ACHIEVING ALARP WITH SAFETY INSTRUMENTED SYSTEMS

C.R. Timms,

MIEE, United Kingdom, Tel: + 44 (0) 1339 886618, Email: c.timms@ifb.co.uk

Keywords: ALARP, hazards, risk, safety, SIS.

Abstract

This paper sets out a methodology for setting tolerable risk levels, for various methods of Safety Integrity Level (SIL) determination, to meet the principles as low as reasonably practicable (ALARP). It makes proposals on how to deal with the tolerable risk concept for safety instrumented systems (SIS) protecting against single hazards.

1 Introduction

Safety Instrumented Systems (SIS) are one of the most commonly used methods of reducing the risks associated with major accident hazards in the process and other sectors. They can be found in various systems such as emergency shut down, fire and gas and machinery protection. A single SIS normally provides protection against a single hazard, and this poses a dilemma for designers when they are trying to fulfil the overall requirements for reducing risk to as low as reasonably practicable (ALARP) (see Figure 1), since the concepts of ALARP are concerned with the total risk from **all** likely hazards to workers or the public.

However, SIL determinations for safety instrumented systems are processed on a hazard by hazard basis. Thus these SIL risk assessments are all based on protective functionality for single hazards and not the total risk posed by all likely hazards corresponding to ALARP.

Obviously the residual risk for any single hazard must still be tolerable but additional allowances must be made to ensure that the combined risk from all likely hazards will meet the ALARP concept. This paper makes proposals for calibrating and using various risk assessment methods to achieve levels of tolerable risk, associated with any single hazard, to try and meet the principles of ALARP.

The paper discusses the determination of safety integrity levels for safety functions where failure results in injury/fatality of workers and also the societal risk aspects when members of the public could be injured due to failure of the safety function. It develops further some of the topics and issues raised in 'Determination of Safety Integrity Levels Taking into Account ALARP-Cost Benefit Analysis', 2006 [1] presented at the Hazards XIX conference earlier in the year. The following methods will be included in the context of setting tolerable risk levels:

- Fault Tree Analysis
- Risk Graphs
- Risk Matrices
- Layers of Protection Analysis (LOPA)

2 Regulation compliance

In the European Community the European Seveso II directive, 1996 [2] requires all hazards to be identified and risks to be reduced in line with the ALARP principles, as does the BS IEC 61511, 2003 standard [3]. The UK HSE offers guidance in their publication 'Reducing Risks Protecting People (R2P2), 2001' [4] for the purposes of the Statutory Instrument 'Control of Major Accident Hazards Regulations 1999 (COMAH)' [5] which is the UK implementation of the European Seveso II Directive. However, in all the guidance, the concept of ALARP is concerned with the total risk from all likely hazards to workers or society. In Figure 1, the thresholds between the intolerable and tolerable region, and the tolerable to broadly acceptable regions, are indicators for the total risk to the most exposed employee or society at large.

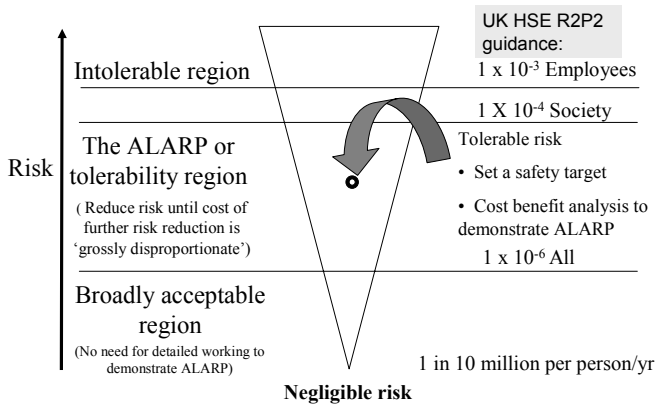

Figure 1 – ALARP Principles

The primary objectives are to ensure that the correct safety functions have been identified from the risk assessment, and that the functionality of these safety functions is specified and implemented correctly to reduce the residual risk posed by any hazard to a tolerable level. The difficulty is determining what the tolerable risk level should be for any single hazard.

In addition, tolerable is not the same as acceptable, it is a level at which there is a willingness to live with the risk in order to obtain the benefits.

