DEVELOPMENT OF LOW-COST RISK ANALYSIS METHODS FOR PROCESS PLANT

R.A. Cox and P.J. Comer *

A detailed study has been made of alternative approaches to the reduction of the amount of effort required in risk assessment work. A minimum specification for the methods was that they should be capable of giving output in the form of iso-risk lines round the plant and F-N curves.

It was concluded that a 'Simplified Classical Method' would be most suitable for examination of particular plants or complexes in a specific location, or for general planning problems. A 'Simplified Parametric Method' would be more suitable for a general 'ranking' of hazards or for generating an extensive survey of a very large number of plants.

1.0 INTRODUCTION

1.1 Background to this Study

The public authority Rijnmond, which is the local government authority for the Rotterdam/Europoort area, has been active in the field of risk assessment of process plant for many years. These activities have included a very detailed pilot study of six selected hazardous installations (the COVO pilot study - Blokker et al., 1980) and a project for the 'Inventorisation of Hazardous Objects in Rijnmond' (the IGOR project) in which the aim is to create a complete inventory of all installations in the area with the potential to cause a major accident. The COVO pilot study used the 'Classical' method of risk analysis, that is, identifying all the failure modes, quantifying their probabilities and consequences and summarising these in terms of risk to life. Following the completion of that study, however, it became clear that the Classical risk analysis method would be much too laborious for application to all of the installations in the IGOR inventory; nevertheless, the general objective of evaluating the level of risk from the whole of the industrial complex was still seen as desirable and therefore it was decided to seek low-cost methods for carrying out the risk analysis. This was implemented in the form of a feasibility study carried out under the auspices of the IGOR project steering committee, and the present paper is based on the resulting final report.

1.2 Forms of Output Required from the Risk Analysis Method

The main requirements for the possible methods were that they should be capable of giving risk values in the form of risk contours, cumulative frequency curves (F-N) and average rates of death for both population and employees, as in the COVO pilot study.

*TECHNICA Ltd., 11 John Street, London WC1N 2EB.

For public risk, both the risk to an individual and the risk to groups of people have to be considered. Commonly, individual risk is represented by risk contour plots, which give the geographical distribution of risk, from which the risk to an individual can be determined if one knows his pattern of movements through the locality. These contour plots depend only on the hazard itself and not on the population density distribution, which leads to certain simplifications in their derivation. Group risk is generally represented by graphs of the frequency F with which accidents exceeding a given size N are expected to occur (the F-N curve, or cumulative frequency curve). The size of the accident is usually measured in terms of numbers of fatalities.

For employees' risk, again both individual and group risk could be considered. However, for individual risk it is unlikely that the risk contour method of presentation would be useful, because for the relatively large failure cases which are the principal interest here, the contours inside the plant would be rather flat. Moreover, to get a good description of damage zones inside the plant would require a more detailed analysis than is practicable in a 'low-cost' method. Similarly, for employees' group risk it was considered to be too difficult to make the analysis sufficiently detailed to be useful for the employees on their own, but there is a strong case for treating employees and population together when calculating group risk, because many of the scenarios would affect both simultaneously.

For the calculation of group risk, it is inherently necessary to consider the population density distribution, which leads to some additional complexity in the analysis, as compared with the individual risk case.

The above remarks concern the principal outputs required from the analysis. In addition, it is useful to have certain details which will help in checking that the analysis has been correctly performed, and in giving an indication of the features of the plant which contribute most to the risk, so that design and layout may be optimised (for new plant) or remedial measures recommended (for existing plant). Examples of these more detailed results include:

(i) extent of damage zone for individual failure cases;

(ii) expected frequency values for individual failure cases;

(iii) risk contours or F-N curves for separate process or storage units.

2.0 OUTLINE OF POSSIBLE APPROACHES

Two main types of approach were identified during the study: first, the simplification of the 'classical' method; and second, a so-called 'parametric' approach in which risk values are related to a limited set of descriptive parameters through a correlation relationship. There are many variations of the parametric method, depending on the number of parameters, the number of output values and the degree of modularisation used in the plant description.

These two approaches have strongly contrasting basic philosophies. The simplified classical method is an attempt to model the reality; it involves therefore the representation of all the possible events and the generation of much detailed information. This mass of information then has to be integrated into a concise form to make it useful. On the other hand, the parametric approach is based on the idea that such a volume of detailed intermediate results cannot really be necessary if, in the end, they are presented in summarised form.

I. CHEM. E. SYMPOSIUM SERIES NO. 71

Therefore, it is assumed, it should be possible to establish a direct relationship between the risk output and the gross observable features of the plant, provided one can determine what the most important factors are.

The 'simplification' approach would involve consideration of individual failure cases, as in the full classical method, but only a limited number of selected cases would be considered. Additionally, the means of identifying and selecting the failure cases and the analysis itself can be simplified so that the quantification of each failure case and the summarisation of the results is efficient. However, it was soon clear that it would be impossible to implement a Simplified Classical Method at sufficiently low labour costs without the use of computers, so this scheme was designed from the start with computer implementation in mind.

In the Parametric Correlation Method (PCM), it is assumed that the characteristics of a plant or section of plant can be summed up in the form of a small set of parameters, such as complexity, scale, intrinsic hazardous properties of materials, and so on, to each of which a value can be given in any particular case by some reasonably objective method.

It is then assumed that the main features of the risk impact of the plant or unit will depend uniquely on these chosen parameters and can be related to them through a correlation function which can be determined. In order to do this, it is first necessary to find ways of expressing the risk impact (both in risk contour and F-N curve form) in a simple way which is itself characterised by a small number of parameters. This may be best explained by a simple example:

Assume that, for a particular type of process unit, the risk contours will be nearly circular, and that the decay of risk with distance is expressed by the function R(D). Assume further that for all units of this type, the form of R(D) belongs to a family of curves

$$R(D) = \frac{R(0)}{2}\left\{1 + \cos\left(\frac{\pi D}{Dmax}\right)\right\} \qquad \text{(see Figure 1)}$$

Any one curve in this family is specified by R(0) and Dmax so that once these parameters have been estimated, the risk impact is defined and the contribution from this unit of plant can be added in with those from other units.

In this simple example, the aim would be to find ways of relating the parameters R(0) and Dmax directly to the observable plant characteristics through a transfer function. In principle, this transfer function could be found by a mixture of multiple regression techniques and more direct theoretical deductions. An example of the latter is that Dmax will depend mainly on the scale of the installation and the hazardous properties of the materials and not at all on the complexity of the plant.

Although the PCM could be used with a great variety of different parameter vectors, the special case of two parameters (representing unreliability and hazard potential respectively) was studied in detail during this project.

3.0 DEVELOPMENT OF A SIMPLIFIED CLASSICAL METHOD

In this section we describe the overall calculation scheme for the SCM, which contains three principal parts: failure case selection and probability estimation; consequence sub-models and their interconnection; and summarisation of the overall risks to life.

3.1 Failure Case Definition and Probability Estimation

The definition of a representative set of failure cases constitutes an important and difficult problem even in the full classical method, and is even more troublesome in the SCM because each of the selected failure cases has to represent a broader range of real failure cases. Consider first the total set of possible actual failures. This comprises certain distinct and unique events (such as unstable sudden crack propagation in a vessel, leading to total failure) and also certain events which may vary in scale in a continuous way (such as a hole in a pipe wall, which could take various sizes within certain limits).

