AUSTRALIAN ATOMIC ENERGY COMMISSION

NUCLEAR REACTOR SAFETY - A REVIEW OF THE RASMUSSEN REPORT (WASH-1400)

by

C.P. GILBERT

ABSTRACT

The US Reactor Safety Study (WASH-1400), often referred to as RSS or the Rasmussen Report, sets out to estimate the risk to which the public is exposed by the possibility of an accident in one of the first 100 nuclear power stations situated in the USA.

This review briefly describes the methods used in RSS, and indicates how the results were obtained. It then discusses the nature of these results, and the comparisons which RSS makes between them and other accident statistics. Finally, the criticisms which have been levelled at RSS are assessed, and the conclusions of a very recent US Nuclear Regulatory Commission review are included in an appendix.

RSS concludes that the risk due to reactor accidents is smaller than that associated with many other aspects of everyday life. While there are shortcomings in RSS, the author has little doubt that the above conclusion is valid.

National Library of Australia card number and ISBN 0 642 59666 2

The following descriptors have been selected from the INIS Thesaurus to describe the subject content of this report for information retrieval purposes. For further details please refer to IAEA-INIS-12 (INIS: Manual for Indexing) and IAEA-INIS-13 (INIS: Thesaurus) published in Vienna by the International Atomic Energy Agency.

HAZARDS; RISK ANALYSIS; SAFETY REPORTS; REACTORS; SAFETY ANALYSIS; HUMAN POPULATIONS; REACTOR SAFETY; REACTOR ACCIDENTS

CONTENTS

(Continued)

CONTENTS (Continued)

1. INTRODUCTION

Since nuclear power first came into use about 20 years ago its public acceptance has been hampered by widespread fears of a nuclear catastrophe, in which a release of radioactive material would kill many thousands of people. Are there substantial grounds for these fears? Just what are the chances of a nuclear accident of some sort?

The Reactor Safety Study [Ref.1], often known as the Rasmussen Report after the director of the study, and usually referred to as RSS, attempts to answer such questions by estimating the average risks of death, injury and property damage to which the public is exposed due to the possibility of an accident in one of the light-water reactor power stations operating or planned to operate at various sites throughout the USA.

A detailed risk analysis of this type had never before been carried out on such a large and complicated plant as a nuclear power station. The leader of the study was Professor Norman C. Rasmussen of the Department of Nuclear Engineering at the Massachusetts Institute of Technology. He directed a group of 60 people including a number of scientists and engineers with extensive academic and industrial experience: he also made use of various consultants, and sub-contracted certain aspects of the study to external bodies. The work took over three years to complete, and cost 4 million dollars. A draft report was issued in 1974 [Ref.2] and criticisms were invited for inclusion in the final report. The latter, which has 11 appendices containing over 2000 pages, was published in October 1975.

RSS concludes that the risks involved in the operation of existing and proposed light-water reactors are very small compared with many other risks accepted by the general public. This information paper describes briefly how the RSS results were obtained, and examines the results and their implications. Finally the main criticisms that have been levelled at RSS are considered, at which stage the present author's views are introduced.

2. DEFINITION OF RISK

RSS defines *risk* in the following way:

RISK = FREQUENCY x CONSEQUENCES

Depending upon the field of investigation the *consequences* may be defined as death, injury or any other undesired event, and the *frequency*

is simply how often the event is likely to occur.

As a specific example consider the prediction of 'road' deaths. In this case:

Risk (expected number of road deaths per year)
= (estimated number of road accidents per year)
x (estimated number of deaths per accident).

One implication of this definition is that the *risk* is the same whether there are ten accidents each killing one person, or one accident killing ten people; this point is discussed in Section 5.2. Since there is, in fact, a range of accident sizes, the total risk to the population, R (expected numbers of deaths per year) can be expressed as the sum:

R = (number of accidents per year killing one person) x one
+ (number of accidents per year killing two people) x two
+ (number of accidents per year killing three people) x three
+ ...

and so on.

As one would expect, in practice the *frequency* of such accidents (number per year) gets smaller as the *consequences* (number of deaths) get larger. The above value for R gives the risk run by society as a whole; if there were 100 000 people exposed to this risk, then the chance that any one of them would die in a car accident in a particular year is R/100 000.

Since there is no reason to suppose that past trends will change significantly, very reliable estimates can be made for future car accidents because of the detailed statistics which have been collected for many years. The risk calculation does not predict who will have an accident, nor where an accident will occur, but it does give a reliable estimate of how many people will be killed in small accidents, how many in large accidents and so on.

When it comes to estimating the risk of being killed by an accident in a nuclear power station, very much the same principle is used.

The method is quite different however, mainly because the risks are very much smaller and the statistics from the past, from which the estimates have to be calculated (known as the *data-base*), are of a very different character from those discussed above. To begin with, there is no record of any accident in a commercial nuclear power station serious enough to kill or injure a member of the public. Further the major

components of these power stations have failed so rarely that they do not provide a reliable guide to the failure rates to be expected in the future. Finally, a nuclear power station is an extremely large and complicated plant, carefully designed to be as safe as possible, and so the analysis of risk is bound to be complicated and involve a very large amount of detail. The steps involved in such an analysis are outlined in the next section.

3. RISK CALCULATIONS FOR REACTOR POWER STATIONS

3.1 General Consideration

A nuclear power station differs from any other power station only in the method of providing steam for the turbo-generators; instead of using a coal- or an oil-fired boiler the water is heated using a nuclear reactor. Since it is the reactor core itself which is the main source of risk a very brief description of such a power station will make the later discussions clearer.

In the diagram of Figure 1 the fuel elements, which constitute the reactor core, heat the water being pumped through the system. (If for any reason this water supply fails there is an *emergency core cooling system* (ECCS) to prevent the core from overheating.) In a *boiling-water reactor* (BWR) the water boils in the core, and the resulting steam is used to drive the generators. In a *pressurised-water reactor* (PWR) no boiling is permitted in the core, but the hot water is pumped through a heat exchanger where it boils the water in a totally separate piping system. This 'secondary' steam is used to drive the generators.

Although RSS treated the two types of reactor quite separately, the results of each are so similar that for the purposes of this brief review only a generalised water reactor need be considered.

