ISO 26262: EXPERIENCE APPLYING PART 3 TO AN IN-WHEEL ELECTRIC MOTOR

M. Ellims, H. Monkhouse, A. Lyon

Protean Electric Ltd, UK, Mike.Ellims/Helen.Monkhouse@proteanelectric.com

Keywords: hazard analysis, risk analysis, functional safety, functional safety concept.

Abstract

This paper presents a case study that applies ISO 26262 Part 3 to the hazard analysis of an in-wheel electric motor. It describes the activities undertaken, their mapping onto the Standard, and discusses the limits and strengths of the analysis and possible alternative approaches.

1 Introduction

ISO 26262 Part 3 [3] mandates a process for evaluating the functional hazards associated with electrical and electronic systems and components fitted to road going vehicles of up to 3500 kg. While the processes detailed within are perhaps well suited to the needs and capabilities of large vehicle manufacturers and established suppliers, the capacity of smaller organisations to apply ISO 26262 and to bring to market a novel device is less well understood. In this paper the authors document the actual process used to develop the Functional Safety Concept (FSC) for an in-wheel electric motor and compares this with the idealised process presented in ISO 26262. We also compare our process with the options that are possibly available to larger organisations.

Using the terminology defined in ISO 26262, the "item" under analysis is an in-wheel electric motor capable of generating over 800Nm of torque for extended periods of time. Unlike a number of similar systems this device incorporates all the high voltage and control electronics within the hub of the wheel, outboard of the vehicle's suspension. It should be noted that even though the development is not being performed by a major OEM the analysis we perform is of an "item" rather than a "Safety Element out of Context" (SEooC). This approach was chosen because although the "item" is not targeted at a specific vehicle programme, it does directly influence the safety of the vehicle; as the item combines elements of a drive-by-wire engine, a brake system and a differential. Indeed it is difficult to conceive how an analysis of the "item" could be performed out of context.

2 Overview of Part 3

Part 3 of ISO 256262 is divided into four major areas: Item Definition, Process Initiation, Hazard and Risk Assessment and Functional Safety Concept.

Item Definition involves defining the attributes of the item under development. Where attributes considered include: functional and non-functional requirements, interactions between the item and other systems, and any assumptions being made.

Process Initiation is essentially deciding whether the item is a new development or a modification to an existing item, previously developed to ISO 26262, and imposes requirements on downstream process.
Hazard and Risk Analysis defines the process for evaluating the Automotive Safety Integrity Level (ASIL) of the item being developed. Of particular interest here, is the defined process and associated technique mandated for assessing and assigning risk, which involves three primary parameters: exposure, controllability and the severity of the outcome.

The Functional Safety Concept involves stating the safety goals for the item and defining the associated safety requirements.

This paper focuses on these last two areas, the hazard analysis and the development of the FSC.

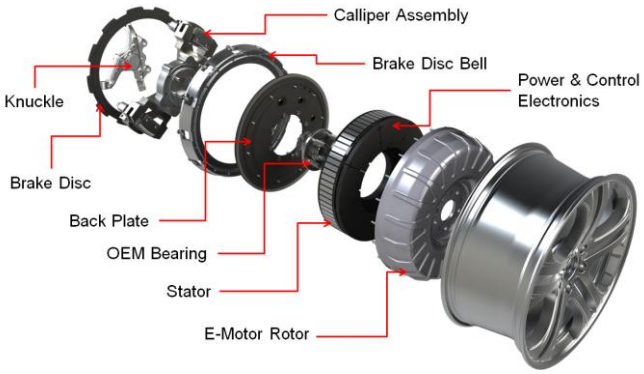

Figure 1: exploded view of the Protean Electric in-wheel motor.

3 Initial Investigations

For the first stages of the safety analysis the item definition was effectively the description of the existing development

motors. The main failure modes of the items were obvious: the motor can provide more torque than commanded or it can provide less torque than commanded, both when acting as a motor and when acting as a generator (i.e. braking).

Given this, the main question then became, "what are the vehicle level effects that will be induced by a failure?" Informal discussions concluded that the major effect would be to induce a yaw moment on a moving vehicle. However, not all yaw moments are created equal, and yaw moments can be induced on a vehicle from a wide variety of external sources: wind gusts, standing water, potholes, tyre deflation etc. Normally a driver can cope with these provided that the magnitude of the disturbance is limited.