It is essential, in any analysis which seeks to evaluate the total risk of a plant, that this full 'spectrum' of failure cases should be represented properly. This means that each of the 'discrete' failure cases selected for use in the analysis must stand for a range of actual failures, in such a way that the integrated probability/consequence impact of the actual failures is well modelled. The method most often used at present is to pick a discrete failure which is typical of the actual failures in respect of its consequences, then to assign the total probability of that particular range of actual failures to the selected discrete failure. It is important to ensure that there are no gaps and no overlaps in the representation of the complete set of failures for the whole plant.

Despite considerable effort to find a short-cut method, it was concluded that a small set of failure cases could not be chosen without first defining the complete set. In order to achieve a sufficient reduction in labour, it was therefore necessary to devise easier methods of engineering appraisal and failure case definition. This can be done by (a) taking care only to collect engineering information which will actually be used in the models and (b) using a computer routine to define the pipework failure cases (numerically the largest class) from simple input data.

3.1.1 Methods of obtaining the 'full' failure case set

The 'full' failure case set, from which the final short list of failures will be derived, is itself already a discretisation of the real spectrum of failure cases. In Classical risk analysis, it is obtained by a combination of techniques: Checklists, Hazard and Operability Study, inventorisation of hazardous materials, and so on. Often, some kind of qualitative engineering appraisal is required, to acquire factual information and to identify any special process hazards (e.g. runaway reactions). These steps can all be simplified to some degree.

3.1.1.1 Engineering Appraisal

The study concluded that this work could be simplified by:

(a) Using available process drawings rather than checking everything on the plant itself - the possibility of unsafe modifications is a risk that is properly allowed for in the 'Management Factor' discussed in Section 5.0 below.

(b) Collecting only the information that will actually be used in the analysis, by use of prepared questionnaires.

(c) Streamlining the description of control and protective systems by (for example) recording the number of independent protective devices but not the details of each one.

(d) Streamlining the description of operating procedures by recording only the proportions of time spent in each operating mode. The adequacy of the procedures has to be assessed more generally and included in the 'Management Factor'.

3.1.1.2 Checklist of containment failures

Containment failures will, in practice, make up the majority of the failure cases. The definition of failures in this class must start with a complete inventorisation of all hazardous materials on the plant, and of all components in the containment system. It is suggested as an initial basis that containment failures should be classified in the following way:

(i) Complete failure of vessel (storage tank, process vessel etc.).

(ii) Partial failure of storage tank (e.g. due to overflow, failure of roof seam, tank fire, etc.).

(iii) Guillotine break in main pipework.

(iv) Small break in main pipework or equipment (e.g. flange leak, pump seal failure, small branch broken off, puncture by external impact, etc.).

3.1.1.3 Definition of pipework failures

The generation of the pipework failure set is a task which can be readily automated using a computer algorithm. Such an algorithm would greatly reduce the effort required in defining this set of failure cases. An outline scheme for a pipework failure algorithm is shown in Figure 2. The data required for this algorithm are summarised in Table 1.

These data provide sufficient information for the construction of a network showing each of the items, their connections, and the location of important valving. The network can then be used to identify inventories of material that could be released following a failure, taking into account vessel connections and shut-off valving. The ability to identify various types of valving will give greater flexibility and a better representation of the full failure case set. A limited number of valve types are defined within the algorithm which are then referenced in the data file. The valve types would include: excess flow valves, check valves and block valves (manual, remote-operated and automatic). The effect of these valves on the failure cases is two-fold: first, if the valve operates as intended the duration of discharge from that side of the break is reduced, and second, the probability of the failure case in which the valve fails to close includes the unreliability of the valve.

3.1.2 Reduction of the failure case set

The number of failure cases in the full set, determined by the methods described above, will be very large for most process plant. They must be reduced to a manageable number of 'equivalent discrete failures' (EDF) before the quantitative analysis is undertaken. The principle which is followed in making this reduction is that of grouping failures which have similar effects, and adding together their frequencies. Two alternative approaches to this were defined: one in which 'clusters' of failures with similar values of their descriptive parameters (e.g. release rate, toxicity of material, etc.) are identified (the 'Clustered EDF' approach) and another in which each failure is compared with a predetermined standard set of failures for which the consequence calculations have already been carried out in detail (the 'Standard' EDF approach). The latter was found to be the more efficient and repeatable method, mainly because of the difficulty of determining criteria for identifying clusters of failures with several different controlling parameters.

In the Standard EDF approach sets of standard failure cases are defined for each release type, and the full range of consequences associated with each Standard EDF are evaluated and recorded as part of the methodology. Each set of Standard EDFs would be characterised by a single parameter, or 'benchmark' result, which would typically be an effect distance such as vapour travel distance in some standard weather condition. Thus use of a benchmark result to characterise the release case rather than the primary release parameters allows combinations of dissimilar parameters (such as material, release rate, temperature) to be combined into a single measure. Thus the standard EDF results are independent of material or other individual release conditions.

The data file of Standard EDF results would contain sets of effect distances for all the consequences possible for that release type. Thus it would include results for different dispersion conditions, and, for flammable releases, the effects of ignition at various set distances away from the source. The data file would contain only the consequences, and would not include any of the probabilities or conditional probabilities for alternative outcomes. These would be added in the summarisation routine (Section 3.3 below) where the effects of location would also be evaluated.

In applying the Standard EDF method a 'benchmark' consequence calculation is performed for each failure in the full set of failure cases, using the sub-models as described in the following sections for the standard weather conditions. The benchmark result is then used to identify which Standard EDF is applicable for each case.

3.1.3 Estimation of frequencies of the EDFs

Because of the practical impossibility of examining each plant item individually for features affecting the failure rate, the frequencies used in the calculations must be based on generic classified data. The most important objective will be to get the relative magnitudes right, so that the method gives a correct impression of the proportion of total risk attributable to different types of plant. The absolute level of the frequency values should also be given close attention, but as a second priority. Possible sources of raw data for this include: previous risk assessment studies, such as WASH 1400, Canvey, COVO study; data banks, such as SRS; and published papers, such as Baldock (1979) on ammonia handling equipment and Welker and Schorr (1979) on LNG.

For the special process hazards, no general rule can be given, however it will often be found that special hazards are protected against by one or more control systems, and so these cases often fall within the general class of "systems hazards" whose failure rate cannot be simply derived from historical data. In the COVO study, these failures were analysed using complicated fault trees, which could not be utilised in a 'low-cost' analysis. However, some credit must be given for the presence of protective systems and it is suggested that this should be done by assuming that the overall failure rate for any one type of special hazard depends only on the number of process

abnormalities that could initiate such a hazard, and the number of independent protective systems. A small fault tree then permits the overall failure rate to be estimated.

3.2 Simplified Procedures for Consequence Calculations

Consequence models include: discharge rate, vaporisation, dispersion, explosion, toxic effects etc. In principle, it would be possible to create correlations of the results of predictions by well-established theoretical models, and to use these correlations in place of a full computation. For certain simple models which have few input parameters this correlation is likely to be quite easy, but some of the models are sensitive to many factors, and therefore the correlations would have to be multi-dimensional. An example is the dense vapour dispersion problem, which is sensitive to release rate, initial density, safe concentration level and surface roughness. It would require a relatively large effort to develop the correlations for this model, but it is clearly quite feasible and the effort, once made, need not ever be repeated.