While many things can go wrong in a nuclear power station most of them would have only trivial consequences. RSS concerned itself only with those accidents that would affect the public, i.e. involving a release of radioactive material. Broadly this material consists of the fission products within the partially-consumed fuel elements. Exposure to these fission products can cause cancer, genetic defects in future generations and (at high levels) illness, possibly followed by death within a few days.

A serious release of fission products can occur only if

(i) the core melts,

(ii) the high-pressure primary coolant circuit is breached, and

(iii) the containment building leaks,

(not necessarily in that order). Thus the bulk of RSS is devoted to studying the ways in which these three events could occur, the chances of their happening and their effect on the public, despite the fact that not one of them has ever occurred so far in any commercial power station. (In fact, some Russian stations do not have a containment building at all because the likelihood of failures in the earlier 'lines of defence' is considered so remote.)

The chance of an event occurring is quantified by its *probability*. Thus a probability of 1 means that an event is certain to occur, and a probability of 0 means that it is certain NOT to occur. When tossing a coin the probability of a head coming up is 1/2, and that of a tail is also 1/2. Similarly, in rolling dice the probability of a three, say, is 1/6 = 0.1667 for a single throw.

More typically for the work in RSS, a probability of 1/100 = 0.01 means that there is 1 chance in 100 of an event occurring in a given period of time. If this period is 2 years, say, then a probability of 0.01 implies that the event can be expected to occur approximately once in 2/0.01 = 200 years, or, if the value is to be used in the expression for risk as defined earlier, with a frequency of 1/200 = 0.005 event per year.

3.2 Description of the Accident - Event Trees

The first step in the RSS analysis was identification of all possible accident sequences which could lead to core melting. Such a sequence would start with the *initiating event*, i.e. the accident itself, and this might be any one of a range of mishaps such as a burst pipe, a power failure or a very severe load transient. The sequence would continue with a series of events such as the availability (or failure) of electric supplies, the operation (or failure) of the ECCS and so on. Provided there were no serious failures in the so-called *engineered safety features* the accident itself would be contained and there would be no risk to the public despite possibly severe damage to the plant. The analysis made use of *event trees*, which are simply convenient pictorial representations of sequences, including every event which could affect the course of the accident, and making sure that no possible combination of the listed events was missed. The very simple event tree of Figure 2 shows a pipe burst as the initiating event. The

two following events lead to four possible sequences (A, B, C and D). Of these C can be ignored because the ECCS will not work correctly without electric power.

While they vary in size, a single tree rarely leads to more than 30 distinct sequences. Each event (such as the starting of a diesel generator to supply electric power, or the operation of a core spray to cool the core and wash out fission products) could either occur, or fail to occur with a certain probability. If the probability is known for all the individual events in a sequence the overall probability that that sequence will occur can be found. Provided the events are independent, or if they are dependent, that conditional probabilities are used, the individual probabilities can be simply multiplied together to give the sequence probability.

Trees were drawn up for a range of initiating events. Also a separate event tree was prepared for the containment to list all of the ways it could be damaged and leak, and the appropriate part of this tree was included with the various accident sequences described above.

The process of identifying accident sequences requires extreme care, a very high degree of technical competence, and complete familiarity with the plant, its method of use, and the maintenance and test procedures. It appears that these conditions were substantially fulfilled in this study, and while there is no certainty that all possible accidents were recognised it is unlikely that a serious one was overlooked, since the conscious search for accident sequences is an essential aspect of reactor design and has been actively pursued for over 25 years.

3.3 Probability Calculations - Fault Trees

To calculate the probability of an accident sequence it is necessary to know the probabilities of all its component events. In RSS some of these could be found directly from the data-base from previous industrial experience. Thus, although pipes have not often burst, it was possible to use accident statistics to make a reasonable estimate of how often they could be expected to do so in the future. Similarly, failure rates could be established for a whole range of common components such as relays, pumps, valves, motors, etc.

However, the accident sequences leading to the more severe consequences involved the failure of very complex systems, such as the *low pressure water injection system*, which are made up of many diverse

interconnected components. Because of the very high reliability built into such systems, the small number of them in use, and their relatively short operational history, the probability that the whole system will fail could not be found directly from past experience. In such cases, *fault trees* were used. These show the logical relationships between failures of individual components (such as pumps and valves), and the failure of the overall system. The known failure rates from the data-base for each component in the system were used in the fault trees to arrive at a probability that the complete system would fail. The probability of failure for each of the systems was then used in the event trees mentioned earlier.

While the data-base used in RSS was probably the best then available, it was not nearly so extensive or reliable as could have been wished. Thus the selection of the failure probabilities required a good deal of engineering judgement and there is room for honest disagreement on a number of the choices made. As more experience is gained the data-base will gradually become more and more reliable. It should be mentioned that errors on the part of operators and maintenance personnel were included as contributing events on the fault trees, and the probabilities of such errors were also estimated from the data-base.

As for the event trees, many skills and a great deal of care are required for the construction of fault trees, which can be very complex indeed. Some of the large trees had several thousand elements, but the formal tree structure enabled many important properties of the various failure modes to be identified.

3.4 The 'Degraded' Core Analysis

In the next stage of the analysis the effect of each accident sequence upon the reactor core itself was calculated. Thus for a given sequence of pipe breaks, pump failures, etc. the analysis evaluated in some detail the sequence of processes leading to core degradation such as steam generation, zirconium-water reaction, hydrogen burning (or explosion), fuel melting, reactor vessel melt-through, containment pressure and leakage, containment melt-through, and even the possible underground movement of the molten mass formed from the core [Figure 1]. A number of simplifying assumptions were made to reduce the amount of work involved, and the technique of making conservative assumptions which is used throughout the study assumes special significance in the degraded core analysis.

In this method, which is widely employed in accident analysis, where doubt exists as to the precise course of events in a complex situation the worst course is chosen, i.e. that giving either the most serious consequences or the largest probability. One can then be sure that whatever does occur, the result cannot be worse than that calculated. For instance, it is not easy to evaluate the precise degree of core melting that could occur in a given accident. Hence it was assumed that as soon as any part of the core reached melting temperature, the *whole* core would melt. Similarly, if the ECCS did not supply its specified quantity of water it was assumed to have failed completely: no benefit was claimed for the lesser amount of water which it would, in fact, have supplied. If such assumptions are made carefully, they not only lead to conservative (i.e. pessimistic) results, but also in most cases permit a very much simpler calculation to be used.

