

ASSESSING THE ROLE OF INSTRUMENTED PROTECTIVE SYSTEMS IN THE BASIS OF SAFETY FOR BATCH REACTORS

Gerry Brennan[1] and Paul Harding[2]
[1]ABB Engineering Services, UK
[2]Baker Petrolite, UK

The publication of recent standards for safety instrumented systems has lead to the need for more explicit understanding of their function and integrity in order to ensure that they are designed and managed correctly. For existing systems this brings about a need for assessment, which poses particular challenges for batch chemical plant which has a flexible operational role. This paper describes the development of a method for determining the required safety integrity levels for safety instrumented systems fitted to batch reactors, and its application to a plant containing a large number of such reactors, of varying design characteristics, and intended for a wide range of possible duties.

KEYWORDS: batch reactors, safety integrity level

INTRODUCTION

Baker Petrolite manufactures chemical products for use in the global hydrocarbon recovery and processing industry. Their site at Kirkby, Merseyside operates multi-purpose batch reactors capable of making a variety of products from a range of raw materials. Some of the reactions carried out are exothermic in nature, capable of thermal runaway, and the hazards resulting from this were assessed in the site's COMAH Safety Report. The Bases of Safety (BOS) for the reactors rely on process control and operational procedures, and a mix of safety features or systems such as vessel containment, pressure relief, and safety instrumented trip systems (SIS). All of the existing SIS were installed prior to the publication of BS EN 61508 Functional Safety of Electrical/Electronic/Programmable Electronic Safety Related Systems in 1998. This standard broadly requires that the function of a SIS should be defined, and that its integrity should be appropriate for this function. While the standard is not part of safety legislation and is not intended to be applied retrospectively, it does now represent good engineering practice for SIS, and there is thus an onus on duty-holders to apply it wherever it is relevant. As is the case in many older plants, these "legacy" SIS posed questions because there was not a sufficiently detailed definition of the role allocated to the instrumented safety layer, in the overall protective arrangements, to enable the appropriate safety integrity level (SIL) to be identified. While in some reactions the role of the SIS is clearly significant, e.g. a low temperature trip on hazardous material additions to prevent accumulation, in other recipes made in the same reactor this SIS would not be a required protection. It is desirable from a production viewpoint to have flexibility of use of the reactors, but having a range of possible duties makes determination of the SIL for the SISs more complex than for reactors with a more dedicated role. The determination of an appropriate SIL is clearly important in

maintaining the BOS, but also because there are lifetime costs associated with SIS which increase as the required SIL increases, and in some cases it may be more appropriate to replace a SIS with another form of protection rather than modify an existing design.

This paper describes how the BOS of this group of reactors was defined more explicitly, so that the role of the SIS could be properly understood, and in so doing a model of the overall protective arrangements was created which can be utilised in future to evaluate the impact of any planned operational changes on the BOS.

PLANT OPERATIONS AND REQUIREMENTS

Baker Petrolite provides quality oilfield chemical additive programmes for all phases of petroleum production, namely drilling, well stimulation, production enhancement, pipeline transportation, refining and maintenance reduction. The company offers tailored, performance based solutions through differentiated chemicals and innovative technology. The products are used to minimise pipeline corrosion, waxing and water in oil emulsions and H2S content, lowering costs, reducing downtime, resolving environmental problems and increasing throughput.

The Kirkby site, based at the Knowsley Industrial Estate is one of two manufacturing facilities in the UK. The site produces a wide variety of chemical products used as part of the Baker Petrolite product portfolio. There are currently 15 batch reactors within the main reactor house of varying size and construction, including glass lined and stainless steel units, making a wide range of products and batch sizes. The operating characteristics are variable with differing operating temperatures and pressures within the processes.

The majority of the current plant was built around 20 years ago and was compliant with best practice at the time. The existing SIS systems have been identified as being part of the overall plant instrumentation and not as safety specific equipment.