In the Standard EDF approach, the consequences of the standard EDFs are calculated once-for-all as a separate initial exercise. Figure 3 gives a flowscheme for this process. This creates a standard data file of information about the consequences of each standard EDF in all possible weather conditions and (for flammable clouds) for all possible ignition locations. This will be used as basic data for the summarisation routine.

The second phase of the Standard EDF analysis concerns the application to a particular plant. The flowscheme is given in Figure 4. Here, the information specific to a particular plant (e.g. unit locations and the frequencies assigned to each Standard EDF) are output as a data file which will be read by the Summarisation Routine. The only use of the sub-models is to decide which Standard EDF most closely corresponds to each of the actual EDFs.

The particular sequence of calculations required to evaluate the consequences of any given EDF, for any of these schemes, will depend on the type of failure case and on the type of hazard associated with the material being released. Six main categories of 'Release Type' have been identified, and these are defined in Table 2 along with the typical types of consequence associated with each, for flammable and toxic hazards. Logic diagrams giving the sequence of calculations and the sub-models required for each were worked out in detail during the course of the study. An example, for instantaneous release of flammable flashing liquid, is given in Figure 5.

A summary of the sub-models required for the release types listed in Table 2 is presented in Table 3. The specifications required for each of these sub-models were examined in detail as part of the feasibility study.

3.3 Summarisation of Results

The final phase of the analysis is the summarisation, in which the results of the submodels for all EDFs are brought together with the background data on weather, population and ignition sources. Up to this point, the analysis has been carried out without any reference to the environment within which the plant is located. The failure cases have been identified and their probabilities estimated; the consequences have been evaluated in terms of their geographical extent but not numbers of casualties; the probabilities of ignition at the various standard distances and of different weather and wind directions have not been calculated. All of these matters have to be dealt with in the summarisation routines.

The overall flow diagram of the summarisation routine is shown in Figure 6. The basic method which is followed in this flow diagram is to consider all the possible ignition source locations, wind directions and weather classes for each EDF in turn. The consequences for each individual scenario are simply looked up from intermediate data files, while the probabilities have to be calculated. Note especially that in this approach the effect of different ignition source distributions is not seen as changing the consequences of the release, but as changing the probabilities of the various possible consequences. This is a fundamental idea which makes it possible to carry out the calculations of the consequence sub-models without reference to plant location.

Much of the apparent complication of Figure 6 comes from the need to structure the analysis along the logical subdivisions of the EDF set: plants, units and individual EDFs. This is necessary in order to be able to prepare the data in an orderly manner, and also to determine the amount of risk associated with particular parts of the total hazardous industry.

Certain activity boxes in Figure 6 are of particular importance, and these are described in the following subsections.

3.3.1 Calculation of local population and ignition source density on polar coordinate basis

It is assumed that descriptions of the population and ignition source density distributions are available on a regular Cartesian grid, with principal coordinates (x, y). From experience, a minimum of twelve wind direction sectors is required in order to describe the risk contours reasonably well, giving a sector of 30° for each wind direction.

The computations necessary for calculating the effects of a release from any particular point are much simpler if the population and ignition source distributions are re-expressed in polar coordinates centred on that point. Algorithms for this conversion have been devised.

3.3.2 Calculation of frequencies of final outcomes

The data for the frequency calculation consists of the following:

- EDF overall frequency F_e
- weather class and wind direction probabilities $P_{w\theta}$
- ignition probabilities at regular intervals dr in the θ direction - $P_\theta(idr)$ (i = 1....n)

For simplicity, $P_\theta(idr)$ is written P_i in what follows. Then the frequency of an event with ignition at the nth possible ignition location in direction θ is given by

$$F_{ew\theta i} = F_e \times P_{w\theta} \times P_n \prod_{i=1}^{n-1} \left(1 - P_i\right)$$

3.3.3 Calculation of area covered by event, and numbers of casualties

For each final outcome, specified in general by an EDF, weather and wind direction, and ignition location, the consequence sub-models have already provided parameters such as extent X and nominal cloud width CW. All that is required to define the zone actually affected is to combine this information with the absolute location (x, y) and the wind direction θ. It is simple in principle to devise an algorithm for incrementing the risk values within this zone by the frequency $F_{ew\theta i}$. Equally, the numbers of casualties caused can be simply calculated, knowing the polar population distribution Pop(r, θ) and the effect zone width W(r), which is assumed to be of a specified shape, normalised to the nominal cloud width value CW.

These routines are not considered to present any fundamental difficulty in the development of the simplified risk analysis method.

4.0 PARAMETRIC CORRELATION METHODS

The main problem with the PCM is that of determining the form of the transfer function. One possible method would be to set up a few key hypothetical units and carry out a full risk analysis on them, subsequently varying the input parameters in order to examine the changes in the results. Provided that the number of parameters is quite small this would be practicable but quite time-consuming, although any well-founded theoretical relationships could be used to cut down the number of variations that have to be tried. The dependence of the output "risk parameter vector" on the input "plant characteristic parameter vector" could then be found by regression.

Alternatively, the use of theoretical relationships could be extended to include functional forms that have only an intuitive basis, agreed by a number of experts. This would introduce some risk of errors of judgement and some chance of unresolved differences of opinion, but would greatly reduce the work involved in deriving the needed transfer functions.

In principle, the parametric method can be used at various levels of complexity, depending on the number of parameters included in the input and output vectors. For example, the Dow Index system can be considered to be a parametric method, having many parameters in the input and just one parameter (the Fire and Explosion Index) in the output.

It is clear, however, that information is lost if the number of independent parameters in the output is less than the number in the input; an alternative way of looking at this is that more data was collected than was really needed. Conversely, if the number of outputs exceeds the number of inputs, spurious detail is being produced and the outputs cannot all be independent. Therefore, a scheme in which the numbers of input and output parameters are equal is favoured as an ideal, although it is recognised that it may be necessary in practice to collect much engineering information in order to define one input parameter correctly.

If more than, say, three or four parameters are used in the analysis, the definition of the transfer functions becomes a very complicated problem. If the input parameters are P_i(i = 1....n) and the output risk parameters are R_i(i = 1....n) then the transfer functions will have the form:

$$R_i = f_i\ (P_1, P_2, P_3 \ldots P_n) \quad \text{(n equations)}$$

and even with the simplest possible choice of the functional forms f_i, the number of arbitrary constants to be determined within the functions f_i will be very large (e.g. at least 32 for n = 4). In view of the obvious difficulty of determining these transfer functions it was decided to consider the feasibility and value of a simple version of the parametric method, in which only two parameters are used. This would provide a more informative output than the Dow Index scheme, but would not be unduly complicated in the transfer functions.

