].

The paper is divided into five main sections as follows:

- Part 1 examines the environment in which the motor operates, including the driver and impacts of regulation.
- Part 2 examines the vehicle level effects of in-wheel motor failure at the vehicle level using a black-box approach.
- Part 3 approaches the problem using analytical methods supplemented by detailed vehicle modeling.
- Part 4 synthesises the results of both investigations.
- Part 5 briefly considers the impact of the results from part 4 on the functional safety concept.

Performing what is essentially the same study twice may at first seem odd. However as we hope will become evident from reading this paper, having that alternative independent

approach has greatly increased our confidence in the final result.

PART 1: ENVIRONMENT

The concept for an electric vehicle powered by in-wheel electric motor is simple. Place an electric machine (i.e. a motor/generator) at two or four corners of the vehicle in order to replace the centralized drive train with an electric transmission.

The potential advantages of this arrangement are numerous; it frees up significant amounts of space within a vehicle, it simplifies the routing of the steering, and it removes the need for a differential. In addition it simplifies the implementation of torque vectoring and provides a mechanism to use the electric machines as a component of the antilock brake system (ABS) and electronic stability control (ESC) [14].

Fitting in-wheel electric motors to hybrid vehicles has similar advantages; for example adding two electric in-wheel motors to the rear axle of a front wheel drive vehicle is relatively straight forward conceptually. If the vehicle is

small enough, one could even conceive of the friction components of the service brake being deleted from the rear of the vehicle.

At the conceptual level, the in-wheel motor is simply considered as a black box with no inputs or outputs, and for the purpose of evaluating the vehicle level affects of in-wheel motor failure this is sufficient.

Clearly this abstraction is not tenable as a realizable device. However, very few connections between the in-wheel motor and the vehicle are required to produce a device that can be realized. The minimum set comprises:

- A means of transferring electric power.
- A means of control.
- A means of removing heat.
- A means of mounting to the vehicle.
- A means of mounting to the wheel and tire.

The internal implementation of the in-wheel motor is far from simple, but aside from its location on the vehicle (e.g. front right), the wheel has no information about the vehicle on which it is being used. Specifically the in-wheel motor has no information regarding the state of the vehicle, of other equipment fitted to the vehicle or of the driver's intentions. Although the possibility exists for the wheel motor to take unilateral failure mitigation, this is not done because the context needed to validate a given action is not known. Also, if the motor were to actively attempt to compensate for faults then additional complexity would be introduced; primarily thought interactions within the control hierarchy as shown in [Figure 1](#).

It is for this reason that we have attempted to treat the motor conceptually as a “dumb actuator” throughout its development. That is, while being a complex electronic device; it will deliver to the best of its ability the torque demanded of it and signal failures to a vehicle level control system as soon as it is able; relying on devices higher in the hierarchy to take coordinated mitigating actions.

The Environment

The environment in which the in-wheel motor is used affects how motor failures will affect the vehicle to which it is mounted. There are two aspects to this, first there are direct effects on vehicle dynamics. Second there is the set of assumptions that can be made about the equipment that will be fitted to the vehicle.

To simplify the analysis undertaken, the vehicles to which the in-wheel motor could be fitted were limited to vehicles in the class M1/N1 [8]. That is passenger and light commercial vehicles less than 3,500kg, with a further restriction being that motors would only be fitted in pairs at the end of any axle.

By limiting the study to vehicles in the M1/N1 class we obviously limit the size and weight of the vehicles under consideration. However this then allows us to make strong

assumption about the equipment that will be fitted to the vehicles.

For example, M1/N1 class vehicles are required to conform to either UNECE Regulation 13H [41] or FMVSS 135 [24]. Both regulations specify braking performance and policies for the fitment of ABS and ESC. Additionally an electric vehicle will be expected to conform to UNECE Regulation 100 [42] or FMVSS 305 [25]. Consequently during the development of the safety concept certain measures can be assumed; such as ground loop detection and automated power shutdown in crash situations. All of these assumptions can be incorporated into the functional safety concept for the motor and associated safety documentation.

Vehicle Architecture

As previously stated the wheel motor is intended for use on only road going vehicles. In general we do not consider hazards that exist in contexts other than this. That said, the context in which the motor is used is complex and when considering hazards one must also consider all other systems with which the motor may interact.

