

Safety Analysis of Vehicle-Based Systems

Peter H Jesty & Keith M Hobley
University of Leeds
Leeds, UK

Richard Evans
Rover Group Ltd
Coventry, UK

Ian Kendall
Jaguar Cars Ltd
Coventry, UK

Abstract

The Motor Industry Software Reliability Association Steering Group is producing guidance on the safety analysis of vehicle-based systems to support its original Development Guidelines for Vehicle Based Software. Using existing generic techniques, these new guidelines will explain how they may be used in the automotive context. Topics will include System Analysis, Hazard Identification, Hazard Analysis, the identification of Safety Integrity Levels, and the uses of Failure Mode and Effects Analysis and Fault Tree Analysis.

1 Introduction

In 1994 the Motor Industry Software Reliability Association (MISRA) published its guidelines for the development of software for vehicle-based systems [MISRA 1994]. A key aspect of these guidelines is that the hazards associated with a system must be both understood, and taken into consideration, from the beginning of its design cycle. It is therefore important to:

- assess the risks associated with the behaviour of a system;
- do this early enough in order to take design actions that can reduce those risks to an acceptable level;
- provide documentary evidence of the reasoning that lies behind the design decisions made.

Whilst the scope of the original MISRA Guidelines is software, system issues are covered to the extent that they influenced software development. Eight specific issues are addressed:

- Integrity
- Software in Control Systems
- Noise, EMC and Real-Time
- Diagnostics and Integrated Vehicle Systems
- Software Metrics
- Verification and Validation
- Sub-Contracting of Automotive Software
- Human Factors in Software Development

In the section on integrity the MISRA Guidelines call for Safety Analysis to be performed, which includes Hazard Analysis and Integrity Assessment. Whilst it does not mandate the use of specific techniques or methods for safety analysis it does refer to Preliminary Safety Analysis (PSA), which is part of the 'PASSPORT' methodology [Hobley 1995a]. MISRA has found that there are a number of approaches that can be used depending on the nature and scope of the system being analysed. A good example being a Preliminary Hazard Analysis (PHA) based on a Hazard and Operability (HAZOP) study [Redmill 1999]. The automotive industry has long been in favour of Failure Mode & Effects Analysis (FMEA) as a means to manage the system and hardware risks associated with its products. However, FMEA is not easy to apply before a design exists, and is therefore often under-utilised.

The MISRA Steering Group therefore decided to produce a document to show how these techniques are related, and to provide detailed and practical recommendations for bringing the techniques together. The document, due in 2000, will explain how to approach PSA for 'blue-sky' conceptual systems, and PHA for systems whose basic design can be already assumed. It will not only expand on the recommendations made in the original MISRA Guidelines on the issue of 'integrity' but will also augment the scope to include system and hardware as well as software considerations. The treatment of safety analysis will also be expanded to include Detailed Safety Analysis (DSA) as well as Preliminary Safety Analysis (PSA). The intention is to offer ideas which complement the automotive industry's quality system standards, which include specific requirements for FMEA [QS 9000 1995 and VDA 1996], in such a way as to bring the concepts covered by PASSPORT, and in standards such as [IEC 61508 1999] and [DEFSTAN 00-56 1996], into a single 'MISRA Safety Analysis' approach.

This paper describes the principal contents of the proposed MISRA approach to safety analysis. It is divided into three main sections as follows:

- A description of how PSA can be applied to vehicle system concepts, and of how PHA can be applied to top-level designs;
- A discussion of how Safety Integrity Levels (SIL) can be assigned to vehicle systems, both those that are related to the moving vehicle scenario and those which are not;
- A description of how FMEA and FTA may be applied to automotive systems as part of DSA.

The responsibility for the safety of a vehicle ultimately resides with the manufacturer, which must therefore be free to choose the most appropriate set of techniques for its circumstances. Thus, following the philosophy of the initial MISRA Guidelines [MISRA 1994], the safety analysis document will not mandate the methodology that should be used; instead the properties that should be

demonstrated are given, together with some systematic techniques that may be used to obtain them.

2 Hazard Analysis of Vehicle Systems

Before one can perform a hazard analysis it is necessary to produce an abstraction, or model, of the system on which to base the analysis. By its very nature a model must be an approximation, since the only 'model' that can be identical to the system is the system itself [Carroll 1939]. It is therefore necessary to choose the approximation so that it highlights the features necessary for the particular task. Although some definitions of a hazard limit themselves to physical situations with a potential for human injury, in practice one should include all situations that can threaten people, property and the natural environment; thus any model that is used as a target for a PSA and/or PHA should highlight the boundary between the system under consideration and those 'things' that might suffer the hazard, and the interaction between them. The model must therefore show:

- Components, especially those close to the boundary, with their attributes;
- The interconnections between the components, with their attributes.
- The boundaries of various kinds;

2.1 Components

In its basic physical sense a system is built up with components. The level of detail at which one considers a component will depend on the nature of the work and the type of the system, but it is wise to choose a consistent level at each stage, e.g. all line replaceable units (LRU), or all electronic components.

The attributes of a component include not only its physical properties, but also the function(s) that it performs (for this reason it is sometimes called a *functional element*). Thus the failure modes associated with a component can be physical, e.g. short circuit, or functional, e.g. brakes do not act on the wheels.

The behaviour of a system, however, is usually more than just the sum of the behaviours of its components. A system may often exhibit *emergent properties* which can only be obtained when the components are working together. One consequence of this is that an analysis based on LRUs may not produce the same results as one based solely on the electrical and mechanical components of which it is comprised. This is particularly true when some of the components are programmable.