4.1 Choice of Parameters in a Two-Parameter Method

For the output functions, an example with two parameters was given in Section 2 above. This example was for the R(D) function (risk as function of distance from the plant) and took the form:

$$R(D) = R(0) \times f(D/Dmax)$$

where R(0) and Dmax are the two parameters describing the output. The value R(0) is the risk in the centre of the plant and is therefore dominated by the total frequency of failure and is not significantly influenced by the size of accidents (hazard potential). It is therefore an output parameter which is mainly dependent on the reliability of the plant. The Dmax parameter is a measure of the hazard potential (maximum consequences) of the plant. It will depend only on the hazardous inventory of the plant and not on the reliability. R(0) and Dmax are therefore completely independent parameters, and moreover we can expect that if the input parameters are chosen to represent unreliability (U) and hazard potential (H) then the corresponding transfer functions will be uncoupled, that is:

$$R(0) = f_1(U)$$

and $Dmax = f_2(H)$

This leads to considerable simplification in the method and offers a possibility for developing a parametric method which has common transfer functions for all types of process or storage unit, because the input parameters U and H can be expressed in the same form for all units. The U parameter automatically has the same meaning and physical units for all plants, by definition; the H parameters for different types of hazard can be related to each other by equivalence of the effect distances.

4.2 Determination of Indices for Unreliability and Hazard Potential

4.2.1 Unreliability index U

The U parameter has to represent the total frequency of events of all sizes sufficient to cause a certain minimum level of damage or injury. It should therefore be possible to devise a scheme for evaluating this parameter by the following actions:

(i) count the number of tanks, vessels, interconnecting pipes in different size ranges, hoses and loading arms;

(ii) apply suitable assumptions, based on sample surveys, for the lengths of pipelines and the amount of in-line equipment found on average;

(iii) for items whose integrity depends on the functioning of protective systems, apply correction factors for the number of independent systems and the number of different kinds of hazard;

(iv) total up the frequencies for all items.

The methods used should be broadly similar to those described for the SCM in Section 3.1.3 above.

4.2.2 Hazard potential index H

What is required for the H parameter is some suitable measure of the magnitude of the damage that could be caused by the plant. It should not depend in any way on the likelihood of failure. The H value will be dependent on the hazardous properties of the material or materials stored and on the quantities and the possible rates of discharge.

It seems important that the concept should be modified slightly from that outlined in Section 2 above, where the maximum conceivable accident was implied, because this would not necessarily be a typical measure of the magnitude of all the possible releases. Hence, we propose that the H index should be some weighted function of the following quantities:

(i) largest extent of the effect zone for sudden release of total content of any one container;

(ii) largest extent of effect zone for 50mm (2 inch) diameter hole in any container;

(iii) largest extent of effect zone for special process hazards such as reactor runaway.

Part of the task of developing the Simplified Parametric Method would be to define the needed weighting functions, using the results of sample risk analyses. The transfer functions f_1 and f_2 can be defined at the same time; indeed, with this method the need for these functions may not exist at all because the input parameters U and H could be defined in such a way that R(0) = U and Dmax = H.

The U and H values are next used to generate the R(D) function for each plant unit (see Section 2). From this, the overall risk contour map can be approximately determined by superimposition. For this purpose, wind direction effects may be incorporated as an extra weighting factor, but it will not be possible to allow for the actual ignition source distribution. Some general assumptions about ignition delays would have to be used instead. With these assumptions, the algorithm for summation of the risk contours is very simple and it should not be very laborious either to develop the method and programs or to execute the analysis.

For the generation of F(N) curves, information about the population distribution P(x, y) has to be used. It is possible to calculate the F(N) function knowing only the values of the functions P(x, y) and R(D), by assuming that all effect zones have similar shapes, and that, for any given wind direction, the population distribution is a function of distance only, that is, neglecting population density changes in the cross-wind direction.

The way in which the analysis would proceed is as follows:

(i) assume that all effect zones have a similar shape, defined by the function $W(\frac{d}{D})$ such that the width at a point d from the source is given by $D \times W(\frac{d}{D})$.

(ii) calculate, from the Cartesian population density distribution P(x, y) and the known location of the source, the population density functions $P_\theta(d)$ for each of a number of directions θ.

(iii) for events with extend D, estimate the average number of people affected as follows:

$$N_\theta(D) = \int_0^D P_\theta(d) \times D \times W(\tfrac{d}{D})\, dd$$

(iv) calculate frequency f(D) with which effect zones occur with an extent exceeding D, using the relationship

$$R(D) = -\frac{1}{2\pi D} \int_D^\infty \frac{df}{dd} \times d \times W(\tfrac{D}{d})\, dd$$

(v) factor the function f(d) by the probability of each wind direction θ, to obtain functions $f_\theta(D)$.

(vi) knowing $f_\theta(D)$ and $N_\theta(D)$ calculate the functions $F_\theta(N)$ and sum up to obtain the overall function F(N).

This approach has the great merit that it allows the whole of the technical analysis of the plant to be done in isolation (i.e. without reference to its location or the population distribution) to generate the function f(D).

As with all parametric methods, it is not possible to obtain detailed results for all the individual failure cases on the plant. However, the effect distances that are used in calculating the H parameter are based on a set of standard loss-of-containment cases which will provide physically meaningful intermediate results should these be desired.

5.0 INCORPORATION OF THE 'MANAGEMENT FACTOR' INTO THE ANALYSIS

It has been generally recognised that the adequacy of operational management plays a major role in determining the level of risk from a hazardous process. It is desirable that this factor should be represented in a risk analysis because this will generate an extra incentive for the improvement of training, maintenance and other 'software' matters relevant to the control of major hazards. Against this, however, it has to be remembered that there can be large changes in the management quality during the life of a plant, so that this aspect has to be reviewed frequently if full credit is to be given for it. Alternatively, with less frequent reviews, the importance weighting given to the management factor could be reduced.

Howland (1980) has pointed out that human error is a dominant factor in most process plant failures and concludes therefore that failure rates calculated from equipment failure statistics are incorrect. This, however, is not a fair criticism in that the statistics do (or should) include failures actually caused through human error. Thus the statistics give a picture of the average standard of equipment and the average standard of human error. In a given instance, the actual standards may be above or below the average and Howland suggests that with some initial research effort this could be assessed in some

way and a correction factor applied to the frequency values. This is the approach that has been adopted in this study.

The influence of good or bad management systems is felt throughout every part of the plant and it is therefore a factor which should be applied on an overall basis, instead of being included within the details of the analysis. The influence of good operational management may have an effect on risk in one of two ways: firstly to reduce the probability of a particular failure case, and secondly to influence the distribution of probability among the various branches of an event tree (by better emergency training, for example).

The first of these has a straightforward effect on the risk, and it is reasonable to assume that any such effect will be uniform throughout the whole range of failure cases, reducing the probabilities by a constant factor.

The second effect is more complex because in principle it could alter the relative magnitudes of risks in different failure cases. However, any risk transferred from a large failure case to a medium sized one can be considered to be offset by further transfers from medium to small and from small to negligible. Therefore, the degree of change in the distribution of risk between failures is not so large as might be thought at first. For this reason, and because of the poor prospects for obtaining any accurate data on this effect, it seems reasonable again to consider the management factor to have a uniform influence on all of the failure rates.

The remaining question is how to quantify the effect of the management factor. Two steps are necessary for this: firstly, to devise an objective way of assessing the quality of management and expressing this in the form of a very small number of parameters; and secondly, to find a means for interpreting this assessment in quantitative risk terms. The latter problem involves a judgement about the sensitivity of risk to operational management quality, and although it is possible to determine this from observable data on the more commonplace plant accidents, this research has not yet been done. It could, however, be important to allow for the fact that plant with a large amount of automatic protection may not be particularly sensitive to operational errors.