Thus the initial question "what are the vehicle level effects" becomes, "what magnitude of vehicle level effects will be seen?". To address this modified question a limited amount of vehicle level simulation was undertaken. At the same time the existing motor design was subject to HAZOP [4] analysis as part of the process to define the next generation motor architecture.

3.1 Simulation

Simulation of vehicle behaviour was performed using the IPG CarMaker [1] suite of tools, using the driver model provided but replacing the built in steer-by-angle control with a custom steer-by-torque algorithm. The simulation also required the vehicle be modified; removing the standard inboard powertrain model and replacing it with four individual motors driving directly at the wheels. This necessitated including models of the motors and their control systems directly into the standard CarMaker vehicle and removing existing standard functionality where necessary.

Tests were performed at 50 kph in a straight line and with turns of 0.4g and 0.8g lateral acceleration with delta torques of +/- 425 and +/- 850 Nm. The intention was to perform the tests at higher speeds, but it was evident that although in very few cases the vehicle did depart from the lane (the failure criteria), there were issues with the approach. For example the 0.8g turn is near the limit of the vehicle's capabilities. In addition +/-850Nm delta values applied to the wheel often resulted in a total loss of grip with little or no effect observable at the vehicle level.

The major insights gained from these simulations were:

- The initial corrective action needs to take place within 1.5 seconds.
- There was a lack of fidelity in the delta torques used.
- The 0.8g turn was too close to the ultimate vehicle limits to provide useful information.
- The criteria for determining whether a control failure occurred needed to be tightened; with the lane exit on its own being too coarse a measure.

3.2 HAZOP

The initial HAZOP process was applied to the motor as a single unit and acting as an electric motor or a generator. The boundary of the item was taken to be at its interface to the outside world i.e. including the serial communications and power supplies for the power electronics. However, excluded from this preliminary analysis were the power supplies for digital electronics and the hardware disable line as these were known to have issues.

The primary aim of the HAZOP exercise was to record the informal analysis of the existing motor that had taken place within the motor design group meetings. The discussion here took the form of a brainstorming exercise and while outcomes and reasoning were captured in the minutes of the meetings, information was not systematically organised.

The other main outcome of the HAZOP was to explicitly state the failure modes and evaluate the vehicle level effects of potential failures.

3.3 Initial SIL

Intertwined with the activities described above, an initial estimate of the Safety Integrity Level (SIL) for the vehicle was performed using both the method proposed by MISRA [5] and that specified in ISO 26262 Part 3 [3]. The primary assumption that was made here regarded the "controllability" of a failure; with the control options available to the driver of the vehicle being limited to steering and braking.

Application of both risk classification techniques arrived at basically the same answer, that the item in foreseeable cases was SIL 3 or ASIL D. The information that leads to these results is consistent. For exposure both risk assignment approaches [3, 5] result in the highest possible ranking and the same is true for severity. For example, consider a single carriageway "A" road. If the failure is an induced yaw then a possible outcome is a head-on collision with another vehicle. Given a potential delta V in excess of 190 kph it can reasonable to expect that there would be fatalities.

It was more problematic to establish a "reliable" metric for the controllability parameter. Where the MISRA guidelines [5] provide a *relatively* straight forward assessment procedure, ISO 26262 provides only a small set of generic examples and information on the assessment process for category C2. From Table B4 in Annex B [3]: the closest appropriate control behaviour in response to a yaw disturbance is "maintains path", which is ranked C2 for "Motor failure at high lateral acceleration". In this case the controllability parameters are in agreement. What is not clear is whether that would be true in general.

3.4 Analysis

At this stage one thing was clear, and that was that nothing was completely clear! The one exception being that the motor has a high ASIL level, which is consistent with limited information gleaned on other systems such as electronic braking [9], torque vectoring differentials etc.

Simulation work suggested that a motor failure was not catastrophic, i.e. only a limited number of situations resulted in a vehicle leaving the lane and those occurred at high torque

level. Simulations also indicated that extreme driver responses did not appear to be required to maintain control. Likewise the limited HAZOP analysis indicated that while prototype motors have undesirable failure modes these could potentially be designed out.

4 Continued Investigation

The first major output from the initial evaluations described above was the creation of an Item Definition. This is the document that defines (broadly) the device and its interfaces that will be analysed and built.

