 © 2003 IChemE

MODELLING HIGH CONSEQUENCE, LOW PROBABILITY SCENARIOS

R.P. Cleaver, A.R. Halford and C.E. Humphreys
Advantica Ltd, Ashby Road, Loughborough, Leicestershire, LE11 3GR

By their very nature, major accidents in the gas and oil or petrochemical industries have a low frequency but potentially high consequences. However, assigning quantitative measures to the risks of such an accident is difficult, as the frequency has a high degree of uncertainty attached to it and the consequences may be sensitive to the assumptions made in representing the scenario within the analysis. It follows that the assessment of such accidents requires a robust methodology to be used. There may be cases in which a worst-case or single representative analysis may distort the analysis compared with a full evaluation. Advantica has developed a methodology that allows the different realisations of scenarios to be analysed, taking into account relevant parameters such as release and wind directions. This paper provides details about the basis of this method. It is shown that the calculations allow FN curves to be built up from the frequency and number of casualties for each realisation. This allows the contributions from different realisations to be analysed and those contributing most to be identified. It is also noted that using this method may reveal that some preconceived ideas about what is the 'worst case' are not always correct. As well as being of use when carrying out qualitative risk assessments for Safety Reports for example, it is noted that these methods are of great help in cost benefit analysis. Using worst-case assumptions can often mask the benefits of protective measures, as they appear to have little effect on the worst-case realisation. However, using the methodology allows a correct identification of suitable measures to be taken to reduce risks, as the impact of the measure on all the different realisations is assessed.

KEYWORDS: Quantified risk assessment, Major accident scenarios, Gas releases.

INTRODUCTION

Within the UK, under the Control of Major Accident Hazards Regulations (COMAH), there is a requirement for owners of certain installations to prepare Safety Reports. Sites storing an aggregate amount of certain dangerous substances fall under the remit of this legislation. Typically, major gas or oil process and storage sites exceed these threshold amounts and so require Safety Reports. Guidance is available from the Health and Safety Executive (HSE) on preparing the Safety Reports[1]. In terms of the analysis of major accident hazards, the HSE describe the three steps to be followed. These are:

- Identify all the possible major accidents and select a subset for analysis.
- Give a realistic estimate of the likelihood of each major accident hazard or an adequate summary of initiating events.
- Produce an adequate assessment of the extent and severity of the consequences for each identified major accident hazard.

In a previous paper[2], an account was given of the difficulties encountered in analysing the so-called low frequency-high consequence scenarios. In summary, it was argued in that

 © 2003 IChemE

paper that selecting a worst or single representative case may distort the analysis of these scenarios compared with a full evaluation. It was also noted that to compound the problem it is not always clear which case should be selected as 'worst' or 'representative'. A remedy was described in the form of a more realistic treatment that can be used for all such scenarios. The next section of this paper provides a more detailed account of this method. Its application to a typical high consequence scenario is described in the following Section. The paper ends with a discussion of the findings from this and other application of the methodology.

METHODOLOGY

In the first stage, information is collated on the distribution of population, both on-site and off-site. This population is then represented on a map that can be interpreted by the computer-based risk assessment package. Two methods are used for the analysis of people in buildings. The first is to represent a building or a group of buildings as a single point receiver located at the building centre. The second type of analysis is to represent the building using a grid of receivers suitably distributed to take account of its shape. The choice of method depends on the size of the building relative to the fire or explosion being considered. For larger buildings a grid of receivers is more appropriate as it is allows a greater spatial resolution in analysing the effects of incident thermal radiation or overpressure. A group of buildings that is a large distance from the release may be represented by a single receiver. An occupancy level is defined for all buildings and this is taken to vary according to the time of day. Office buildings, for example, are likely to have a higher population during the day than at night. The opposite is likely to be true for domestic buildings. The population who are outside can be represented on a similar grid of points. For on-site locations, the distribution of the points can be selected to reflect working patterns and common activities. In order to account for normal working patterns, the occupancy level at each point is taken to vary according to the time of day. For example, it is usual to consider occupancy during a normal working day (e.g. 8 to 5 Monday to Friday), at weekends and during the night. Special events, such as loading or discharge, may require particular analysis to ensure that hazards that can only arise in connection with certain activities have a population distribution appropriate to that activity.

