ELSEVIER Journal of Hazardous Materials 44 (1995) 231-251

JOURNAL OF HAZARDOUS MATERIALS

A risk perspective for emergency pressure relief system design

Michael A. Grolmes[a,*], Jeff Gabor[b], Marc Kenton[b], Richard Wachowiak[b]

[a] *Centaurus Technology, Inc., 2438 Oak Hill Drive, Lisle, IL 60532, USA*
[b] *Dames & Moore, Westmont, IL 60559, USA*

Received 10 September 1994; accepted 7 April 1995

Abstract

Improperly sized pressure relief devices have historically been a factor in a significant fraction of serious accidents involving process reactions. The Design Institute for Emergency Relief Systems (DIERS) efforts have provided a methodology which minimizes the potential for inadequate pressure relief capacity. Unfortunately, this methodology can lead to requirements for vent sizes which appear to be unreasonably large. The difficulty may be traceable directly to the selection of the design basis. It is often difficult to justify credible bounds for scenarios to be considered in the design by qualitative methods alone. In these instances, quantitative risk assessment methods can be effective in providing criteria for the elimination of excessive conservatism and thus can lead to a satisfactory design. This paper describes a hypothetical but realistic example which effectively combines the DIERS methodology with quantitative risk assessment methods. The result is a more satisfactory basis for evaluating the adequacy of the vent design.

Keywords: Risk assessment; Pressure relief; Process safety; Vent design

1. Introduction

Significant advances in the technology of emergency pressure relief system design have been made in the last 10 years. These advances have been especially important in applications considering liquid phase chemical reaction processes. The methods embodied in the Design Institute for Emergency Relief Systems (DIERS) technology base have been widely accepted by industry as representing the state-of-the art practice [1, 2].

Three key concepts of the DIERS methodology are:

- Establish a design basis upset scenario(s) to be considered in the design, i.e. the design basis.

[*] Corresponding author.

0304-3894/95/$09.50 © 1995 Elsevier Science B.V. All rights reserved

SSDI 0304-3894(95)00056-9

– Determine the reaction rates and other necessary system characteristics at the relief conditions corresponding to these scenarios.
– Use of vent sizing methods which account for both gas and liquid flow in the relief system.

The definition of the design basis upset condition is typically derived from a hazard evaluation. The determination of the corresponding reaction rates and system characteristics requires, in most instances, test data obtained from appropriate calorimetry apparatus. The use of the appropriate reaction rate information in vent sizing relations which consider two-phase discharge has been facilitated by a variety of analytic methods. Much of this technology is summarized in the DIERS Project Manual [1].

Over the past ten years, as this technology has been disseminated to the industry at large, specialists and working groups, have continued to refine test techniques and vent sizing formulations [3]. The one facet of the DIERS methodology that has received relatively little attention is the selection of the system design basis. This would seem to be a conceptually straightforward process.

Various reviews of industrial incidents involving exothermic batch and semi-batch process reactions have identified suspected classes of potential causes of runaway reactions [4–6]. The following scenarios, broadly grouped, contribute about equally to initiation of runaway reactions which challenge the emergency relief system: (i) improper charging of reagents; (ii) loss of agitation; (iii) loss of temperature control.

Accidents involving vessel damage and more serious consequences result when the relief system fails or is inadequately designed for the initiating event, or when the initiating event was not anticipated. Ref. [7] makes clear that the record of incidents shows that relief system failure attributed to inadequate capacity is of substantially greater significance than failure because of mechanical or installation faults of an otherwise adequately designed relief system.

The record of incidents and causes reflected in Refs. [4, 7] can be viewed as the rationale for the keen interest in, and rapid pace of safety technology developments. Recent US regulations [8, 9] have mandated formality and documentation of hazards evaluations. These regulations address the issue of identifying possible causes of runaway reactions. The prevalent techniques employed by industry are HAZOP, checklist and what-if type formats [10, 11]. These methods can be highly effective as discovery techniques for accident initiators.

The DIERS methodologies have been accepted by industry and regulatory agencies as representing state of the art design methods for pressure relief devices for reactive systems. It is fair to say that a conservative DIERS emergency pressure relief evaluation will be adequate for a specified event. The purpose and intent of this study is to demonstrate how risk methodologies can be utilized to define credible design basis.

It commonly arises in the evaluation of a reaction process by the usual hazard evaluation procedures that multiple causes of runaway reaction events are identified. Each of these scenarios can have its own corresponding adequate vent size. In those instances where the bounding (largest) vent size represents a practical installation, the process can be terminated with a design based on the bounding case.

However, it frequently occurs that with due consideration of plant experience, and without postulating physically or chemically impossible events, a scenario can result from hazard evaluation procedures which leads to vent size requirements which are too large for practical implementation. This is often a consequence of compounding events which cause reaction rates to increase in a highly non-linear manner. Now the challenge is to deal with very rare postulated initiating events in a reasonable fashion. At this point the hazard evaluation and vent sizing process can become polarized by conflicting subjective judgements. Without further technical support, responsible parties may be unable to reach a consensus.

It is in this situation that risk assessment methods can play a key role in the design or evaluation of emergency relief devices. There are three main benefits to the introduction of risk assessment into the vent design process. The first is to simplify and formalize the application of judgement at the component and procedural level, rather than to global scenarios. The second is to determine the frequencies of various scenarios and to identify the key contributing elements to each event. Finally, the resulting design (or the acceptability of an existing installation) can be thoroughly documented.

In this paper, a hypothetical process reactor system is considered. With as much realism as possible for a hypothetical case study, the range of resulting vent sizes is related to a corresponding range of scenarios. The selection of a vent size will ultimately hinge on rejecting certain scenarios which are possible but highly unlikely. The purpose of this paper is to illustrate the use of risk assessment methods for accomplishing this.

2. System description

The example system is a catalyzed polymerization reaction shown in the skeleton diagram of Fig. 1. The key features of the system are a 12 000 gal process reaction vessel with monomer and solvent supply tanks. The reaction vessel is agitated and has provision for jacket heating and cooling. The process is controlled by operator actions which are based on recorded indication of process parameters. Further assumptions related to operations will be identified as required.

2.1. Possible upset conditions

Even for this relatively common system, it will become clear that the range of possible accident scenarios is very broad. It is quite reasonable that as a result of a simplified hazard evaluation, the list of events which could initiate the accident scenarios listed in Table 1 would be considered possible.

The list of possible upset conditions in Table 1 is given without concern for their likelihood.