This is an interesting problem, because while it may be obvious that the in-wheel motor constitutes all or part of the vehicle driveline, less obvious is its participation in several other vehicle systems. Specifically the in-wheel motor is a component of:

- The drive train.
- The service brake system (potentially)
- The transmission.
- The steering system (potentially).

The need to obtain type approval (in Europe) for the vehicle is assumed, which allows us to make strong assumptions about the equipment fitted to the vehicle. For example, as the in-wheel motor may form part of the service brake system it is subject to UNECE Regulation 13H [41] and legally:

- It must be able to be controlled by the ABS unit.
- It must tolerate single point failures, including failure of the batteries.
- Must meet specific performance criteria, which in turn means that friction brakes must be able to be fitted.

Thus to form a coherent functional safety concept we make the assumption that the wheel motor will interface with other vehicle based systems (ABS, ESC etc.) through a dedicated controller. This has the advantage of simplifying the motor design, and removes the need for the motor to have any knowledge of what is happening at a vehicle level. This hierarchy of control is shown in [Figure 1](#).

The Hazards

Hazard identification was performed utilizing standard methods - initially “brain storming” supported by a structured HazOP process [21]. The analysis resulted in the following set of moving vehicle hazards being associated with in-wheel motor failures:

- Increased and Reduced Braking.
- Induced Yaw, including Intermittent Yaw.
- High Acceleration and Uncommanded Acceleration.
- Low Acceleration and Uncommanded Deceleration.




Figure 1. Conceptual representation of the vehicle control hierarchy.

The following sets of hazards were associated with stationary vehicles:

- Uncommanded Vehicle Movement.
- Uncommanded Rotation of a wheel.
- Incorrect Rotation of a wheel.
- Low Acceleration.
- High Acceleration.
- Magnetic Field (to persons fitted with pacemakers).

The majority of the hazards described above have the same root cause; either the in-wheel motor produces too much torque (MORE) or too little torque (LESS, NO etc.). This is true of both the moving and the stationary vehicle cases.

Given that all vehicle level hazards are derived from two main root causes, these two classes of faults characterize the primary concerns associated with the in-wheel motor, that of Uncommanded Acceleration and Induced Yaw.

Uncommanded Acceleration

Uncommanded or unintended acceleration has been of concern to the automotive industry for many years. Equation 1 provides the relationship between the torque generated at the wheel and the radius (including the tire) of that wheel.

$$F = ma = \frac{T}{r} \quad (1)$$

Given that the maximum torque (T) that a single in-wheel motor can deliver, the radius of the wheel (r) and the target vehicle's mass, we can calculate the maximum acceleration that will result, which for a vehicle of 1700kg is just under 0.2g.

Guidance on what may be considered tolerable is provided by several sources, but the original NHTSA “Silver Book” [22] suggests that values of less than 0.3g are generally not considered uncontrollable. This suggests that a single motor level failure would be controllable under the majority of situations (C1), especially given the structure of functional safety concept which is outlined in Part 5.

Induced Yaw

Perhaps the more interesting hazard is Induced Yaw. On a moving vehicle an asymmetrical torque across an axle will in general induce a turning force; causing the vehicle to turn about its vertical axis. This induced yaw has the potential to disturb the behavior of the vehicle far more significantly than a longitudinal disturbance such as Uncommanded Acceleration, but is more difficult to quantify.

Because of this we chose to use vehicle simulation to better understand the vehicle level effects of Induced Yaw, and to establish an upper bound for torque or power delivered by the motor. That knowledge could then be used to shape the functional safety concept for, and the design of, the in-wheel motor. The approaches taken to address this problem are the focus of the next two major sections.

PART 2: SIMULATIONS

Overview

Vehicle simulation at Protean was conducted using the IPG CarMaker tool. A default library vehicle model was used. This model was essentially a new model VW Beetle.

The default vehicle model was modified by removing the standard drive train and replacing it with an all electric four wheel drive train. Control was achieved by directly applying the torque requested by the standard driver model (supplied by the package) to all four wheels of the vehicle.

The standard vehicle model was modified by adding a block that simulated the Vehicle Control Unit to independently process the driver inputs to the brake and accelerator pedals and pass commands directly to the models of the in-wheel motors and perform functions such as brake blending.