Each high consequence scenario is analysed in detail, considering the different ways in which the scenario could occur. A representative number of different locations for the release are selected and a range of possible temperature and pressures for the process fluid are selected, as appropriate. For jetted releases, the number of different directions in which the release can point initially is considered. It may also be necessary to consider transient behaviour following the initial release, in order to account for a finite capacity of the system under consideration. The weather conditions that are assumed for each evaluation of the release are chosen to represent meteorological data for that particular location. Given that the scenario occurs, the probability with which each of these particular combinations of conditions arises is determined using a combination of the meteorological data and site information. These combinations are then simulated in the computer package.

The vulnerability at the specific grid points used in the package is evaluated for each realisation of the scenario. For flammable gases, this involves a consideration of the harm arising from any fires or explosions. For gas/vapour dispersion, the main hazard distance is taken to be when the mean in-plume concentration has decayed to the lower flammable limit (LFL). The distance to LFL is considered to represent the maximum distance within which an ignition source could ignite a release, leading to a flash fire throughout all of the source cloud or an explosion in a congested or confined region engulfed by the cloud. In principle, persons and property within this range could be affected in the event of ignition occurring, although in practice the occupants of most buildings would be afforded protection, unless gas ingress had occurred. Ignition of isolated pockets of gas beyond this LFL contour may be a possibility, but it is considered that this would not lead to general cloud ignition[3]. However, in recognition of the potential for such ignitions to cause harm, distances to half of the LFL (½ LFL) are also evaluated and used in the assessment, as noted below.

The effects of thermal radiation are determined by the dose of thermal radiation received as a function of time[4]. A typical hazard range criterion for personnel exposed to thermal radiation is taken as the distance from the fire from which persons can be expected to escape without receiving a defined dose of 1060 thermal dose units $[(kW/m^2)^{4/3}s]$, which is fatal in approximately 1% of cases. A secondary criterion is the distance from the fire from which persons can be expected to escape without sustaining injury in the form of second-degree burns (skin blistering). The dose of thermal radiation required to cause the onset of skin blistering depends on the thermal radiation flux level and the time of exposure, but for the scenarios analysed in this report it is typically of the order of 250 dose units. In calculating the "escape distance" using either criterion, a lower threshold of 1 kW/m^2 is used, to which it is assumed a person can be exposed for an indefinite time without injury. For any assumed escape speed, the "escape distance" is calculated neglecting the possibility of obtaining shelter.

Ignition of combustible material on buildings or structures can also be caused by intense thermal radiation, although this is again dependent on the thermal radiation flux level and the time of exposure. The threshold for buildings exposed to thermal radiation is taken as the flux level at which secondary fires may be started by piloted ignition of combustible materials (minimum 12 kW/m^2).

The effects of an explosion depend on the strength and duration of the overpressure wave that is generated. In the assessment of hazard distances for people inside buildings, a blast wave overpressure of 40 mbar is sometimes used. An overpressure of 40 mbar is estimated to cause 90% window glass breakage. The guidance from the Chemical Industries Association (CIA) on the safety of occupied buildings associated with major-hazard installations[5] indicates that an overpressure of 40 mbar would cause approximately 1% fatality within a population in a typical domestic building. Different levels of overpressure may be required for other buildings, such as supermarkets or sports halls. The overpressure required to cause this level of fatalities rises to 100 mbar for a typical office building. An overpressure value of 180 mbar is considered to be capable of producing fatality in 1% of the population for people who are outside.

The strategy that would normally be adopted in analysing the vulnerability for the high consequence scenarios is summarised in Table 1.