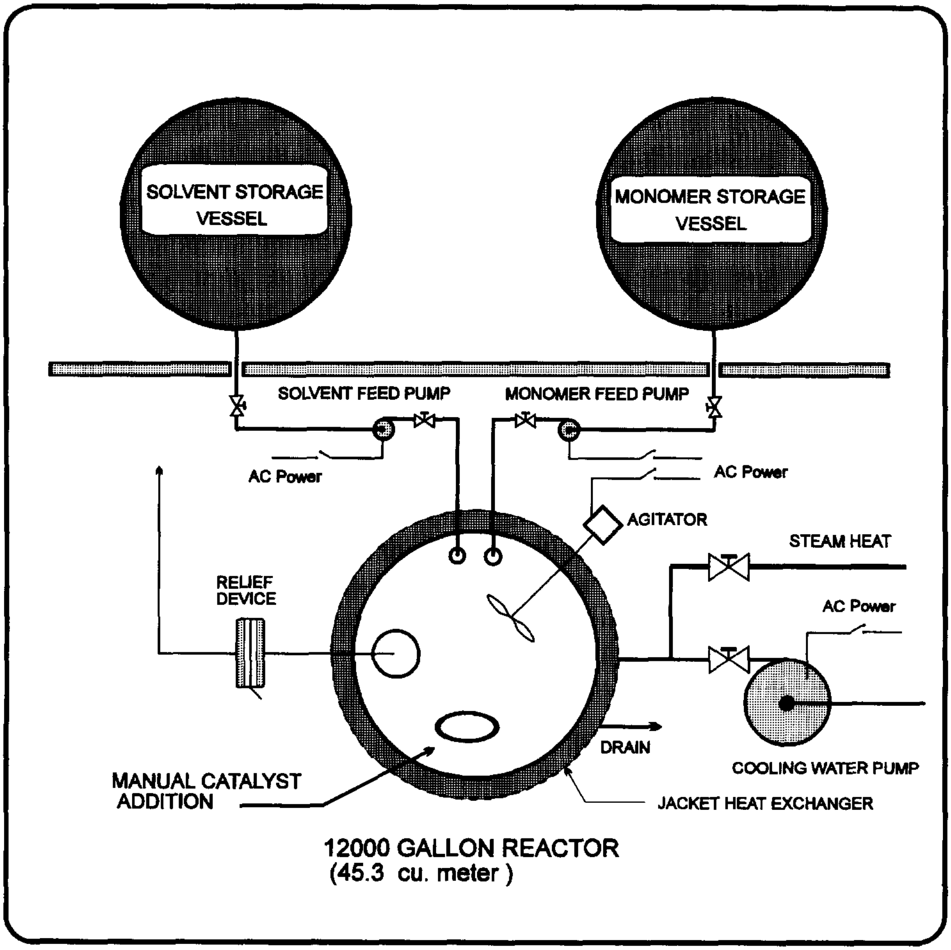

Fig. 1. Ilustration of polymerization reactor system.

Table 1
List of possible process upset conditions

1. *Normal batch charge*
   (a) loss of system power leading to loss of agitation and loss of cooling
   (b) equipment failure leading to loss of cooling, agitator failure or coolant circulation pump failure

2. *Batch mischarge*
   (a) normal batch, wrong catalyst or excess catalyst
   (b) no solvent, normal monomer charge
   (c) no solvent, monomer overfill

3. *External fire*
   (a) on normal batch which leads to loss of agitation and cooling and also adds heat to system

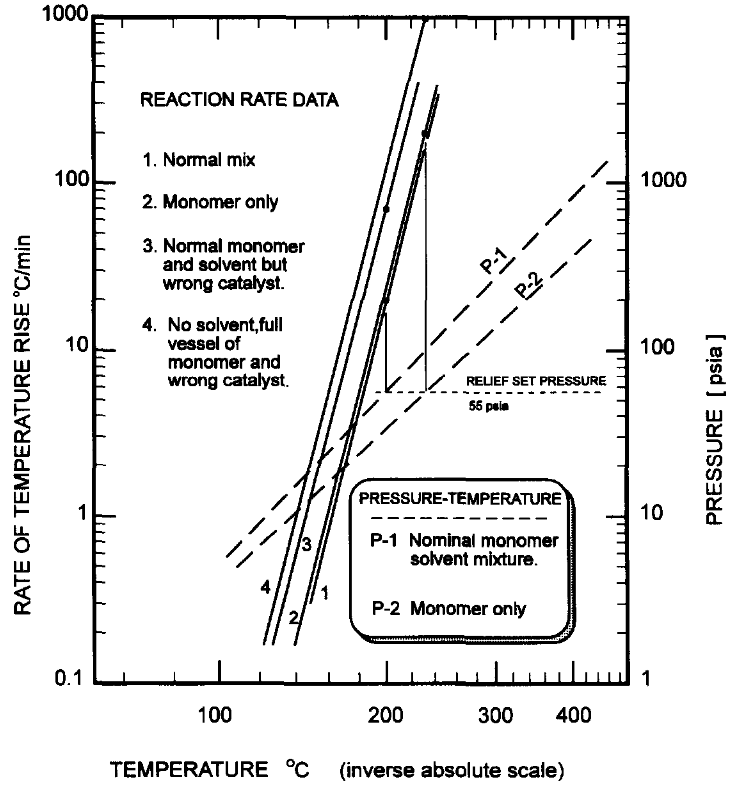

Fig. 2. Reaction self-heat rate and pressure temperature characteristics for process upset scenarios.

2.2. Vent size implication

Determining the minimum vent size for the various upset scenarios requires a knowledge of the underlying reaction kinetics and pressure–temperature relations for the system. For the purpose of this study, it will be assumed that the requisite data have been obtained by an appropriate combination of test and analytic methods, and that the data can be represented in a form similar to that shown in Fig. 2. The figure presents runaway reaction data expressed as self-heat rate (dT/dt), and system vapor pressure as a function of temperature (inverse absolute scale). The curves in Fig. 2 corresponds to the following four cases:

1. Self-heat rate for a runaway reaction with a normal batch loading.
2. Self-heat rate for normal monomer, normal catalyst but no solvent.
3. Self-heat rate for normal monomer/solvent charge but excess or wrong catalyst.
4. Self-heat rate for pure monomer, no solvent and excess or wrong catalyst.

Two vapor pressure curves are shown, one corresponding to a normal mix of monomer and solvent (Case P-1), and the other for a case with pure monomer (Case P-2).

Table 2
Association of initiating events with reaction rate and required vent size

Case	Case description	$\mathrm{d}T/\mathrm{d}t$ (°C/min)	T (°C)	Required vent diameter (in)
0	Vent size based on API-520 fire exposure all-vapor vent flow with no reaction heat	–	200	4
1	Nominal charge loss of cooling – all-vapor vent flow	20	200	10
2	Nominal charge loss of cooling DIERS two-phase vent flow	20	200	16
3	Case 2 with external fire added	25	200	18
4	Nominal charge, wrong catalyst loss of cooling DIERS two-phase vent flow	70	200	28
5	No solvent, normal monomer and catalyst with loss of cooling. DIERS two-phase vent flow	200	230	32
6	No solvent, monomer overcharge with loss of coolant DIERS two-phase vent flow	200	230	48
7	No solvent, monomer overcharge wrong catalyst – Independent of loss of coolant	1000	230	108

The relief set pressure is assumed to be 40 psig (55 psia). The reaction rate at the temperature corresponding to the set pressure is a key parameter for the vent size evaluation. This is shown for two examples in Fig. 2. The illustration shows how the factors which affect reaction kinetics and system volatility can have profound effects on the reaction rate at the relief set pressure, and thus on the required vent size.