Products are manufactured to customer order, with the number and size of batches varying according to demand. The product profile and number of batches can vary from year to year. As part of work to define the BOS an exercise was carried out on the site to identify clearly the number of each type of operation in past years, and an assessment was made as to whether this was representative of the production profile for the reactors or could change significantly in the future.

A previous project had examined the basis of safety for each reaction chemistry and had identified those which could potentially experience strong exothermic reactions capable of becoming uncontrolled. It is essential to understand the characteristics of the reactants being used, and in particular the limits of safe operation, and thus it was identified that there was a need for further experimental testing. This was carried out by an external test house prior to commencement of this project. Having in excess of 500 possible reactions to characterise is a significant problem, both in costs and in the quantity of information to be processed. To overcome this, the products were split into 13 representative chemistry groups, which encompassed all products manufactured. For each representative group a “worst case” recipe was identified by the Baker Petrolite Production and Process Chemistry staff, on grounds such as heat of reaction and complexity of the

process. Where necessary further experimental testing was commissioned until the characteristics of the representative group was thoroughly understood.

METHOD FOR SIL DETERMINATION

Recognising the need for a systematic and thorough review of the role of the safety systems, capable of taking account of the wide range of operations, and assessing performance, a method for SIL determination was identified and developed during the project. This had the following steps.

- Identify the potential causes of reaction runaway
- Identify the protective systems which prevent runaway, developing a fault tree model to represent the logical relationship between the potential causes of failure and the protection which prevents or controls these,
- Decide the maximum tolerable frequency of occurrence of runaway reaction
- Quantify the fault tree using representative failure frequencies
- Take one representative recipe, and adapt the generic fault tree to the process steps; start with the first reactor allocated to this duty, then check for any differences in any other reactor;
- Repeat the previous step for each representative recipe
- Create a failure estimation spreadsheet for each reactor; this takes account of a number of operations for each reaction type, the likelihood of failure from all relevant causes, the effect of the relevant protective systems, and yields a frequency of failure for the reactor
- For each reactor compare the frequency of failure against the maximum tolerable frequency to determine the required SIL for the SIS.

APPLICATION OF THE METHOD

IDENTIFY THE POTENTIAL CAUSES OF REACTION RUNAWAY

Reaction runaway is the situation where the heat generated in a reaction exceeds the capacity of the available cooling means to remove this heat, so that the temperature of the material rises in an uncontrolled way. This rise in temperature causes an increase in pressure in a closed reactor vessel, and if unchecked leads to a vent of the material through the relief system, and ultimately, if this is not effective, to a rupture of the vessel. The energy release on rupture depends on volume and burst pressure, and can cause significant injury and damage.

A team of Baker Petrolite staff comprised of engineering, production, process chemistry, operations and safety was assembled, facilitated by ABB as hazard study leader. The team began by examining records of past incidents or near misses within the company, and relevant external cases (EPA 1999). Next the batch operations themselves were reviewed. All batch recipes are described in an operating procedure, which gives a practical description of the activities involved in production. However it was impractical to consider each

of these individually, because of the large number of recipes in use and hence the team developed a generic definition of the main batch reaction steps and some key permissives. These were:

1. Charge reactor with quantity W1 of raw material A.
2. Start Agitator.
3. Charge reactor with quantity W2 of raw material(s) B (&C).
4. Heat to temperature T1.
5. Add catalyst and confirm addition.
6. Confirm cooling available.
7. Add reactive raw material M at rate Q1, (maximum quantity W3).
8. Confirm exotherm has started.
9. Continue to add M under cooling, rate not to exceed Q2, determined by Temperature between T2 to T3.