2.2 Inter-connections

Components work together through their inter-connections. Mechanical components use their connections to pass forces between each other, electrical

components pass signals, and programmable components may pass signals and/or messages (information). The attributes associated with an interconnection are the 'thing' that is being passed, the rate at which it is doing so, and its value, accuracy, phase etc.

2.3 Boundaries

A knowledge of the boundary is essential in order to define the scope of any analysis. There are four different types of 'boundary', not necessarily distinct, which are relevant when performing a safety analysis on vehicle based systems:

- **System Boundary** - This defines the scope of the 'system of interest' to the development team. This is unlikely to be an entire vehicle, but the automotive industry does refer to its major vehicle units as 'systems' rather than 'sub-systems'.
- **Boundary of the Target of Evaluation (TOE)** - This term can be used to define the scope of the item being considered during a specific safety analysis. It will often be coincident with the system boundary, but may only take in a sub-set of the system of interest, e.g. the team developing the power train may wish to analysis the engine and gear box separately. A TOE is unlikely to cross the system boundary.
- **Zone of Responsibility** - This defines the scope of the authority held by the development team, and relates to the degree to which they can control, or influence, changes in both their own design, and in the design of other systems. This zone is likely to be defined according to the business and organisational requirements of the company, e.g. by office, building, department or country. Whilst it may be normal for the development team to have full responsibility for their own system of interest, it is essential that the company has a mechanism in place whereby, should the development team identify changes that are both necessary and outside their own zone of responsibility, then it is possible for this information to be acted upon in a proper manner.

Figure 1 shows a possible relationship between the System Boundary, the Boundary of the TOE and the Zone of Responsibility, and an example is given in Section 2.6.

- **Moving Vehicle Boundary or Vehicle Hazard Boundary** - Whilst a few vehicle systems may be able to cause harm by themselves, most hazards will be related to the motion of the whole vehicle. In this situation it is necessary to consider the vehicle boundary and this is discussed below.




Figure 1 - Relationship between Three of the Boundaries

2.3.1 Vehicle hazard Boundary

This is the interface between the (emergent properties of all the) systems that make up the moving vehicle, and all the things with which the vehicle may interact, or influence, otherwise known as its environment. This environment includes the driver and passengers, the road, road-side furniture, other vehicles and the atmosphere. The mechanisms by which, and the interfaces through which, a vehicle may interact with its environment are varied and some of them are shown in Figure 2.

Figure 2 is not complete, deliberately, because new systems may introduce new interactions, and it is the responsibility of the team performing the hazard analysis to ensure that their version reflects the reality of their situation. The possible interactions with the environment are of three basic types:

- **Inputs** - These relate to all the devices that a driver can use to provide control signals to a system in the vehicle, e.g. throttle, brake pedal, switches;
- **Outputs** - These relate to all the displays and warning systems for the driver, e.g. speedometer, ABS failure lamp, vision enhancement systems;
- **Physical properties** - These relate to the nature of the physical materials used, e.g. flammability, sharpness, and the emissions that are expelled from the vehicle, e.g. exhaust, compressed air;

There are two main classes of emergent properties that are due to the combined effect of a number of individual systems:

- **Movement** - These relate to the basic longitudinal and lateral movements of the vehicle, e.g. acceleration, deceleration, and steering;
- **Stability** - These interactions relate to the ease with which a driver can control the Movement, e.g. yaw, pitch and roll;

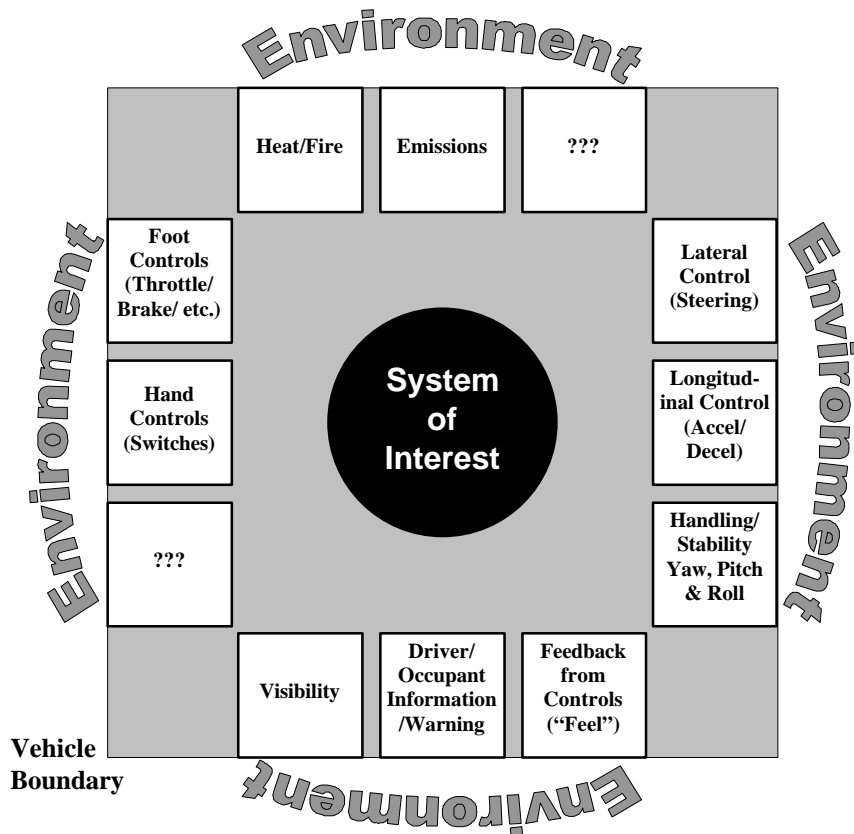


Figure 2 - Interactions and Emergent Properties at the Vehicle Hazard Boundary

2.4 Choice of Model

The choice as to which model should be used depends both on the type of system under investigation, and on how much is known at the concept stage of the design. There are two basic types of TOE associated with vehicles:

- A TOE with interfaces to people and/or other vehicles etc., e.g. autonomous cruise control with radar detection (see Figure 3a);
- A TOE with very well defined interfaces to other vehicle systems (see Figure 3b).