Aside from the event which initiates an accident, subsequent events or system failures can compound the accident further and must also be therefore considered. For example, a scenario could be made worse by a coincident loss of cooling to the reactor. Table 2 summarizes potential scenarios, their reaction rate parameters and resulting vent sizes. The table contains two cases where inappropriate vent sizing methods have been used. Case 0 in Table 2 indicates that a 4 in vent size would be calculated to be adequate based on all vapor vent flow associated with an API-520 [12], external fire heat load, and no accounting for the system reaction. Case 1 uses the appropriate reaction heat rate data for a loss of cooling on a normal batch, but the vent size is again based on all vapor vent flow. The assumption of all vapor flow would not be valid for this polymerization reaction system. Cases 0 and 1 in Table 2 therefore represent inadequate design from a purely methodological stand point. The remaining vent size estimates in Table 2 are based on Fig. 1 data and the DIERS short form vent size relation [1, 13] for vent area

$$A = \frac{2.5\,M\,(\mathrm{d}T/\mathrm{d}t)_{\mathrm{set}}}{C_{\mathrm{d}}\,P_{\mathrm{set}}\,(T_{\mathrm{set}}/C)^{1/2}}, \tag{1}$$

where M is the batch inventory, $\mathrm{d}T/\mathrm{d}t$ is the reaction self-heat rate at the relief set pressure, C_d is the vent discharge coefficient, T_set is the batch temperature corresponding to the relief set pressure, and C is the liquid heat capacity.

In these calculations, the normal batch inventory is assumed to be 31 000 kg of 50% monomer. The vent discharge coefficient is taken as 0.8 and the liquid specific heat capacity is taken as 2200 J/kg K. Homogeneous two-phase flow is assumed through the vent. Eq. (1) includes allowances for two-phase flow and 20% maximum pressure rise after relief activation.

The assumptions listed above and Eq. (1) are used throughout the evaluation for consistency. More sophisticated DIERS methods can in some scenarios lead to somewhat smaller vent sizes, but generally not by large factors when the same reaction rate parameters are assumed. Therefore, Cases 2–7 represent valid vent size evaluations for the upset scenarios being considered.

While the results do not represent a specific plant example, in many ways they are quite representative of experience. The equivalence of the 16 in vent for Case 2 using the DIERS methods with a 10 in vent for the corresponding reaction heat rate assuming all vapor vent flow is fairly typical: adequate pressure relief of homogeneous two-phase vent flow typically requires an area increase of between 2 and 3 times larger than that required for relief of the vapor production only. Table 2 results are also typical in that impracticably large vents (48 and 108 in) can result if initiating events are compounded by multiple concurrent failures.

Installation of a 48 in or larger vent would not normally be recommended. The real difficulty arises in the justification for the selection of a vent in the 16–32 in size range. The differences in the corresponding scenarios are not easy to differentiate on the basis of subjective judgements. It is particularly in this regard that the results of Table 2 are quite typical. How can the vent design process be brought to a satisfactory conclusion?

2.3. Assessment of consequence severity

The consequence of an under-sized emergency pressure relief vent can vary widely from simply exceeding the recommended maximum of 10% over design pressure, up to vessel failure. We assume a vessel design pressure of 46 psig which is reasonable for a 40 psig relief set pressure. The vent design basis assumption of 20% over pressure on an absolute basis corresponds to a maximum venting pressure of 51 psig. This is equivalent to 10% over the 46 psig design pressure on a gage basis. Because of conservatism in the ASME code, one would not normally expect permanent deformation at pressures below $\approx$95 psig. We would expect catastrophic vessel failure at pressures in excess of about 185 psig.

With the reaction kinetics and vapor pressure temperature data shown in Fig. 2, and the assumption of homogeneous flow, one can with the aid of simple computer simulations estimate the maximum venting pressure for vent sizes. Consequence severity categories can then be defined to group these results into a manageable

Table 3
Consequence severity categories

Severity category	Case simulation outcome
0	Normal relief vent activation P_{max} less than 1.1 MAWP
1	Vessel overpressure greater than 1.1 MAWP but P_{max} less than 2 MAWP, no component failure
2	Vessel overpressure in range of 2 MAWP. Vessel requires rework
3	Vessel overpressure greater than 2 MAWP. Permanent strain. Vessel beyond repair
4	Vessel overpressure greater than 4 MAWP. Damage to surrounding equipment
5	Vessel overpressure much greater than 4 MAWP at high rate of pressure rise. Vessel explosion with maximum potential damage

Table 4
Vent size and severity category

	Scenario	Severity category (0–5) vent size inch									
		4	10	12	16	18	24	28	32	48	108
1	Loss of cooling with normal charge (16)[a]	5	2	1	0						
2	Loss of cooling with										
	(a) no solvent, normal monomer (32)	5	5	4	3	3	2	1	0	0	
	(b) no solvent, overcharge monomer (48)	5	5	5	4	4	4	3	2	0	
	(c) normal monomer and solvent with wrong catalyst (28)	5	4	4	3	2	1	0			
	(d) Case 2a with wrong catalyst (96)	5	5	5	5	5	5	5	5	4	0
3	External fire and loss of coolant with										
	(a) normal charge (18)	5	3	3	1	0					
	(b) normal charge with wrong catalyst (32)	5	5	4	3	3	1	1	0		
	(c) no solvent normal monomer (48)	5	5	5	4	4	4	3	2	0	
	(d) Case 3c with wrong catalyst (96)	5	5	5	5	5	5	5	5	4	0
4	Mischarge without loss of coolant										
	(a) no solvent correct monomer (18)	5	3	2	1	0					
	(b) no solvent, overcharge monomer (28)	5	4	4	3	2	1	0			
	(c) case 4b with wrong catalyst (108)	5	5	5	5	5	5	5	5	4	0
	(d) case 4a with wrong catalyst (48)	5	5	5	5	5	4	3	2	0	

[a] Numbers in parentheses represent nominal adequate vent size in inches for the upset scenario, with allowance for 20% overpressure and corresponds to an indicated severity level of 0.

number of cases as shown in Table 3. While still somewhat broad, the severity categories in Table 3 tie reaction characteristics and vessel pressure rating together in a meaningful, quantitative way. Finally, one can associate the consequence severity categories with the vent size and upset scenarios as illustrated in Table 4.

3. Risk analysis

As previously noted, in order to protect against all postulated incidents a relief vent on the reactor vessel would need to be larger than the vessel itself. This is clearly not acceptable. At this point, we are motivated to consider the likelihood of the various scenarios to determine a reasonable vent design.

Risk assessment methods can be used to estimate the frequencies of overpressure events of each severity category as a function of the reactor vent size. These frequencies can then be used to eliminate from considering small vents that allow high consequence events to occur at high frequency and large vents which would only be needed under exceedingly rare circumstances. There is also the possibility of identifying other modifications to the reactor system (such as changes to operating procedures or electrical configuration) that would provide enhanced protection from the high consequence incidents at a lower cost.

3.1. Define the undesired plant condition

The first step in the probabilistic analysis is to define carefully the results of the accidents being studied. These definitions should allow a quantitative comparison of the consequences of different accidents (such as cost of repair and lost revenue) so that overall risk can be determined. In our hypothetical process reactor, the accidents are defined as the degree of vessel overpressurization of a single, independent batch reactor, Table 3.