These steps were found very useful, along with the lessons from past incidents and other information on the possible outcomes of failures, in identifying the relevant causes of runaway reaction, which can be summarised as shown in Figure 1. A principal cause of unplanned exotherm was seen as accumulation of reactive material, due to causes such as too high an addition rate, addition at too low a temperature, loss of agitation in the reactor, or a delay in the start of the exotherm. Loss of cooling to the reactor during the reaction stage could also give rise to unplanned exotherm. The team identified that for a small number of recipes there is the potential that if the temperature of the batch went outside the permitted range, due to not following the operating procedure correctly, contamination, or possibly external fire, then a secondary exothermic reaction could begin with similar consequences to loss of control of the primary reaction. Finally there was the potential that unintended materials could be added in error which might react exothermically with the intended ingredients.

IDENTIFY THE PROTECTIVE SYSTEMS

Each of the causes of unplanned exotherm was now considered in detail to identify which of the hierarchy of protective systems available would be effective in stopping the undesirable consequences from being realised. The logical structure is important here because some protective systems guard against a specific failure, e.g. high flow rate trip system, or loss of agitation trip system. These and other trip systems could stop some unplanned exotherm events from proceeding, whilst other protective layers such as the strength of the reaction vessel or the pressure relief system would protect against catastrophic failure from unplanned exotherm whatever the cause, upto the limiting strength or capacity.

Accumulation of a certain amount of reactive material is a normal part of the process, and accumulation becomes problematic only if it reaches the excess condition where the reaction could no longer be controlled by the cooling system. Thus excess accumulation and failure to have sufficient cooling capacity available are linked, and could conveniently be considered together in the fault tree analysis. Secondary decomposition was considered

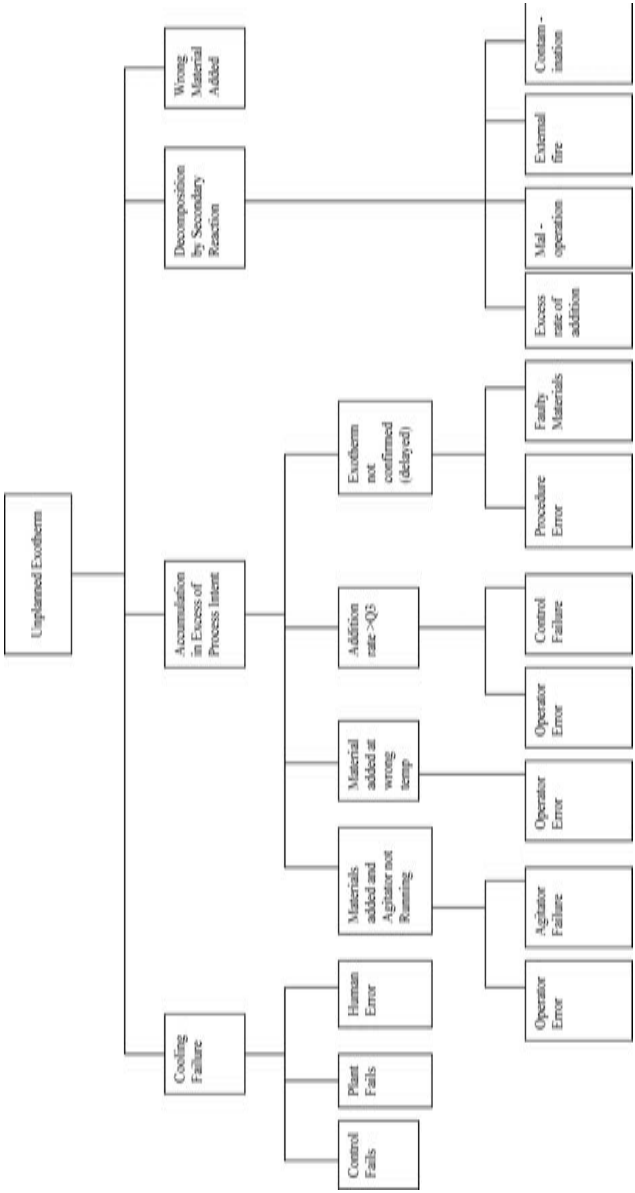

Figure 1. Causes of unplanned exotherm

separately. Addition of the wrong material was only feasible through operator error, and a separate study was begun to consider in detail in which reactions this was possible, and what were the safeguards against it.