Based on this, the analysis work outlined above was carried forward on a number of different fronts. First a much better understanding of what the driver was capable of and likely to do was required in order to gain more confidence in assigning controllability ratings. Second, a more comprehensive set of vehicle simulations were needed to a) evaluate the ability of drivers to deal with induced yaw and b) to quantify the maximum torque error that could be tolerated. Third, the Item Definition needed to be reanalysed for potential faults and associated failures. Lastly, information from the above needed to be woven into a Functional Safety Concept. There was also the minor complication of ensuring that legal requirements were fully considered in the FSC.

4.1 Legal requirements

The general perception is that the automotive industry compared with, for example, the aerospace industry works in a relatively regulation free environment. This is not the case. Currently there are 126 automotive regulations published under the UN Economic Commission for Europe (UNECE) and a further eleven Global Technical Regulations (GTR). Most of these are matched by similar Federal Motor Vehicle Safety Standards (FMVSS) in North America.

The relationship between ISO 26262 and regulation is interesting. On reading the brake regulations (Regulation 13H) [7] it is clear that the existence of Annex 8, covering complex electronic vehicle control systems, forms a major driver for developing the ISO 26262 standard. This relationship is two way, with clause 7.4.2.8 of Part 3 [3] stating that "Class C0 may also be assigned if dedicated regulations exist that specify the functional performance with respect to a defined hazard". Clause 7.4.2.8 goes on to state that if under these conditions C0 is selected, then "no ASIL assignment is required".

All legislative requirements needed to be examined and those found to be applicable were analysed to identify the functional and procedural requirements to be incorporated in the FSC. An interesting aspect of this analysis is that it is not only about finding technical requirements that apply directly to an item, but also about developing the context in which the item will be embedded.

4.2 Human factors

In section 3.3 it was noted that consideration of driver actions was limited to steering and braking; functions with which drivers could reasonably be expected to be conversant with.
An examination of the literature on driver braking suggested that it would not be a viable mitigation on its own. This is primarily because of the relatively long times involved; with brake reaction times ranging from 0.5 of a second to over two seconds and total braking time ranging from 0.5 to over six seconds. Expected (as opposed to possible) reaction times are especially slow; Young and Stanton [18] suggesting that two seconds is appropriate if drivers are inattentive to external stimuli and Triggs and Harris [16] suggest that "response times can be expected to exceed the commonly accepted design value of 2.5s relatively frequently". Given the constraint of 1.5 seconds given in section 3.1, application of brakes was not considered a viable primary control mechanism.

A limited survey of the steering literature proved more fruitful. Here literature suggests that response to steering disturbances is both natural [6, 17] and quick with reaction times on the order of 0.3s [14, 15, 17]. Of particular interest was the work of Neukum *et al.* [14] on steering superposition errors which gives yaw rate and lateral acceleration limits for which the driver can be expected to maintain control.
More problematic was locating publically available data that defines "normal" driving behaviour. Only a single paper by Lechner and Perrin [12] directly addresses the issue, though literature on drive cycle development provides useful information as does a National Highway Traffic Safety Administration (NHTSA) study on real world driver behaviour [11]. The NHTSA report also demonstrates that high acceleration manoeuvres while not common, need to be considered. These sets of information broadly support each other and are consistent with proprietary information the authors are aware of.

A number of other constraints on driver performance were also taken into consideration, for example maximum hand wheel velocity and acceleration limits, limits on maximum torque that could be applied by the driver etc. This set of human factor criteria, which together with the lane exit criteria formed the basis for determining whether the virtual driver had "lost control" during simulation.

4.3 Further simulation

The human factors work described above was the primary input into a second, more detailed round of vehicle simulation where the information gathered on how the driver could be expected to react and limits on the mitigating actions were explicitly built into a monitoring function within the vehicle model to identify driver control failures.

Information on expected normal driving behaviour was built into the scenarios that were evaluated during simulation. Driving scenarios included constant speed, acceleration and deceleration, on straights and curves at 0.4g and 0.6g, with each scenario repeated at speeds of 50, 100 and 150 kph.

These parameters were expected to encompass the majority of actual driving.