3.2. Determine the operating modes at risk

The next step is to identify all operating conditions from which upsets could lead to the defined accidents. The reactor can be either idle or processing a batch. It can be assumed that the idle state poses no threat of overpressurization. The batch mode can be broken into discrete time phases, each with a distinct potential for overpressurization incidents. In the loading phase, the reagents are assumed to be below the reaction temperature, therefore events during this phase are not considered. In the heat-up phase, there is some potential for overpressurization, however, it would occur only after the temperature of the batch is brought up to the reaction temperature. Events occurring in the heat-up phase can therefore be conservatively grouped with beginning of cycle events. At the beginning of the batch cycle, there is a maximum potential for overpressurization because of the large inventory of unreacted material in the vessel.

After mid-cycle, it is assumed that the amount of unreacted material in the vessel has been sufficiently reduced that events such as loss of cooling or external fires will not challenge a reasonable vent design. This assumption should be verified after a vent selection is made. The final batch phase is the unloading of the finished product. It is assumed that there is no threat of overpressurization in this phase. Based on this reasoning, this study will consider only the phase from the beginning of cycle to the mid-cycle of a batch.

3.3. Functional failures leading to the accidents

The hazard evaluation described previously produced a list of possible process initiating events, Table 1. These conditions can be defined in terms of functional failures. The functions are solvent charge, monomer charge, catalyst charge, external fire, and cooling. Note that agitation and loss of power are not listed as functions. This is because they are systems that support the cooling function.

Once an upset condition is present, there may be functions available to mitigate the event. The pressure relief device falls in this category. If the pressure relief device is adequately sized and activates as designed, there will be no challenge to vessel integrity. Some upset conditions, however, require relief capacity larger than the vent installed. In these cases, successful actuation of the vent may merely reduce the severity of the overpressurization. It is also possible that the venting function could itself fail to perform as designed. Historic data on vent failures indicate that vents can actuate at a higher pressure then their set point or only partially open, thereby failing to provide full relief capacity [5]. For our purposes, a failed relief device is considered to behave like a functional vent with a smaller capacity.

3.4. Event sequence analysis

The various events which determine the severity of the accident are arranged in a logical manner on an event tree, Fig. 3. The tree starts at an initiating event (i.e. the beginning of a batch cycle) on the left-hand side, and proceeds through a series of events to the end states on the right. End states can either have no consequence or can be categorized as one of the accidents. Each path through the event tree from the initiating event to an end state is known as a sequence.

The first heading on the tree signifies that 225 batch cycles per year occur. Each subsequent heading is associated with the success or failure of an important function. By convention, an upward branch under an event heading implies success of the function, while a downward branch indicates failure. A separate event tree would generally be constructed for each initial condition of the analysis. As discussed previously, the only tree necessary in this example is for the beginning of the batch cycle.

The success criteria for each branch point on the event tree are defined in terms of the operational states of the plant systems that provide the associated function. As an example, consider COOLING in Fig. 3. This heading determines the success or failure of the cooling function during a batch process. For simplicity, it is assumed that loss of jacket cooling, loss of agitation, and failure to switch from steam heat to cooling all have the same effect on the reaction rate in the vessel. Success of COOLING indicates that the agitator is running throughout the heat up and batch cycle and that the jacket cooling system is properly removing heat during the batch cycle. In addition, success implies that the switch over from steam heating to jacket cooling has taken place. Therefore, failure of COOLING occurs if either agitation or jacket cooling fails, or if the steam heat remains on line after the mixture has reached the reaction temperature. Even though the inventory of reactants in the batch are decreasing with time during

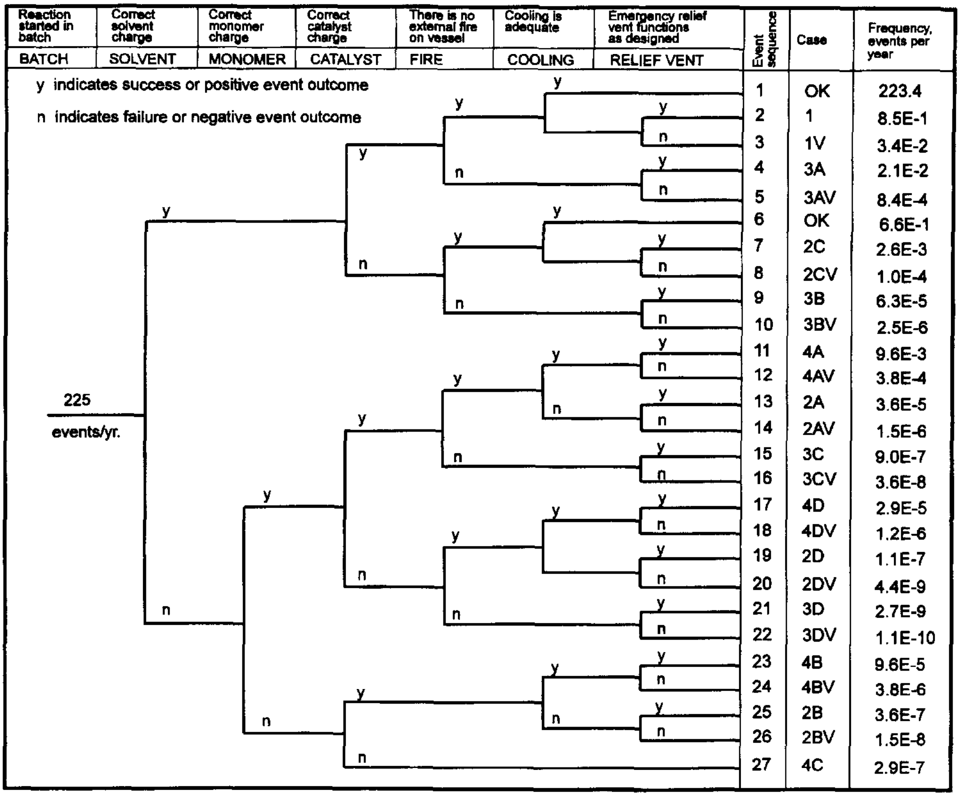
Reaction started in batch
BATCH
Correct solvent charge
SOLVENT
Correct monomer charge
MONOMER
Correct catalyst charge
CATALYST
There is no external fire on vessel
FIRE
Cooling is adequate
COOLING
Emergency relief vent functions as designed
RELIEF VENT
y indicates success or positive event outcome
n indicates failure or negative event outcome
225
events/yr.
y
n

Event sequence	Case	Frequency, events per year
1	OK	223.4
2	1	8.5E-1
3	1V	3.4E-2
4	3A	2.1E-2
5	3AV	8.4E-4
6	OK	6.6E-1
7	2C	2.6E-3
8	2CV	1.0E-4
9	3B	6.3E-5
10	3BV	2.5E-6
11	4A	9.6E-3
12	4AV	3.8E-4
13	2A	3.6E-5
14	2AV	1.5E-6
15	3C	9.0E-7
16	3CV	3.6E-8
17	4D	2.9E-5
18	4DV	1.2E-6
19	2D	1.1E-7
20	2DV	4.4E-9
21	3D	2.7E-9
22	3DV	1.1E-10
23	4B	9.6E-5
24	4BV	3.8E-6
25	2B	3.6E-7
26	2BV	1.5E-8
27	4C	2.9E-7

Fig. 3. Batch process event tree.

the cycle, loss of cooling at any time prior to mid-cycle is conservatively assumed to cause reaction rates that correspond to beginning of cycle conditions.