MAXIMUM TOLERABLE FREQUENCY OF OCCURRENCE OF RUNAWAY REACTION

Criteria for the tolerable level of industrial risk in the UK today are published in the UK by the HSE (HSE 2001), and distinction is made between risk to the workforce, and risk to the general public. The boundary between the intolerable and the ALARP categories of risk is regarded as an individual risk of death of 10^{-3} per year for members of the work force and 10^{-4} per year for members of the public. The boundary between the ALARP and broadly acceptable categories is regarded as an individual risk of death of 10^{-6} per year for both workers and the public. Lesser consequences than fatal injury are not subject to published criteria.

The first step in deciding how to use these criteria in determining the SIL required of a SIS is to understand the worst credible consequence which the SIS guards against. The site Safety Report considered uncontrolled exotherm as having a number of possible outcomes. Runaway events which are contained in the reactor, and which do not lead to rupture or release through the relief system, are regarded as having negligible severity consequences. Runaway events culminating in a release through the relief venting system would have at most low severity consequences. However reactor failure at elevated pressure under runaway conditions, if it occurred, could be a catastrophic event, capable of causing offsite fatalities by virtue of the overpressure created by the reactor explosion. This conclusion was considered very pessimistic but for consistency with the safety report it was decided to adopt it, rather than re-visit the consequence modelling carried out previously.

For the purposes of this study, any reactor vessel where the maximum foreseeable pressure from a runaway event, during a reaction operation for which it was allocated, ignoring the presence of the relief system, could exceed the design pressure of the reactor was considered capable of catastrophic failure. The actual likelihood of such a failure is reduced by safety factors incorporated into vessels built to established codes, as well as safety systems such as pressure relief, and these can be taken into account in the assessment of predicted failure frequency. Failure at the pressure corresponding to the design rating would also have less damaging consequences than the failure at elevated pressure examined in the safety report, but this was not taken into account in the assessment.

For any reaction chemistry the maximum vessel pressure which could arise is considered to be the same for any reactor in which it is performed. The rupture of any reactor under runaway conditions was again pessimistically regarded as catastrophic, regardless of size.

As noted above the overall individual risk to a member of the public caused by a process industry site should never exceed 10^{-4} per year risk of death. If this is taken to

be a total risk due to the site, then clearly the risk contribution from any one of the hazards on the site must not exceed some fraction of this total. The allocation of this fraction is an important issue on which there is currently no published guidance from the authorities. In the absence of guidance it is felt that a reasonable approach is to say that the maximum tolerable frequency of occurrence for a single hazardous event should be such that the sum of all hazardous events should not approach the threshold, and this will normally lead to an overall risk level well within the ALARP region. In this case there are 15 reactors to consider, which each contribute to the risks posed by the site. Although it is in fact unlikely, for operational reasons, that all of the reactors would be at the active reaction stage at the same time, it was decided that each reactor should have a maximum frequency of occurrence of catastrophic failure which is not more than 1/100th of the threshold value, or 1×10^{-6} per year. This will mean that in total, over all reactors, the total frequency of catastrophic runaway events will be 15×10^{-6} per year which represents less than one fifth of the threshold of tolerable risk for catastrophic events.

Allocation of the reactors to chemistry groups showed that every reactor could have at least one reaction duty where the chemistry was such that the reaction characteristics met the criterion for credible failure. This allocation of duties provides a high degree of operational flexibility, and this situation was seen as presenting an opportunity whereby some reactors could be restricted in duty, rather than upgraded or improved, should this be shown to be necessary.

Where an instrumented protective system guards against hazardous events with a severity outcome which is less than major injury (fractures, unconsciousness, loss of a limb) it is assumed that an "ungraded" system will provide adequate integrity, and definition of a maximum tolerable frequency is not necessary. Exceptions to this are where it is considered that the demand rate may be high, or where the system is required for asset protection requirements, where SIL 1 is considered appropriate.