3.5 The Fission Product Source

If the series of physical states through which the core material will pass is known for a given accident, it is possible to calculate the amount and character of the fission products that will be released from the fuel, the way in which this material will become dispersed throughout the containment, and the rate, height and temperature at which it will eventually escape into the environment (if it does escape). In RSS the effect of water sprays in the containment, filters, etc. was taken into account.

In Section 4.1, a description is given of the way in which the whole range of possible radioactive releases from the containment was represented by a limited number of *release categories*.

3.6 Accident Consequences

The fission product release would emanate from the containment as a *plume* i.e. a cloud of hot gas with solid particles suspended in it. Accident consequences were found by estimating the number of people within a 500 mile radius who would be exposed to this radioactive plume or to any material deposited from it. Crude estimates were made for the exposed population beyond this radius.

The two major considerations were:

- weather conditions: the wind blows the plume along in a particular direction, and tends to break it up, or dilute it; also rain may wash some of the suspended particles out onto the ground.

- population density: the plume may be blown over a densely populated area, irradiating many people, or over unoccupied land or sea. In the former case, the affected population was assumed to be reduced, to some extent, by evacuation.

To account for all these factors six 'composite' sites, typical of geographical locations used for reactors, were chosen such that their weather patterns were representative of the 68 actual sites of existing or proposed reactors. For each composite site, 90 different sets of representative weather conditions were selected and used in the consequence calculations.

To find the population distribution to be used with any given composite site, census data for the populations surrounding the corresponding actual sites were collected, within equal radial sectors. These population densities were ranked and graded into 16 different levels (including one for the highest density found), and this enabled the probability of any one density of population being irradiated by the plume to be found for that particular composite site.

The consideration of all possible combinations of *release category*, composite site, weather conditions and population density for both PWRs and BWRs resulted in over 120 000 separate calculations. For each, a unique set of consequences in terms of death, injury, and the cost of damage was calculated. Associated with each set of consequences was the probability of it occurring, found from the probability of that particular release combined with the probabilities associated with the particular weather conditions and population density. The results are discussed in Section 5.

4. REVIEW OF THE SCHEME OF CALCULATIONS

The above summary of the techniques and calculations used in RSS, of necessity, has omitted certain aspects; the more important of these are discussed below.

4.1 Volume of Work

Although 120 000 separate results seems a very large number, in fact the number was only kept as small as this because great efforts were made to reduce the bulk of the work without compromising its validity. One way in which this was done was the use of the six composite sites to represent the 68 actual sites. Other techniques are now described.

Initially the calculations were performed very simply with very

conservative assumptions (i.e. assumptions which tended to make the risk appear worse, but made it easier to calculate). The crude results obtained were adequate to enable a large number of cases (which did not add significantly to the risk), to be neglected. In these cases either the probability or the consequences were vanishingly small compared with those of other similar accidents. Next, the remaining sequences were considered much more carefully, but at every stage in the subsequent calculations, sequences of very low risk were rejected.,

Using these methods the number of accident sequences requiring further study was reduced to about 1000, but instead of performing detailed consequence calculations on each one they were separated into *release categories*. Nine release categories were chosen for the PWR and five for the BWR, and they were selected to be representative of the whole range of accidents and types of fission product release. Each of the 1000 sequences was then included in a release category having a similar (or worse) release, and within each release category the sequence probabilities described in Section 3.2 were added to give the overall category probability. In each category, a few large probabilities dominated the total, and provided that all these so-called *dominant accident sequences* were finally identified, all other accidents were of secondary importance. There were less than 150 of these dominant sequences, and when 'category smoothing' was introduced — a very conservative technique which allowed for the possibility that a sequence had been allocated to the wrong release category — the number of sequences making a substantial contribution to the total risk was reduced even further. The probabilities of the dominant sequences were then carefully checked, and the consequences of each release category were calculated as described in Section 3.6 to give the major contributions to the total risk.

4.2 Reliability of Data-base

Values for equipment failure rates were obtained from industrial records relating to both nuclear and other situations. Much of this information was necessarily vague and incomplete, and it was frequently unclear to what extent the data were applicable in a given case. Rather than choosing a single failure probability for a given component, a range of values representative of the basic data was used. This range of values for a probability was assumed to have the form of what is known as a log-normal distribution, which may be described approximately

by a median, or 'middle' value, together with a lower boundary which is smaller than 95 per cent of the values in the range, and an upper boundary which is larger than 95 per cent of the values. When the complete distributions were used, the process of combining probabilities became rather complicated, and so for most of the conservative calculations, only the upper 95 per cent boundary was used.

Throughout the analysis, many assumptions had to be made to simplify the calculations, and many choices had to be based mainly on engineering judgement. It is not surprising that a number of these assumptions were challenged by critics of the draft report, and their criticisms led to some important changes in the final report. However, because the assumptions were made in what was judged to be a 'safe' direction, it has been claimed that a good deal of conservatism is built into the RSS results.

It should be remembered that only the dominant accident sequences have much influence on the total risk. Thus if the data for the other low-risk accidents were found to be wrong, even by a large amount, it would hardly affect the total risk at all.

4.3 Non-independent Events

Suppose that it were possible, in the sequences outlined in Figure 2, for the pipe-burst to destroy all the electric power supplies, by forcing high pressure water into the supply cubicles. These two events would no longer be independent, and the original accident would prevent the ECCS from working, i.e. the equipment designed to mitigate the effects of an accident would have been put out of action by the accident itself. In such a case the simple probability calculations outlined above would grossly underestimate the risk because they assume that such events are independent of one another.

Naturally, power stations are very carefully designed to avoid such situations, but interactions of this sort can occur in very subtle and devious ways which are extremely difficult to foresee.