The simulation runs were performed in two phases. During the first phase of simulations a single motor was forced to fail silent (zero torque), from various torque levels. This proved unsatisfactory for a number of reasons. Foremost of which was the necessity of driving on slopes to force the motor to produce sufficient torque to achieve the desired delta torque.
In the second round of simulations rather than forcing a motor to zero torque, a delta torque value was imposed on it. These simulations were then performed using the same set of scenarios as the fail silent tests. The results were then collated from both sets of simulations to estimate the maximum torque delta which could be tolerated without a control failure occurring.

In addition to the two major threads described above, a number of minor investigations were undertaken on a more ad-hoc basis. Examples of this type of investigation included: examining the effects of different drive configurations i.e. four-wheel drive vs. front wheel drive; comparison of different tyre models and the effect of weight distribution and so on; the main purpose being to examine the sensitivity of the simulation model.

The major output from the work was an estimate of the delta torque at a wheel that a vehicle could tolerate before the limits on driver capability were exceeded. This establishes the maximum permissible size (in terms of torque) that a fault, or faults within a motor can produce while remaining "normally controllable" at a vehicle level. This torque limit in turn directly influenced the architecture of the motors internal electronics and interface.

4.4 Item definition HAZOP

The HAZOP work outlined in section 3.2 was repeated at two different levels. The highest level considered the motor to be a single torque producing device; effectively ignoring all the design decisions made in order to reduce the likelihood of a single failure causing the maximum torque delta threshold being exceeded. This established the major failure modes that would likely be present. This process benefited from being able to observe and evaluate the major failure modes during simulation.

A second pass examined faults that could be induced at the interface to the motor, e.g. the serial communications links (CAN), the power supply for the digital electronics (12V), hardware enable lines, high voltage supply and so on. This second phase matches the HAZOP activity described in section 3.2. A specific activity performed here was to explicitly merge the two sets of requirements developed from performing the HAZOP at the interface level. Not unexpectedly the overlap between the two sets was not complete but there no major omissions found.

A potential weakness of the two initial phases is that they focus on the physical realisation of the motor and the analysis process requires the mapping from artefact to function to be performed as each interface is considered. A potential downside of this is that interactions between interface elements could be missed. The obvious next step is to repeat the analysis, this time looking at detailed functions delivered by the motor. To date this has not been performed as a HAZOP due in part to resource and timing constrains, but rather has been incorporated into a system Design Failure Mode and Effects Analysis (DFMEA) activity.

The work described to this point is encompassed by the first two sub-clauses of 7.4.2.2 "Hazard identification" [3] which comprises only nine lines of text. For less novel systems, i.e. those already in widespread use, this part of the hazard identification activity could reasonably be expected to be less involved.

4.5 Scenario identification

The second part of clause 7.4.2.2 [3] involves determining what hazardous events can result from the hazards in "relevant combinations of operational situations", which the standard defines as a "scenario that can occur during a vehicle's life".

This immediately introduces two problems: firstly what is a relevant situation and secondly, almost any scenario could occur during a vehicles life.

The problem was approached in the following way. For driving situations the Road Accident Data [8] was used as a starting point. Situations were built up by starting with the road class (motorway, A road etc.) and extended for both the road type (e.g. roundabout, slip road etc.) and the junction detail. Modifiers were also considered for traffic condition (light, normal, heavy, crawling, surging etc.). Modifiers identified, but not fully applied include the driving activity being undertaken (accelerating, braking) and road surface conditions (dry, wet, snow etc.). This set in turn was correlated with information in Annex B of Part 3 [3] to construct a set of unified exposure ratings.

Also considered at this stage were the "actors" that could be involved with a vehicle in a scenario. The inclusion of the driver is obligatory. Passengers are obvious, as are occupants of other vehicles and other road users. Perhaps less obvious are persons who become involved after an accident i.e. emergency service personnel and good Samaritans, along with service and maintenance personnel.

To evaluate all scenarios derived from the situations, modifiers and actors identified above would result in an enormous number of evaluations being performed. What ISO 26262 *requires* is unclear as there is no specific guidance on when the hazard classification process should or could be terminated. However clues can be gleaned from the objectives of clause 7 that the purpose of the analysis is:

- to identify hazards,
- to categorise hazards, and
- to define safety goals to prevent or mitigate hazardous events

Thus what we consider to be a workable solution was to terminate the process when it was concluded that all three

high level objectives had been met. Thus it is not adequate to simply state that "there is a situation where the item is ASIL X"; where X is the highest ASIL reasonably expected. Rather it is a case of determining whether all safety goals have been discovered. Simply, the termination criteria for this analysis appear to be the verification criteria defined in clause 7.4.5.