QUANTIFY THE FAULT TREE USING REPRESENTATIVE FAILURE FREQUENCIES

Having identified the logical relationship between causes of failure and protective systems the team then proceeded to discuss the likelihood of failure, and ascribe preliminary values to each of the events in the fault tree analysis. External sources of data provided the starting point for this discussion, and useful references include Kirwan 1994, and Gertman 1994 for human error, and Lees 1996 and CPR 12E1997 for equipment failure data. Failure data estimated from actual site experience was used wherever it was felt that justifiable conclusions could be drawn. For example, one possible cause of delayed exothermic reaction was identified as "contaminated material". Past experience, where a supplier provided a substance of greater strength than required causing an unplanned temperature rise that was nevertheless within the range of the cooling system to control was used to derive a probability of this failure arising in the operations concerned.

APPLICATION OF THE FAULT TREE TO A REPRESENTATIVE RECIPE/REACTOR COMBINATION

The generic fault tree was compared with the recipe for the representative product from the first chemistry group, being carried out in the first reactor of the four possible reactors allocated to this duty, and it was checked that the failure and protective system assumptions were appropriate. Where necessary the fault tree was tailored to suit the details of the operation, and this fault tree variant was saved as a characteristic fault tree for this group. The team also brainstormed whether any recipe-products of the group were significantly different from the representative case as far the logic of the fault tree analysis was concerned, and concluded they were not. The comparison was then extended to the remaining three allocated reactors. In this first case there were no significant differences between the operations in the four reactors and so one characteristic fault tree was sufficient to cover all cases.

APPLICATION OF THE FAULT TREE TO ALL REPRESENTATIVE RECIPE/REACTOR COMBINATIONS

The previous step was repeated for each chemistry group, across all the relevant reactors in which they could be made. In all it was found that 16 characteristic fault trees were sufficient to cover all cases.

BASIS OF SAFETY SPREADSHEET

A basis of safety spreadsheet could now be created for each reactor. An illustration is shown in figure 2. The contributions to the calculation are chosen to represent the reactor being considered. An explanation of the columns is given below.

Operations per year is the representative number of operations of each group per year. This can be used to explore the impact of different production profiles on the risk calculation.

Rate or Quantity Excess is the probability of a failure leading to the addition of the reactive ingredient at too high a rate. The range of values here illustrates the differing circumstances which apply. Some chemistry groups are not susceptible to runaway through this fault. Other groups, where the reactive ingredient is added through a dedicated line, benefit from multiple independent layers of protection so that the likelihood of failure is extremely small. For one chemistry group however the reactive ingredient is added by manually controlled weighed addition, and this was seen as having a much larger error potential.

Proceeds without cooling is the probability that a failure of cooling occurs and is undetected by the operators and the trip system. It can be seen that in one group of reactions cooling is not in fact needed.

Materials Fault is the probability that the exotherm start is delayed due to either contaminated materials or a fault in the quantity of catalyst added, where this applies.