These different types of TOE tend to require different approaches to modelling at their respective boundaries. Figure 4 shows the situation at the boundary of any TOE. The TOE interacts with its environment through its Boundary Elements; these are situated immediately inside the boundary of the TOE. The 'thing'

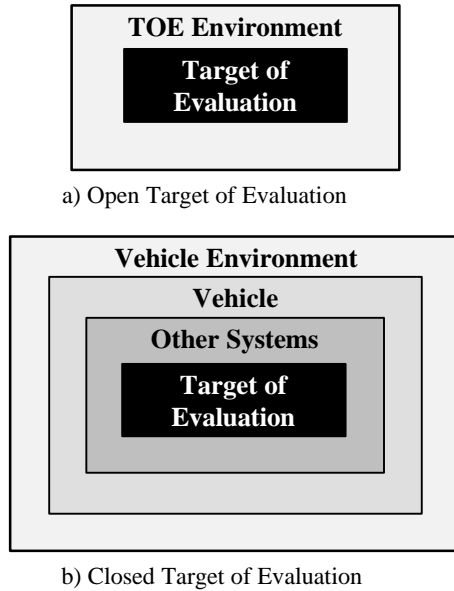


Figure 3 - Different types of TOE

immediately outside the boundary of the TOE, and with which the Boundary Element interacts, is a Terminator.




Figure 4 - The Boundary of the TOE

For systems of the form shown in Figure 3a the Boundary Element tends to be of three basic types:

- If the TOE is communicating with another system X, then the Boundary Element will normally be an 'Interface to X', or its name;
- If the Terminator is a person the Boundary Element will normally be the human machine interaction device;
- If the TOE is a control system then the Boundary Elements will normally include sensors and actuators.

Note that for systems of the form shown in Figure 3a the Terminators may not be well defined, however, for those shown in Figure 3b both the Terminators and the Boundary Elements are normally well defined and understood.

2.5 Preliminary Safety Analysis

For TOEs of the form shown in Figure 3a the PASSPORT methodology for PSA has been shown to be particularly effective as a systematic way of identifying hazards at the concept phase of the design [Hobley 1995b]. Hazard analysis is performed as part of PSA, along with system analysis and system decomposition, to the extent necessary to be able to allocate SILs. The recommended model is known as the PASSPORT Diagram, a very simplified version of which is shown in Figure 5. Note that the model only includes items that are situated within the boundary of the TOE, the description of the Terminators will be included in the textual description of each Boundary Element that must accompany the Diagram. Note also that the 'Data' is that which passes from the Boundary Element to/from the Kernel of the TOE, and not the data that passes from the Boundary Element to/from the Terminator, which may be different in nature due to some transformation that may take place inside the Boundary Element.



Figure 5 - Simplified Outline PASSPORT Diagram

2.5.1 Hazard Identification

Hazard identification is undertaken using an adapted version of FMEA that is sometimes called 'What If?' analysis. Each boundary element is considered in turn and the hazards to the environment are considered systematically (see Section 2.7), both for when the TOE is working normally, and when there is a failure in the operation of the Boundary Element, or there is an error in the data being transferred from/to it.

2.5.2 High-Level Safety Requirements

It is also possible to identify the top-level safety requirements for the TOE by analysing the information that does exist about the design using an adapted version of FTA called 'What Causes?' analysis. This will identify the (high level) faults that may result in a hazard, for which suitable safety requirements may then be formulated to mitigate their effect.

2.6 Preliminary Hazard Analysis

Whilst the PASSPORT methodology can be used for TOEs of the form shown in Figure 3b, especially if they are novel, the more normal situation is when an existing system is being replaced by more modern technology, or being enhanced in functionality. In this situation an outline design of the TOE and its Terminators can be drawn and the system analysis and system decomposition steps of PSA are not necessary.

At Rover Group Ltd, a Preliminary Hazard Analysis process has been defined in order to meet the MISRA requirements for Safety Analysis. It was strongly influenced by Defence Standard 00-56 [DEFSTAN 00-56 1996] but has been developed specifically for use with embedded automotive controllers. It has been applied during the development of engine management and transmission systems, and of chassis controllers.

The first phase of the Rover process involves hazard identification using Hazard and Operability (HAZOP) studies. The process used is very similar to that described in Defence Standard 00-58 [DEFSTAN 00-58 1996] but with specific deviations and enhancements where necessary.

The most important concepts defined in Defence Standard 00-58 are:

- **Entity** - A label associated with an interconnection between components;
- **Attribute** - A property of an entity.
- **Guide word** - A word that describes a deviation from the design intent.

The HAZOP study involves the construction of HAZOP cases through the combination of entities, attributes and guide words in order to describe some deviation from design intent. Each HAZOP case is then used to identify potential hazards e.g.

What if [Entity] [Attribute] is [Guide word]?