The distinction made in the previous paragraph is critical. The safety goals and requirements derived from them are what ensures that hazards are as far as possible removed from the item by design. In contrast, the ASIL associated with each of the safety goals are requirements on the design rigour necessary to ensure those functions are delivered reliably.

On a more practical level, to deal with the potential volume of scenarios, the scenarios were divided into two primary categories of moving and stationary vehicles.

For moving vehicles road classes were listed and expanded by inserting road types and traffic conditions. In addition each of these scenarios was further split into two classes, one where failures would be normally controllable and one where a failure is not expected to be controllable (i.e. C3) using the torque limits derived from the simulation work.

4.6 Hazard classification: risk analysis

The statement in the previous section about the "highest ASIL reasonably expected" may seem odd; however in section 3.3 we noted that one of the first activities performed was to estimate the SIL using the process defined in [5]. Thus the approximate ASIL expected can reasonably be obtained before the ISO 26262 compliant hazard classification process is performed in detail.

The hazard classification process defined in Part 3 is straight forward, if somewhat involved and in places somewhat subjective.

In ISO 26262 hazard classification has three primary components, exposure to a scenario, controllability of the situation and severity of the outcome. As defined in ISO 26262 each of these components has issues associated with them.

Exposure: for moving vehicle exposure the rating can be estimated from the information provided in Annex B [3]. However, for stationary vehicles only minimal guidance is present which is mostly related to drivers and driving situations. What for example is the exposure rating for Bob the mechanic in the workshop with the large spanner? Annex B suggests that a classification of E1 or E2 is appropriate based on operating time for "the vehicle". However, mechanics don't work on single vehicles (working on vehicles is their profession), consequently we have assumed that the exposure rating for this "actor" is E4.

Controllability: from the simulation studies we have *a priori* knowledge that the magnitude of the failure, in terms of torque delta, influences the control retained by the driver on the moving vehicle and the expected controllability rating [14]. This influences the safe states that are required for different levels of failure and the safety goals required to reach those states. In this situation it is relatively straight forward to assign the controllability rating. However, in general controllability ratings as defined in ISO 26262 cannot be reliably established for any class except C2 and C3. The only feasible option is use the suggested procedure for establishing C2 to disprove a classification of C0 or C1.

Severity: this component has been divided into four categories according to severity estimated relative to the Abbreviated Injury Scale (AIS) [10]. If however you browse the accident investigation literature fewer divisions are usually used; typically only two, AIS 0+ and AIS2+ or 3+. Of the severity examples presented in Annex B [3], it is stated that "no generally valid conclusions can be derived" and that "accident statistics can be used to determine the distribution of injuries that can be expected to occur in different types of accidents". However the complication here is that accident statics are generally compiled using a different scale (e.g. STATS19 is used in the UK) and mapping between scales is at best problematic [13].

4.7 Functional safety concept

Development of the FSC from the safety goals formulated during the Scenario Identification and Hazard Classification process is not particularly straight forward. A major part of the complexity (possibly self inflicted) is that the FSC "should" be represented as a tree or graph that maps safety goals to functional safety requirements in a process of hierarchical decomposition as shown in Figure 2 of clause 8 [3].

For the wheel motor this logical structure was represented as a tree within a spreadsheet, where high level goals were decomposed into sub-goals and then into functional safety requirements. The hierarchy of this structure was created following the derivation of the safety requirements. This process is quite straight forward if only the top level safety goals and requirements directly derived from those goals are considered.

The complexity we discovered arises from the fact that not all functional safety requirements are derived via this mechanism. In our case functional safety requirements were also derived from legislative requirements and from the customer specification (developed in house), which is not formally considered as an input by ISO 26262. Safety requirements are also found in the item definition and were naturally derived as part of the hazard analysis (HAZOP). One interesting facet of this process is that while all the requirements are functional safety requirement, many are also what ISO 26262 refers to as technical safety requirements, most notably those derived from the HAZOP activities.