Generic Chemistry Groups	Operations per year	Rate or quantity excess	Proceeds without cooling	Materials fault	No Agitation	Low emperature addition proceeds	Delayed exotherm trip fails	Proceeds with exotherm delay	Heat generated exceeds natural cooling	Wrong material added	Secondary decomposition	Frequency of unplanned exotherm	Relief fails to open	Vessel strength not sufficient	Injuries result	Total frequency per year
1	0						1.0E-01		0.0	1.7E-04	0.0E+00	0.0E+00	0.01		0.1	0.0E+00
2	0						1.0E-01		1.0	0.0E+00	0.0E+00	0.0E+00	0.01		0.1	0.0E+00
3	0						1.0E-01		1.0	0.0E+00	0.0E+00	0.0E+00	0.01	0.001	0.1	0.0E+00
3	87	9.5E-06	8.3E-05	3.1E-05	1.1E-04	1.7E-04	1.0E-01	3.1E-05	1.0	0.0E+00	0.0E+00	9.9E-03	0.01	0.010	0.1	9.9E-08
4	12	0.0E+00	8.3E-05	5.0E-03	0.0E+00	1.7E-04	1.0E-01	5.2E-04	1.0	0.0E+00	0.0E+00	7.2E-03	0.01	0.010	0.1	7.2E-08
5	0						1.0E-01	0.0E+00	1.0	1.7E-04	0.0E+00	0.0E+00	0.01		0.1	0.0E+00
5	5	0.0E+00	8.3E-05	3.0E-05	1.4E-04	1.7E-04	1.0E-01	3.4E-05	1.0	0.0E+00	0.0E+00	5.9E-04	0.01	0.100	0.1	5.9E-08
7	0						1.0E-01	0.0E+00	0.0	0.0E+00	0.0E+00	0.0E+00	0.01		0.1	0.0E+00
8	0						1.0E-01	0.0E+00	0.0	0.0E+00	0.0E+00	0.0E+00	0.01		0.1	0.0E+00
9	25	6.2E-03	8.3E-05	3.0E-05	0.0E+00	0.0E+00	1.0E-01	3.0E-06	0.0	0.0E+00	0.0E+00	9.4E-04	0.01	0.100	0.1	9.4E-09
11	6	9.1E-06	8.3E-05	3.8E-05	1.1E-04	0.0E+00	1.0E-01	1.5E-05	1.0	1.7E-06	4.8E-05	9.4E-04	0.01	0.010	0.1	9.4E-09
12	5	9.1E-06	8.3E-05	3.0E-05	0.0E+00	0.0E+00	1.0E-01	3.0E-06	1.0	1.7E-04	4.8E-05	1.6E-03	0.10	0.001	0.1	1.6E-08
13	1						1.0E-01	0.0E+00	0.0	0.0E+00	0.0E+00	0.0E+00	0.01	0.001	0.1	0.0E+00
															Total	2.5E-07

Figure 2. Basis of safety example

No agitation is the probability that the agitation fails or is not effective and the reaction is allowed to continue. In some cases agitation is not significant.

Low temperature addition proceeds is the probability that the reactive ingredient is added at a temperature below the specified range, allowing a possible accumulation of unreacted material, and is only a concern in certain cases.

Delayed exotherm trip function was still in design at the time of assessment and hence a probability of failure on demand at ungraded level was assumed.

Heat generated exceeds natural cooling allows credit to be taken for the fact that in some cases the natural cooling of the reactor will be sufficient to prevent runaway even if an exotherm occurs.

Wrong material added is only possible in certain chemistry groups, which were identified through a separate human error study. An assessment of the error probability was made using the HEART technique (Williams 1985) which gave the probability of adding the wrong bulk chemical to a reactor and not detecting the error as 0.00017 per batch. In the group referred to as the resins the probability of error is reduced further because these recipes have a laboratory test of the mixed materials, for quality control purposes, before addition of the reactive material, and this would also have to be in error for the operation to continue.

Secondary decomposition is only possible in a small number of chemistry groups, and in certain reactors. Decomposition arising from accumulation was felt to be included in the previous failure contributions, so a separate fault tree analysis of the non-accumulation cases was carried out.

Frequency of Unplanned exotherm sums the previous contributions for calculation purposes.

Relief fails to open recognises that there is a probability that this protection (using bursting discs) will not work as intended, estimated as a probability of failure on demand of 0.01 (Smith, 1997).

Vessel strength not sufficient quantifies the probability that the reactor will actually fail under pressure if an uncontrolled exotherm is initiated, and the safety systems and relief fail to curtail it. In this case the temperature within the reactor will rise to a level determined by the heat of reaction, volume of material present, and the natural heat loss from the system, with corresponding increase in pressure. A reactor vessel built to an appropriate standard and in good condition is likely to withstand several times its rated design pressure before it fails catastrophically. Pressure vessels are normally tested in use to 1.5 times the design pressure, and it would be expected that deformation would occur rather than rupture, up to certain elevated pressures. The multiple of design pressure at which catastrophic failure is actually likely to happen is not well established. The Dutch authorities calculation method (CPR 14E 1997) expects a safety factor of around 2.5, and suggests that rupture consequences be calculated at this multiple of the design rating. However some practical pressure system design guides require a higher safety factor to be applied, for example a factor of 4 (Borzileri 1999). In this assessment it was decided that good construction standards and inspection regime made it reasonable to assume a safety factor of at least 3.