The failure states of functions in the event tree can affect the success of branches later in the tree. For example, if the solvent load is successful, it is not possible to fill the vessel with monomer. Therefore MONOMER is always successful following the success of SOLVENT. Conversely, if SOLVENT, MONOMER, and CATALYST are all failed, any relief device is assumed to be insufficient to prevent overpressurization, so the function VENT is assumed to fail. These assumptions are directly reflected in the way the Event Tree is drawn. So, for example, there is no branch in Fig. 3 under the MONOMER loading if SOLVENT is a success.

3.5. Assignment of end states

Next, each of the event sequences needs to be associated with one of the predefined overpressurization severity categories. In order to determine a severity of the end states, the vent size of the reactor system must be known. Thus, the severity categories associated with end states of the event tree are different for each proposed vent size. In our case, however, the functional failures are independent of the relief vent size. It is therefore possible and convenient to calculate the frequency of the end states once for the reactor system regardless of the vent size. For each vent size, the end states are then associated with a severity category.

The 'Case' column in Fig. 3 indicates which upset scenario from Table 4 is assigned to each of the end states of the event tree. As an example, sequence number 1 has the correct solvent load (therefore the correct monomer as well), correct catalyst, no fire near the vessel, and cooling available. This would be considered a normal, uneventful batch run, and therefore does not result in overpressurization. It is assigned the 'OK' end state designation. Sequence number 6 is similar except that the catalyst is wrong (i.e. too much catalyst or a more reactive catalyst is loaded). The deterministic analysis discussed earlier shows that the wrong catalyst alone is not sufficient to cause a relief vent challenge, so this end state is also assigned the 'OK' designator. Sequence number 27 has no solvent, an over-fill of monomer, and the wrong catalyst. This combination of events corresponds to the scenario labelled '4C' in Table 4, so the end state of sequence number 27 is therefore given this label.

The case designations for sequences with a failed vent have a 'V' following the upset scenario identifier. This is because the deterministic analyses assumed that the installed vent worked properly. We assume here that end states with a failed vent can be assigned the severity category for a reactor with a relief capacity of approximately one-half of the installed vent size.

3.6. Systems analysis

To calculate the probability that an event tree function, such a COOLING, will fail, we must consider all combinations of component failures, human actions, and external events that are sufficient to cause failure of the function. This is done by constructing fault trees. In a typical risk analysis, a fault tree is drawn for each function displayed on the event tree. Many systems, in turn, require support systems, such as electric power, in order to perform their designated function. These may also be modeled by fault trees.

A fault tree is made up of a top event, corresponding to an event tree heading or a support system function, and a logic structure that models all of the combinations of events that must take place to cause the top event. All of the elements in the fault tree typically represent failures. Fig. 4 presents the major components used in the fault tree modeling illustrated here; some analysts use a few additional symbols. An AND gate represents a failure if all of its inputs are failures, while an OR gate represents a failure if any of its inputs are failures. A basic event represents a failure that is not developed with any further logic. Basic events can be any of the following: human actions (such as operator improperly adjusts cooling flow bypass), component failures (such as cooling pump fails to run), or external events (such as loss of offsite power). The probabilities of basic events are input directly by the analyst. A transfer gate represents logic that is modeled in another fault tree.

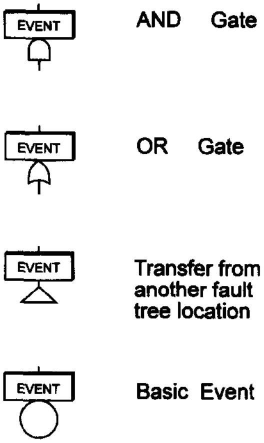

Fig. 4. Basic fault tree elements.

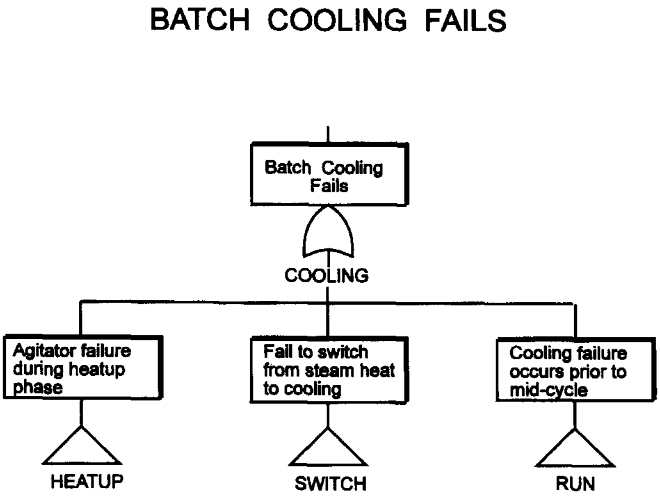
BATCH COOLING FAILS
Batch Cooling Fails
COOLING
Agitator failure during heatup phase
Fail to switch from steam heat to cooling
Cooling failure occurs prior to mid-cycle
HEATUP
SWITCH
RUN

Fig. 5. Cooling function fault tree.

Figs. 5–7 illustrate the fault tree modeling of the COOLING function for the batch-reactor. The success criteria defined for this function states that batch cooling requires the agitator to start and run during the heat-up phase, the switch from steam heating to jacket cooling must be made when the reactants are at the proper temperature, and the cooling system and agitator system must successfully run from the beginning of cycle to mid-cycle. The fault tree that represents the failure of any of these will cause failure of the function. This fault tree contains transfers to supporting