Similarly a single cause may result in the failure of many similar pieces of equipment. Thus, in some cases, three or four identical sets of safety equipment are used so that if one fails the others remain to protect the plant. A multiple failure would result if, for instance, the maintenance technician misunderstood his instructions and adjusted all the replicated equipment incorrectly, with the result that none of it would work when called upon in an emergency.

Situations of these types are not easy to identify, and much skill and care are needed to evaluate and combine the failure probabilities correctly. RSS used several novel techniques to help find where such effects could be significant, but there can be no absolute certainty that these efforts were completely successful.

4.4 Equipment Quality and Design Adequacy

The data used in RSS were based on the assumption that the plant had been competently designed, that all components were suitable for their task and had been subjected to rigorous testing and quality control and that, when in use, the plant was regularly tested and maintained.

It is the function of the US Nuclear Regulatory Commission to ensure that there are no major deficiencies in these respects. Minor deficiencies, such as an occasional faulty valve, pipe leak or maintenance error are accounted for by studies of the type included in RSS.

5. THE RESULTS OF RSS

5.1 Total Risks

RSS calculated the risks associated with 100 nuclear power stations. The reason for doing this was that when RSS was written there were 29 PWR and 33 BWR power stations operating in the USA and many more were under construction or planned. The issue was complicated by the fact that power station designs are changing all the time - the later versions are bigger than the earlier ones, comply with more stringent safety codes, and benefit from the experience gained in the previous stations. The RSS analysis was based on two 'average' stations (one BWR, one PWR), but it was considered that by the time the 100th plant comes into operation (probably in the 1980s) the above two stations will no longer be typical and the whole topic should be re-examined.

For the 100 power stations, the consequences of accidents were calculated in terms of:

- early fatalities (the number of people who die within a short time of the postulated accident);
- early illnesses;
- delayed health effects (including delayed fatalities); and
- property damage (e.g. the cost of crops rendered unfit for consumption, of lost industrial production, and of the eventual decontamination of land and buildings).

Let us look first at the results pertaining to early fatalities: comparable statistics are available for a wide range of other classes of accident, and in general the RSS results for early deaths follow trends similar to those for injuries and damage. Although RSS did not stress the point, it showed that a much larger number of deaths would be caused by cancer between 10 and 40 years after a nuclear accident. Because few other types of accident give large numbers of delayed fatalities, the latter are discussed separately in the next section.

The total risk of early death, caused by all types of reactor accidents, can be summarised by combining the risk associated with each release category, as in Table 1.

TABLE 1

Accident Type	Total Number	Individual Chance per year
Motor vehicle	55 791	1 in 4 000
Falls	17 827	1 in 10 000
Fires and hot substances	7 451	1 in 25 000
Drowning	6 181	1 in 30 000
Firearms	2 309	1 in 100 000
Air travel	1 778	1 in 100 000
Falling objects	1 271	1 in 160 000
Electrocution	1 148	1 in 160 000
Lightning	160	1 in 2 000 000
Tornadoes	91	1 in 2 500 000
Hurricanes	93	1 in 2 500 000
All accidents	111 992	1 in 1 600
Nuclear reactor accidents (100 plants)	-	1 in 5 000 000 000

Table 1, a copy of Table 1-1 from RSS, lists the risks run by an average individual in the USA of being killed by a range of possible accidents.

Thus RSS claims that the risk of early death associated with an accident occurring in any of 100 nuclear power stations is at least a thousand times smaller than that of being killed by lightning, and a

million times smaller than that of dying in a car accident. Note, however, that the value for reactor accidents is only an *estimate* whereas the other values are based on historical facts.

The question to be asked is not whether the estimate for reactor power stations is accurate, because until there is a long history of fatal reactor accidents this can never be established. What has to be considered is: what are the chances of the RSS result being wrong by such a large amount that the risk becomes appreciable? We will return to this point in the next section.

5.2 The Importance of Accident Size

It follows from the definition of risk as the product of frequency and consequences, that a risk of, say, 10 deaths/year, could have the alternative meanings given in Table 2. The community appears to accept a continuous stream of small accidents (on the roads, say) more readily than it does a few very large accidents (aeroplane crashes), even though the total number of people killed by the small accidents may be much larger. Whether this attitude is reasonable or not it must be taken into account. Thus the 'equal-risk' alternatives listed in Table 2 become increasingly less acceptable to the community from top to bottom of the table.

TABLE 2

Average interval between events	Average frequency (events/year)	CCF*	Consequences (deaths/event)	Risk (deaths/ year)	Type of event
36 days	10	11.111	1	10	Frequent, trivial
1 year	1	1.111	10	10	
10 years	0.1	0.111	100	10	
100 years	0.01	0.011	1 000	10	
1000 years	0.001	0.001	10 000	10	Rare, catastrophic

** CCF, the complementary cumulative frequency, is found by adding to the frequency on a given line all frequencies below it in the table. Thus for the 10-year event the CCF is 0.1 + 0.01 + 0.001 = 0.111.*

It follows then that it is not sufficient to examine total risks alone, as in Table 1, but that the risk of accidents of a *range of sizes* must also be examined. One method of displaying the extent to which a total risk is made up of frequent, small events as opposed to rare, large ones is illustrated in Figure 3, using the data of Table 2. For illustration purposes, these values are no longer regarded as five equal alternatives but as five separate components characterising a total risk of 5 x 10 = 50 deaths per year. The five plotted points are joined by a line whose shape illustrates the nature of this total risk - in this case an equal chance of being killed in a small, intermediate or large accident.

Suppose we wish to know how often an event resulting in 200 deaths is likely to occur. Starting at 200 on the horizontal axis leads us, in this hypothetical example, to a frequency of 0.05, or once in 20 years. (Statistics from RSS show that this is roughly the situation for dam failures or earthquakes.) A second way of displaying this information, which has important technical advantages, uses the complementary cumulative frequency (CCF) in Table 2. These values do not differ much from the frequency values, and when plotted they lie within the circles shown in Figure 3, and so the curve joining them is insignificantly higher than the line shown. However, proceeding as before, we can now estimate the average interval between events which result in 200 *or more* deaths; the result is slightly less than 20 years.