3 Risk Assessment Methods

The risk assessment method may be either qualitative or quantitative and it does not have to be quantified if it can demonstrate that the residual risk is better than, or close to, the 'broadly acceptable' region.

However quantified or semi quantified methods are usually used when the levels of risk are quite high and include:

- Fault Tree Analysis
- Semi quantified Risk Graphs
- Semi quantified Risk Matrices
- Layers of Protection Analysis (LOPA)

Experience shows that a semi quantified approach such as a risk graph or risk matrix is most appropriate to undertake the initial risk assessment for large installations. They are simple to interpret, relatively fast to use and they also produce conservative results. This is due to their rigid framework, as this limits the number of parameters for which risk reduction credit can be claimed. A more detailed analysis of the higher SIL results can then be undertaken by methods such as Fault Tree Analysis or LOPA. Fully quantified methods such a Fault Tree Analysis tend to be time consuming and when a whole process plant is to be analysed then it become an impracticable exercise.

Whether quantified or semi quantified methods are used, they require a numerical value for the target or objective residual risk. This then allows the analysis to determine if sufficient risk reduction is achieved. This tolerable risk objective will be referred to as the Safety Target.

4 Setting Safety Targets

Although most major organisations will probably have established corporate values for 'tolerable risk', this is often a new experience for the majority of end users of Safety Instrumented Systems (SIS), and an area which causes considerable problems.

None of the UK regulations offer a 'typical' value for a tolerable risk and this is understandable since it will be dependant on the disproportional cost of further risk reduction, and this will never have any 'typical' value.

The HSE document 'Reducing Risks, Protecting People, 200' (R2P2) (paragraph 128) indicates the upper limit boundary between Tolerable and Unacceptable risk, for workers, would be 1 in 1,000 (1.0E-03) per annum, as shown in Figure 1.

In (Paragraph 130) of R2P2 the boundary between Tolerable and Broadly Acceptable risk, for risks entertaining fatalities, for both workers and public is indicated as 1 in 1,000,000 (1.0E-06) per annum as this corresponds to a very low level of risk, as shown in Figure 1.

These boundaries are based on total risk to the most exposed individual and the dilemma faced by engineers involved in SIS is to understand and/or calculate where the tolerable risk, or safety target, should sit within the upper and lower boundaries for a single hazard.

Where a major accident hazard could impact on both workers and the public then it is necessary to establish tolerable risk levels for both. It is generally expected that the tolerable risk for the public, often referred to as societal risk, will be at least ten times less, i.e. more restrictive than for workers.

Risk assessment by SIL determination only makes an assessment of the risk associated with a single hazard whilst ALARP is concerned with the most exposed person to ALL risks per annum. It is practically impossible for the SIS designer to estimate how many simultaneous risks an individual might be exposed to when considering a design to protect against a specific hazard in a particular area, as 'all risks' embrace a wide spectrum of hazards including slips, trips and falls. It is simply not practicable to try and resolve such a complex problem for all workers and society for every individual safety instrumented function (SIF).

Thus some rationale needs to be established to make allowances for the fact that only single hazards are analysed during SIL determination. It may be possible to try and estimate all the hazards that might be encountered by the most exposed individual, under normal routines and occupancy, and then take an average for the Safety Target.

Another option is to use judgement to set the single hazard Safety Target a factor more sensitive than a corporate all risks per annum tolerable risk target. This would make allowances for these uncertain multiple risk conditions; e.g. if a corporate all risks per annum was set at of 1 in 1000 (1.0E-03), for the most exposed individual, then if a factor of ten was used the single hazard tolerable risk would be reduced to 1 in 10,000 (1.0E-04). The HSE do not currently provide guidance on the scale of this factor, and opinion ranges between 10 and 100, but there is general agreement in the industry that a factor of at least 10 is sufficient.

Where the consequences of a hazardous event impact upon the public then a Societal Risk assessment is required. Each operating organisation must set their societal risk criteria and this is quite commonly set to be a factor of ten times lower than for on site workers; e.g. If the tolerable risk for workers was set at 1.0E-04 then the tolerable societal risk could be 1.0E-05. Thus the suggested factor of ten risk reduction for single hazards can be applied to both the tolerable risk for workers and society.

