Derivation of Safety Targets for the Random Failure of Programmable Vehicle Based Systems

Richard Evans[1] and Jonathan Moffett[2]

[1]Jaguar Cars Limited, W/1/014, Engineering Centre, Abbey Road,
Whitley, Coventry, CV3 4LF, United Kingdom.
revans52@jaguar.com
[2]Department of Computed Science, The University of York, Heslington,
York, YO10 5DD, United Kingdom.
jdm@cs.york.ac.uk

Abstract Increasingly, the dependability of vehicle based programmable systems is becoming a key feature in ensuring the safety of those in and around the vehicle. The goal of those responsible for the design and manufacture of such systems must be to control adequately the associated risks so that the potential of the technology may be exploited fully. The Motor Industry Software Reliability Association (MISRA) has provided guidance for the management of the safety risks associated with software, but there is no comparable guidance for the management of the risks associated with the random failure of electronic hardware. This paper describes the development of an automotive industry specific risk model and goes on to derive safety targets for the random failure of programmable vehicle based systems. In addition the work provides a basis for comparison between the MISRA Guidelines and related national and international standards.

1 Introduction

Safety engineering provides many processes, methods, and techniques for use during the design and development of safety related and safety critical systems - irrespective of application domain. However, most if not all approaches rely on there being some defined safety targets against which the system will be measured.

In 1994 the Motor Industry Software Reliability Association (MISRA) published the “Development Guidelines for Vehicle Based Software” [1] which addressed the management of the safety risks associated with vehicle based software. However, there is no comparable guidance for the management of the corresponding risks associated with random failure. The MISRA approach uses the concept of Safety Integrity Levels (SIL) but these are automotive industry specific and have no defined mapping to other industry specific or generic SILs.

It is the belief of the authors that targets for the random failure of electronic systems are desirable in order to act as design drivers during system development. The purpose of these targets would be to provide a sound basis for decisions which affect the dependability of embedded systems, as for example, the introduction of redundancy.

This paper results from a project done as part of the MSc in Safety Critical Systems Engineering at the University of York, and sponsored by Rover Group Ltd [6]. It is structured as follows. Section 2 describes important characteristics of the automotive industry. Sections 3 & 4 describe the approach to deriving a safety target, with section 3 describing the method, section 4 creating a risk model. Section 5 & 6 derive safety targets, first for accidents and then for systems. Finally section 7 discusses the results and reaches some conclusions.

2 Characteristics of the Automotive Industry

2.1 Introduction

Historically there has been a trend to deploy programmable systems which enhance the functionality of existing mechanical systems e.g.

- engine control systems,
- gearbox control systems,
- chassis control systems.

The hazards associated with failure of these systems are largely independent of the introduction of programmable technology but the likelihood of the hazards occurring is very much dependent on the programmable system. The implication is that there is an existing level of safety performance by virtue of mechanical components which can be used as a baseline for risk.

In the future the programmable systems fitted to vehicles will perform more novel functions which to a greater or lesser extent can be characterised by automatic interaction between vehicles. As a result these will give rise to new hazards. Examples of these systems are:

- adaptive cruise control,
- road trains,
- collision avoidance.

Since these hazards are new there is no existing baseline against which to measure risk.

2.2 Hazard Classification

The concept of controllability was developed during the project DRIVE Safely [3]. Controllability is derived from the usual representation of risk i.e.:

$$R = P * E$$

Where R is the risk, P is the probability of a system failure which could result in an accident (not the probability of an accident given a failure) and E is the effect of the system failure given that the failure has occurred.

In the automotive context the environment in which vehicles are used is extremely variable due to the large number of factors that influence the safe passage of road users. A small subset of these variables are:

- Weather
- Traffic density
- Road type
- Driver behaviour
- Lighting levels

The consequence of this situation is that it is not possible to predict with any certainty the likely ‘effect’ of a system failure.

The DRIVE approach was to assign integrity levels based on the classification of the ‘effect’ in terms of five controllability categories. These categories had textual definitions which describe qualitatively the degree of control of the safety of the situation given that a failure has occurred.

3 Safety Target Derivation Method

In order to make a decision on which approach is most suitable it is necessary to bear in mind some general requirements for safety targets:

1. They should be justifiable and defensible.
2. They should be capable of adapting to changing perceptions of risk.
3. They should be reasonable.
4. They should cope with rapidly changing technology.