The process used by Rover can be characterised as follows:

- It is based on Defence Standard 00-58 [DEFSTAN 00-58 1996];
- The entities for the analysis are only associated with the outputs of the control unit;
- The choice of attributes is derived from behaviour at the electro-mechanical boundary;
- An application specific set of guide words has been defined;
- The causes of hazards are not investigated.

In the Rover PHA process entities are labels given to the outputs of embedded controllers. The restriction of the analysis to outputs is a pragmatic decision based on the assumption that an analysis of outputs will identify all those hazards that

would have been identified during an analysis of the inputs. In addition, the behaviour that results due to deviations from design intent at the inputs depends on the functionality that has been/will be implemented. This may not be well defined at the earlier stages of the development, which is when the PHA should take place.

A common problem when performing a HAZOP study has been the choice of attributes. An output from a controller can be considered from many perspectives e.g. voltage/current levels, frequency, output state, position of an actuator, etc. Experience has shown that the best attributes to choose are those associated with the behaviour of components being controlled on the electro-mechanical interface.

Consider the example of the embedded controller depicted in Figure 6 which maintains the vehicle ride height by controlling an air suspension system. The output controls a valve that regulates the flow of compressed air into an air spring. The opening of this valve is managed by a solenoid, the current to which is regulated by varying the mark space ratio of a pulse width modulated signal at the output of the controller. Whilst the voltage at the output pin of the controller could be the focus of the HAZOP study, experience suggests that it is more meaningful and productive to consider the effects directly due to the behaviour of the actuator, rather than as a consequence of following the cause and effect sequence arising from some behaviour described at the output pin of the embedded controller. In other words, the HAZOP cases are mainly used to describe behaviour of an electro-mechanical actuator rather than that associated with the immediate boundary of the TOE.

The Rover PHA process defines its own set of guide words which have been chosen for relevance to the actuator technologies currently being used. If the nature of the system being controlled changed then the list of guide words may need to be revised.

Figure 6 depicts an embedded controller with one output signal that controls a pneumatic valve, and one input signal which is used to provide vehicle height feedback from a height sensor. The height sensor input will play no part in the HAZOP study because inputs are ignored. The interconnection between the controller and the valve becomes an entity for the HAZOP study. In this example, a good choice of attribute would be valve position, since it is associated with a component on the electro-mechanical boundary. It is then necessary to choose a set of guide words that are meaningful in the context of this system. One possible choice would be the guide word MAXIMUM so, for example, the HAZOP case would become:

What if [VALVE] [POSITION] is [MAXIMUM]?

This entity, attribute, guide word combination describes some deviation from the design intent, however, it does not have a defined meaning in the context of this




Figure 6 - Example air suspension system

system due to the generic nature of the guide words. It is therefore necessary to assign a meaning to this question. A possible interpretation of its meaning is that the valve is fully open. The effects of this behaviour would then be determined and traced to the relevant hazards. This would be achieved using the combined knowledge and experience of a carefully selected group of engineers.

Figure 6 also shows the relevant boundaries in the context of this example as follows:

- **System Boundary** - This denotes the vehicle system of interest. It includes all those components associated with the air suspension system - irrespective of technology.
- **Boundary of the TOE** - This only includes the controller because the purpose of the analysis in this example is to determine the SIL for the controller.
- **Zone of Responsibility** - This boundary exists by virtue of the company organisation. In this example this boundary encompasses the electrical/electronic and programmable electronic components within the air suspension system.
- **Vehicle Hazard Boundary** - This is primarily associated with the interface between the road wheels and the road, since an air suspension system may cause hazards related to the lateral control, handling and stability of the vehicle.

2.7 Vehicle Hazard Identification

The [IEC 61508 1999] definitions of Fault, Error and Failure highlight the fact that, in complex systems, there may be a chain of events between an initial fault and the final hazard. This is shown in Figure 7 for a vehicle, using the concepts introduced in Figure 1 and Figure 3b.

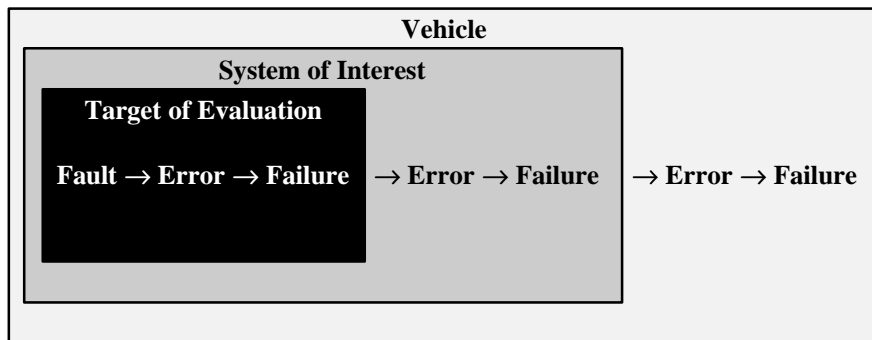


Figure 7 - Fault, Error, Failure Chain of Events within a Vehicle

Thus in order to identify most of the hazards associated with a vehicle, especially when the vehicle is moving, consideration must be given to the relationship of the TOE to Figure 2, noting in particular that the outputs of the TOE may be combined with those of other systems to produce an emergent property. Figure 2 may be used as a form of checklist whereby each interaction between the vehicle and its environment is considered in turn as to whether a property of, or a failure in, the TOE may produce a hazard at the boundary of the moving vehicle.