If risk is being assessed for a multiple proximity plant site, then the target for society/public needs to be proportionally further reduced, i.e. if there are ten plants in close proximity

then the public risk target needs to be a factor of another ten times lower. This would not be the case for workers as they usually work on one plant with exposure limited to hazards from that plant. However, consideration needs to be given too any potential for domino escalation between sites.

4.1 Safety Targets and Risk Reduction

Risk analysis determines if existing risk reduction measures are sufficient. Risk reduction is generally made up of a number of different layers and some typical examples are shown in Figure 2.

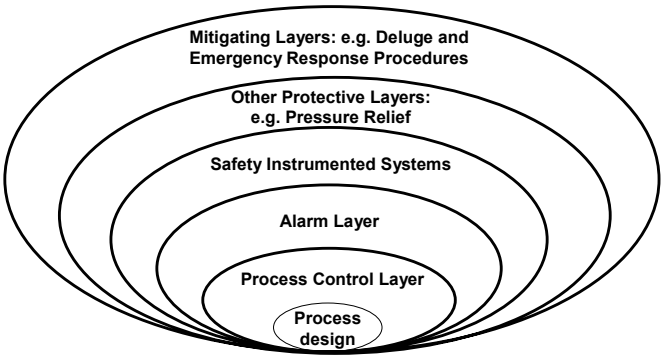

Figure 2 – Risk Reduction Layers

Total risk reduction is a product of the individual risk reduction measures. Thus each risk reduction layer normally has a quantified value attributable to it. Having established a Safety Target or tolerable risk for any single hazard, the objective of the risk assessment is to estimate the risk reduction achieved by existing measures, in numerical terms, and to compare this with the Safety Target such that:

$$\frac{\text{Safety Target}}{\text{Product of Existing Risk Reduction Measures}}$$

This will indicate if additional risk reduction measures are needed and the value of required risk reduction required. The relationship between the required risk reduction and the probability of failure on demand (PFD) used in SIS design is:

1 / Risk Reduction = PFD

The Relationship between SIL, risk reduction and PFD is provided in BS IEC 61511 and is shown in Table 1.

SIL	Risk Reduction	PFD
1	10 - 100	1.0E-1 – 1.0E-2
2	100 - 1000	1.0E-2 – 1.0E-3
3	1000 - 10000	1.0E-3 – 1.0E-4
4	10000 - 100000	1.0E-4 – 1.0E-5

Table 1 – SIL, Risk Reduction and PFD

The PFD range of values, for a specific SIL, is often referred to as the 'SIL Band'.

5 Setting Safety Targets

This section will discuss ways in which Safety targets can be set for different methods of risk assessment.

5.1 Risk Graphs

Risk graphs are far more useful and consistent when all the parameters are semi quantified. The general arrangement of the BS IEC 61508 risk graph used for a Personnel Safety is shown in Figure 3.

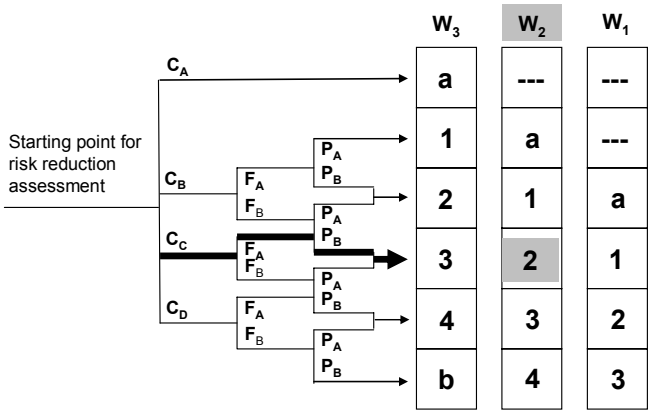

Figure 3 – Risk Graph General Arrangement

The graph uses the following parameters for making the risk assessment:

- C = Consequence severity
  Where C = Number of people x Vulnerability Factor
- F = Occupancy of the hazardous area being considered
- P = Alternatives to avoid the hazard
- W = Frequency of demand on the SIS or likelihood
- 1,2,3,4 = Safety integrity level (SIL)

