Session Five:

Functional Safety and Engineering Judgement

Harvey T.Dearden
Time Domain Solutions Ltd

Abstract

Discussion of the role of professional judgement in the context of the functional safety standards IEC 61508 and IEC 61511.

Introduction

It is the role of a professional engineer, having acquired the appropriate competencies, to exercise professional judgement with due regard to pertinent guidance. In terms of the functional safety standards, engineers should recognise that we approach compliance asymptotically along a curve of diminishing return; we may approach closer and closer to full compliance, but it requires ever increasing effort and investment. There is a point where the marginal increase in compliance does not warrant the additional effort, which may be more gainfully employed on other safety concerns. Professional judgement must be exercised to identify when this point has been reached.

If engineers do not exercise appropriate judgements, there is a real danger of people tying themselves up in allsorts of metaphorical knots in an attempt to establish total, rigorous compliance with the standards, with the result that wrongheaded approaches may be adopted and resources and efforts misdirected from where they may well yield a better safety return.

Note that ‘exercising judgement’ does not mean the same as ‘going out on a limb’, which would imply accepting a significantly higher degree of risk than would otherwise arise; it will often simply come down to a matter of employing some common sense rather than ‘going through the motions’ of strict compliance even when there is little or no benefit (or even conceivably negative ‘benefit’).

The wish of some engineers to avoid any accountability is recognised by some as a marketing opportunity. Occasionally you will find interested parties use ‘scare tactics’ to promote their product or service. It is all very well picking out individual clauses from the standards to ‘demonstrate’ a particular requirement, but without proper consideration of the context and the underpinning philosophy of the standard, it is easy to end up with a wrong headed approach.

If ever you are being told that you need to ‘rip it all out and start again’ whether metaphorically or literally, you need to take a step back and consider whether the ‘logic’ that is being deployed has carried you to a place that is no longer sensible. Do not fall into the trap of assuming a logical proposition is necessarily a sensible one. (Which is itself an error in logic!) Careful consideration of the starting point and the explicit and implicit assumptions being made may reveal that a logical conclusion may be far from sensible.

As an example, consider the following: A company had been hired to do a review of existing safety system provisions against the requirements of IEC 61508. A typical argument deployed in their report ran as follows (here distilled to its essence):

A legacy system has a number of components. None of these is certified as being suitable for deployment in a SIL rated system and the safe failure fractions are not identified. Therefore conservatively assume that the safe failure fractions are less than 60%. But the standard specifies a safe failure fraction of less than 60% is 'Not Allowed' for the specified SIL target. Therefore replace all the components.

Logical, but very far from sensible: The equipment had been procured and installed in accordance with good practice and there was no suggestion that it was not fit for purpose. An intelligent, responsible review, exercising appropriate judgement, would have identified this.

The important thing is to bring a considered, systematic and responsible approach to these matters; Insistence on an entirely rigorous approach where not only the spirit of the standard is met, but also the letter, may well produce unwarranted distortion in deployment of resources, resulting in a net loss of safety. Engineers should recognise that there is often a trade off between rigour and robustness; a more rigorous approach may look good on paper, but may well prove more fragile in operation, with a tendency to suffer inadvertent corruption and with people prompted to make short cuts. This can in turn undermine the wider safety culture within an operation.

Many organisations have been prompted to generate their own guidance, particularly on matters of SIL determination, with a variety of risk matrix and risk graph offerings. These may well be useful as an aid to assessment, but these tools have limited resolution and their 'calibration' may not always be appropriate for a given context. It would be a mistake to simply accept the output of these tools as definitive. These tools should rather be considered as a means of probing plant design and provisions and identifying potential anomalies where provisions appear inconsistent with the risks. If the output of the tool is at odds with your judgement, be prepared to examine the situation more deeply. From where does the apparent discrepancy arise? Is the assessment flawed or should your judgement itself be 'recalibrated' by newly identified considerations? There may well be qualifying circumstances or provisions that have not been properly recognised in the application of the tool. If you are persuaded that there is a real shortfall in safety provisions, then there is a corresponding obligation to address this.