3 Safety Integrity Levels

In order to take advantage of being able to vary the degree of rigour with which a system may be specified and designed, as specified by [IEC 61508 1999] and followed by [MISRA 1994], it is first necessary to classify the hazards according to their severity, after which the system can be allocated a Safety Integrity Level (SIL). There are currently four discrete levels for specifying the safety integrity requirements to be allocated to the safety-related system, in order to reduce the risk (*probability of occurrence* × *degree of severity of harm*) of a hazard to an acceptable level. [IEC 61508 1999] suggests a number of different techniques for the identification of SILs but, because they are targeted at specific types of protection system, they are not universally applicable and each industry sector needs to produce its own interpretation [Redmill 1998].

Whilst there are a few protection systems within a vehicle, e.g. a system to stop an electric window trapping a limb, most moving vehicle systems exhibit different properties. Until now all the effort on this subject has been on providing a

technique for categorising moving vehicle hazards (see Section 3.1.2), but MISRA felt that this should be enhanced to cover all vehicle system hazards.

3.1.1 Vehicle Protection Systems

It has been found possible to use a risk graph, similar to that shown in Figure 8, to allocate a SIL to a vehicle protection system. The exact meaning of each parameter and the final allocation of SILs must, however, gain consensus within the industry before the final version of the graph can be placed in the public domain.

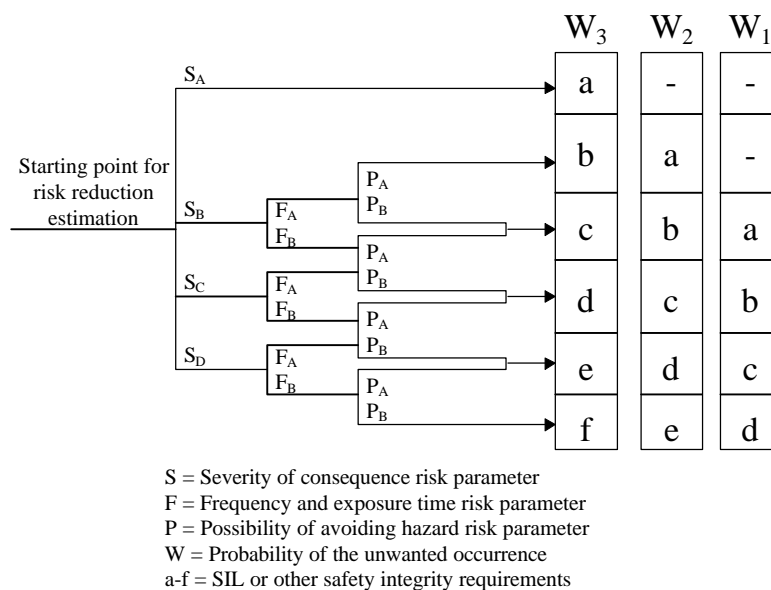


Figure 8 - Risk Graph based on [IEC 61508 1999]

There is an interesting issue with regard to certain vehicle protection systems whose objective is to reduce the risk of a hazard that is currently considered acceptable, e.g. a system to warn the driver of a high vehicle that it is about to pass under a low bridge. In this situation it should be noted that a SIL only defines the degree of risk reduction desired, whether or not the original hazard was considered acceptable. However, before one is tempted to reason that an anti-bridge-bashing system is not safety-related, because a driver should always check the signs, it must be born in mind that, once drivers are aware of the system's existence, they will be tempted to rely on it alone to indicate whether they can pass safely under a bridge.

It should be noted that it is not in the remit of MISRA to define SILs for any particular system, this is the responsibility of the designers and the manufacturers of that system. MISRA can only provide advice as to how the task may be done.

3.1.2 Controllability

There are, however, few vehicle protection systems and therefore few systems to which a risk graph can be applied; most of them are control systems. In addition vehicle based systems are obviously not static, and the road traffic environment within which they operate varies continuously. The consequence of this is that when analysing the “severity of harm”, one usually ends up with “it depends on the situation at the time”. Taking the worst case scenario is not much help either, because it is usually possible to dream up a situation where even what would normally be considered to be a ‘mild’ hazard could produce the worst possible conceivable event.

We therefore need a technique which is independent of the traffic conditions, and that can take into account the possible reactions of the road user to the occurrence of the hazard. It also needs to be independent of the number of units deployed so that, say, a high volume vehicle manufacturer would use the same SIL for the same system in the same application, as a low volume manufacturer.

A technique that satisfies all these conditions is known as *controllability* [Hobley 1995a & Jesty 1996]. This provides a qualitative assessment of the:

controllability of the safety of the situation (after a failure)

and is similar, though not identical, to the categorisation of aircraft hazards used in [DO-178B/ED-12B 1992]. The technique makes no attempt is made to classify the final effect of the hazard but uses the fact that during its occurrence, e.g. between a failure and a final event, there is a loss of control. The degree of loss of control is assessed by considering:

- The degree of control that the sub-system has on the safety of the system when it is working normally, and that therefore might be lost;
- The number, and type, of other sub-system(s) available to mitigate the loss of control (caused by the failure);
- The speed with which it is necessary for a user to react with the back-up sub-system(s) in order to mitigate the loss of control;

after which an assessment is made of the Controllability Category according to Table 1.

The distinction between the risk graph and the controllability approaches is that controllability does not consider the frequency of exposure to the hazard nor the probability of its occurrence. Travelling in a vehicle along a road is, by its very nature, a hazardous undertaking, and thus the ‘frequency and exposure time’ parameter is effectively a constant. The ‘probability of the unwanted occurrence’ parameter only applies to protection systems. Controllability therefore classifies the risk solely in terms of the immediate potential consequence to the vehicle occupants.