For automotive safety targets to be of most use they must be reviewed by all relevant parties and agreed as the definition of ‘best practice’. If such targets can be agreed upon they could be used as a means to bridge the gap between industry specific standards (e.g. MISRA) and generic standards (e.g. IEC 61508).

In the light of these considerations the chosen solution was to:

1. Derive a model which enables a risk expressed in terms of transport accidents to be related to vehicle based systems.
2. Establish a level of risk, associated with accidents resulting from technological failure in vehicle based systems, that would be regarded as broadly acceptable in the sense used by the Health and Safety Executive [4].
3. Use the target risk in terms of accidents and the derived risk model to set safety targets for vehicle based systems.

4 Risk Model

4.1 Analysis of Accident Data

The analysis described in this section seeks to answer the questions:

1. What kind of vehicle mode best determines the severity of an accident?
2. What kind of vehicle mode best determines the probability of an accident?

The source of the data was the annual report from the UK Department of the Environment, Transport and the Regions. It is entitled “Road Accidents Great Britain: 1997 - The casualty report” (known as RAGB). The underlying data which provides the basis for this report is collected by the police, either as a result of an officer attending the scene of an accident, or through reports made to the police at a later time.

There are very many measures that can be used to describe the severity and probability characteristics of accidents. The Road Accidents Great Britain report employs a consistent approach to describing the severity of casualties and accidents:

- Fatal (i.e. killed)
- Serious (i.e. seriously injured)
- Slight (i.e. slightly injured)

The data has been split into a number of modes. These modes have been grouped into classes as shown in Table 1.

Intuitively it was expected that the mode classes would influence both the probability and severity of an accident due to the differing factors involved.

A statistical analysis of a sample of the data shows that the severity distribution is largely independent of driving condition, their mean being: 1.52% fatal; 14.73% serious; 83.75% slight [6]. This is seen as an important result in this paper since it suggests that the overall distribution of outcomes is fixed and that the significant factor becomes the probability of an accident, independent of the level of uncontrollability of a hazard.

It should be noted that the accident data recorded in RAGB are only a subset of the relevant outcomes of accidents. This is because the accident data are only collected if the accident involves personal injury. However, the RAGB report states that the cost-benefit value of prevention of road accidents in 1997 was estimated to be £14,814 million, of which £10,453 million is attributable to personal injury accidents and the remainder being associated with damage only accidents. The report also states that the average cost per 'damage only' accident is £1,210 and, for all accidents involving casualties, the average cost was £43,550. Hence the number of damage only accidents, and personal injury accidents, can be estimated as:

$$\text{1997 damage only accidents} = (£14.8 \times 10^9 - £10.4 \times 10^9) / £1{,}210 = 3.6 \times 10^6$$

$$\text{1997 personal injury accidents} = £10.4 \times 10^9 / £43{,}550 = 2.4 \times 10^5$$

This means that the ratio between 'personal injury' accidents and 'damage only' accidents is approximately 15:1. Using this result we can scale the personal injury accident data with respect to the damage only accidents and hence derive the severity distribution in Table 2.

4.2 Resulting Risk Model

Table 3 defines terms used in the derived risk model, which itself is presented in Fig. 1. Relevant probabilities are defined as:

- $P(A|H_x)$ is the probability of an accident, given a hazard with classification denoted by H_x.

- $P(I|H_x)$ is the probability of an incident, given a hazard with classification denoted by H_x.
- $P(C|H_x)$ is the probability of control being maintained, given a hazard with classification denoted by $H_{x.}$
- $P(A_x|A)$ is the probability of an accident of severity denoted by A_x given that *an* accident has occurred.

In the absence of data generated from the results of an experiment such as the use of a driving simulator, we make the following intuitive assumptions:

- A hazard classified as uncontrollable will result in either an incident or an accident with a probability of 1.
- A hazard classified as nuisance only will not result in an accident or incident.
- There is likely to be a logarithmic relationship linking the likelihood of an accident for the remaining three categories of controllability.

If the above hold then we can derive a table which maps controllability categories to probability of an accident or incident. This is shown graphically in Fig. 2.