Because the tree structure was explicitly encoded in the spreadsheet it was also possible to display the data (with a little programming) as a tree in a PDF document more or less as shown in ISO 26262. This proved valuable for determining whether the organisation of the FSC was rational and allowed a pictorial view of the safety requirements to be displayed. However given the size of the FSC tree it requires an A0

plotter to allow the printed tree to be read. The general outline can be seen in figure 2 where the first two layers of safety goal 07 are shown; the full tree having seven layers.

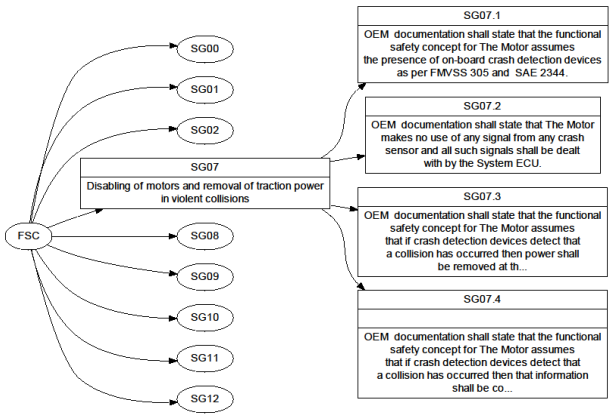

Figure 2: a small section of the FSC tree showing the first two levels for SG07.

The top level safety goals associated with a moving vehicle are as follows;

- SG00: no single fault shall prevent the driver from retaining control.
- SG01: the types of vehicle and any external equipment required will be documented.
- SG02: external failures that cannot be mitigated by the motor shall be protected by external means.

Safety goal SG02 is required because there are failures which cannot be controlled by the item, such as the provision of traction power.

The second part of the tree comprises the safety gaols associated with stationary vehicle and motor;

- SG07: making the motors safe if a vehicle crash event is detected.
- SG08: prevention of unintended vehicle movement or wheel rotation.
- SG09: minimising expose to risk during low speed manoeuvres.
- SG10: protection of susceptible persons from strong magnetic fields.
- SG11: provision of end user information.

The purpose of most of the safety goals listed above should be obvious, however the need for SG01 and SG11 are perhaps less so. These two safety goals are directly aimed at communicating information to the OEM about what assumptions we have when conceiving the system and this is discussed further in section 5.2.

4.8 Limits of Simulation

The simulation work conducted to date has some obvious limitations that should be stated. Perhaps the most obvious is that only the driver's initial reaction, i.e. the corrective steering, has been considered. In the real world, a "normal" driver is unlikely to try to maintain speed and would probably take further actions (e.g. reduce speed) depending on the severity and the disturbance.

A further restriction which became apparent was a lack of fidelity in the modelled driver behaviour, particularly in terms of reaction and muscular/torsional response. It became clear that the level of human behaviour response required in order to accurately determine "controllability" was beyond the capabilities of the existing models and is possibly something which had not been attempted before. This necessitated a significant amount of work to understand the limitations in the simulated driver behaviour and to adapt the analysis of the results accordingly. This is an area that offers much scope for improvement.

Simulations are also limited in the range of situations examined; the most notable omission being the lack of any stationary or low speed work.

It is our intention to carry forward the simulation work by generating faults directly within the MatLab/Simulink model comprising the in-wheel powertrain and CarMaker vehicle models. It is then intended to "drive" the vehicle around virtual test tracks to obtain data from a larger range of situations.

This will be a large undertaking as it necessitates a more rigorous integration of the wheel motors into the existent vehicle model. Other areas that will require attention include more detailed integration of electric and friction brakes, more accurate control behaviour of the wheel motors and automated evaluation of success/failure criteria that does not cause the model to be prematurely halted.

In addition, in the near future we expect to start the limited validation, or correlation of the simulation results using a vehicle on a test track. This will not be a full validation but is rather aimed at establishing correlation data between the simulation models and real life and supports the confirmation of software tools activity.

5 Discussion

5.1 Comparison with the OEM process

A large vehicle manufacturer (OEM) would most probably have taken a different approach for a number of reasons. Firstly they can reasonably be expected to have the required human factors expertise on *driver behaviour* in house. In addition larger OEMs would either have access to their own vehicle simulators or arrangements for access. The net result being that the work to establish what was reasonable driver behaviour would have been shorter. Secondly an OEM could reasonably be expected to proceed to testing the effects of failure on a prototype vehicle more quickly than we have been able to do.