As noted in the Table, the methodology distinguishes between people who are indoors and those who are outdoors. It is assumed that the people who are outdoors at the time of the event attempt to escape to a safe distance. In the case of fires, the thermal dosage that they receive in doing this is evaluated and used in a probit relationship to infer their vulnerability. People who are indoors initially are assumed to stay inside unless the building is predicted to start to burn because of piloted ignition. A certain proportion (10%) of the people inside such buildings are then assumed to become 'trapped' fatalities. It is assumed that the remaining people try to escape from the building. The size of the fire relative to the size of the building is taken into account in this analysis and also the possibility that people could use the different exits from the building. For example, people within the piloted ignition distance within buildings are assumed to attempt escape form the nearest exit, whereas anyone outside of this distance, but still within the building, is assumed to attempt escape from the most favourable exit.

Table 1. Levels of harm used for the impact assessment

Hazard	Hazard range/dose/ metric	Effect of people within range		
		Outside	Inside "normal" building	Inside "hardened" building
Free field overpressure (mbar)	Received blast loading	Calculated based on overpressure - correlation given by Baker[3] for percentage fatalities arising for specified free field overpressure (464 mbar taken as producing 100% fatality, 300 mbar 50% fatality, 180 mbar producing 1% fatality)	Calculated based on overpressure - correlation given in CIA guidelines[5] (600 mbar taken as producing 100% fatality, 250 mbar 50% fatality, 40 mbar producing 1% fatality)	Case specific analysis based on specification for building – typically greater than 1000 mbar required to produce fatalities.
Jet fire, pool fire or fireball	Secondary fire range (based on piloted ignition of wood)	-	10% of residents of buildings completely engulfed without any fire proofing are assumed to be fatalities– the remaining 90% of residents seek to escape at the time of piloted ignition and vulnerability is calculated as for people outdoors within escape distance	

	Within escape distance (calculated for person outside attempting escape at uniform speed from event)	Percentage vulnerability calculated from Probit relationship, based on received dosage in attempting escape.	People within buildings that are located outside of the secondary fire distance are assumed to remain in buildings and to be safe.
Flash fire	Within LFL contour	Assumed fatalities 100%	Protected by building
	Between LFL contour and 0.5 LFL contour	Vulnerability reduces from 100% at edge of LFL contour to zero at edge of 0.5 LFL contour	Protected by buildings

The vulnerability information at each grid point and the frequency information for each realisation of the scenario can then be combined to determine the overall risks posed by each scenario. That is, the location specific individual risk is defined at each grid point. Further analysis of the output from the calculations enables values of the maximum number of fatalities, the average number of casualties and the societal risk (the combination of the frequency of each event with the expected number of fatalities) to be determined. This also enables the contribution to the risk from different realisations to be analysed and those contributing most to be identified.

ILLUSTRATIVE EXAMPLE

In order to demonstrate the principles of the methodology, a hypothetical LNG storage site has been created. The population distribution surrounding the site is based on one specific location in the UK, in order to give realistic numbers for subsequent analysis. The LNG Storage site and the accident scenarios have been ‘superposed’ for demonstration purposes and are not intended to reflect the situation at any real site.

The Site has been split into four main areas, with associated population distributions, as follows:

- Inside the administration building, with entrance building
- Inside the workshop
- Inside the control room
- Outside in the process area (within the site boundary)

It is assumed that ‘night’ occurs for 50% of the time and peak daytime hours and off-peak daytime hours each occur 25% of the time. A typical working day (peak hours) therefore occurs for 42 hours per week. The control and Administration building, the Workshop and the Stores were each represented using a grid of receivers. The remainder of the site was divided into 117 rectangular areas, each of which was represented by a receiver at its centre. Seven local business or commercial properties are also assumed to be situated close to the site, with a number of isolated farms of hamlets and smaller villages, and a large industrial estate nearby.

Figure 1 shows the locations of the populations assumed for this study.

It would normally be assumed that the population of domestic properties is lower during the working day, when some people are at work. However, in order to account for a number of small businesses in the settlements, it has been assumed in this analysis that the population is constant.