For the assessment of management quality, a scheme is favoured in which the observable characteristics are noted in a site visit, such as:

- existence of a safety policy, and a clear management structure for safety responsibility;
- duration and content of training at all levels;
- existence of adequate maintenance schedules and records;
- extent of remote instrumentation and control facilities;
- general state of repair of plant and instruments;
- communication channels;
- degree of automation in process control and safety systems;
- use of safety audits and inspections;
- investigation and analysis of incidents; communication of results;

- procedures for permit-to-work, plant modifications etc.

These and other characteristics would be reported in a standard format. Credit points could then be allocated, with importance weightings to reflect the adjudged sensitivity of the risk functions to each characteristic. This would result in a single rating index, and the scheme would be so drawn up that it would not be possible to attain a high rating without scoring reasonably well in every critical subject area.

6.0 EVALUATION OF THE MERITS OF THE ALTERNATIVE METHODS

6.1 Accuracy

The final risk values calculated in the COVO Pilot Study (Blokker et al, 1980) were judged to be generally accurate only to within one order of magnitude in each direction, and possibly two orders of magnitude in some cases. This level of accuracy is not high by the standards of normal engineering calculations, but it is sufficient to discriminate between plants of different types and degrees of hazard, and to make a general comparison with other existing risks, bearing in mind that risk is measured on a logarithmic, rather than linear, scale.

For the low-cost methods of risk analysis, there is a presumption that the accuracy must be lower than for a full risk analysis, because of the smaller amount of analytical effort. In practice, however, this effect will probably not cause insuperable problems because of the 'law of diminishing returns' which applies to analyses of this kind. In the case of risk analysis, the underlying mathematical cause of the law of diminishing returns is that as the analysis becomes more detailed there are two opposing effects on the accuracy:

(i) the modelling becomes closer to reality, and

(ii) the number of procedural steps becomes larger.

Turning to the specific example of the SCM, it must be the case that the accuracy of the analysis will be somewhat less than in the full classical method, because the technique is exactly the same while the amount of detail involved is much reduced. There will be greater problems in the definition of release cases, and it may be difficult to represent the full range of variations properly. However, the simplicity of the analysis, and in particular the feasibility of checking a sample of the intermediate results by hand, should ensure that there is less likelihood of computational error than in a full risk analysis. Therefore, it is felt that the loss of accuracy because of the simplification will be small.

For the PCM, the accuracy will depend on the method of 'calibrating' the transfer functions. If this is done by reference to a full classical analysis, then the accuracy would obviously be no better than that analysis. If it is done by reference to actual experience (which is only feasible for certain common types of unit and for smaller accidents) then it could be more accurate. There are other sources of error due to inappropriate choice of the parameter vector, and due to inaccurate assessment of the parameters for a particular plant. The latter, however, should be avoidable by making the parameters deliberately unambiguous and easily observed.

I. CHEM. E. SYMPOSIUM SERIES NO. 71

However, the method does not lend itself to checking and therefore there is some danger of data preparation or computation errors going unnoticed. Also, the correlation of risk analysis results which is required for generation of transfer functions will still leave a considerable scatter, whose magnitude is not known at present.

For the Simplified Parametric Method, using two parameters for unreliability and hazard potential as described in Section 4, the accuracy will probably be significantly lower than for the SCM.

6.2 Costs

The costs of implementing the risk analysis methods comprise two parts: non-recurring 'set-up' costs and application costs. In this study, the application costs were estimated on a 'per plant unit' basis, where a 'unit' within a plant is taken to mean one distinct functional section, such as a group of storage spheres, the furnace section of a cracker, cell room of a chlorine plant and so on. An individual distillation column would not usually constitute a 'unit' by this definition, but a complete fractionation train would do so. Naturally, some units would involve much more work than others, and the figures given in the table are intended to be average values covering plant ranging from the simple or trivial cases to large and complex units with various possible hazards.

The labour costs were estimated to be as follows:

person-days

Method:	SCM		PCM		Typical Full Classical Risk Analysis
	Clustered EDF	Standard EDF	Full	2 Parameter	
Set-up costs:	180	190	350	180*	0*
Per unit costs:	3.5	2.6	2.2	1.3	25

*these figures, not given in the final report of the study, are the author's own current estimates.

It is emphasised that the above figures for application costs should be regarded as asymptotic values which would only be approached when the analysts had gained experience in the use of the systems. On this basis, the simplified methods did appear to offer some substantial advantages over 'Classical" methods when applied to a very large number of installations. However, much of the gain was due to the fact that simplification or parameterisation both make possible an analysis which is completely computerised; the residual labour costs consist almost entirely of information gathering and data preparation. There is little prospect of reducing these costs any further.

7.0 CONCLUSIONS

(i) In this study, three possible approaches to the development of a low-cost risk analysis method, capable of producing risk contours and F-N curves, have been examined: a Simplified Classical Method (SCM); a Parametric Correlation Method (PCM) and a two-parameter version of the PCM known as the Simplified Parametric Method (SPM).

(ii) The SCM has been examined in considerable detail and a specification for the method has been produced. It is concluded that this method would not be much less accurate than a full risk analysis but the amount of effort involved in routine application, assuming that use is made of computer methods, would be only about a tenth of that of the full analysis.

(iii) The SPM was found to be a practicable method and could be used either for ranking hazards or for production of approximate iso-risk contours or F(N) curves. Its accuracy is significantly lower than that of the SCM or PCM but its labour cost is the lowest of all. It is concluded that it would be impossible to reduce this cost still further and still produce an output which gives a measure of the risk impact of the plants.

(iv) The overall accuracy of the SCM and PCM need not be much lower than that of a full risk analysis. However, the SPM has a lower accuracy than the other two.

(v) The PCM uses a method which is difficult to understand because it does not model the real mechanisms in a realistic way; therefore, the results of this method may not be credible to non-technical people. The cost of developing a full PCM method is very high and it is not certain whether the method will be feasible if more than two parameters are used. For these reasons, the full version of the PCM was not recommended for implementation.

(vi) It has been shown to be feasible to incorporate the effects of safety management quality and the existence of protective and preventive systems into the analysis, whichever risk analysis method is used.

Acknowledgement

The study to which this paper relates was performed under a contract from the public authority Rijnmond and is published with their permission. The authors wish to acknowledge the valuable contributions made by members of the IGOR Steering Committee during discussions on this project. The views expressed in this paper are, however, those of the authors alone and do not necessarily represent those of the public authority Rijnmond, the IGOR Steering Committee or any other participant in the IGOR project.

REFERENCES

Baldock, P.J. (1979) Accidental Releases of Ammonia - an Analysis of Reported Incidents. Proc. AIChE Loss Prevention Meeting, Houston, Texas. April 1979.

Blokker, E.F., Cox, R.A., Lans, H.J.D. and Meppelder, (1980). Evaluation of risks associated with process industries in the Rijnmond area - a pilot study, Proc. 3rd Int. Loss Prevention Symp., Basel, Sept. 15-19, 1980.

Howland, A.H. (1980) Hazard analysis and the human element. Proc. 3rd Int. Loss Prevention Symp. Basel, Sept. 15-19, 1980.