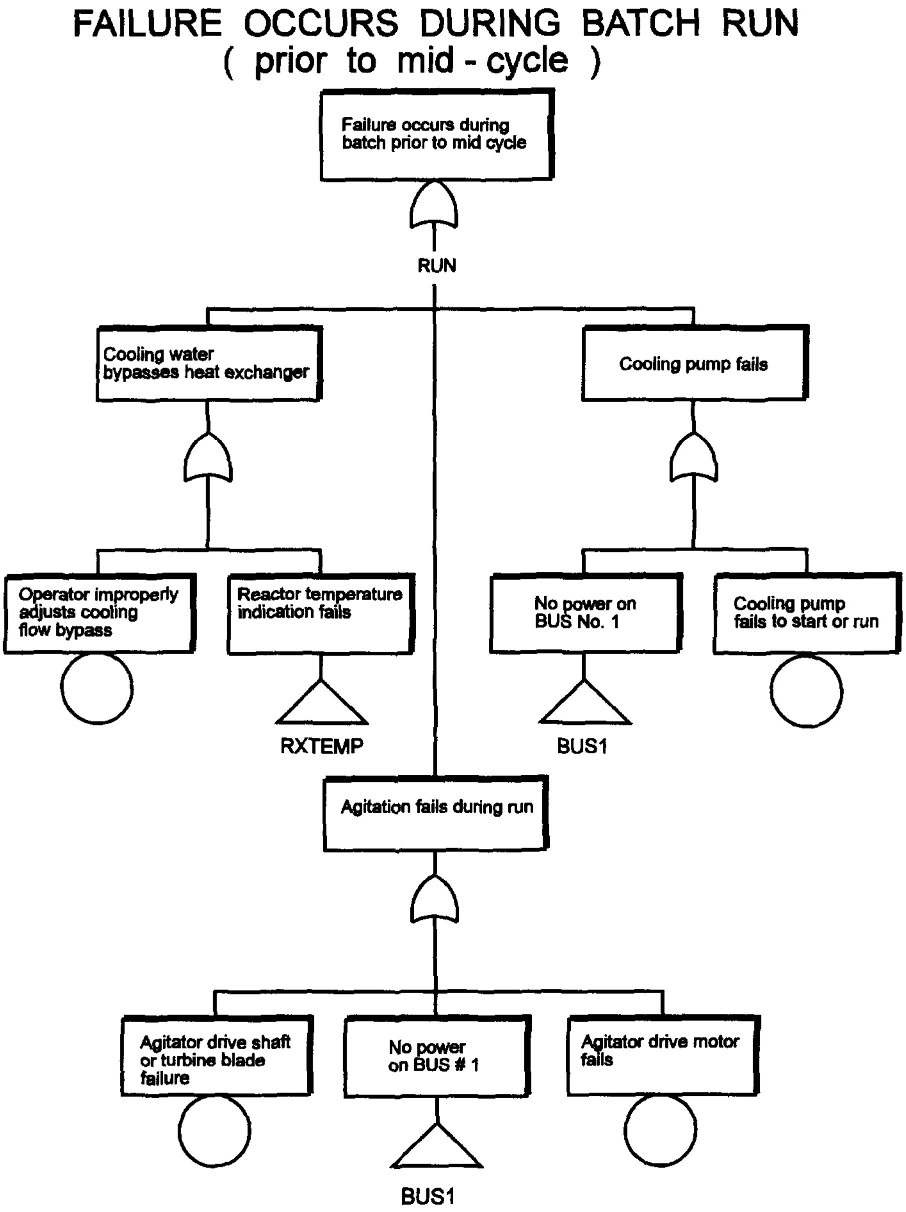
FAILURE OCCURS DURING BATCH RUN
(prior to mid - cycle)
Failure occurs during batch prior to mid cycle
RUN
Cooling water bypasses heat exchanger
Cooling pump fails
Operator improperly adjusts cooling flow bypass
Reactor temperature indication fails
RXTEMP
No power on BUS No. 1
BUS1
Cooling pump fails to start or run
Agitation fails during run
Agitator drive shaft or turbine blade failure
No power on BUS # 1
BUS1
Agitator drive motor fails

Fig. 6. Fault tree for cooling failure during batch run.

fault trees (e.g. failure occurs during the batch run, Fig. 6) that represent the failure of these individual systems.

Following the logic one level down, failure of cooling during the batch (Fig. 6) occurs if there is insufficient cooling flow to the jacket heat exchanger, the cooling

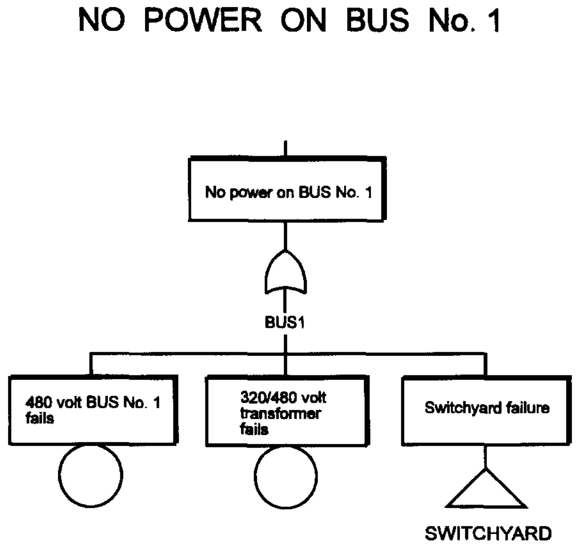

Fig. 7. AC power fault tree.

pump fails to run, or the agitator fails to run. This model assumes that both the coolant pump and the agitator have been successfully started and are running at the beginning of the batch cycle. (These failures are modeled in one of the other supporting fault tress, not shown.) As in the previous model, this tree contains both basic events and transfers to support system fault trees. Electric power and reactor temperature indication are support systems for this model.

In general, a single support system can be used by several systems. In this example, electric power from bus #1 is required for the operation of both the agitator and the cooling pump. The failure of this power system (Fig. 7) contains the basic events representing the failure of the components unique to that bus and a transfer to the main electrical system fault tree. This development of the fault tree logic proceeds until only basic events exist at the ends of the trees.

To summarize, an event tree is used to define systematically the functional failures which can lead to accidents. For each function represented on the event tree, a fault tree is constructed to define systematically the equipment or human failures necessary for the function to fail. In this way, the operation of very complicated systems can be represented in a clear, documentable fashion.

3.7. Component analysis

Once all of the basic events are identified, each must be assigned a probability of occurrence. These can come from several sources. For equipment failures, there are published sources of generic data that can be directly used in the analysis, e.g. Ref. [14]. Alternatively, failure for specific components at a site can be generated from information contained in operation logs and maintenance records. The probability of operator errors are generally determined from interviewing plant personnel, reviewing procedures and evaluating past incidents.

As with all statistical analyses, there are uncertainties associated with the estimates of the probabilities of the basic events. Note, however, that the same values are used to compare the risk of the same reactor process for the different vent sizes. While their absolute values will change, the relative frequencies of the upset scenarios will generally be similar even if different data are used. Secondly, to ensure that the absolute frequencies of the scenarios do not violate our acceptance criteria, uncertainty distributions are provided for each basic event probability used in the analysis. The computer codes used to solve the logic models are able to propagate the distributions through the solution and to provide an overall uncertainty distribution for the frequency of the scenarios.

3.8. Quantification

While simple trees can be quantified by hand or with a spreadsheet, it is common to use a dedicated computer program. The code used in this analysis is the Integrated Reliability and Risk Analysis System (IRRAS) [15]. This code was selected because of its wide availability, ease of use, and because it operates on a personal computer. IRRAS provides the capability for creating and graphically displaying the logic models, quantifying the frequencies of the undesired events, and reporting the combinations of individual failures that lead to the various end states. In addition, the code can perform uncertainty analysis on the results and indicate the relative importance of the various basic events used in the model.

Since the model was built to determine the frequencies of specific upset scenarios rather than the frequencies of the severity categories, the end states are binned into severity categories for each proposed vent size. For a given vent size, the frequencies of all of the end states in a category are combined to determine the overall frequency of the severity level. Table 4 is used to assign severity levels to the upset scenario-based end states. Fig. 8 presents the mean frequencies of accidents within each of the severity level categories for all of the postulated vent sizes.

