Invited Paper

The problems of assessing the safety of ATP systems which have been developed under different national standards

J. Catmur

Arthur D. Little, Science Park, Milton Road, Cambridge, CB4 4DW, UK

Abstract

As railways broaden their potential suppliers of ATC systems they are faced with selecting between systems that have been built at different times and in different countries. The techniques used to manage the development of safety-critical computer systems have evolved over the last 20 years and the way in which they have been applied has varied from country to country and supplier to supplier. This paper examines the difficulty that exists in assessing the safety of computer based ATP systems which have been developed using distinct, and sometimes conflicting, national standards. It is based on extensive practical work in reviewing the protection against systematic failures that the design processes, used in a variety of ATP, systems has given.

1. Introduction

When selecting computer based signalling systems railways must, if they are to deliver the service required by their clients, ensure that the system installed will meet a series of operational, reliability and safety targets. When assessing the possibility of the new signalling system meeting the required targets the purchaser must take into account the methods used by the supplier to ensure and demonstrate that the targets are achieved. During the design, construction, installation and commissioning process the system's manufacturer will need to use a series of techniques and processes which are aimed at ensuring that their client's requirements are met.

The causes of the reliability, availability and safety requirements not being met can be divided into those caused by random failures of the system's hardware and those caused by systematic errors in the design and installation process. If the supplier misinterprets an operational requirement a systematic error will be introduced They can also design a system which is prone to random hardware failures and may thus not fulfil the reliability or service targets. Systematic

failures can occur at all stages of the project, from requirements identification through to final commissioning.

Many of the purchaser's requirements can be demonstrated to have been achieved through extensive testing during the design through commissioning process, and this testing may also help remove some of the systematic errors that could have occurred. The possibility of random hardware failures occurring will often be analysed as part of the availability, reliability and maintainability studies that the manufacturer produces. For those systematic errors where a significant number are found during testing and use, reliability growth modelling can be used to predict the probability of further undetected errors. There are, however, no methods yet available for analysing or identifying the classes of systematic errors which are very unlikely to occur. In general systematic errors which could cause safety targets to not be achieved are in these latter classes.

The safety targets for signalling systems are often defined in terms of the mean time between unsafe failures of the system, and are in the range of 100,000,000 - 1,000,000,000 hours. For the random failures that could cause such unsafe failures analysis can be carried out to give a reasonably accurate demonstration that such figures are attained. For systematic failures it is impossible to carry out enough testing to demonstrate this (one might require 100,000,000,000 hours of testing), so we must rely upon the quality of the process that is used to develop the signalling system as a method for avoiding such critical systematic failures.

When assessing a supplier's ability to produce a safe system we need to judge the degree to which his processes reduce the chance of an undetected safety-critical systematic error being present in the final system. During the assessment no proof can be made that such errors do not exist , but a degree of confidence can be obtained. Over the last few years standards have begun to appear which define the processes a manufacturer may use to help build up confidence that they have avoided allowing safety critical systematic errors to slip through their development process of computer based signalling systems. These standards may in the future allow easier assessment of the design to commissioning process, once they are in wide use.

For a railway which wishes to purchase an existing computer based signalling system the problem of assessing a manufacturer's ability to deliver a safe system is not so easy to resolve. Existing systems may have been designed over the last 25 years in different countries, by different companies, using techniques that have evolved with time and varied substantially.

Over the last three years we have been involved in carrying out pre-tender evaluation and post-tender safety review of numerous computer based signalling systems in a variety of countries and have had to struggle with these differences, in order to reach a useful conclusion. This paper reviews the problems we have found in carrying out this work. While it cannot use specific examples, for confidentiality reasons, we hope that it can help shed light on how such assessments can be carried out and the results they can give.

There were many areas where we found problems in assessing the systems and, therefore, we had to use our judgement in order to reach a useful conclusion as to their safety. The five main areas of concern were:

- Different standards currently exist in each country
- The systems were often built before these standards existed or as they evolved
- The philosophy behind achieving a safe system varies substantially
- For older systems original documentation may no longer exist
- There is no clear definition of how to develop a safe system

This paper gives an overview of how we assessed the systems concerned and reached a useful conclusion.

Aim of the reviews carried out

The reviews carried out attempted to answer the basic question "Was the system developed in such a way as to give reasonable confidence about its basic level of safety?" for each of the systems under consideration. When we started the work we knew that no correct answer existed, and as the studies we carried out progressed it became clearer and clearer that existing systems could not meet emerging standards. We therefore aimed to measure the quality of each system's development process and the general soundness of its architecture and to thus identify those areas where either could have been improved.

Once this was done we evaluated the results in order to come to a conclusion as to our confidence in the system's safety. We were clear from the start that it is impossible to identify the level of protection against systematic failures a system would give, so we would have to be content with an evaluation based on the quality of the development process. We also recognised that this meant that we had to approach the evaluation with an open mind and not to pre-judge the correct solution.