The actual Safety Target delivered by this kind of graph is best demonstrated by an example shown in Table 2

An example risk graph has been assigned parameters with the values opposite: A 'typical SIL determination has resulted in a risk path with the highlighted parameters. Each parameter has a range value with a high and low end values	**Parameter**	**Low Value**	**High Value**
	W3	0.3	3.0
	W2	**0.03**	**0.3**
	W1	0.003	0.03
	CA	0.001	0.01
	CB	0.01	0.1
	CC	**0.1**	**1.0**
	CD	>1.0	
	FA	**0.01**	**0.01**
	FB	0.1	1.0
	PA	0.1	0.1
	PB	**1.0**	**1.0**
	SIL 1	0.01	0.1
	SIL 2	**0.001**	**0.01**
	SIL 3	0.0001	0.001
	SIL 4	0.00001	0.0001

Table 2 – Example Risk Graph Calibration

The calibration assigned to each parameter has a range of values with high and low end value. An example SIL determination is highlighted on the risk graph in Figure 3 and in the parameters of Table 2. These result in a specific calibration for the risk graph.

The 'high' end and 'low' end values of any parameter impact on the overall risk assessment e.g. From the example in Table 2, the frequency of demand parameter of W2 ranges from 0.03 years to 0.3 years. Obviously the risk is less if the demand is 0.03/year than if it was 0.3/year.

If all the parameters were at their 'low' end of range this would result in the highest level of achievable risk reduction of 3.0E-08, or best case as shown in Case A:

(Low CC) * (Low FA) * (PB) * (Low W2) * (Low SIL 2)
Thus:
Case A = (0.1) * (0.01) * (1.0) * (0.03) * (0.001)
Case A = 3.0E-08

If all the parameters were at their 'high' end of range this would result in the lowest level of achievable risk reduction of 3.0E-04, or worst case as shown in Case B:

(High CC) * (High FA) * (PB) * (High W2) * (High SIL 2)
Thus:
Case B = (1.0) * (0.1) * (1.0) * (0.3) * (0.01)
Case B = 3.0E-04

Taking the two extreme cases the average risk reduction can be calculated as follows:

The risk reduction afforded by the best case:
Case A = 3.0E-08

and

The risk reduction afforded by the worst case:
Case B = 3.0E-04

Since the example risk graph is based on a logarithmic scaling the average risk figure is the logarithmic average of the best and worst risk values:

Ln(Average risk) = (Ln A +Ln B)/2
Ln(Average risk) = ((Ln(3.0E-08) +Ln(3.0E-04))/2

Average risk for this example = 3.0E-06 for any single hazard and represents the Safety Target achieved by this risk graph example.

Similar risk graphs with specific tolerable risk targets for societal consequences, asset loss and environmental consequences usually form a complete risk assessment.

The Safety Integrity Level (SIL) which is delivered by this type of risk graph represents the additional risk reduction that is required to meet the Tolerable Risk for a single hazard i.e. the Safety Target.

5.2 Risk Matrix

Risk matrices afford a far more limited risk assessment than risk graphs, but they can still be semi quantified to help with assessment consistency. Figure 4 shows a basic 5 x 4 matrix.

Consequence		Likelihood (Demands Per Year)			
		W3 3.0-0.3	W2 0.3-0.03	W1 0.03-0.003	W0 <0.003
CA	None	–	–	–	–
CB	0.01- 0.1	1	a	–	–
CC	0.1- 1.0	2	1	a	–
CD	1.0-10.0	3	2	1	a
CE	>10	4	3	2	1

Figure 4–Example Risk Matrix

There are now only three parameters to be considered:

- W = Likelihood
- C = Consequence severity
  Where C = Number of people x Vulnerability Factor
- 1,2,3,4 = Safety Integrity Level (SIL)

As with the risk graph example, each parameter has a range of values. To understand how the calibration works it is best demonstrated by an example as shown in Figure 4, where a SIL assessment has resulted in W2, CC and a SIL 1. Then if all the parameters were at their 'low' end of range this would result in the highest level of achievable risk reduction of 3.0E-05, or the best case as shown in Case A:

Case A = (Low CC) * (Low W2) * (Low SIL 1)
Thus:
Case A = (0.1) * (0.03) * (0.01)
Case A = 3.0E-05

If all the parameters were at their 'high' end this would result in the lowest level of achievable risk reduction of 3.0E-02, or the worst case as shown in Case B:

Case B = (High CC)* (High W2) * (High SIL 1)
Thus:
Case B = (1.0) * (0.3) * (0.1)
Case B = 3.0E-02

Since the risk example matrix is based on a logarithmic scaling the average risk figure is the logarithmic average of the best and worst case risk values:

Ln(Average risk) = (Ln A +Ln B)/2
Ln(Average risk) = ((Ln(3.0E-05) +Ln(3.0E-02))/2

Average risk for this example = 9.0E-04 for any single hazard and represents the Safety Target achieved by this risk matrix example.

5.3 Fault Tree Analysis

Fault Tree Analysis is probably the most straightforward of all risk assessment methods for determining the difference between any existing risk reduction and a Safety Target.

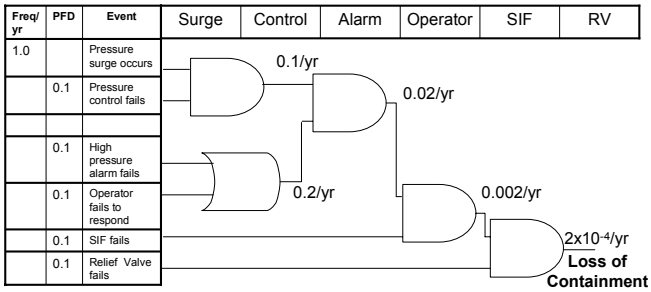

Figure 5–Example Fault Tree Analysis

Taking a very simple example as shown in Figure 5 the likelihood of an undesirable loss of containment event is calculated at 2.0E-04/year. This can be compared with the required tolerable risk, and adjustments can then be made to the SIF PFD, or other measures taken as required.

5.4 Layers of Protection Analysis (LOPA)

LOPA is rapidly becoming a popular method for risk assessment and SIL determination and guidance on application can be founding the American Institute of Chemical Engineers Centre for Chemical Process Safety document -'Layer of Protection Analysis Simplified Process Assessment', 2001[6]. A simple example of a safety based LOPA worksheet is shown in Figure 6.

Impact Event		**Initiating Cause**	**Initiating Cause**
Overpressure and loss of containment from First Stage Separator		Blocked Outlet	Process Control
	Event Frequency	0.1/year	0.2/year
Protection & Mitigation Layers	Process design	0.5	0.5
	Relief valve	0.1	0.1
	Process Control	0.1	-
	Occupancy	0.1	0.1
	Independent alarm	0.5	0.5
	Intermediate Event Frequency	2.5E-04	5.0E-04
	Total Mitigated Event Frequency	7.5E-04	
	Tolerable Event Frequency	1.0E-05	
	Required SIS risk reduction	**1.0E-05/7.5E-04 = 1.3E-02 (SIL 1)**	

Figure 6– Example Layers of Protection Analysis

As with risk graphs LOPA must be quantified to provide meaningful and consistent results. A tolerable risk or Safety Target must also be set. The consequence severity and likelihood of a hazardous event are quantified for every cause/consequence pair, in the absence of any protection or mitigation measures. Credit is taken for risk reduction afforded by independent protection layers and also factors which are considered to make a contribution towards mitigating the consequences. Each layer, for which risk reducing credit is taken, is assigned an appropriate quantified risk reduction factor. Conditional modifiers are also used to credit further risk reduction which might only be specific to workers, society, the asset or the environment; e.g. a bund round a tank or closed drains will reduce the environmental impact of a spillage but have no effect on the asset loss. Some typical risk reduction layers are shown in Figure 2 and might include:

- Inherent process design factors;
- Process control systems;
- Alarm systems if independent from the process control;
- Existing safety instrumented functionality;
- Personnel occupancy of the area;
- Mechanical protection such a pressure relief;
- Mitigating measures such as fire and gas systems;
- Safety procedures;
- Emergency response procedures.

The total mitigated event frequency is compared with the tolerable risk frequency target that has been set for workers, society, the asset and environment and the difference equates to any additional risk reduction factor to be contributed by the SIS.