Another area where a large OEM has an advantage is access to supplier systems. A critical system here is the brake control system which embodies both the anti-lock braking (ABS) and the electronic stability program (ESP) functions. The vehicle

we will initially test has neither of these and the functions are not expected to be present on a test vehicle for at least another six months.

However aside from these three fairly major points the overall process applied would probably have been very much the same. The hazard analysis activity would be similar, likewise the scenario identification (but possibly already in existence) and vehicle simulation would have almost certainly been performed prior to vehicle testing.

5.2 Evaluation of ISO 26262

ISO 26262 appears to have the underlying assumption that it will be the OEM who performs the functional safety analysis. If it is not the OEM who performs that analysis, then the other option embodied in the standard is to perform the functional safety activities considering the item as a "Safety Element out of Context". This may be suitable for a high integrity real-time operating system for example, but this is certainly not possible for our motor, as the context in which it is to be used is known as are the vehicle level functions it will provide.

When considering the functional safety analysis of the motor, our approach to assume the role of the OEM, seems to be the only workable option. The down side of this approach is that we now have a number of assumptions about the environment in which the motor operates that need to be discharged by the OEM when fitting the motors to a vehicle. This in turn implies a need to communicate those assumptions to the OEM so that they know of and can discharge the associated obligations. This on its own makes up a substantial part of the FSC and in no way removes our obligation to ensure that the assumptions can in fact be discharged. Note that while ISO 26262 requires that assumptions be documented, there seem to be no explicit requirements to discharge those assumptions.

ISO 26262 also assumes a linear flow of activities. That is not usually possible in the real world and the activities described in this paper have overlapped at all stages as our knowledge of the system and the process has grown. Strictly speaking this implies that we do not meet the requirements of ISO 26262 but in the development of a novel item this is unavoidable.

From the discussion on hazard classification we find several general issues. Perhaps most obvious is a potential issue with the calibration of the risk matrix as required by IEC 61508 [2]. Currently the derivation of the calibration process could perhaps be at best described as "opaque".

Another issue that arises comes from the necessity to derive severity parameters from accident statistics. Aside from the issue where different scales are used by different groups, there appears to be an un-intended interaction between the use of accident statistics and what external mitigations can be applied. For example guardrails and air-bags are mentioned as examples of external measures in Part 10. However, these examples will already have been factored into the severity rating via accident statistics. The unintended effect of this being that it may be possible to build in credit for an external mitigation twice. Our solution has been to ignore any mitigation that does not directly affect the vehicle dynamics.
However the authors' view of ISO 26262 is not all negative. At first glance the requirement to perform the semi-formal scenario identification activity seemed unduly pedantic, especially given that no explicate stopping criteria are provided. However in practice the exercise has proved reasonably worthwhile. Not as might be expected because of the associated hazard classification process, but rather because it forced us to consider just what may comprise a scenario and what (if anything) would be the safe state associated with it.

The "requirement" that the FSC be derived as a directed graph, though not formally stated as such, also produced some unexpected benefits. As stated above, it forced us to more carefully consider the structure of the requirements and gave us a picture of their organisation. This, to a small extent allowed some holes to be identified.

As well as identifying a few holes, the directed graph also forms the beginnings of an "evidence tree" for the final safety case. For some sub-goals it was obvious that there were requirements on the provision of supporting evidence that needed to be provided to satisfy the top level goal. As a safety argument represented in Goal Structured Notation (GSN) this "evidence tree" could be thought of as part of the "solution" that supports a "strategy" that argues that safety obligations have been effectively communicated and discharged. To date components of this "solution" have been incorporated in the FSC tree directly, but no detailed work to develop this into GSN for the safety case has been carried out.

5.3 The relationship to regulations

In section 4.1 it was noted that for hazards where dedicated regulation specify the functional performance with respect to that hazard, ISO 26262 allows a C0 controllability rating to be assumed. What is less clear is how to deal with the safety requirements supporting a safety goal that has an unassigned ASIL rating. Should the design rigour applied to these requirements be for that of a safety goal assigned ASIL QM or higher? To bypass this question we have opted to assign a specific ASIL of NC (not classified) to explicitly distinguish these from QM rated requirements.