Controllability Categories	Definition	SIL
Uncontrollable	This relates to failures whose effects are not controllable by the vehicle occupants, and which are most likely to lead to extremely severe outcomes. The outcome cannot be influenced by a human response.	4
Difficult to Control	This relates to failures whose effects are not normally controllable by the vehicle occupants but could, under favourable circumstances, be influenced by a mature human response. They are likely to lead to very severe outcomes.	3
Debilitating	This relates to failures whose effects are usually controllable by a sensible human response and, whilst there is a reduction in the safety margin, can usually be expected to lead to outcomes which are at worst severe.	2
Distracting	This relates to failures which produce operational limitations, but a normal human response will limit the outcome to no worse than minor.	1
Nuisance Only	This relates to failures where safety is not normally considered to be affected, and where customer satisfaction is the main consideration.	0

Table 1 - Definition of Controllability Categories [MISRA 1994]

4 Detailed Safety Analysis

Once the design phase begins a detailed safety analysis (DSA) may be started. The objectives of the DSA are to:

- Confirm the findings of the PSA or PHA;
- Identify any additional hazards that may have been introduced as a result of the design used;
- Identify the possible causes of each hazard;
- Confirm the allocation of SILs;
- Predict the frequency with which a particular failure may occur;
- Identify the degree to which the system can accommodate any fault.

Whilst these processes should be undertaken for each iteration of the design, in practice they are usually only performed once, due to the time and resources required. This means that although issues may be identified, their solution may not be confirmed.

The two principal techniques that are used during DSA are FMEA and FTA. However, although they are in common use, MISRA discovered that the details of how they were done, and the objectives for performing them differed both between companies, and sometimes within companies. This is not surprising because, for example, the purpose of FMEA may include [IEC 812 1985]:

- Evaluation of the effects and the sequences of events caused by each identified item's failure mode, from whatever cause, at various levels of a system's functional hierarchy;
- Determination of the significance or criticality of each failure mode as to the system's correct function or performance and the impact on the reliability and/or safety of the related process;
- Classification of identified failure modes according to their detectability, diagnosability, testability, item replaceability, compensating and operating provisions (repair, maintenance and logistics, etc.) and other relevant characteristics;
- Estimation of measures of the significance and probability of failure, subject to the availability of data.

When FMEA is used to determine the criticality and probability of occurrence of the failure modes it is sometimes given the title Failure Mode Effects and Criticality Analysis (FMECA), but this is not applied universally.

The QS 9000 document set [QS 9000 1995], which includes a section detailing the unified approach to FMEA, was defined by the 'big 3' in the US (Ford, GM, Chrysler), and is now widely adopted by many of their subsidiaries and suppliers. Similarly the VDA (an association of German automotive manufacturers and suppliers) have published FMEA Guidelines as part of their Quality Assurance Guidelines [VDA 1996]. These publications include both 'Design FMEA', which is actually FME(C)A as defined above, and 'Process FMEA', which applies the same technique to the manufacturing process for a component.

FMEA is a bottom-up inductive technique that works up from a fault, or failure mode, to the resultant effect. A complementary technique is FTA which is a top-down deductive technique which starts from an event and seeks to determine its causes. The PASSPORT Methodology recommends the use of both FMEA and FTA as a means of providing an independent check, using FMEA to find the effect, and then using FTA to check whether the possible causes of those effects coincide with the original faults, or failure modes [Hobley 1995b]. The PASSPORT Methodology also recommends that, before either an FMEA or an FTA is begun there should be a check on the design for both its completeness and its consistency, since there is little point in analysing something that is inherently wrong.

4.1 Failure Mode and Effects Analysis

FMEA can be applied at many different levels in the hierarchy of a vehicle. Figure 9, which is one way of decomposing a vehicle, shows that there are a number of different levels at which it may be sensible to undertake an FMEA. However, as shown in Figure 3b, the only level at which the end effects produce hazards is at the vehicle boundary. It is therefore necessary to maintain traceability between the fault, or failure mode, and the final hazard through the ‘intermediate effects’ at each level.

4.1.1 Scoring

When performing an FMECA, the usual situation, the automotive industry tends to use three parameters for each failure, namely ‘severity’, ‘occurrence’ and ‘detection’, and to give each one a score between 1 and 10 (not critical - extremely critical). Whilst the generic definition of these scores are not tailored to the specific considerations relevant to programmable systems, both [QS9000 1995] and [VDA 1996] guidelines for FMEA allow for application specific scoring schemes to be defined. The MISRA Safety Analysis guidelines will therefore propose a scoring scheme that is tailored to programmable vehicle-based systems along the following lines:

- **Severity** - The hazards associated with each failure could be categorised by extending controllability (see Section 3.1.2) or the risk graph (see Section 3.1.1) into ten different levels;
- **Occurrence** - Whilst it is often possible to model the probability of random faults in hardware, this is not usually possible for systematic software faults. The occurrence of random faults should ideally be scored objectively using reliability data. The occurrence of systematic faults cannot be scored using probability, however, since FMECA is more of a management tool, rather than a pure engineering tool, a score could be based on the level of integrity associated with the SIL of its production process.
- **Detection** - A meaningful interpretation of this is a measure of the degree to which a programmable system can accommodate faults which manifest themselves during the operation of the system while in the hands of the customer. It could be made up of the probability of detecting the error, and the degree to which the risk associated with that error can be reduced by the system taking some beneficial action.