5 Accident Safety Targets

In this section we derive accident safety targets that could reasonably be considered to represent a 'broadly acceptable' level of risk. The method used is to evaluate the overall level of risk to which drivers are exposed so as to act as a baseline. In addition, the HSE work on the definition of a broadly acceptable level of risk, and the MEM (Minimum Endogenous Mortality) criterion are evaluated. Finally an accident safety target is proposed.

The risk of death arising directly from travelling by car is approximately 10^{-4} per person.year based on the following data:

1. Total number of car road traffic fatalities in 1997 = 1,934 {[2] p54}.
2. Total car & taxi traffic in 1997 = 3,678 x 10^8 km {[2] p53}.
3. It is assumed that the average distance travelled per annum = 20,000 km.

The HSE (Heath and Safety Executive) ToR (Tolerability of Risk) framework involves the definition of the upper and lower limits of tolerable risk. The lower of the two limits corresponds to the point below which the risk is considered to be 'broadly acceptable'. The following quote from the HSE gives a definition of the limit between tolerable and broadly acceptable risk: "This level might be taken to be 1 in a million (1 in 10^6) per annum bearing in mind the very small addition this would involve to the ordinary risks of life… "{[4] p31}.

Endogenous mortality is the rate at which a particular age group of a population die due to technological causes. In central Europe and well developed countries in general, the Minimum Endogenous Mortality (MEM) rate corresponds to the age group 5 to 15 years and has a value of $2x10^{-4}$ fatalities/person.year [5]. A feature of the MEM criteria is that it is assumed that there is a maximum of 20 technological systems that affect an individual at any one time, therefore they each share a proportion of this failure rate with the result that the acceptable risk to an individual from a given system is 10^{-5} fatalities/person.year.

Therefore the target risk is defined by the MEM criteria at the upper limit and the HSE ‘broadly acceptable’ value for risk at the lower limit. Therefore the derived accident safety target becomes:

$$10^{-6} < P(A_3)_{target} < 10^{-5} \text{ per person.year}$$

Since there are approximately 10^4 hours in a year this target can be expressed as:

$$10^{-10} < P(A_3)_{target} < 10^{-9} \text{ per person.hour}$$

6 System Safety Targets

6.1 Single Vehicle Systems

The derivation of safety targets for the single vehicle context relies on the following characteristics:

1. Vehicles operate independently of each other.
2. The upper limit on the number of fatalities that could reasonably be expected to result from an accident is close to unity.
3. The severity ratios that were derived through an analysis of present day accidents is valid for these types of system.

Table 4 presents accident probability versus controllability category and suggests that an uncontrollable hazard is certain to result in a loss of control and hence either an incident (near miss) or an accident. Therefore, using the accident safety target derived in section 5, and since:

$$P(H_4)_{vehicle.target} = P(A_3)_{target} / P(A_3|A)$$

the safety target for all the uncontrollable hazards for an entire vehicle can be derived as:

$$10^{-7} < P(H_4)_{vehicle.target} < 10^{-6} \text{ per hour}$$

As each vehicle contains a number of systems it is necessary to reason about the collective influence that these systems can contribute to causing a hazard of a given controllability classification.

Assume that the various programmable systems on the vehicle are of similar designs, and use similar technology, and each have a probability of failure, P(F), per hour. If it is further assumed that these systems have high levels of reliability, then for N of these systems the probability of failure is N x P(F). At the current time there are very few systems fitted to vehicles that can cause the most severe hazards i.e. those which would be considered uncontrollable. For the purposes of deriving these safety targets it is conservatively estimated that there are 10 such systems fitted to the vehicle.

Based on the previous assumption the system safety target for an uncontrollable system hazard, $P(H_4)_{system.target}$ can be calculated as follows:

$$P(H_4)_{system.target} = P(H_4)_{vehicle.target} / 10$$

Giving the result:

$$10^{-8} < P(H_4)_{system.target} < 10^{-7} \text{ per hour}$$

The equivalent results for the remaining controllability categories may be calculated in the same way but with reference to Table 4 for the corresponding accident probability. These results have been calculated and are recorded in Table 5.