Welker, J.R. and Schorr, H.P. (1979) LNG Plant Experience Data Base, Proc. AGA Operating Section Transmission Conference, New Orleans, 21-23 May, 1979.

TABLE 1 - Data for pipework failure algorithm

A:	Schedule of Item Interconnections Each item to item connection defined by: - identification of items of either end of pipe; - pipe diameter; - location of connection at either end of pipe - in vapour space (V) or below liquid level (L); - Presence of particular types of valving along pipe (excess flow, non-return and shut-off valves).
B:	Schedule of Items For each plant item; - item identification - material and inventory contained - temperature and pressure The item location is specified by the item identification which identifies the plant unit on which the item is located.
C:	Schedule of Plant Units For each plant unit: - Set of parameters to define failure frequencies - location coordinates.

TABLE 2 : RELEASE TYPES WITH TYPICAL CONSEQUENCES FOR FLAMMABLE AND TOXIC HAZARDS

	RELEASE TYPE	FLAMMABLE	TOXIC
1.	LIQ. T > Tbp. INST Instantaneous Release of Flashing Liquid from Pressurised Storage	BLEVE - from failure due to fire exposure Fireball from early ignition Flash Fire/Explosion	Toxic Cloud
2.	LIQ. T > Tbp. CONT Continuous Release of Flashing Liquid from Pressurised Storage	Jet Flame Explosion	"
3.	LIQ. T ≤ Tbp. INST Instantaneous Release of Liquid at its boiling point (Tbp < Tambient)	Pool Fire Flash Fire Explosion	"
4.	LIQ. T ≤ Tbp. CONT Continuous Release of Liquid at its boiling point (Tbp < Tambient)	Pool Fire Flash Fire Explosion	"
5.	GAS. INST Instantaneous release of gas from pressure	Fireball Flash Fire Explosion	"
6.	GAS. CONT Continuous release of gas from pressure	Jet Flame Explosion	"

Note: Tbp = Boiling Point

TABLE 3 : SUMMARY OF SUB-MODELS WITH CROSS REFERENCE MAP FOR RELEASE TYPE

SUB-MODEL	BRIEF DESCRIPTION	RELEASE TYPE											
		1		2		3		4		5		6	
		a	b	a	b	a	b	a	b	a	b	a	b
DISCHARGE RATE	Calculates discharge rates for liquid flows, two phase critical flows and sonic gas flows from containment failures.			X	X			X	X			X	X
JET DISPERSION	Calculates distances to the LFL due to dispersion in a momentum jet.			X								X	
DENSE VAPOUR DISPERSION	Calculates distances to the LFL and cloud width for both instantaneous and continuous release cases using a correlation to dense vapour cloud model results.	X				X		X		X			
TOXIC DISPERSION	Calculates distances to specified toxic load levels using Gaussian dispersion theory for both instantaneous and continuous release cases.		X		X		X		X		X		X
VAPORISATION	Evaluates an equivalent constant vaporisation rate from instantaneous and continuous spills of refrigerated liquids.					X	X	X	X				
RADIATION	Calculates distances to specified radiation intensity levels from Fireballs, Pool Fires and Jet Flames.												
- Fireball		X								X			
- Pool Fire						X		X					
- Jet Flame				X								X	
EXPLOSION	Calculates distances to three damage effect levels from unconfined vapour cloud explosions	X		X		X		X		X		X	
EXPANSION	Predicts fraction of liquid lost from cloud after initial expansion phase for liquid releases, and cloud density for the gas release case.	X	X	X	X					X			

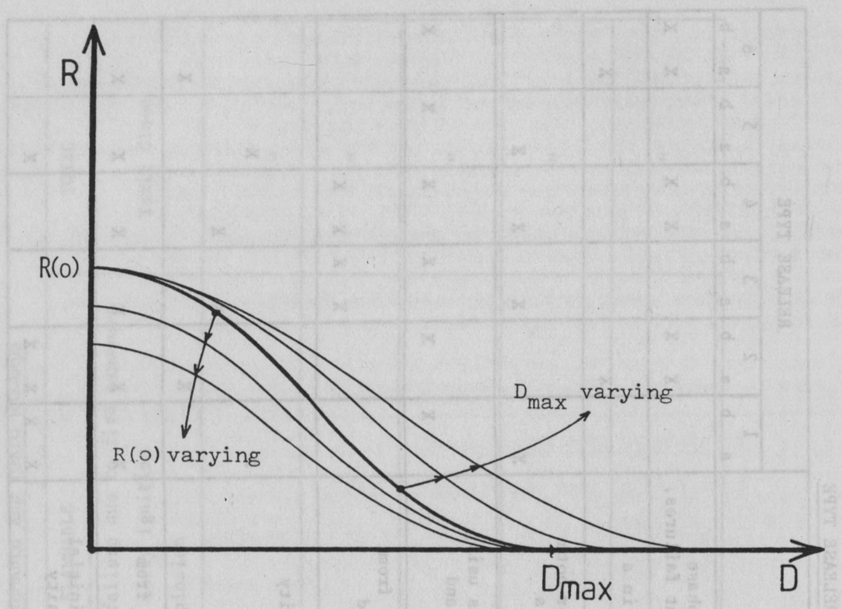

FIGURE 1 : Families of curves for the function R(D)