Historically the industry has multiplied these three scores to give a Risk Priority Number (RPN), and then concentrated on dealing with those faults, or failure modes, with a high RPN. In fact, however, the RPN should be used with great care, since it is a one dimensional assessment of three dimensions of information, i.e. information is lost during its creation. Of particular concern is the case where only one of the parameters has a high value; this may result in a low RPN but be highlighting a specific major problem. It is for this reason that some companies are now only using the ‘severity’ parameter.

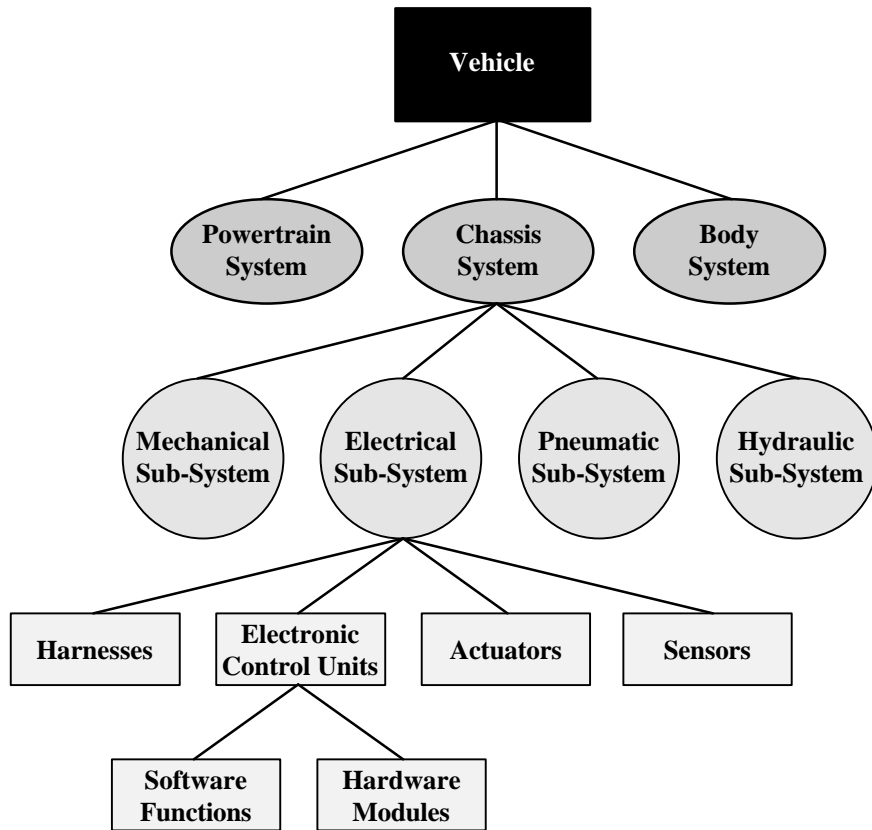


Figure 9 - Decomposition of a vehicle

4.2 Fault Tree Analysis

FTA is not as widely applied in the automotive industry as FMEA (e.g. there is no equivalent section of QS9000 covering FTA requirements). However there are two areas in which its use has been shown to be appropriate.

Firstly fault trees are often used by design engineers to communicate to those writing service repair documentation, or automated off-board diagnostic tools, the possible causes of typical vehicle faults which may be reported by a customer. In such cases the starting point will be, for example, “warning light is on” and the fault tree will be a means of identifying the possible causes, and therefore possible repair actions a technician may need to perform.

Secondly, and more relevant to this paper, they are beginning to become recognised as a valuable technique in the analysis of safety-related automotive systems. According to [IEC 1025 1990] FTA may be used to:

- Identify the causes, or combination of causes, that lead to the top event;
- Determine whether a particular system reliability measure meets a stated requirement;
- Demonstrate that the assumptions made in other analyses, regarding the independence of systems and non-relevance of failures, are not violated;
- Determine the factor(s) which most seriously affect a particular reliability measure, and the changes required to improve that measure;
- Identify common events, or common cause failures.

In the analysis of a safety-related system FTA is a complementary technique to FME(C)A and helps to give additional confidence that nothing has been missed. By starting with a hazard as the top event, it is possible to work down towards failures of LRUs, or basic events, which can give rise to the top event. This is done via a hierarchy which gradually refines the hazard through a series of intermediate (system and sub-system level failure modes) down to the failures of individual LRUs. It is possible to analyse the AND-OR logic of a fault tree to identify the combinations of basic events which can cause the top event. These are known as 'minimal cut-sets'. Where a minimal cut-set is shown to be a single basic event, this reveals a single point of failure which can give rise to a hazard. Depending on the severity of the hazard, and the known reliability of the LRU concerned, this may or may not be acceptable, depending on the level of risk. In addition, where there are multiple points of failure (i.e. more than one event in a minimal cut-set), it may be possible to make a subjective assessment as to how likely these events are to occur together, again in accordance with the acceptable level of risk.

Further analysis which can be performed on a fault-tree is the quantification of hazard occurrence rates based on the reliabilities of the LRU components of the system. Such quantification can range from a simple combination of probabilities (i.e. multiply for an AND gate, add for an OR gate), to complex time-dependent failure rate calculations ($R=1-e^{-\lambda t}$). Allowance can also be made for repairable systems, where planned maintenance and service replacements are involved, although this complicates the calculations still further. Of course, software packages are available which can assist this process.