The plotted points lie substantially on a straight line in this simple example, but as discussed above the larger accidents are less acceptable to the public even though they are less frequent. To meet this community attitude it is reasonable to expect that in an acceptable situation the risk curve should drop away more rapidly at the right-hand (high casualty) end as shown dotted. On this basis, 200 or more deaths might now be expected to occur in the example only once in about 75 years (which was roughly the situation for US air crashes when the statistics for RSS were compiled).

5.3 Comparison of RSS Results

RSS uses the CCF method of presentation to compare its results with other risks to which the community is already exposed: Figures 4 and 5 are based on data taken from RSS.

Curve A of Figure 4 shows the total risk which can be considered as being a result of everyday human activities, i.e. aircraft crashes,

explosions and dam failures. Notice that although road accidents give rise to far more deaths than all other causes combined (Table 1), they hardly influence Curve A at all because any given road accident very rarely causes 10 or more deaths, the smallest consequence plotted.

Curve B shows the total risk due to natural events, such as hurricanes and earthquakes. Curves A and B are both supported by statistics from actual incidents, but naturally the data for the frequent events are much more reliable than those for the rare events. Curve C is the sum of curves A and B, and thus represents the total risk to the community (in the USA): this curve is re-plotted on Figure 5.

The risk of early death from the potential accidents associated with 100 nuclear power stations as estimated by RSS is given in Figure 5 as curve D, which was drawn through a large number of points calculated for the 14 release categories under a range of weather conditions and a range of population densities (see Section 3.6). Compared with curve C, this risk is many, many times smaller than the total of other risks that are already accepted by the public for both small and large consequences. The components making up curve C represent risks such as those due to fires, air crashes, earthquakes, etc., and, had all these been shown, it would have been clear that curve D is smaller than any one of them taken on its own, with the exception of the risk of being killed by a meteor; this is roughly the same as curve D.

6. DISCUSSION OF RSS RESULTS

6.1 Uncertainty in Results

The conclusions were summarised in RSS as:

"(a) The possible consequences of potential reactor accidents are predicted to be no larger, and in many cases much smaller, than those of non-nuclear accidents. The consequences are predicted to be smaller than people have been led to believe by previous studies which deliberately maximised estimates of these consequences.

"(b) The likelihood of reactor accidents is much smaller than that of many non-nuclear accidents having similar consequences. All non-nuclear accidents examined in this study, including fires, explosions, toxic chemical releases, dam failures, airplane crashes, earthquakes, hurricanes, and tornadoes, are much more likely to occur and can have consequences comparable

to, or larger than, those of nuclear accidents."

Table 1 and Figure 5 appear to support these conclusions completely.

However, since some uncertainty is inevitable about the data used and the assumptions made, it is most unlikely that the RSS results as presented in curve D are exactly correct. To account for this, RSS provided factors of 1/5 and 5 to give the tolerances on probability, and 1/4 and 4 on the consequences, giving the shaded band in Figure 5 which would enclose many of the 120 000 results referred to previously.

While is is not difficult to find a number of shortcomings, the evidence suggests that, taken as a whole, RSS was carried out conscientiously and competently. The point at issue is: what are the chances of RSS being too optimistic by a very large factor? For instance, in Table 1, even if the RSS result were to be increased by a factor of 1000, it would still be the smallest entry in the table, and if it were to be increased a *further* 100 times, the risk would be similar to that for death due to drowning. While the original result could quite possibly be too small by a factor of 10, or even 100, a factor of 1000 let alone 100 000 is not really very likely.

6.2 Delayed Deaths Due to Cancer

In addition to causing the risk of immediate deaths considered above, excessive radiation may cause cancer in some members of an exposed population, after a delay of up to 40 years. While the number of early deaths can be calculated with reasonable certainty, deaths due to these cancers and other possible 'latent' effects are very difficult to estimate because the dose/effect relationship is progressively less well-known as the dose gets smaller.

The usual method of estimating the number of latent cancers induced by exposure to radiation is to use the 'linear hypothesis'. This states that the carcinogenic effect of a radiation dose is in simple proportion to its size, however small the dose, and is independent of the rate at which it is given. Most competent authorities believe that these assumptions considerably overestimate the effects of low-level radiation which may even be zero [Ref.3]. In accidents of the type considered by RSS, more than 90 per cent of the total population dose would be at low or very low individual doses, and, to moderate the degree of conservatism, dose-effectiveness factors were used to reduce the estimates when low doses were considered. When combined with other smaller sources of delayed fatalities these cancers give a total risk of delayed

death illustrated by curve E of Figure 5 (which is similar to the risk of early death due to chlorine releases as reported in RSS). Curve E shows that the number of people who die 10 to 40 years after a nuclear accident would be many times larger than the number of deaths at the time of the accident. In Table 1 an entry for this might make total nuclear deaths comparable with those due to lightning.

In the case of an extremely serious radioactive release (expected once in 10 million years) RSS estimates that there would be 3300 prompt deaths, followed later by a death rate of 1830 per year for 30 years. (For comparison, air travel in the USA actually causes over 1700 deaths per year - Table 1.) There is, however, a further complicating factor; of the ten million people potentially at risk from such a release, an average of about 17 000 die annually from cancer in any case, and the additional contribution of 10 per cent due to the postulated reactor accident would probably be indistinguishable from the statistical fluctuations which normally occur in this total.

Thus the causes of most of these deaths could never be positively identified, and while this does not make them any less serious it does indicate some of the problems involved in assessing long-term accident consequences.

7. A SURVEY OF THE CRITICISMS OF RSS

7.1 Published Commentaries

When the draft RSS was made available in 1974, 61 organisations and 25 individuals submitted comments for consideration before the final report was written. The major documented commentaries were prepared by:

- The US Environmental Protection Agency [Ref. 4];
- The US Atomic Energy Commission [Ref. 5];
- The Union of Concerned Scientists combined with the Sierra Club [Ref. 6];
- The University of California, Los Angeles [Ref. 7];
- A Euratom Working Group at Ispra [Ref. 8];

and, in addition, the American Physical Society made its own independent review of reactor safety [Ref. 9] which contains a number of references to RSS. Staff of the Australian Atomic Energy Commission also made a detailed examination of the draft RSS.