A typical scenario that could have onsite and offsite effects is a large spill of a flammable, volatile liquid. A release of LNG from a storage tank has been selected as an illustrative example. The failure of an LNG tank is extremely unlikely. Studies have been carried out to estimate an appropriate failure frequency for this event (see, Lees[6] for examples). For the purposes of illustration, however, it is assumed that the tank fails catastrophically leading to a release of its contents (cryogenic LNG). A number of different cases are modelled as follows:

- Case 1 Total tank failure, simultaneous failure of surrounding bund – release spreads as though on flat terrain.
- Case 2 Total tank failure, surrounding bund undamaged, some LNG overtops depending on spread calculation.
- Case 3 Total tank failure, liquid retained in bund.

Cases of immediate and delayed ignition are considered, as appropriate. This is typical of a high consequence-low frequency event of the type considered in Safety Report for an LNG Storage site (see discussion on Glass and Johnson[7]).

For immediate ignition, a running pool fire model was used to calculate the maximum diameter of the pool fire. Radiation predictions for a steady state pool fire with this maximum diameter were used to calculate the hazard to people. The pool fire was assumed to exist until the mass of LNG burned by the fire exceeded the mass of LNG released. It is likely that after the bulk of the LNG pool had burned away, a much smaller pool fire would persist close to the release point. However, sensitivity studies showed that the additional effects of the radiation from this fire were negligible.

If no immediate ignition occurs then the dispersing clouds may ignite at a later time. There is also a possibility that local pockets of flammable gas may exist between the location of the mean in-plume LFL and half LFL contours. For the purposes of this analysis, it has been assumed that when ignition occurs, all of the gas in the cloud inside the mean in-plume LFL contour is burnt in the flash fire and that all of the local pockets of flammable gas between the LFL and half LFL are ignited. It is also assumed that the remaining LNG in the liquid pool ignites at the same time.

For delayed ignition, a running pool fire model was used to calculate the maximum diameter of a pool fire ignited in the appropriate time interval. Radiation predictions for a steady state pool fire with this maximum diameter were used to calculate the hazard to people. The pool fire was assumed to exist until the mass of LNG burned by the fire exceeded the mass of LNG released. A transient dispersion model was used to predict the maximum extent of the mean in-plume LFL and half LFL contours in the appropriate time interval.

In summary, the following scenarios were modelled:

- 5 different ignition times
- 3 pool diameters
- 11 different combinations of wind speed and atmospheric stability
- 12 different wind directions

To calculate the average fatalities, the probability of each wind speed and direction was taken from the wind rose data. Each appropriate atmospheric stability was assumed to be equally likely, that is, in a 2 m/s wind speed there was a 20% probability of each atmospheric stability between B and F, whereas a stability category of D was always used for a 10 m/s wind. The overall average number of fatalities for this scenario is calculated by assuming that immediate ignition occurs 30% of the time, delayed ignition 60% of the time and no ignition 10% of the time. It has also been assumed that case 3 happens 90% of the time, case 2 9% of the time and case 1 the remaining 1% of the time. If delayed ignition occurs, then ignition in each of the separate time intervals is assumed to be equally likely.

Figure 2 shows how the average number of fatalities varies with time of ignition. As can be seen, for this transient event, it would not be obvious in advance which set of conditions would lead to a 'worst' case. Figure 3 complements this information by showing how the relative vulnerability (or relative individual risk) varies with distance from the source of the release along a particular 'ray' emerging from the source. The vulnerabilities are shown for 5 different times after the start of the scenario. This shows that vulnerability at each distance increases and then decreases throughout the event. However, the time of maximum vulnerability is different for different distances from the source.

The FN curve for this scenario is shown in Figure 4, assuming an appropriate event frequency. This is compared with a line that can be inferred from HSE Guidance as to acceptable societal risks for any single event from a large installation[8].

DISCUSSION

An analysis of the type described above allows the calculation of individual risk at specific locations, for instance in the control room, of individual risk contours and of societal risk. The structured way in which the scenarios are handled removes a degree of subjectivity from the analysis and allows a consistent and auditable approach to be used. Using a realistic range of parameter values, when exact values cannot be know in advance, reduces the sensitivity of the analysis to the input conditions. However, like all methods of analysis, the answers that are obtained are only as good as the quality of the input data used to describe the scenario.