The second major modification to the standard vehicle model was to introduce both failure triggering and monitoring functionality. This functionality allowed delta torque values to be applied directly of a single wheel at a specific point on a virtual test track and to automatically capture and evaluate data on performance of the vehicle (e.g. yaw rate, lateral

acceleration) and the driver model (steering wheel angle, torque etc) from large numbers of vehicle runs.

The Human Driver

The primarily objective of the vehicle simulation was to investigate the controllability limits for a vehicle subjected to a torque error occurring in a motor. The torque error can result in either an increase or decrease in torque, leading to a torque delta or asymmetric torque across the axle of the vehicle.

To better understand both the performance limits of the driver, and how they can be expected to react when a failure occurs, the following areas of human behavior and performance were investigated:

- Limits in driver performance; e.g. how fast and how far the driver could be expected to steer.
- Limits of driver behavior. Given that the vehicle response is “quick”, then that necessitates a controlling action that is equally quick; braking and steering being the two obvious examples.
- Examining what are the limits of “reasonable” behavior for drivers on the road. For example very few drivers corner at the limit of adhesion, but what is a typical usage profile?

Performance Limits

Data on human performance came from a number of sources including the test procedures for braking regulations [41]. The set of results is summarized in Table 1.

Table 1. Selected criteria for human performance limits and the sources they were derived from

Factor	Limit	Source
Steering Torque	7Nm	Neukum <i>et al.</i> [26]
	4 Nm	Switkes <i>et al.</i> [35]
	15 Nm	Pick <i>et al.</i> [28]
Steering Rate	100°/s	Schmidt <i>et al.</i> [29] [30]
	120° in 2s	UNECE Reg. 13H [41]
Steering Angle	240°	UNECE Reg. 13H [41]

Reaction Limits

A second factor considered was the driver's speed of response. That is how fast the driver can cope with a disturbance and whether that reaction could be classed as a “normal” response. Two candidates suggest themselves, braking and steering and Table 2 summarizes results from several sources.

Regarding controllability the most interesting work is by Neukum *et al.* [26], Neukum [27] and Wierwille *et al.* [37]. The paper by Wierwille *et al.* on driver response to wind gusts was located first and appears to be the most complete study to date. This strongly suggests that for impulse wind gusts applied toward the front of the vehicle the driver response time is low, in the order of 0.3 seconds. This is

supported by the studies from Neukum *et al.* which looked at driver responses to steering errors on a range of vehicles, which again gives a response time in the order of 0.3s (data is summarized in Table 2).

The fact that the steering response times given in Table 2 are under half a second is significant, as it indicates an automatic response (process) [31, 10, 32] is being triggered and that conscious cognition (i.e. controlled processing) is not involved. What is being observed is a learnt, associative response. It should be noted that for an accurate automatic response to be ensured, there must be adequate statistical or associative training. We suggest that this is provided in normal vehicle use, as drivers must respond to lateral disturbances caused by wind gusts, standing water and even road surfaces imperfections such as potholes.

The work by Neukum *et al.* is important for this study as it presents limits for acceptable vehicle yaw rates, specifically:

- 4.0 °/s at 50 km/h
- 3.0 °/s at 100 km/h
- 2.5 °/s at 150 km/h

In addition a maximum lateral acceleration of 1.25 m/s² is given. In the context of the ISO 26262 these values equate to a C2 controllability level. That is, we can expect 90% of drivers to retain control following this level of *additional* yaw being induced by a failure.

The data outlined above provide a set of tangible criteria that can be used in the evaluation process. Intuitively some of the values presented seem low, particularly at higher speeds. However it should be remembered that these values are a delta. The values for the highest speeds are in accordance with work conducted on wind gusts by McAdam *et al.* [16], which gives a value of just under 2.5 °/s using a subjective rating system similar that that used in [26].

Why the driver can react so quickly is also of interest. Neukum [27] suggests that it is due to steering wheel torque feedback. However track data shows no such feedback from rear wheel disturbances. Given the importance of the vestibular senses in driving, these are probably also significant. MacAdam [17] suggests that they will both contribute and that disentangling their individual influences may not be feasible.

Reasonable Behavior

Driver reactions do not occur in isolation and are relevant only in the context in which the reaction becomes necessary, or to use the ISO 26262 terminology, “scenarios”.