It is fair to say that while there were many criticisms of detailed aspects of RSS, most of them were constructive. Most reviewers regarded

the draft RSS as a major pioneering achievement and supported its methods and its general conclusions even though some had major reservations about certain individual numerical results. The only exception to this general statement was the contribution of the Union of Concerned Scientists, whose major objections are considered below.

Some of the criticisms made of the draft report have been incorporated into the final version of RSS: the major modifications involved the calculation of latent effects and the number of people evacuated from an affected area, and the study of large electrical fires. Although these changes made RSS more comprehensive they gave rise to only minor changes in the numerical results and no changes in the general conclusions.

When this information paper was in the final stages of preparation a copy of a report prepared for the US Nuclear Regulatory Commission became available: this report, which set out to review RSS, is discussed briefly in the appendix.

7.2 Objections by the Union of Concerned Scientists

In addition to their comments on the draft, the Union of Concerned Scientists has since published a review of the final RSS [Ref. 10]. This is an updated version of their earlier critique, and is referred to here as UCS.

UCS contains a number of responsible criticisms, but also has a number of flaws and is not considered the most informed of the published reviews. However, it is certainly the most critical of the reviews, and so by examining its objections to RSS most of the points raised by other reviewers will be covered.

(1) UCS strongly criticises the uses to which RSS has been put and the way in which the results have been interpreted by others to support nuclear power. This claim may be justified, but even if others have misused it, RSS itself is not necessarily at fault, except as suggested in the following criticism.

(2) UCS claims that the RSS main report, and particularly the Executive Summary, gives a seriously misleading impression of both the total consequences of a nuclear accident and the uncertainties involved in the estimates.

To the extent that RSS fails to make a clear statement of the estimated magnitude of the delayed deaths, this criticism is fully justified. The data are available but are not presented with the clarity or force that their importance deserves.

UCS goes on to point out that when results are calculated in the RSS Appendices they are rightly accompanied by qualifying notes listing their limitations and uncertainties, but that when the results are presented in the Executive Summary (which is all that most people will read), the comments limiting their applicability are omitted. This claim also has some truth in it, but it is open to conjecture whether the omissions are in the interests of simplicity, or are intended to deceive (as UCS implies).

(3) RSS is criticised by UCS for omitting certain topics which it was never intended to examine. The two most important topics are sabotage and design adequacy. (RSS assumed in most cases that, provided equipment worked *as designed*, it was adequate to perform its intended task.)

Both are undeniably important matters with a fundamental bearing on risk. However, apart from other difficulties the first could hardly be studied without violating the secrecy of the security measures, and the second would be a monumental task, well beyond the resources of the study group.

While it is important that these matters should be pursued, RSS was hardly the place to do so.

(4) UCS found RSS difficult to follow: it claims that the reasoning is often obscure and because of numerous omissions it is impossible to follow the calculations in detail from stage to stage.

This was certainly true of the draft RSS, and while the final report is better, there is still room for much further improvement.

(5) The validity of much of the data used by RSS was criticised by UCS. All reviewers agree that the basic information on equipment failure rates was not nearly as extensive nor as detailed as would be desired.

One point to be remembered, however, is that RSS set out to find what could reasonably be expected to occur on average, not the worst possible result. In the past, safety analyses traditionally calculated the effects of the worst conceivable catastrophe, and RSS, which used conservative assumptions only when it was forced to do so by lack of data, introduced a long overdue note of realism into such studies by relating the seriousness of an accident to the chances of its occurrence. In a number of cases UCS seems to be

expecting RSS to adopt the older approach when claiming that it used optimistic data which resulted in the calculated risks being much too small. Among the other reviewers there are some who support the UCS view, though to a much lesser degree in most cases, whereas others claim that RSS used pessimistic data.

Unfortunately many choices had to be made on the basis of engineering judgement, and some disagreement was inevitable. Because of this it is considered that most reviewers would agree that the tolerances or error bands suggested by RSS are too small, possibly much too small. Thus it is not unreasonable to suggest that the areas of uncertainty surrounding curves D and E of Figure 5 should be perhaps twice their present width to account for the various uncertainties.

Bearing in mind the suggested range of the error bands one might be forgiven for asking what is the justification for a long and costly analysis of the RSS type. Probably the most important point is that it directs attention to the areas where reactor systems are least safe and can most benefit from further work, e.g. it highlights the dominant accident sequences. (In fact, one dominant sequence has been eliminated from later reactor designs.) Also the RSS techniques permit studies to be carried out in which the effect of individual factors is examined. Thus where an uncertainty exists in the data, the results can be recalculated for a range of values of the particular number to see how big it must be before it becomes important. For example, it is not possible to *demonstrate* that the ECCS would be effective in the case of a large loss-of-coolant accident (LOCA) even if all the equipment worked. However, RSS was able to show that the ECCS would have to be unsuccessful once in every ten attempts for this to affect the overall risk significantly. By implication then, if the ECCS *never* worked during a large LOCA, the risk would be increased by a factor less than 10, which would still be within the tolerance band shown on Figure 5. Such understanding can be very valuable for future reactor design.

Also, as time goes on and more reliable data are accumulated, the error bands associated with the predictions made by this type of analysis will become progressively narrower.

(6) UCS claims that the event tree/fault tree techniques are quite

unsuited for quantitative analysis, and that better methods have been developed in the aerospace industry in the past 15 years. USC cites cases in which these techniques gave unrealistic failure rates, and implies that because the methods have sometimes given wrong answers they can never give correct ones. No references to the 'better methods' were provided by UCS, and no evidence of such methods has become apparent among the growing literature devoted to accident and reliability analysis.

While other critics have attacked the validity of some of the absolute values of the RSS numerical results, there has been no substantial support for the UCS view, even when comment was sought by the USNRC from the National Aeronautics and Space Administration, who could be expected to be familiar with the latest aerospace techniques.

The main basis for the UCS criticisms is that event and fault trees rely on the skill of the user to obtain a correct result; i.e. that the accident sequences and the logical relationships have to be selected by fallible human beings. This is true, but for UCS to imply that other methods are free from corresponding objections is quite unjustified. UCS also claimed that there is no certainty that RSS identified *all* the possible accident sequences. There can never be absolute certainty, but although RSS has been subjected to widespread expert criticisms, no significant omissions in the trees have been suggested.