The results obtained can be compared with risk criteria in order to aid decision-making on the acceptability of the risks. A typical comparison is shown in Figure 4, where the FN curve is compared with a particular acceptability criterion that has been applied to determine if a scenario is totally unacceptable in all but exceptional circumstances. One of the main purposes of carrying out the assessment of major accidents is to aid the decision making process, as to whether risks are broadly acceptable or whether further safety improving measures need to be taken. The realistic analysis gives a much better representation of a major accident, as opposed to using a simple worst case or representative analysis. This in turn gives a more accurate assessment of risk, thus enabling decisions to be taken with a higher degree of confidence.

The analysis method illustrated above has been applied in practice to real sites. For example, it has been used in the preparation of COMAH reports in the UK and it is currently being used in updates to some of these reports. The method has also been used in quantitative risk assessments for sites elsewhere in the world, both for existing and new facilities. There are number of important benefits in using the methodology. It allows a better judgement to be made of the distribution of risks. For example, cases have been found where the maximum number of fatalities and the average number of fatalities peak at different times after the event initiation. Depending on the area of most concern, different mitigation measures may be proposed to tackle what is the largest contributor to the overall risk. Again, use of this approach is less subjective and gives consistency between assessments.

The analysis also allows the benefits of any specific mitigation measures to be analysed, particularly if linked to cost benefit analysis. This was illustrated elsewhere[2], where a particular liquid spill was analysed, with or without flow limiters installed to reduce the spill rate. It was shown there that the realistic analysis allowed the effects of the flow limiters to be assessed more rigorously by not only comparing the maximum number of fatalities but also the average number of fatalities. Using a simple analysis, the results of a cost benefit analysis showed little benefit in installing the suggested safety improving measures, whilst a realistic analysis gave a much better measure of the improvement. The implications of this are significant, as using a worst case analysis may mean that safety improvement measures may not be installed when they should be or, if alternatives are being assessed, a less effective option may be selected. Having automated the method it is efficient to use on a standard desktop PC to investigate such issues. Despite its complexity, changes can be assessed quickly, which is of particular use in cost benefit analysis or if the project is at a design stage. The speed of modern computers means that time taken to carry out calculations is not significant. The most time consuming stage of the process is in obtaining the appropriate input conditions to represent the scenario and discussing with site personnel the way in which the plant will respond to different types of release. Experience suggests, however, that time spent in this way is well spent.

Finally, it is noted that, in principle, the vulnerability of the environment could also be evaluated on a similar array of grid points. For example, the amount of any toxic fluid calculated to reach a sensitive area or waterway or that is calculated to percolate into the subsoil could be used in evaluating the harm to the environment in a similar way to using thermal radiation or over pressure in evaluating the harm to people. In this way the environmental consequences of major accidents could also be assessed.

REFERENCES

1. Environmental Agency, HSE and Scottish Environment Protection Agency, 1999, Preparing safety reports – control of major accident hazards regulations 1999, HSG 190*, HSE Books.*
2. Cleaver, RP., Robsinson, CG., Halford, AR., July 2002, Analysing high consequence, low frequency accidents on process and storage plants, *Chemical Engineer.*
3. Birch, AD., Brown, DR., Fairweather, M., Hargrave GK., 1989, An experimental study of a turbulent natural gas jet in a cross-flow, 1989, *Comb. Sci. & Tech,* 66; 217–232.
4. Hymes, I., 1983, The physiological and pathological effects of thermal radiation, *SRD* R275.
5. Chemical Industries Association. Guidance for the location and design of occupied buildings on chemical Manufacturing sites, February 1998, *Chemical Industries Association (CIA) publication.*
6. Lees, F., Loss prevention in process industries, 1995, *2nd Edition butterworth Heinemann*
7. Glass, D., Johnson, M., Demonstrating the tolerability of risk from major accidents, *March 2003,* Submitted to Hazards XVII, Manchester March 2003.
8. HSE, Reducing risk, protecting people, 2001, *HSE Books*.