4. Evaluation of results

The main purpose of a properly designed relief vent is to ensure that the most severe accident categories have little chance of occurring during the expected lifetime of the process reactor. To screen out unacceptable designs, we define three acceptance criteria: events of severity level 5 must not occur more than once in 10 000 reactor years, severity level 4 must not occur more than once in 1000 reactor years, and events of severity level 3 must not occur more than once in 100 reactor years. Judged by these criteria, Fig. 8 clearly indicates that the 4 and 10 in relief devices are unacceptable.

These criteria are obviously somewhat arbitrary, but serve here to illustrate the methodology. Acceptance criteria can be defined in a somewhat more formal manner.

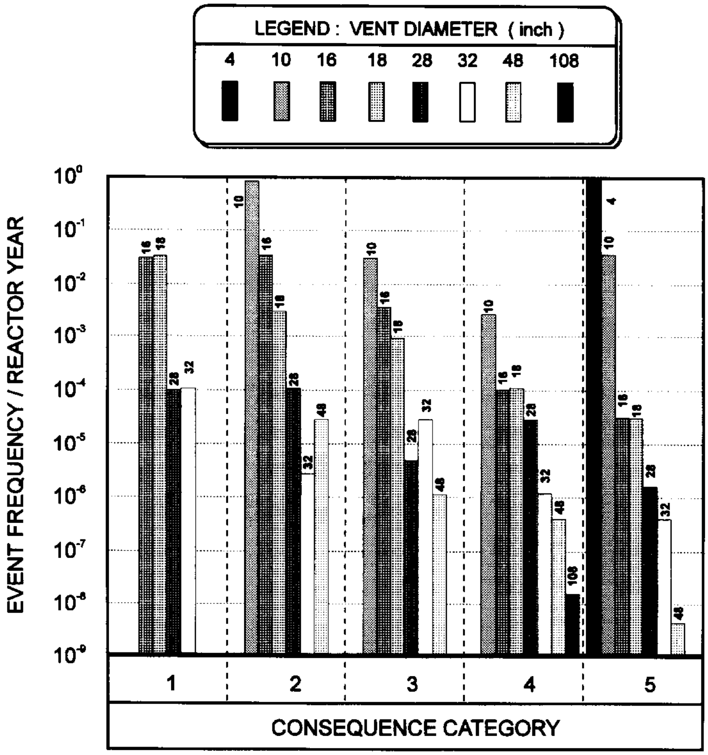

Fig. 8. Frequency distribution of runaway reaction events.

For example, one could insist that the reactor increase the risk of death or injury to near-by personnel by an amount that is small compared to other risks that are implicitly accepted. When the expected frequency of an accident is near the acceptance criteria, an uncertainty analysis should be performed for that result. Note that in Fig. 8 the estimated frequency of severity level 3 events for the 16 in vent is within an order of magnitude of our acceptance criteria.

A Monte-Carlo sampling procedure is used to compute the frequency distribution of this accident class based on the probability distributions input for each of the basic events in the risk model. Fig. 9 shows the cumulative distribution for the frequency of severity level 3 accidents for a reactor with a 16 in relief. We can see that our acceptance criterion is met at about the 95% confidence limit. This makes the 16 in vent marginally acceptable.

To discriminate between the larger vent designs, the concept of risk is introduced. Risk is defined as the product of consequence and frequency. The preceding discussion has considered the frequency of occurrence of the undesired events. If a consequence is assigned to each severity level based on the cost of cleanup, replacement, and litigation, the total risk of operating a reactor with different vent sizes can be directly

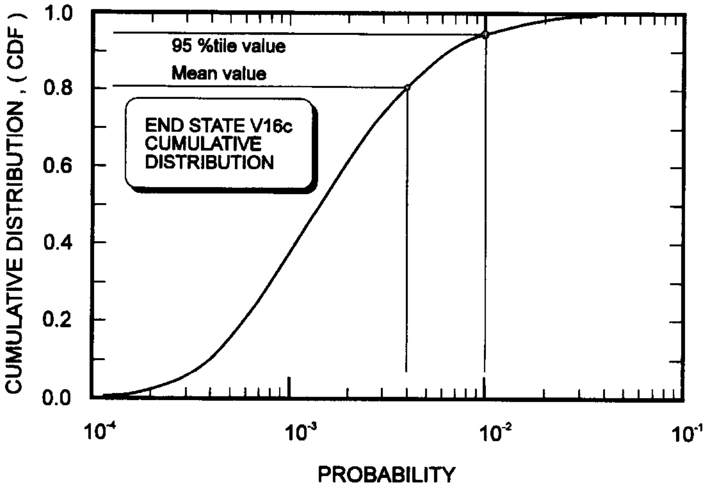

Fig. 9. Cumulative distribution of severity level frequencies for a reactor with a 16 in vent.

compared. Note that in doing this we are not explicitly considering the 'cost' of injuries or fatalities to personnel or the public other than through the frequency-based criteria. Since we are using this figure of merit to compare the same severity levels for slightly different configurations of the same reactor system, the results are generally not sensitive to the absolute dollar values assigned. It is also possible to assign uncertainty distributions to the consequence measures, as was done for the frequencies of the basic events, and propagate these distributions through the analysis.

Fig. 10 presents the risk, in terms of cost incurred over the life of the reactor, due to operational overpressurizations as a function of vent size. Here it can be seen that there is essentially no added benefit of installing a vent device larger than 28 in. Further, when the cost of installation of the particular vent device is added to the operational risk (Fig. 11), we can see that the total cost (which includes both installation costs and the weighted cost of accidents) is minimized for vents in the 16–18 in range. In other words, even though the 28 in vent provides additional protection against all categories of overpressurization incidents, the increase in the cost of installation is much greater than the cost savings due to accident avoidance.

As we can see, it is difficult to determine precisely whether to choose the 16 in vent versus the 18 in vent based only on the mean value of the cost (installation + risk) of the system. The estimation of the frequencies and the associated cost of overpressurization events can be subject to large uncertainties, and considering uncertainty can be helpful in distinguishing two similar alternatives. Fig. 11 shows that even though the mean value of the total cost is approximately the same for both of these vent sizes, the maximum expense (95% confidence limit) that would be expected from the 18 in vent is significantly smaller than that of the 16 in vent. This is mainly because the 18 in vent provides greater protection from the more frequent, less severe accidents.