Contrary to the UCS view that the RSS techniques are outdated, they are coming into wider use in the aircraft and chemical industries. For instance similar methods were used to investigate the risk associated with a petro-chemical complex on Canvey Island [Ref. 11] and it has been indicated by the UK Health and Safety Commission that similar analyses will be expected as a matter of course when future installations are considered.

(7) UCS cites the Browns Ferry fire as evidence that the RSS results cannot be relied upon.

Briefly, during the commissioning of the third BWR at Browns Ferry power station a fire was started in a cable duct to the station control room and burnt for a number of hours. For Reactor No. 1 the fire destroyed the control connections to the ECCS and those to the *core isolation cooling system*, which were both in the same duct (a possible breach of the design codes). However,

Reactors Nos. 1 and 2 were safely shut down from full power, and depressurised.

Many people regard this incident as supporting rather than invalidating the general philosophy of safety design. Although the effects of the fire would have prevented much of the safety equipment from working had it been required, they did not result in a single injury, nor in any damage to the reactor fuel, because of the existence of multiple lines of safety defence. And this degree of protection was achieved despite the fact that (in hindsight) it can be seen that several errors of judgement were made by the station personnel during their attempts to extinguish the fire.

The final version of RSS calculates that if there were three hundred fires of the Browns Ferry type, in *one* of them some core melting could be expected. UCS claims that one out of thirty fires would cause core melting, and that other factors ought to have been considered.

Even if the lower figure is correct, the chances of another fire under circumstances similar to those at Browns Ferry are small: and even if some core melting did occur, a significant quantity of toxic material could be released only if both the primary circuit and the containment building were seriously damaged.

7.3 Summary

UCS contains many valid points and a number of comments and recommendations which deserve serious consideration. It is a pity that attention should be diverted from these aspects of the review by a number of more extravagant criticisms and exaggerated claims, and that some of the descriptive material was written in an emotive manner unsuited to a basically technical document.

However, the following quotations put the UCS authors' considered opinion of the risk of nuclear power into perspective: first, from p.93

> "... the risks of death and illness from a reactor accident in a country populated by 100 to 1000 reactors may be significant relative to other major classes of accidents";

and from p.126

> "We conclude that the expected average fatal impact of a nuclear plant could be within, and under pessimistic upper bound assumptions could exceed, the range of impact of a comparable coal plant".

UCS qualifies the latter statement with the distinction that the coal plant deaths result from normal operation, whereas those from the nuclear power station would be expected to result from a small number of accidents. One might add the comment that the former are certain to occur - they have been occurring for years - whereas the latter are based upon estimates, and have not occurred and may not occur at all. On p.136 of UCS appears the statement,

> "While we cannot exclude the possibility that the US *could* escape such an event in the foreseeable future ... we are far from believing that it *will*".

Bearing in mind the highly critical nature of UCS, the above quotations hardly constitute a serious indictment of the safety of nuclear power.

8. CONCLUSIONS

Few people who are qualified to make an informed judgement would deny that RSS is a major landmark in safety analysis, and will form a starting point for many future risk assessments in both the nuclear field and in others. The many detailed shortcomings and uncertainties which have been identified will gradually be reduced as more experience is obtained and as the data-base improves, and this will lead eventually to greater confidence in the numerical results obtained by the various methods outlined.

However, the question of whether the RSS results are correct has still not been answered: nor can it be. It must be realised that the type of nuclear power station considered will never be used in sufficient numbers for long enough to show if the RSS numerical results are correct. On the other hand, these estimates can very easily be proved wrong if casualties or damage arise at rates higher than those predicted. So far this has not occurred, and the very small amount of experience that exists is compatible with the RSS results.

For it to be comparable with other risks, the nuclear risk would have to be about 100 times worse than is estimated by RSS. It is very unlikely that the estimates are wrong by such a large factor, and in the author's view the major RSS conclusions may be accepted as being essentially valid, i.e. the risk due to 100 nuclear power stations is less than other risks that the public in the USA now accepts, even for very large, rare, catastrophes.

However, one must always keep in mind the nature of statistical results such as those of RSS: no matter how small the probability of an event may be, there is no guarantee that it will not happen. A serious reactor accident can occur at any time. All that can be said is that this is most unlikely, much more unlikely than a dam bursting with similarly catastrophic results, or than a large aircraft crashing on a crowded sportsground and causing many thousands of deaths.

9. REFERENCES

[1] Reactor Safety Study: An Assessment of Accident Risks in US Commercial Nuclear Power Plant. US NRC WASH-1400 (NUREG-75/014), October 1975.

[2] Draft Reactor Safety Study: An Assessment of Accident Risks in US Commercial Nuclear Power Plants. USAEC WASH-1400, August 1974.

[3] Watson, G.M. [1975] - The Effects of Ionizing Radiation on Man. AAEC/IP1 also At. Energy Aust., 18 (4) 1-11.

[4] Comments by the Environmental Protection Agency on RSS. EPA Washington, DC, November 1974.

[5] Review of RSS: Comments by the AEC Regulatory Staff. USAEC Washington DC, November 1974.

[6] Preliminary Review of the AEC Reactor Safety Study: Joint Review Committee, Sierra Club - Union of Concerned Scientists, San Francisco-Cambridge, December 1974.

[7] What Did WASH-1400 Prove? UCLA Course 239D, E.N. Cranier, December 1974.

[8] Comments on the RSS WASH-1400, Euratom JRC Working Group, Ispra, January 1975.

[9] Report to the American Physical Society by the Study Group on Light-water Reactor Safety. Reviews of Modern Physics, Vol. 47, Supplement No.1. Summer 1975.

[10] The Risks of Nuclear Power Reactors: A Review of the NRC Reactor Safety Study WASH-1400; Union of Concerned Scientists, Massachusetts, August 1977.

[11] Canvey: An Investigation of Potential Hazards from Operations in the Canvey Island/Thurrock area: HMSO 1978.

[12] Risk Assessment Review Group Report to the US Nuclear Regulatory Commission: Ad Hoc Review Group. NUREG/CR-0400, September 1978.