It is surprisingly difficult to obtain information on what constitutes “normal” driver behavior on the road. However this information links directly to the information needed to examine the simulations results.

The classic paper in this area is by Lechner and Perrin [15] which suggests that the majority of driving takes place within the following bounds:

- Longitudinal acceleration: $\pm 0.2g$
- Lateral acceleration: $\pm 0.3g$

Table 2. Reaction times for braking and steering

Factor	Limit	Source
Brake Reaction Time (BRT)	0.36s to 6.3s - range	Young [36]
	0.7 to 0.75s - best	Green [9]
	1.2 to 1.35s - common	
	1.5s - surprise	Mehmood [19]
	0.89 to 1.93 s - normal	
Total Braking Time (TBT)	1.5s to 3s	Triggs[39]
	0.5s to 2.4s	Young [36]
	1.2s (mean)	
Steering reaction Time	0.25s	Neukum <i>et al.</i> [26]
	0.3 to 0.45s	Schmidt <i>et al.</i> [4]
	0.3 to 0.45s	Wierwille <i>et al.</i>
	0.29s	Switkes <i>et al.</i> [35]

The best source of supporting evidence for these general figures come from drive cycle studies [13]. Lechner and Perrin also note that there is a strong correlation between vehicle speed and lateral acceleration, giving the following approximate limits: for acceptable lateral acceleration “roughly up to 0.5g below 60 km/h, 0.3g at approximately 120 km/h and 0.2g at approximately 160 km/h.”

ISO 26262 also has a requirement to consider “reasonable misuse”, however it is not immediately obvious what may constitute misuse when considering the drive train. One interpretation is that misuse involves the driver pushing the vehicle beyond the bounds suggested by Lechner and Perrin.

Blundell and Harty [1] when discussing active control characterize the driving envelope in three lateral acceleration categories, where “normal” matches the definitions above, “spirited” is defined as 0.3 to 0.6g and “emergency” as 0.6g or greater. These values correlate with a large study conducted by the National Highway Traffic Safety Administration (NHTSA) [23]. In their study NHTSA observed that “lateral accelerations of greater than 0.60 g were so infrequent that differences among these groups of drivers could not be determined”. An examination of the summary data presented in [23] suggests that activity above 0.5g is insignificant and that 0.4g possibly represents the bound of normal usage (figures 38 and 39). However it must be noted that one weakness in this set of data is that it is not correlated with speed.

The Driver Model

With the information above we were able to adjust the driver model accordingly, with two changes being required to the steering behavior: The first was to modify the standard

model parameters so that the steering correction was not dependent on the distance the vehicle had deviated from its intended path (tolerated deviation is zero). The second was to enforce a delay so that the driver model did not immediately react to the torque delta. That is, explicitly adding a delay of 0.3s before the driver model reacts. Similar modifications were then made to the driver's accelerator behavior.

The monitoring subroutine was then modified to include the driver parameters from the parameters listed in Table 1 and the yaw rate and lateral acceleration limits from [26]. This allows a snapshot of driver and vehicle performance to be captured for the duration of the torque disturbance. Prior to the torque delta being applied, the monitoring subroutine latches all values from the time step before the disturbance is applied. Over the course of the torque deviation event, the maximum deviations from the initial conditions are logged. If the defined limits (e.g. yaw rate) are exceeded then a failure is logged.

Driving Scenarios

The base set of driving scenarios examined comprise three stable curves: 0.0g (i.e. a straight line) 0.4g and 0.6g; with 0.4g being selected rather than 0.3g as it covers all “normal” driving, with an amount of leeway. The above scenarios were expanded to consider road gradients (uphill and downhill) of between 0 and 20¹ degrees. Each scenario was then run at the same speeds at which Neukum *et al.* [26] established yaw rate limits.

Two different approaches were used to test the vehicle. The first approach involved reducing the torque on a single motor to zero, which is representative of a failure such as the loss of traction power. The second approach introduced a torque offset of between ± 200 and ± 700 Nm in increments of 100Nm. Additional work (i.e. binary search) was then performed to more accurately determine the point at which any of the constraints failed.