6.2 Multiple Vehicle Systems

The derivation of safety targets for the ‘multiple vehicle’ context relies on the following characteristics:

1. Vehicles operate in systems of more than one vehicle.
2. The upper limit on the number of fatalities that could reasonably be expected to result from an accident is of the order of 10.
3. The severity ratios that were derived through an analysis of present day accidents are valid for these types of system.

The concept of Differential Risk Aversion (DRA) [5] can be used to modify the safety target for single vehicle systems to be applicable in multiple vehicle systems. This results in an order of magnitude decrease in the allowable frequency of an accident for each order of magnitude increase in the upper limit of the number of fatalities per accident. Therefore, the system safety target for the ‘multiple vehicle’ scenario becomes:

$$10^{-9} < P(H_4)_{system.target} < 10^{-8} \text{ per hour}$$

This result is then used to derive the system safety targets for the remaining hazard classifications as in section 6.1, which are summarised in Table 7.

7 Conclusions and Discussion

The following is a summary of the main conclusions of the paper:

1. It is considered possible to define a reasoned approach to the definition of safety targets for automotive systems.
2. An analysis of UK road accident data suggests that the proportion of fatalities in road accidents is effectively independent of vehicle operating conditions and influencing factors. This fact resulted in the hypothesis that vehicle hazards have a bearing only on the probability of *an* accident and not the likely severity. This result has an impact on the definition of controllability and suggests that severity should be removed.
3. A safety target for automotive systems should be set to a level that gives rise to a negligible level of risk. This is despite the fact that vehicle users are exposed to a relatively high level of risk, due mainly, from human error rather than technological failures.

4. Safety targets based on a risk to an individual derived from the HSE definition of ‘broadly acceptable’ risk and the MEM criterion are considered appropriate and reasonable.

We observe that the conclusions are based on a number of assumptions, the most important of which are:

1. That the statistical analysis, which leads to the conclusion that accident severity is independent of the level of controllability of the hazard leading to the accident, is confirmed.
2. That the logarithmic relationship between the level of controllability and the probability of an incident or an accident is confirmed.

Further assumptions are detailed in [6], together with proposals of how to test them. In particular, the relationship between controllability and severity could be tested out by means of experiments in a driving simulator, which would enable realistic testing of driver behaviour, and its results, in the presence of hazards. Additionally, further investigation is needed into how the risk model can be affected by human factors, some of which could be addressed by experiments using a driving simulator. It is also uncertain how the structure and granularity of the data affect the risk model and the safety targets. If more detailed data were available, the results could change.

References

1. “Development Guidelines for Vehicle Based Software”, MISRA, 1994
2. “Road Accidents Great Britain: 1997 The Casualty Report”, Department of the Environment, Transport and the Regions, August 1998, ISBN 011 552068 6.
3. DRIVE Safely, Towards a European Standard: The development of Safe Road Transport Informatic Systems (Draft 2), DRIVE Project V1051, 1992.
4. “The Tolerability of Risk From Nuclear Power Stations”, Health and Safety Executive, 1992, ISBN 0 11 886368 1.
5. “Generalised Assessment Method, Part 2:Guidelines”, ESPRIT P9032, CASCADE, 1997.
6. "Derivation of Safety Targets for the Random Failure of Programmable Vehicle Based Systems", Richard Evans. MSc Thesis, Department of Computer Science, University of York. September 1999

Appendix: Tables

Table 1. Driving Modes Relevant to Accidents

Mode class	Modes
Infrastructure	Road type, Junction type, Street lighting
Driver	Sex, Age
Environment	Daylight/Darkness, Weather conditions, Road conditions
Temporal	Day of the week, Time of day
Other issues	Special conditions, Region, Object hit, Carriageway hazards

Table 2. Severity distribution for all accident severities

Fatal	Serious	Slight	Damage only
0.09%	0.92%	5.23%	93.76%

Table 3. Risk Model Definitions

Term	Definition	Classification
Accident	An unintended event or sequence of events that causes death, injury, environmental or material damage.	Slight, Serious, Fatal
Incident	An unintentional event or sequence of events that does not result in loss, but, under different circumstances, has the potential to do so.	Boolean
Hazard	A hazard is a situation in which there is actual or potential danger to people or the environment.	Nuisance, Distracting, Debilitating, Difficult to Control, Uncontrollable