6.0 Use of Cost Benefit Analysis (CBA)

When a safety target is set for any risk assessment method then, by achieving the determined SIL for the SIS, the safety target will be met.

However, there still remains the question about whether sufficient risk reduction has been made for satisfying the 'as low as reasonably practicable' requirement i.e. should additional risk reduction measures be applied?

6.1 CBA Background

The UK guidance, on the relationship between 'as low as reasonably practicable' (ALARP) and the use of cost benefit analysis (CBA), can be found in the HSE R2P2 document.

The CBA is based on an estimate of the costs of risk reduction measures and the number of casualties saved by implementation. Thus the cost of preventing a fatality (CPF) takes the form:

$$\text{CPF} = \frac{\text{Total cost of the risk reduction measures}}{\text{Total fatalities prevented}} \quad (1)$$

Then by comparing this with the value of preventing a fatality (VPF) an estimate can be made of the proportion factor:

Proportion Factor = CPF / VPF (2)

When the Proportion Factor is 1 or less (or even 2 or less) then R2P2 advises that additional measures should be implemented.

* *R2P2 Appendix 3 Para graph 13: 'VPF is often misunderstood to mean that a value is being placed on a life. This is not the case. It is simply another way of saying what people are prepared to pay to secure an average risk reduction. A VPF of £1,000,000 corresponds to a risk reduction of 1 in 100,000 being worth £10 to an average individual. VPF is therefore not to be confused with what society, or the courts, might put on the life of a real person or the compensation appropriate to its loss.' 'VPF will vary depending on the particular hazardous situation.'*

This is fine for assessing the benefit where no risk reduction has been previously specified. The CBA is a little more complex when an operator already has specified certain risk reduction measures, but needs to demonstrate whether further risk reduction would be cost effective. This is often the situation when designing SIS. In many cases the cost of existing measures is not known, particularly on legacy or brown field installations.

6.2 Developing CBA for an Existing Safety Target

The difference between the current risk reduction measures and the additional risk reduction, achieved by implementation of further measures, has to be analysed. This involves the additional costs of implementation, the difference in risk reduction achieved and the value for all fatalities prevented over the predicted life time operation of the facility.

In this case from (2):

Proportion Factor = CPF/VPF

Or:

P_f = CPF/VPF

Then:

CPF = P_f * VPF (3)

And since (1)

$$\text{CPF} = \frac{\text{Total cost of the risk reduction measures}}{\text{Total fatalities prevented}}$$

Then :

CPF* Total fatalities prevented = Total cost of risk reduction measures (4)

Substituting (3) in (4) for CPF:

(Pf * VPF) * Total fatalities prevented = Total cost of risk reduction measures (5)

And since the total fatalities prevented is represented by a product of the frequency of demand (F) on the SIS, the probability of failure on a demand (PFD), the operating life of the plant (PL) and the number of fatalities (N) resulting from the hazardous event:

Total fatalities prevented = F x PFD x PL x N (6)

Then substituting for (6) in (5):

(Pf * VPF) * (F * PFD * PL * N) = Total cost of risk reduction measures (7)

Where:

Pf = Proportion factor
VPF = Value of preventing a fatality
F = Frequency of demand on the SIF (for a range use high frequency value)
PFD = Probability of failure of the SIF
PL = Plant operating life
N = Number of fatalities per hazardous event

Where an existing risk reduction proposal has been made through a risk assessment (such as by use of a risk graph or LOPA), then the additional fatalities prevented will be proportional to the difference between the PFD of the existing solution and the PFD with further risk reduction measures.

If

pfd_1 = PFD of existing proposal from risk assessment
pfd_2 = PFD with additional risk reduction measures

Total ADDITIONAL fatalities prevented
= F * (pfd_1 - pfd_2) * PL * N (8)

Then by taking equation (5):

P_f * VPF * Total fatalities prevented = Total cost of risk reduction measures

Substituting (8) for 'total additional fatalities prevented':

Total justified cost of FURTHER risk reduction measures = P_f * VPF * (F * (pfd_1 - pfd_2) * PL * N) (9)

But at what value should the objective pfd_2 be set?