In theory, to perform all test runs would require approximately 1700 tests. However, in practice we curtailed test runs when failures were observed, or when the driving situation was clearly unrealistic as judged by watching the animated vehicle. This significantly reduced the number of simulation runs performed.

Vehicle Simulations

Vehicle simulations with the Carmaker tool were essentially performed in two separate passes. The first pass was essentially a full pass though the set of scenarios as defined above, using the standard CarMaker vehicle as described. The second pass was initially intended as a check that the results of the first pass were “solid” and to extend the work to different situations, most notably to low speed scenarios not investigate in the first round.

¹Baldwin St. in Dunedin, New Zealand has a slope of 35% (19 degrees).

First Pass

Reprise: the purpose of performing the simulations was to establish the relationship between a torque disturbance resulting from a fault in the motor and the performance criteria listed in [Table 1](#) and [Table 2](#); in particular the yaw rates and lateral acceleration criteria given in Neukum *et al.* [26].

In the event it was only the yaw rate limit criteria that was found to be of significance. None of the other criteria came close to being exceeded. There was a correlation between steering wheel angle and yaw found, but the steering wheel angles involved were relatively small; typically 20 to 30 degrees depending on the scenario.

In any test of this kind, determining where to set the failure point is problematic. Should one apply the worst case or is it better to use some other criteria? At this point we chose to use the worst case value. This resulted in a delta torque of 280Nm to meet the C2 criteria developed by Neukum *et al.* [10].

Second Pass

On completion of the first pass it had been our intention to revisit the problem for a number of reasons: the primary one being that by simply selecting the worst case was not very satisfactory. While “safe”, it represents a single point on a multi-dimensional surface, whose shape was not clearly defined and was potentially unnecessarily conservative.

Reaction Times

One aspect that we wanted to examine was the effect of different reaction times on the scenario outcome, that is, how dependent was the controllability of the vehicle on the 0.3s reaction time?

Over the three speeds tested, using values of between 0.1s and 0.6s had no significant effect on the outcome observed. That is while there were small increases in the yaw rates and lateral accelerations observed, the delta torque values where the vehicle exceeded the limits were not significantly different.

Low Speeds

In previous simulations the lowest speed investigated was 50 km/h. A small series of simulations were conducted to examine vehicle behavior at low speeds, with test being run at 40, 30, 20 and 10 km/h.

As expected, the results were that as speed decreased, induced yaw was supplanted by acceleration as the dominant vehicle level effect. This is true for all but the highest torque delta values, i.e. 600 Nm or greater. Limited amounts of track testing in the EV1 vehicle ([Figure 2](#)) support these results.

Internal Consistency

As a check we also ran the modified vehicle model against the generic vehicle model for two different scenarios; wind gusts and braking. Here we were checking to see if any

changes we had made to the driver model significantly affected vehicle behavior.

CarMaker's ability to generate wind gusts made a back-to-back comparison of the driver/vehicle behavior possible for the standard models and the modified models; excluding effects due to torque disturbances at the wheel. While there were small differences in the level of induced yaw performance of the two systems was similar.

The driver model was also tested against the performance of the standard vehicle under braking on a single wheel. It was found that while there is a small shift in the torque delta where the standard model exceeds the Neukum limits, the data from runs where the standard vehicle brakes were used showed the same behavioral pattern.

Reevaluation

While the simulation model was being rebuilt the focus for the target vehicle on which the in-wheel motor would/could be used changed. With the donor vehicle for the first production intent motors having changed, the primary focus of this round of simulation also changed. The question moved from one of confirming and building on the previous results, to evaluating whether the proposed changes to the in-wheel motor architecture would remain safe.

To this end the evaluation stopped being about whether or not the motor could be made “safe” for any generic vehicle above the mass and dimensions of the default CarMaker vehicle, to whether the motor in its new form would be “safe” for a particular sub-set of vehicles that would likely be the basis for early deployment on the road.



Figure 2. The Brabus 4WD vehicle (EV1) fitted with Protean in-wheel Motors during fault injection testing (Photo D. Harty)

[Figure 3](#) shows how changing vehicle parameters effects the behavior of the vehicle in relation to the yaw induced, for any particular value of delta torque on a wheel at 100 km/h. The following are shown in descending order:

- Grey squares are data points from EV1 track testing.
- The simple dashed line at the top is simulation data for the standard CarMaker vehicle.