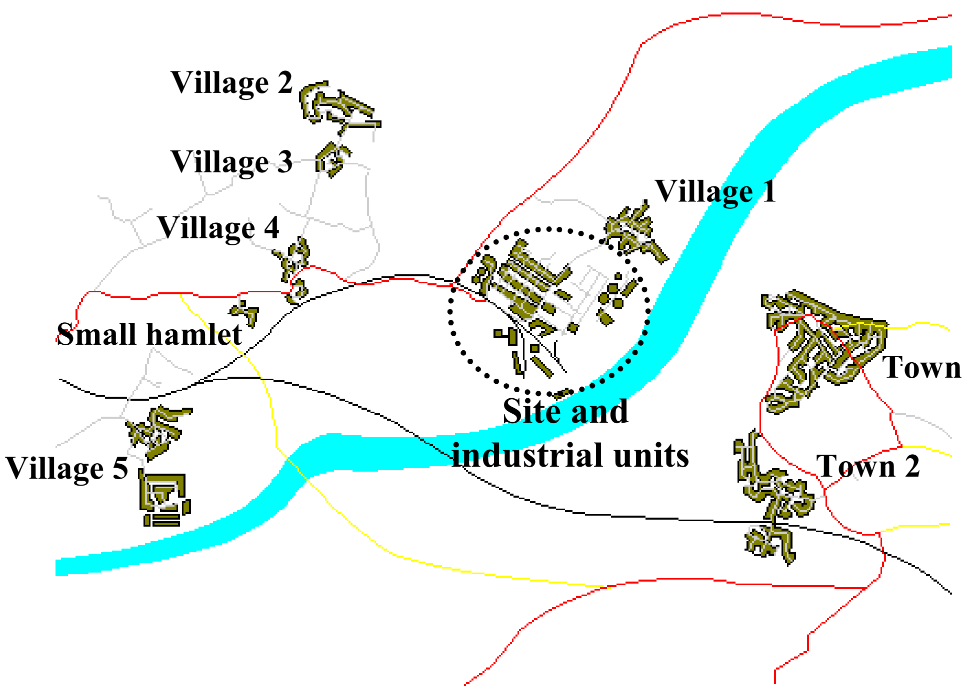

Figure 1. Location of the off-site population assumed in this hypothetical example