Differing standards

As we carried out this work we encountered a series of different standards which have been applied in the past, and also a series of emerging standards. As stated above, we found that almost all the systems we examined could not comply with the latest standards. The problem was how to judge the quality of systems which had not been developed, and could not have been developed, to all comply with all these standards. The fact that they did not comply with a standard did not make them unsafe. It may have reduced our confidence in the system's level of safety, if we believed that straight forward compliance with the standard would make the system safer, which we did not.

Compliance with a standard can mean many things and for any standard which is not explicit compliance is difficult to evaluate as a black and white issue. This means that compliance with the spirit of a standard is possible, while the actual practise might not be as robust as would be liked. So we could not simply check

compliance but had to test the depth and quality of the compliance.

The systems were often built before these standards existed

In the case of systems which were built to no clear standard it became more difficult still to evaluate them. Systems may well have been developed to a very high quality for a computer based system in the '80s, but obviously would not comply to the new standards in this field. For these systems it was easy to judge them against the new standards but very difficult to say whether their non-compliance made them unsafe. We could also have judged them against the state of the art for when they were built and see how well the matched that, but we knew that the state of the art has moved on.

The philosophy behind achieving a safe system

The actual philosophy used by a company in achieving a safe system varies substantially, with differences in the basic design architecture and development process as well as greater or lesser emphasis on certain parts of the design process. A company may not have done one activity because they compensated through greater usage of another, or they may have used an architecture that could not comply with some of the new or existing standards. This difference in philosophy was very noticeable when reviewing systems from different countries (we have reviewed systems in the UK, France, Belgium, Germany, Italy and Canada), and it is impossible at the moment to say that any one philosophy is correct. In the end one has to understand the philosophy applied and the benefits and problems it brings with it and to evaluate those.

For older systems original documentation may no longer exist

The lack of some of the early documentation in a project which has been going on for five or ten years makes review more difficult. It increases the level of uncertainty as it is impossible to fully understand and check the early work on the system. The quality of the documentation that still existed at the time of the audits helped give an idea of what the quality of the former documentation may have been like, but did not allow full traceability through the design process. The latter was our preferred route to assessing the quality of the development process. If the lack of early documentation did not prevent us answering all our questions about the system and fully understanding how the sample we examined worked, then we could feel more confident, but if we could not fully understand the design without it this led more cause for concern.

There is no clear definition of how to develop a safe system

There are numerous opinions on how to develop a safe system and there is no correct answer. We do not believe that one perfect solution exists, so had to study the solution a manufacturer had chosen and evaluate the merits of it. If a perfect method were to have been defined we would almost certainly have found no system which complied with it, but as no such definition exists we had to

work with each manufacturer to understand their views and then to weigh up the pluses and minuses against all the ideas which exist on how to develop a safe system.

Resolution of these problems

During the studies we conducted on nine different computer based signalling systems we encountered all these problems, and yet still had to advise our client as the confidence they could have in the safety of each system. To do this we first spent two weeks carrying out an audit of the system development. This audit reviewed:

- How and when the system was developed
- How the development process was managed
- The techniques involved
- The appropriateness of the techniques and management process used
- Compliance with the stated techniques and processes

During each audit we also conducted a small hazard analysis session within our team in order to understand the degree to which the system in general and individual components were susceptible to systematic errors. As part of this we also reviewed all the existing safety studies on each system. Our tests of compliance with the stated techniques were based on taking a vertical slice through the system and testing all the documentation from the specification to the final system. This was normally done for both hardware and software.

At the end of each audit we then reviewed the overall results, which consisted of a set of findings which reflected the concerns we had identified. For each concern we then considered whether correcting it would have a substantial improvement on our confidence in the system's safety. To do this we had to consider the quality of the design team, the design of the system and the way potential hazards were managed by the team/system. This review was based on the expertise of the team, and the informal risk analysis we conducted for each finding.

It would have been easy to simply say that a certain system did not, for example, "meet RIA23 in the following areas:", but that would not have helped us to reach a final decision, as we never found a single system that met all the standards or ideas in this area. In the end our final conclusions were reached through the team building up a thorough understanding of the system and its development and then reviewing how confident they were in the system's safety. One thing we never did was to compare the systems, because the difference in their basic design philosophies and development histories would have made that a very difficult exercise.

Over the next few years standards will continue to appear, and methods for assessment will improve, but the problems of assessing and comparing systems

which have been developed over a disperse period of time and geography will always exist and will always require a more pragmatic approach to system evaluation.

References

1. Catmur, J.R., Anderson, P.. Developing a Safety case for signalling systems which employ safety critical software.