This paper suggests that the ALARP threshold of 'broadly acceptable' is the ultimate objective i.e. 1.0E-06 for both workers and public for **all** risks.

Thus using the factor of ten times more sensitive for any single hazard this would be a pfd_2 of 1.0E-07.

Note.
R2P2 indicates the Proportion Factor 'P_f' should be:

- 10 when working close to tolerable/unacceptable boundary;
- 1-2 when working close to the broadly acceptable boundary.

A value of 1 will be used for P_f as pdf_2 is at the 'broadly acceptable' level of risk.

6.3 Example of CBA for a given Safety Target

The justifiable additional cost for achieving the principle of ALARP is best demonstrated by way of an example:

Where:

- The tolerable risk will be based on the example risk graph calibration described in Section 5.1 i.e. 3.0E-06;
- Value of preventing a fatality (VPF) = £2,000,000 for voluntary (workers);
- The boundary between Tolerable and Broadly Acceptable = 1 in 1,000,000 for both workers and public – *HSE R2P2*.
- The number of **onsite** fatalities estimated due to a major toxic release (N) = 10.
- The frequency of demand (F) = 0.1 (1 in ten years).
- The operational plant life expectancy (PL) = 30 years.

The PFD of current proposal (pfd_1) = 3.0E-06 (Tolerable risk calibration of the risk graph).
The PFD with additional measures (pfd_2) = 1.0E-07 (The 'broadly acceptable' value of 1.0E-06 for public and workers increased by a factor of ten for the single hazard analysis).

Then by using equation (9):

Total cost of risk reduction measures (C_t)
$= P_f * VPF * (F * (pfd_1 - pfd_2) * PL * N)$

$C_t = 1.0 * 2{,}000{,}000 * (0.1 * (3.0E\text{-}6 – 1.0E\text{-}7) * 30 * 10)$

$C_t = 174$

Thus the total discounted cost of further risk reduction measures would need to be below £174. Therefore, in this example, additional measures would be implemented if the total cumulative discounted cost over the plant/project life of 30 years was below £174. This also demonstrates that the calibration of the safety target, for the risk graph used in this example, is very close to ALARP.

7.0 Finding the Optimal Safety Target

The SIS engineer really needs to know where the optimal Safety Target is before setting out on the risk assessment process. If this could be determined then calibrating the chosen risk assessment method to the optimal Safety Target value, and designing the SIS to achieve the target, would ensure compliance with the ALARP principles.

This cannot be done by simply making a 'stab in the dark' but making multiple iterations for the CBA calculations could be prohibitively time consuming. The optimal Safety target can be found by using equation (9) to plot a range of initial PFD options (i.e. pfd_1) to a 'broadly acceptable' PFD value (i.e. pfd_2). The VPF can be set to the company value and the number of fatalities per hazardous event can be assessed by the normal risk assessment methods.

This is again best demonstrated by an example:

- VPF = £2,000,000
- N = 10
- PL = 30 years
- F = 0.1/year

If a range of pfd_1 values is plotted for 10 fatalities against a 'broadly acceptable' risk for pdf_2 of 1.0E-07 (for a single hazard) then Figure 7 indicates that further risk reduction measures of many millions of pounds would be justified if the current risk reduction measures achieved less than 1.0E-2. It also indicates that there is a distinct flattening of the curve at around 1.0E-04 indicating the region of the optimal Safety Target and achieving greater risk reduction than this value is likely to satisfy ALARP principles.

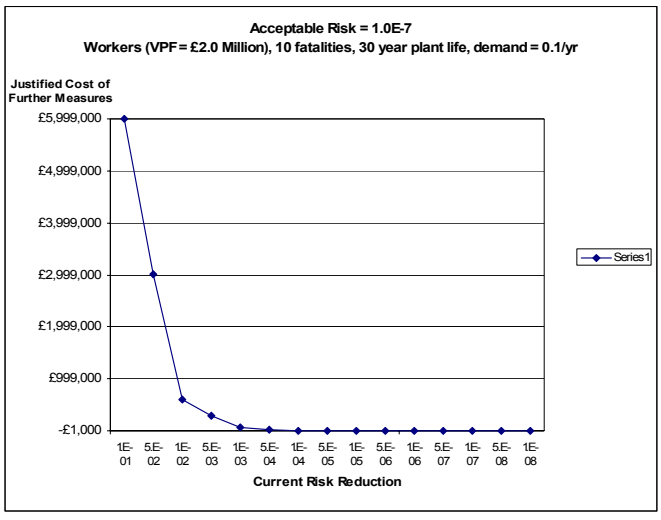

Figure 7a – Justified Cost of Further Measures / Current Risk Reduction

If the curve is plotted for a single fatality over the same range as in Figure 7b the scale of the justified cost of further risk reduction measures decreases, as would be expected, but the knee of the curve is identical.