Uncertainty in estimation

The following statement, drawn from the 'Guidance on Risk for the Engineering Profession' published by the Engineering Council [1], highlights this concern:

• bear in mind that risk assessment should be used as an aid to professional judgement and not as a substitute for it

Uncertainty is a feature of many aspects of risk management. We should be particularly aware of this in the context of the functional safety standards. In order to identify the appropriate Safety Integrity Level (SIL) and implement

appropriate design and test provisions, we need to employ a variety of guesses about, for example; hazard consequences, demand rates, failure rates, and safe failure fractions. Some of these guesses may be your own; some may be imported from colleagues. Some may be sourced from vendors or other third party sources. I use the term 'guesses' for dramatic effect, to bring home the point that functional safety is NOT an exact science. Being engineers, we do not simply guess however, we make considered and, where appropriate, substantiated judgements. It is by this means that our guesses earn the superior designation of 'estimates'. They remain estimates however; and there really is no point troubling over the 3rd significant figure, or even possibly the 2nd. Why should the SILs be separated by a factor 10 in terms of PFD/PFH? Why not a factor 2 or 3? It is an implicit acknowledgement of the inexact nature of functional safety; we look to implement risk reduction with an appropriate order of magnitude. Remember that whereas random failures admit of analysis through statistical calculations, the same cannot be said of systematic failures. A particularly pertinent note in 61511 says "The safety integrity level is defined numerically so as to provide an objective target to compare alternative designs and solutions. However, it is recognised that, given the current state of knowledge, many systematic causes of failure can only be assessed qualitatively."

Do not be fooled by the nature of the standards (in particular their bulk!), the way they are written or the apparent rigour with which the requirements are formulated, into thinking that functional safety is anything like an exact science. The standards are a challenging read and it can be difficult to see the wood for the trees. The real thrust (wood) of the functional safety standards is the requirement that you approach your assessment of protection provisions in a systematic, considered, risk-based manner. Prior to the introduction of the standards, process plant operations were not giving rise to exceptional numbers of fatalities; established good practice has not been rendered obsolete, although on seeing the scope of the standards you could be forgiven for thinking previous practice must have proved woefully inadequate.

The standards do not discuss the uncertainties of estimates and their propagation through dependent analyses. The implication is that estimates should have a suitably narrow uncertainty at a suitably high confidence level. But the same can be said of any engineering calculation. This is an area that is rightly left to professional judgement. Any attempt to place this on a more rigorous footing would complicate matters enormously without bringing corresponding benefits. From a practical perspective, the uncertainty in estimates employed in functional safety studies can be very broad; we are essentially looking to establish an appropriate order of magnitude of risk reduction. An estimate may typically be out by a factor 2, or 3 or more without critically affecting the analysis. Periodic reviews against operational experience (as called for in the standards) protect against unduly optimistic estimates.

There is a note in the standards that says estimates for failure rate should have a confidence level of at least 70%. The implication here is that there should be a less than 30% chance that the actual failure rate is higher than your estimate. But without a population of estimates and some notion of their distribution this is not very helpful. And simply specifying a confidence limit does not constrain the factor by which our estimate could be exceeded, it only constrains the probability of its being exceeded. None of this need trouble us however. This

is the modern equivalent of debating how many angels can dance on the head of a pin. Attempts to introduce such mathematical rigour are failures to recognise the nature of the subject and the real aim of the standard; a considered, systematic, risk based evaluation of protection provisions.

Given the tendency to conservatism in both the assessment of the hazard consequences (and therefore their tolerable interval), and demand rates placed upon protection systems, there is often a corresponding tolerance in under estimation of failure rates. If tolerable interval and demand rate are both conservatively high by a factor of only 1.7, we may use a failure rate figure that is low by a factor 3.0 (being $\approx 1.7^2$) and still deliver the risk reduction required for the hazard.