Table 4. Probability of an incident/accident by controllability category

Controllability category	Probability of incident/accident
Uncontrollable	1
Difficult to control	1 in 10
Debilitating	1 in 100
Distracting	1 in 1000
Nuisance only	0

Table 5. Safety targets in single vehicle systems

Category of controllability	$P(A \mid H_x)$	$P(H_x)_{system.target} = P(H_4)_{system.target}$ x $P(A \mid H_x)$
Uncontrollable	1	$10^{-8} < P(H_4)_{system.target} < 10^{-7}$
Difficult to control	10^{-1}	$10^{-7} < P(H_3)_{system.target} < 10^{-6}$
Debilitating	10^{-2}	$10^{-6} < P(H_2)_{system.target} < 10^{-5}$
Distracting	10^{-3}	$10^{-5} < P(H_1)_{system.target} < 10^{-4}$
Nuisance only	0	N/A

Table 6. Safety targets for multiple vehicle systems

Category of controllability	$P(A \mid H_x)$	$P(H_x)_{system.target} = P(H_4)_{system.target}$ x $P(A \mid H_x)$
Uncontrollable	1	$10^{-9} < P(H_4)_{system.target} < 10^{-8}$
Difficult to control	10^{-1}	$10^{-8} < P(H_3)_{system.target} < 10^{-7}$
Debilitating	10^{-2}	$10^{-7} < P(H_2)_{system.target} < 10^{-6}$
Distracting	10^{-3}	$10^{-6} < P(H_1)_{system.target} < 10^{-5}$
Nuisance only	0	N/A

Appendix: Figures

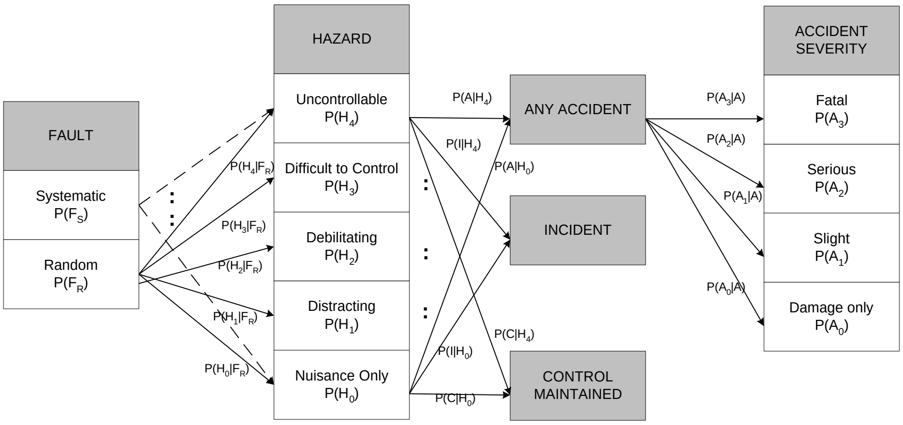

Fig. 1 **Risk Model**

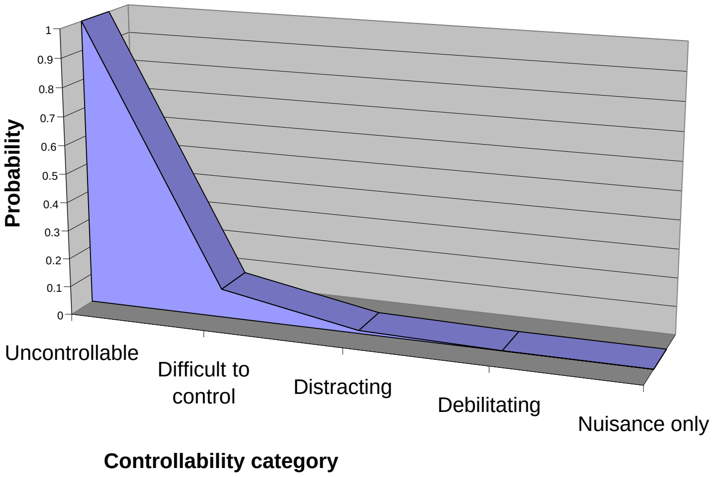

Fig. 2 **Probability of incident/accident by controllability category**