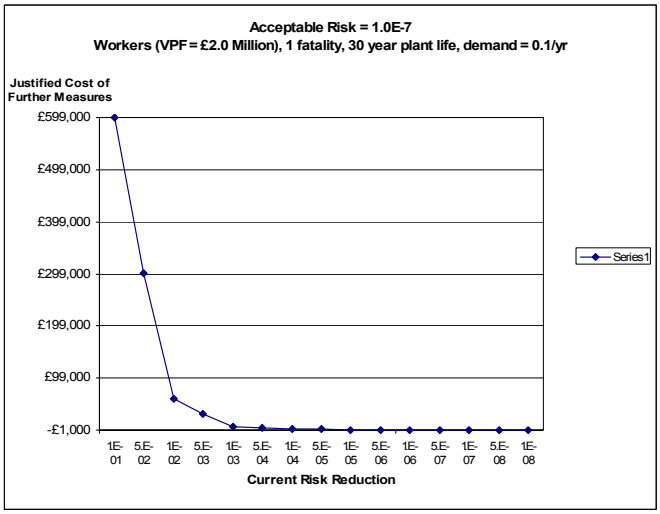

Figure 7b – Justified Cost of Further Measures / Current Risk Reduction

It would appear to make a great deal of sense to plot a curve like Figures 7a and 7b, in preparation for a risk assessment exercise. It then only requires a simple cross check to determine the justifiable cost of further risk reduction measures to comply with the ALARP principles. Ideally, the Safety Target would be set to a value low down in the knee of the curve so that the outcome of the risk assessment would then deliver SIL requirements meeting ALARP principles.

Similar curves can be plotted to take account of the appropriate level of VPF for hazards that put the public at risk.

8 Discussion

There is a considerable gap between the current guidance on ALARP, which is concerned with total annual risk from all hazards, and the situation facing designers of functions within safety instrumented systems protecting against single hazards.

Designers are currently forced to develop their own rationale for implementing SIS to meet the ALARP principles for protection against single hazards.

This paper has made suggestions for setting tolerable risk targets, or Safety Targets, for single hazards, and it has also indicated how these can then be developed to demonstrate ALARP through cost benefit analysis. Methodology has been proposed for assessing the difference between the proposed, or current, risk reduction measures, and what would the justifiable cost of additional risk reduction measures, in ALARP terms.

It has also been demonstrated that there is no need for extensive cost benefit analysis providing the risk assessment, and subsequent design is undertaken to suitably set tolerable risk levels.

However, the approaches and methodology that have been described are only one practitioner's perspective. The purpose of this paper is to stimulate discussion and promote the need for further guidance on ALARP from a Safety Instrumented Systems perspective.

References

[6] American Institute of Chemical Engineers Centre for Chemical Process Safety, Layer of Protection Analysis Simplified Process Assessment, 2001.

[3] BS IEC 61511-3: "Functional safety - Safety Instrumented Systems for the process industry sector –Part 3": *Guidance for the determination of the required safety integrity levels*. (2003).

[2] European Seveso Directive, 1982 (Council Directive 81/501/EEC) reviewed 1996 and adopted as Seveso II Directive.

[5] HMSO Statutory Instruments 1999 No. 743. "The Control of Major Accident Hazards Regulations 1999 (COMAH)" - ISBN 0 11 082192 0.

[4] HSE document: "Reducing Risks, Protecting People (R2P2)". (2001) - ISBN: 07176 21 51 0.

[1] C Timms – "Determination of Safety Integrity Levels Taking into Account ALARP- Cost Benefit Analysis" presented at the Hazards XIX conference March 2006.