Where the maximum foreseeable operational pressure (Pmax) was less than the design pressure rating (Pdesign) of the vessel, the probability of rupture in use was assumed to be very small. Where Pmax exceeds Pdesign, the probability of rupture was assumed to increase, reaching 1 when Pmax/Pdesign is three. The allocation is shown in the following table.

Pmax/Pdesign (ignoring pressure relief)	Probability of failure
≤ 1	0.001
Between 1 and 2	0.01
Between 2 and 3	0.1
3 or above	1

Injuries Result was set at a probability of 0.1 recognising that even in the event of a catastrophic rupture it is quite likely that injuries could be avoided. Contributing factors here include that the onset of such a failure would be expected to be relatively slow, allowing time for escape and evacuation, and there would be some shielding provided by the building and other plant.

COMPARISON AGAINST MAXIMUM TOLERABLE FREQUENCY

The predicted frequency of failure for each reactor can now be compared against the maximum tolerable frequency of occurrence. For each of the reactor scenarios examined it could be seen that the maximum frequency was not breached, with each of the SIS assigned to be a SIL 1. Thus this safety integrity level would be an adequate specification for the systems to meet. The need for a higher SIL was considered, so as to show that risk was ALARP. However in view of the low level of residual risk achieved with SIL 1 systems, the small reduction in risk which could be achieved through a higher SIL would not be a reasonable requirement.

If the future operating regime is required to change, for example in the numbers of operations, or modification of equipment, the impact on the likelihood of failure can be estimated readily, and if there is a need for improvements this can be seen quickly. The operator should be able to see the choices are available for these improvements, whether by enhancing the SIS, adding additional non-SIS protective systems, or limiting the duties of the reactors.

CONCLUSIONS

A strategy for evaluating the role of SIS in reactor safety has been developed which allows the appropriate SIL to be determined in accordance with current standards. The highly interactive nature of this project provided an invaluable learning opportunity for all

concerned and raised the awareness of a large cross section of site personnel. The method takes account of the flexibility of production required in a realistic business scenario, and allows the contribution to safety from a range of protective layers to be included. The resulting model should be of benefit in the future in assessing the safety implications of any planned operational change, and determining the best combination of protective measures.

REFERENCES

Borzileri, C., Pressure Safety Standard UCRL-AR-128970 Rev 1, May 1999.

Committee for the Prevention of Disasters, Netherlands, CPR 12 E 'Red Book', Methods for Determining and Processing Probabilities, 1997.

Committee for the Prevention of Disasters, Netherlands, CPR 14E, 'Yellow Book', Methods for the Calculation of Physical Effects, Part 2, 1997.

Environmental Protection Agency, CEPP, How to prevent Runaway Reactions, 1999.

Gertman, G I, Blackman, H S, Human Reliability & Safety Analysis Data Handbook, Wiley Interscience, 1994.

HSE, Reducing Risks, Protecting People, HSE Books, 2001 (R2P2).

Kirwan, B, A Guide to Human Reliability Assessment, Taylor & Francis, 1994.

Lees, F P, Loss Prevention in the Process Industries, Butterworth Heinemann, 1996.

Smith D J, Reliability, Maintainability and Risk, Butterworth-Heinemann, 1997.

Williams J C: HEART – A proposed method for achieving high reliability in process operation by means of human factors engineering technology, Proceeding of symposium 'Achievement of reliability in operating plant', The Safety and Reliability Society (Manchester, UK), 1985.