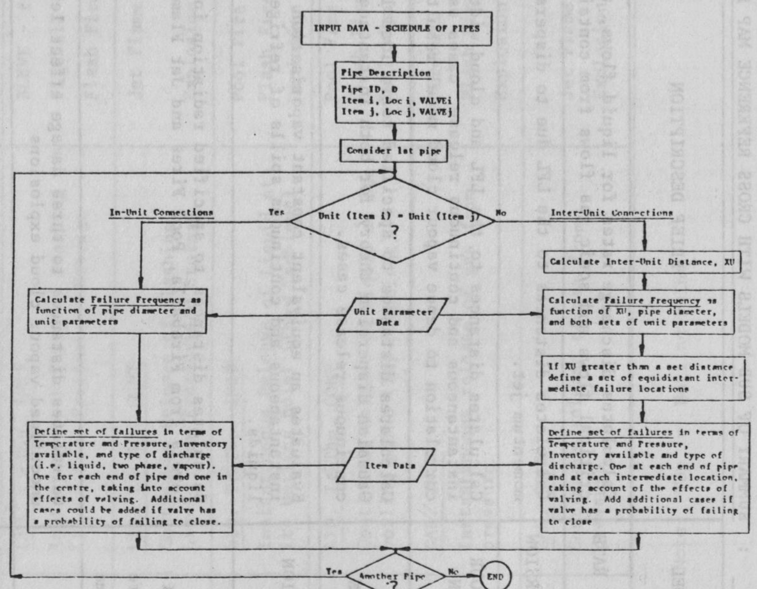

FIGURE 2 : Flowscheme for pipe failure algorithm

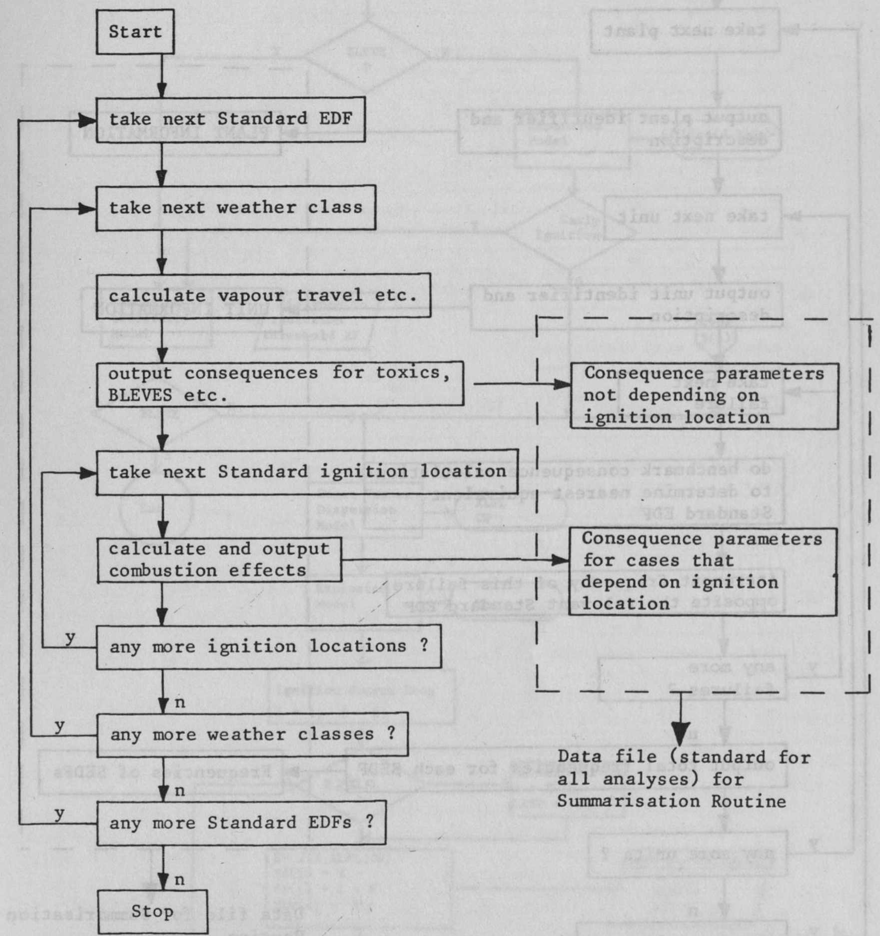

Figure 3 Flowscheme for consequence analysis of standard EDFs

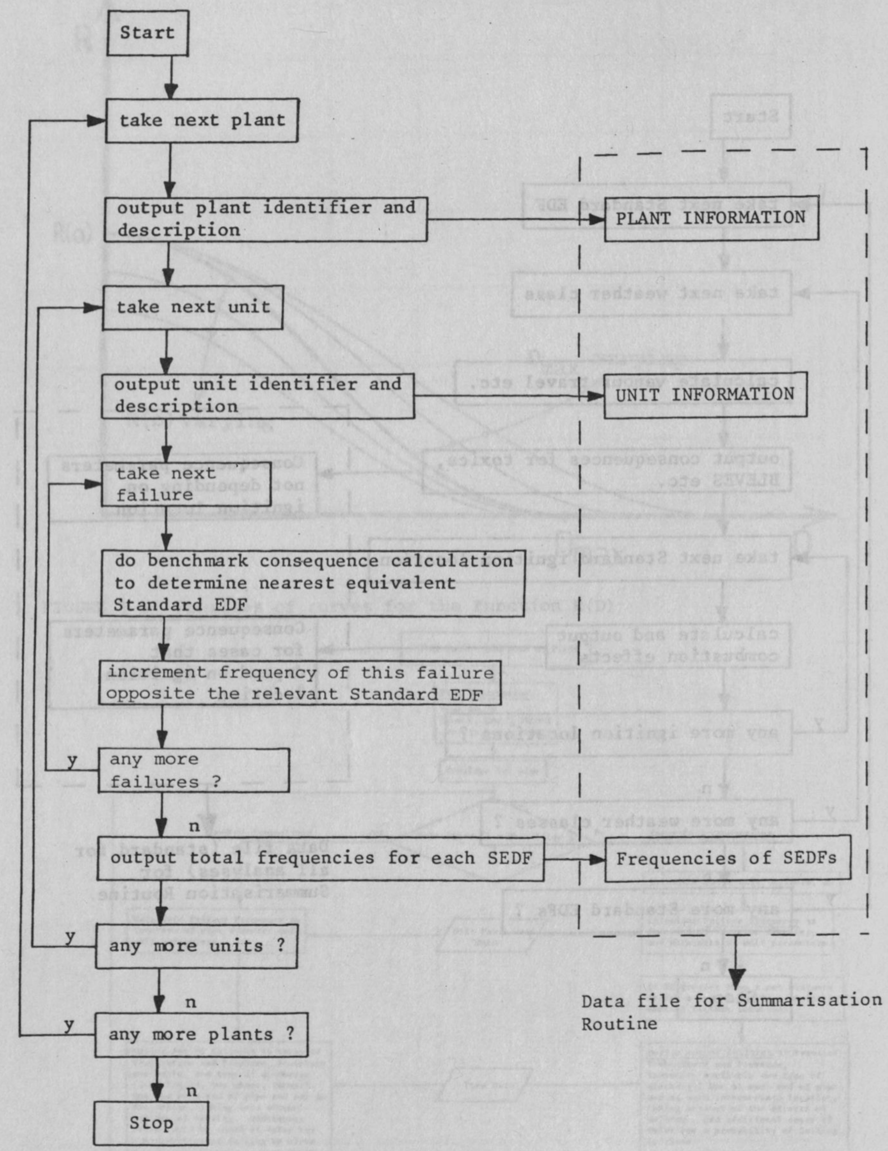

Figure 4 Flowscheme for analysis of actual plants, using standard EDF approach

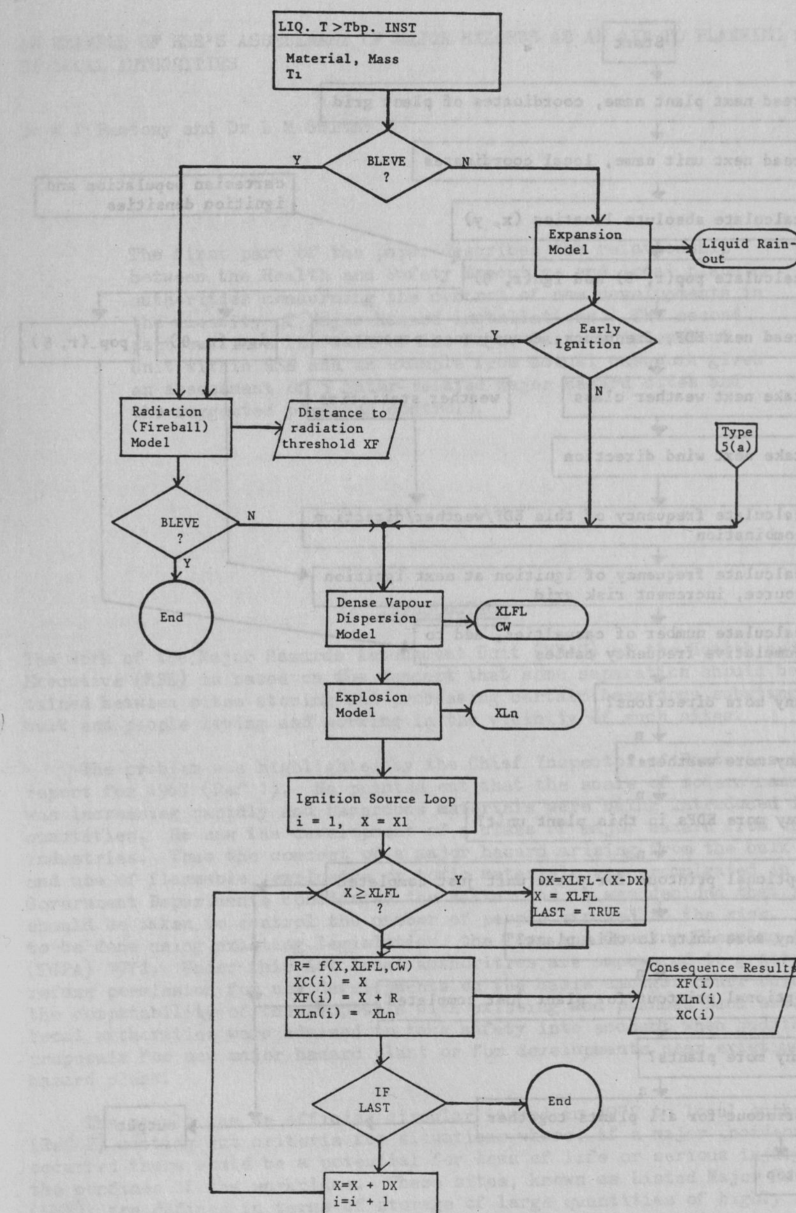

Figure 5 Release Type 1(a) instantaneous release of flammable flashing liquid

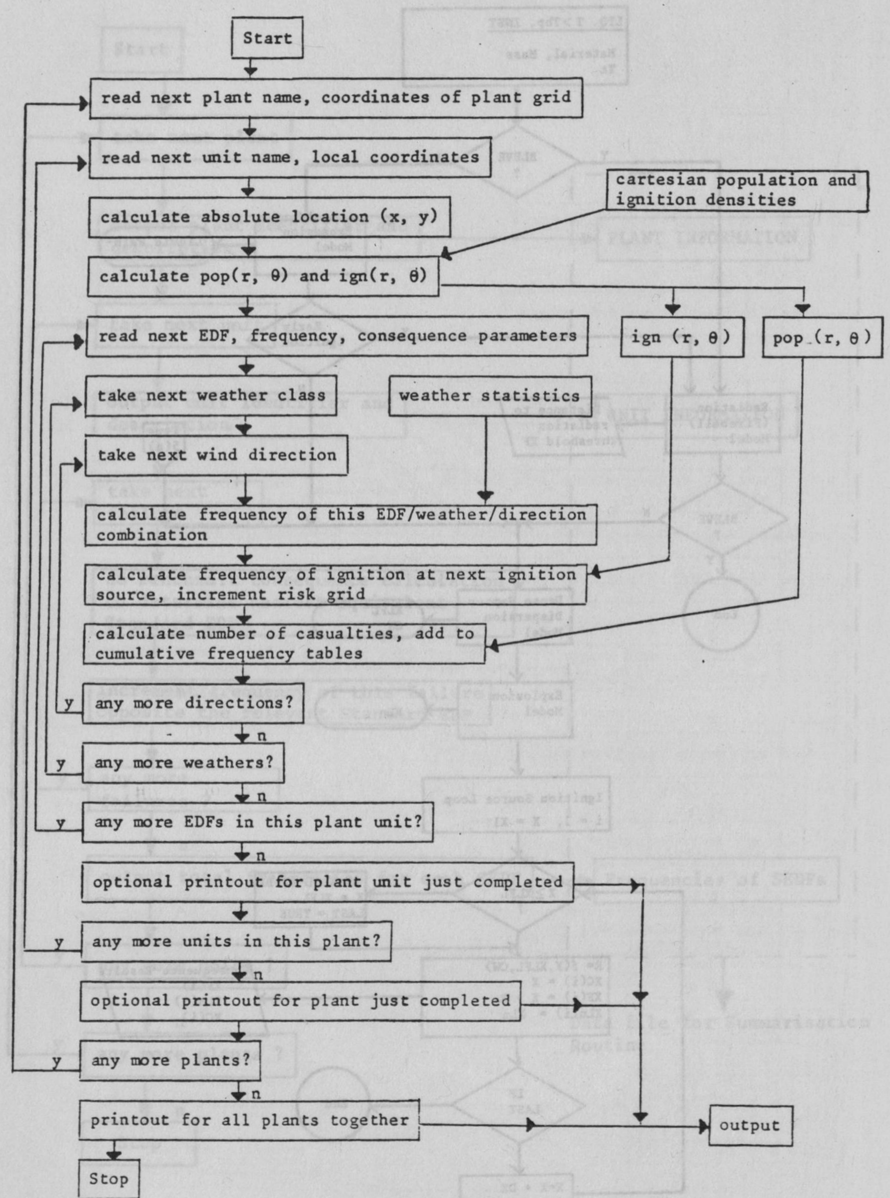

FIGURE 6 : OVERALL FLOWSCHEME FOR SUMMARISATION ROUTINE

AN EXAMPLE OF HSE'S ASSESSMENT OF MAJOR HAZARDS AS AN AID TO PLANNING CONTROL BY LOCAL AUTHORITIES

Dr M F Pantony and Dr L M Smith*

The first part of the paper describes the relationship between the Health and Safety Executive and local planning authorities concerning the control of new developments in the vicinity of Major Hazard installations. The second part outlines the work of the Major Hazards Assessment Unit within HSE and an example from actual casework gives an assessment of 5 inter-related major hazard sites and the suggested planning controls.

INTRODUCTION

The work of the Major Hazards Assessment Unit (MHAU) of the Health and Safety Executive (HSE) is based on the concept that some separation should be maintained between sites storing and processing certain hazardous substances in bulk and people living and working in the vicinity of such sites.

The problem was highlighted by the Chief Inspector of Factories in his report for 1967 (Ref 1). He pointed out that the scale of modern manufacture was increasing rapidly and hazardous materials were being introduced in large quantities. He saw the development of a class of major hazard from these new industries. Thus the concept of a major hazard arising from the bulk storage and use of flammable, explosive or toxic materials was promulgated in the UK. Government Departments considered the matter and it was decided that steps should be taken to control the number of people exposed to the risk. This was to be done using existing legislation, the Town and Country Planning Act (TCPA) 1971. Under this Act local authorities are empowered to grant or refuse permission for new developments on the basis amongst other things of the compatability of the proposals with existing and planned land use. The local authorities were advised to take safety into account when judging proposals for new major hazard plant or for developments near existing major hazard plant.

The result was an official circular issued in 1972 to local authorities (Ref 2) setting out criteria for situations where, if a major incident occurred there would be a potential for loss of life or serious injury outside the confines of the workplace. These sites, known as Listed Major Hazards (LMH), are defined in terms of storage of large quantities of highly flammable

*Major Hazards Assessment Unit, Health and Safety Executive, 25 Chapel Street, London NW1