Quantifying fault trees is standard practice in some industry sectors, but is not common in the automotive industry, and a number of notes of caution must be sounded. One key reason for this is the non-availability of valid figures for hardware component failures. Although failure rates can be measured or obtained from tables, e.g. [ROME Lab 1993], they tend to render any calculations excessively pessimistic. This is because a fault tree often works down towards very specific failure modes of an LRU, however the figures which are available tend to 'lump together' several failure modes. Therefore using these figures in an un-revised manner will lead to pessimistic calculations, which are of limited, if any, practical use for a vehicle development programme.

Where software is identified as a basic event, the question arises as to what figure should be used for the failure rate. The concept of software failure rate is not well defined, and there are a number of equally (in)valid approaches, e.g.:

- If available, use measured Mean Time Between Failure figures of software from a similar application, of similar complexity, developed to a similar process as the software for the system being analysed, by similar people in a similar organisation, using similar tools etc. etc.
- Assume that the software is 100% reliable (i.e. failure rate = 0). This approach is supported by the assertion that developing the software to a given integrity level seeks to provide sufficient confidence that this is the case.
- Apply the target failure rate figures from IEC 61508 for the appropriate integrity level.
- Estimate the software failure rate based on an algorithmic argument, for example based on the number of instructions executed, processor clock rates, and typical defect rates etc. [ROME Lab 1993]
- Guess, and attempt to justify its validity.

Despite these shortcomings, quantified fault trees can still be very useful in gaining an understanding of the failure mechanisms within a vehicle system. For example, even though the absolute validity of the calculated occurrence rates of the top event is questionable, the tree can be used to test assumptions and to perform sensitivity analysis on the effects of the uncertain component failure rates on system failure rates (sensitivity analysis). It may also be possible, accepting the pessimistic outcomes, to see how close a particular system design gets to the target figures in [IEC 61508 1999], for example when choosing between different vehicle system architectures.

5 Conclusion

The MISRA Steering Group has identified a need to provide advice on how to undertake safety analysis on automotive systems, and has decided to produce some guidelines based on the experience of its various members. Whilst many generic techniques exist, they do need interpreting for the specific sets of circumstances that are found in this industry sector. These guidelines will cover the various techniques that may be used to identify hazards, classify hazards and to analyse designs to ensure that the risk associated with each hazard has been reduced to an acceptable level.

References

[Carroll 1939] Carroll L: Sylvie and Bruno Concluded, in The Complete Works of Lewis Carroll. The Nonesuch Press, 1939 (First Published 1893)

[DEFSTAN 00-56 1996] DEFSTAN 00-56: Defence Standard 00-56 - Safety Management Requirements for Defence Systems, Issue 2. Ministry of Defence, 1996

[DEFSTAN 00-58 1996] DEFSTAN 00-58: Interim Defence Standard 00-58 - HAZOP Studies on Systems Containing Programmable Electronics, Issue 1. Ministry of Defence, 1996

[DO-178B/ED-12B 1992] DO-178B/ED-12B: Software Considerations in Airborne Systems and Equipment Certification. RTCA-EUROCAE, 1992

[Hobley 1995a] Hobley K M, et al.: Framework for Prospective System Safety Analysis Volume 1 - Preliminary Safety Analysis. Deliverable N° 9a, V2058 PASSPORT project of the Advanced Transport Telematics (ATT/DRIVE II) sector of the TELEMATICS APPLICATIONS Programme, Third Framework Programme (1991-94), 1995

[Hobley 1995b] Hobley K M and Jesty P H: Analysis and Assessment of Advanced road Transport Telematic Systems. Proceedings of the 14th International Conference on Computer Safety, Reliability and Security (SafeComp '95), Belgirate, Italy, G. Rabe, Ed., Springer-Verlag, 1995, ISBN 3-540-19962-4

[IEC 812 1985] IEC 812: Analysis Techniques for System Reliability - Procedure for Failure Mode and Effects Analysis (FMEA). International Electrotechnical Commission, 1985

[IEC 1025 1990] IEC 1025: Fault Tree Analysis (FTA). International Electrotechnical Commission, 1990

[IEC 61508 1999] IEC 61508: Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems. International Electrotechnical Commission, 1999

[Jesty 1996] Jesty P H and Hobley K M: Integrity Levels and their Application to Road Transport Systems. Proceedings of the 15th International Conference on Computer Safety, Reliability and Security (SafeComp '96), Vienna, Austria, E. Schoitsch, Ed., Springer-Verlag, 1996, ISBN 3-540-76070-9

[MISRA 1994] MISRA: Development Guidelines for Vehicle Based Software. MIRA, CV10 0TU, 1994 {<http://www.misra.org.uk/>}

[QS 9000 1995] QS 9000: Quality System Requirements. Chrysler Corporation, Ford Motor Company and General Motors Corporation, 1995. Available from Carwin Continuous Ltd, Thurrock, Essex, UK

[Redmill 1998] Redmill F, IEC 61508 - Principles and Use in the Management of Safety. Computing and Control Engineering, vol. 9, No. 5, Institute of Electrical Engineers, 1998.

[Redmill 1999] Redmill F, Chudleigh M, and Catmur J: System Safety: HAZOP and Software HAZOP. John Wiley & Sons Ltd, 1999, ISBN 0 471 98280 6

[ROME Lab 1993] ROME Lab: Reliability Engineer's Toolkit. The ROME Laboratory, US Airforce Material Command, Griffiss Air Base, NY, 1993

[VDA 1996] VDA: Quality Management in the Automotive Industry, Quality Assurance before Series Production: Volume 4 Part 2: System FMEA, Failure Mode and Effects Analysis. Verband der Automobilindustrie e.V., 1996