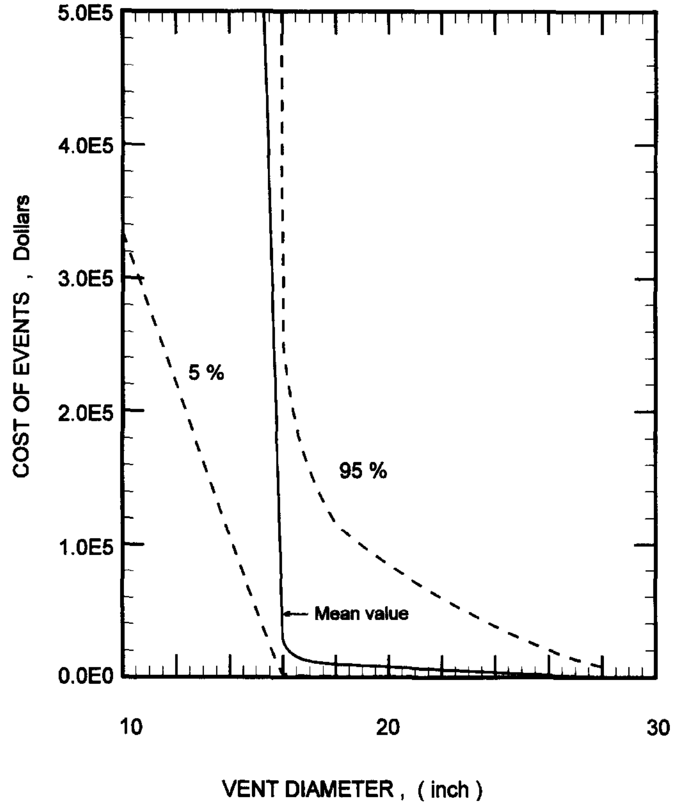

Fig. 10. Expected cost due to overpressurization events.

IRRAS also reports a measure of the relative importance of each basic event used in the calculation of the overpressurization frequencies. Important measures can then be used to determine which events and functions provide the greatest contribution to the frequency of an undesired outcome. For this purpose, the undesired outcome can be the failure of a particular system, a single overpressurization event sequence, an accident with a given severity level, or the combination of overpressurizations of all severity levels.

If there is a way to reduce the frequency of the more severe overpressurization events in a manner that is less costly than installing a larger vent, a less expensive design could be chosen. Improving the reliability of the most important basic events usually provides the most simple and effective means of reducing the frequency of the undesired events. It may also be possible to modify the reactor system (e.g. by installing redundant equipment) so that highly important failure modes present in the original design no longer exist.

Inspection of the important measures reveals that a significant difference between the 16 and 18 in vent designs is found for severity category 3 incidents. For the 16 in relief, improper catalyst loading is the prime contributor to these events. If the

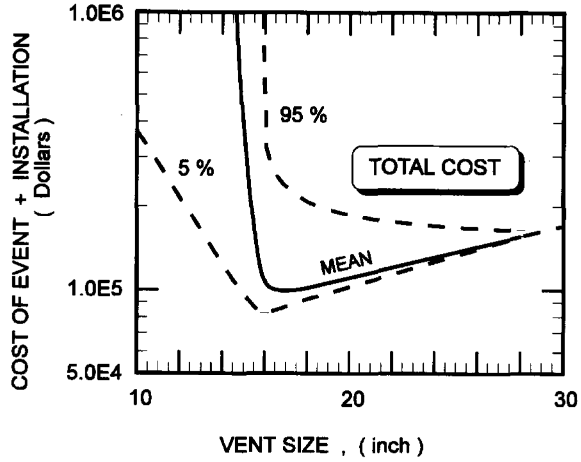

Fig. 11. Total cost.

reliability of this function is increased by an order of magnitude (perhaps by implementing procedures to provide independent verification of the type and amount of catalyst loaded into a batch), the frequency of level 3 incidents is reduced by about 65%, and the uncertainty distribution of the frequency is well within our acceptance criteria. In addition, the frequencies of severity level 4 and 5 accidents will be reduced by 85%. Therefore, a 16 in vent with additional controls on catalyst loading would provide protection against the most severe overpressurization incidents that is similar to an 18 in vent without the additional controls.

When we look at this modification in terms of overall risk savings, however, the picture is quite different. The reactor system with the 16 in vent incurs most of its risk from the severity category 2 incidents, and controlling the catalyst does little to mitigate these. For this reason, the order of magnitude increase in catalyst reliability translates into a mere 20% reduction in the cost of overpressurization events, and the uncertainty band remains large. Here, the controlling factor in overall risk is the reliability of the vent itself, which is a more difficult parameter to control for vents of this size.

The 16 in vent meets the requirements that the severity category 3 incidents do not occur more than once per 100 reactor years, category 4 events do not occur more than once per 1000 reactor years and category 5 events occur less than once per 10 000 reactor years. When the uncertainties in the total cost are considered, we see that an 18 in relief could be a better choice. Even though it is a little more expensive to install, it provides additional protection against the less severe category 2 incidents.

5. Conclusions

This study has subjected a hypothetical process reactor to a vent sizing evaluation. On the basis of a typical hazard evaluation, required vent sizes ranging from 4 in to larger than the vessel diameter are obtained. It is a relatively easy matter to reject the smallest vent sizes (4 and 10 in) which are based on inappropriate evaluation formulae. Similarly, it is easy to recognize the impractical nature of the largest vent sizes (48 and 108 in). This leaves the engineer with a range from 12 to 32 in which is more difficult to deal with on qualitative grounds.

The risk assessment methods illustrated here, clearly narrowed the decision process to the selection of a 16 or 18 in vent. Further precision in the selection, based only on numerical ranking is not advocated. However, various sensitivities are made apparent so than an informed final selection can be made. Another result of the analysis is the identification of other measures for preventing runaway reactions.

It is estimated that the cost of the additional effort for the quantitative risk evaluations would not add more than 50% to the cost of a thorough evaluation of a large process reactor. In this hypothetical example, one can consider the avoided cost of a 32 in vent installation relative to a 16 or 18 in vent as a measure of the benefit to be realized.

Acknowledgements

The authors would like to thank the various reviewers for the helpful comments on the draft of this paper.

References

[1] H.G. Fisher et al., The Design Institute for Emergency Relief Systems (DIERS) Project Manual, AIChE Publication, 1992.

[2] J.E. Huff, Chem. Eng. Progress, (1988) 44.

[3] H.G. Fisher, Paper No. 46a AIChE Summer Meeting, Denver, 15–17 August 1994.

[4] P.F. Nolan, J. Hazard. Mater., 14 (1987) 233.

[5] B. Rasmussen, J. Loss Prevention Process Ind., 1 (1988) 92.

[6] G. Drogaris, Safety Sci., 16 (1993) 89.

[7] G.P. Marrs, F.P. Lees, J. Barton and N. Scilly, Chem. Eng. Res. Des., 67 (1989) 381.

[8] OSHA Process Safety Management Guidelines for Highly Hazardous Chemicals, 29 CFR Part 1910.119, 24 February 1992.

[9] EPA Risk Management Programs for Chemical Accidental Release, 40 CFR Part 68, 20 October 1993.

[10] T.A. Kletz, Chem. Eng., 92(7) (1985) 48.

[11] Guidelines for Hazard Evaluation Procedures, 2nd. Edn., AIChE/CCPS, 1992.

[12] API-520, 5th edn., 1990 American Petroleum Institute, 1220 L Street Northwest, Washington, DC 20005, 1990.

[13] H.K. Fauske, Plant/Oper. Prog., 3 (1984) 213.

[14] Guidelines for Process Equipment Reliability Data with Data Tables, AIChE/CCPS, 1989.

[15] K.D. Russel et al., Reference Manual, NUREG/CR-5813, EGG-2664 vol. 1, 1992.