6 Conclusions

From the outset it was clear that developing a safety related product, such as the in-wheel motor, to meet the requirements of ISO 26262 was going to be a challenge for a small engineering team and a number of unexpected issues only served to exacerbate this. Ambiguities in the pre-released standard and a lack of unified industrial opinion on its interpretation made determining the correct direction for activities a more complex process than anticipated. Limitations in supporting knowledge (such as normal and expected driver behaviour) and the tools used became apparent. Attempting to incorporate the driver as part of the control loop within the vehicle simulation, in order to

evaluate the controllability of failures, highlighted 'short-falls' in the existing modelling tools. Finally, such a small team prevent any significant level of independence from being brought to bear without the use of external and costly consultancy.

Despite all of these points, Part 3 of ISO 26262 was successfully applied at a level which, on the whole helped the product design process rather than hindered it. With the experience gained throughout this process and the ongoing improvement in industrial expertise we feel that repeating this process for a new product will be a much quicker and less frustrating exercise. Despite its FDIS status, ISO 26262 has a number of issues remaining and its application is clearly not ideal for small engineering suppliers. That said, this project has shown that its use can be adapted with reasonable success to even the smallest business.

Acknowledgements

Thanks to Prof. N. Leveson for an interesting lunch time discussion on the philosophies of system safety standards.

References

[1] CarMaker, http://www.ipg.de/ accessed 28 June 2011.
[2] IEC 61508-5, Part 5: Examples of methods for the determination of safety integrity levels. Edition 2.0
[3] ISO/CD 26262-3, Road vehicles — Functional safety — Part 3: Concept phase, 2008-02-29
[4] Ministry of Defence, "Defence Standard 00-58 HAZOP Studies on Systems Containing Programmable Electronics", Issue 2, 19 May 2000.
[5] MISRA "Guidelines for Safety Analysis of Vehicle Bases Programmable Systems", Motor Industry Software Reliability Association, Nov. 2007.
[6] MISRA "The use of Controllability for the Classification of Automotive Vehicle Hazards" MISRA Technical Report, Version 1, Jan. 2007.
[7] Regulation No 13-H of the Economic Commission for Europe of the United Nations (UN/ECE) — Uniform provisions concerning the approval of passenger cars with regard to braking.
[8] Road Accident Data – GB, Variables and Values and Export Record Layouts, UK Data Archive Study Number 6254 - Road Accident Data, 2008.
[9] T. Fülep, "Design Methods of Safety Critical Systems and their Application in Electronic Brake Systems", Section 6.4. PhD Thesis, 2007.
[10] T. A. Gennarelli, E. Wodzin, "AIS 2005: A contemporary injury scale", Injury, Volume 37, Issue 12pp. 1083-1091 (2006)
[11] S. G. Klauer, T. A. Dingus, V. L. Neale, J. D. Sudweeks, and D. J. Ramsey, "Comparing Real-World Behaviors of Drivers With High versus Low Rates of Crashes and Near-Crashes", DOT HS 811 091, Feb 2009.
[12] D. Lechner, C. Perrin, "The actual use of the dynamic performances of vehicles" Proceedings of the Institution of Mechanical Engineers. Pt.D. Journal of Automobile Engineering. Vol. 207, no. D4, pp. 249-56. 1993
[13] A. Morris., M. Mackay, E. Wodzin, J. Barnes, "Some Injury Scaling Issues in UK Crash Research", Proc. Ircobi Conf., Lisbon, Portugal Sept. pp. 283–292 (2003)
[14] A. Neukum, E. Ufer, J. Paulig, H-P. Kruger, "Controllability of Superposition Steering System Failures", Steering Tech 2008, Munchen.
[15] P. J. Switkes, J. C. Gerdes,. G. F. Schmidt, M. Kiss, "Driver Response to Steering Torque Disturbances: a User Study an Assisted Lanekeeping", Advances in Automotive Control, Volume 5 Part 1.
[16] T. J. Triggs, W. G. Harris, "Reaction Time of Drivers to Road Stimuli", Human Factors Report No. HFR-12, Monash University, June 1982.
[17] W. W. Wierwille, J. G. Casali, B. S. Repa, "Driver Steering Reaction Time to Abrupt-Onset Crosswinds, as Measured in a Moving-Base Driving Simulator", Human Factors, 1983, 25(1), pg. 103-116.
[18] M. S. Young and N. A. Stanton, "Back to the future: Brake reaction times for manual and automated vehicles", Ergonomics, Vol. 50, No. 1, 15 January 2007, 46–58.