Confidence in equipment compliance

Some will place their faith in component design certificates as a means of gaining assurance, but a certificate is simply a declaration of someone else's estimates. They are one possible means of acquiring confidence that equipment is fit for purpose, but they are not the only way, and not necessarily to be preferred. Certification is very variable in its quality. Some certificates only address the hardware and do not cover any software deployed, which rather limits the usefulness of the certificate to the user (it may help substantiate an estimate of random hardware failure rate). Some may certify the equipment under conditions that are not representative of expected use. A critical review of a certificate (or its underpinning report) may reveal deficiencies in the certification process, particularly in respect of the software. The certified declarations may be a starting point, but may need qualifying in respect of the anticipated deployment of the equipment.

The alternative approach is to use a 'prior-use' argument in accordance with IEC 61511. There are any number of people that will queue up to tell you how hard the prior use (PU), argument is to deploy. (Most with an interest in certification.) Some shameless scaremongering has been employed. Is it really so difficult? Consider; why would the standard committee include the PU provision and then make it next to impossible for the user to adopt as a practicable option? The compilation of PU substantiation is not necessarily that onerous, particularly for SIL1 & 2.

Some people seem to regard prior use as the poor relation of design certification, but this is not correct; good experiential evidence is to be preferred to theoretical prediction and is ultimately the only means of validating reliability modelling. FMEA studies are routinely employed in certification of equipment failure rates, but these assessments will be for operation under reference conditions and will not include; scaffold pole strikes, being stepped upon, being leaked upon, process/environment excursions, process connection failures, maintenance failings etc. An estimate based upon established experience, even if not offering the same rigour in methodology, may well be superior to the 'rigorous' theoretical assessment. You may also find that suitably certified equipment is simply not available; your only recourse will be to a prior use argument.

In terms of a system (as distinct from a component), it is questionable whether you should look to a certificate (or audit report) to gain complete assurance of compliance of a system with a given SIL. The scope of the standards is so broad, the lifecycle so extensive, the detailed provisions so pervasive (including project procedures remember), that total compliance is hardly to be believed. It is an aspiration that we approach asymptotically. Look for compliance in all essentials. Look for compliance to the point of being fit for purpose. Exercise your professional judgement.

Do not attempt to acquire a 'get out of jail free' card by writing '…must comply with IEC 61508/61511' into your procurement specifications. This supposedly catch all clause is not helpful and is an abdication of responsibility. You should explicitly include appropriate provisions from your identified system safety requirements.

Sometimes it may be appropriate to refine your estimates with Quantified Risk Assessment, using tools such as fault tree studies. But do not kid yourself; they remain estimates, albeit of a higher order. The point is that instead of simply, for example, judging hazard rate say as belonging to a certain band of frequencies (possibly based on some back-of-envelope calculation), you compute the rate for the top hazard event as a function of a larger number of component estimates at the bottom of your fault tree.

Compliance with lifecycle model

The standards identify a lifecycle model for the design and implementation of safety functions which can be used to inform the approaches to be adopted with project procedures. But complete compliance may well be an unrealistic proposition; the model represents an ideal which may not be realisable on a practicable basis in all aspects. A clause within part 1 explicitly recognises the difficulty*: '7.1.1.4 The overall, E/E/PE system and software safety lifecycle figures (Figures 2 to 4) are simplified views of reality and as such do not show all the iterations relating to specific phases or between phases. Iteration, however, is an essential and vital part of development though the overall, E/E/PE system and software safety lifecycles.'* Common sense should prevail. Actual project documentation and execution may be perfectly reasonable without necessarily mapping directly to the model.

Conclusion

Before you undertake a program of work to introduce more rigorous compliance, or refine your calculations or enhance the substantiation for your estimates, be careful to consider whether the additional efforts are warranted and whether you will be adding useful value. You may spend a great deal of time and effort agonising over the possible refinements, but the apparent improvement in precision may be illusory, or even if real, may be largely irrelevant in the face of other uncertainties. Recognise that functional safety is not an exact science. Do not look for 'bombproof', absolute assurances of compliance in every particular. Once you have acquired the appropriate competencies, be willing to exercise your professional judgement. It is then possible to meet the spirit of the standards in a relatively straightforward manner. The functional safety standards are so comprehensive in their scope

and so detailed in their provisions that absolute compliance in every particular is hardly to be believed, nor necessarily to be required.

References

[1] Guidance on Risk for the Engineering Profession, Engineering Council, (March 2011).