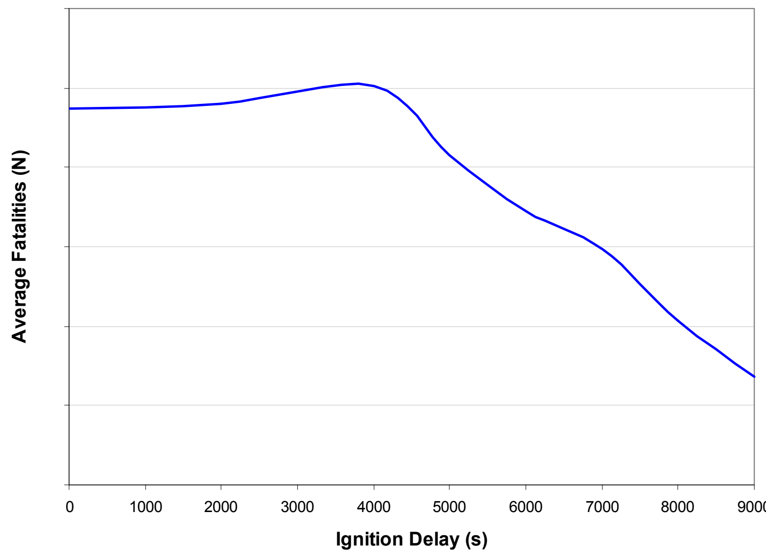

Figure 2. Average number of fatalities for LNG tank failure

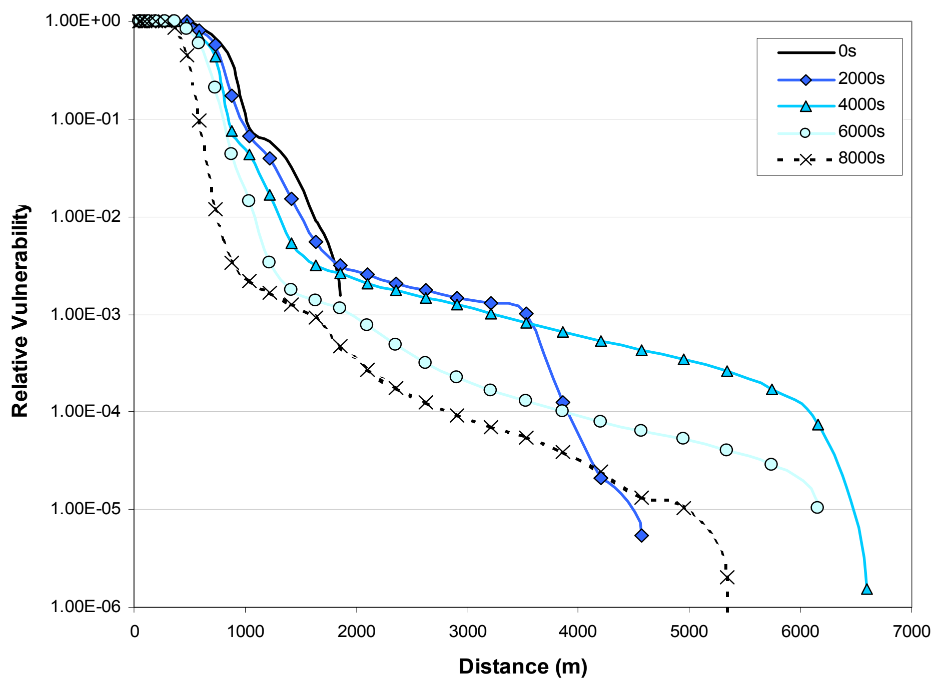

Figure 3. Vulnerability of individuals along a line from the centre of the release for persons inside 90% and outside 10% of the time.

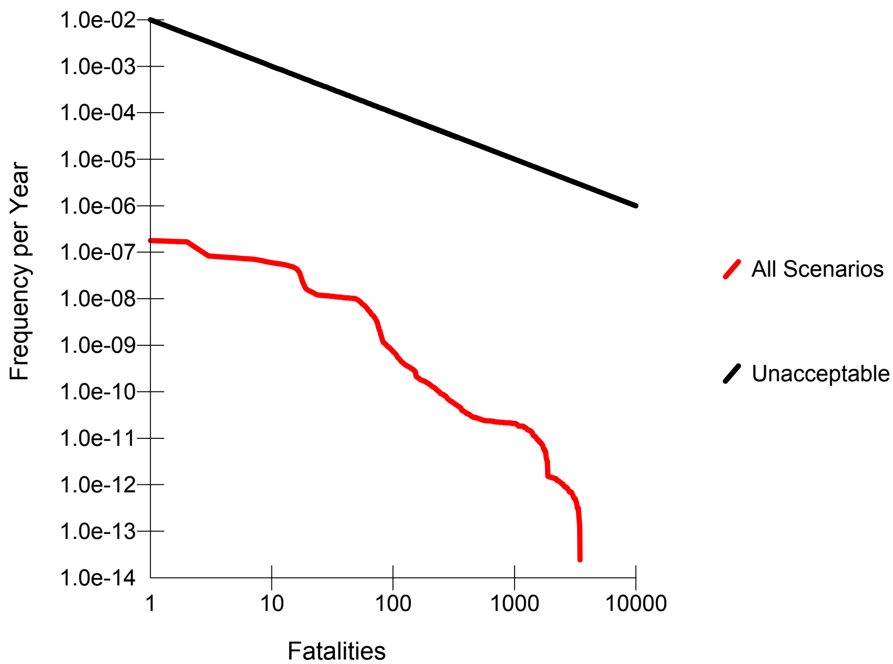

Figure 4. F-N curve for a failure of the LNG tank