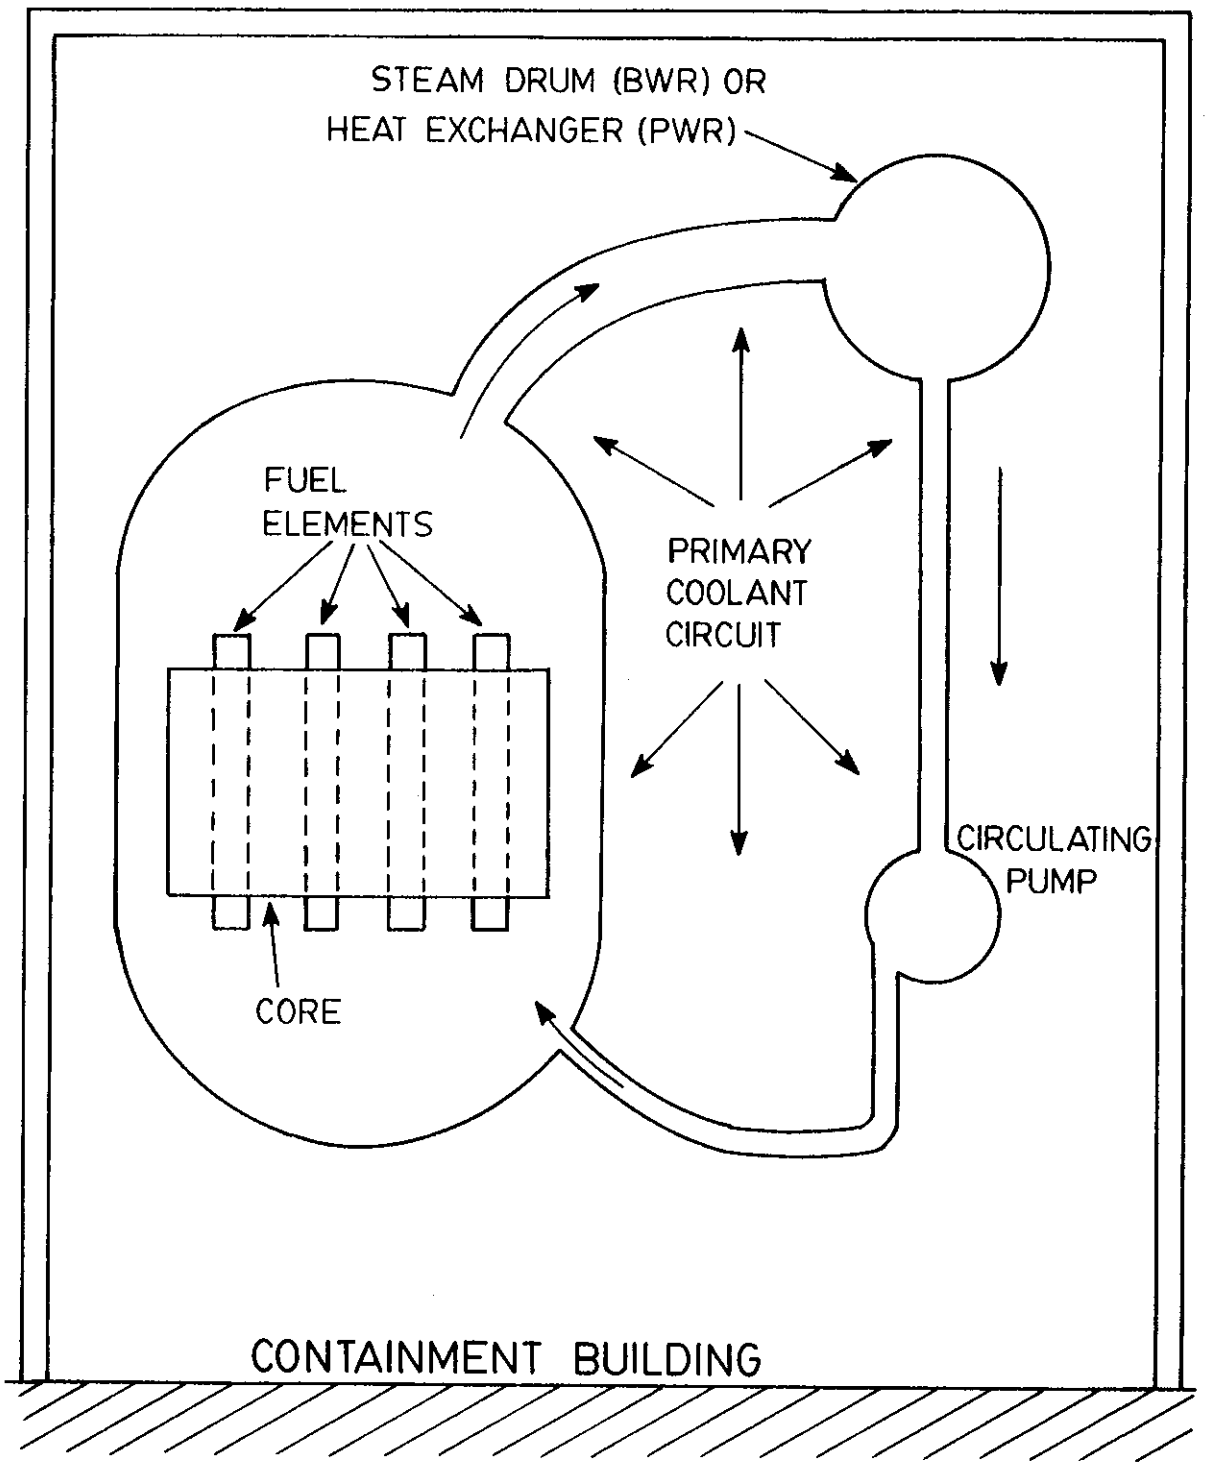

FIGURE 1. THE WATER REACTOR PRINCIPLE, SHOWING THE FUEL ELEMENTS MOUNTED IN THE CORE, COOLED BY THE HIGH PRESSURE WATER (AND POSSIBLY STEAM) IN THE PRIMARY COOLANT CIRCUIT, TOTALLY SURROUNDED BY THE CONTAINMENT BUILDING

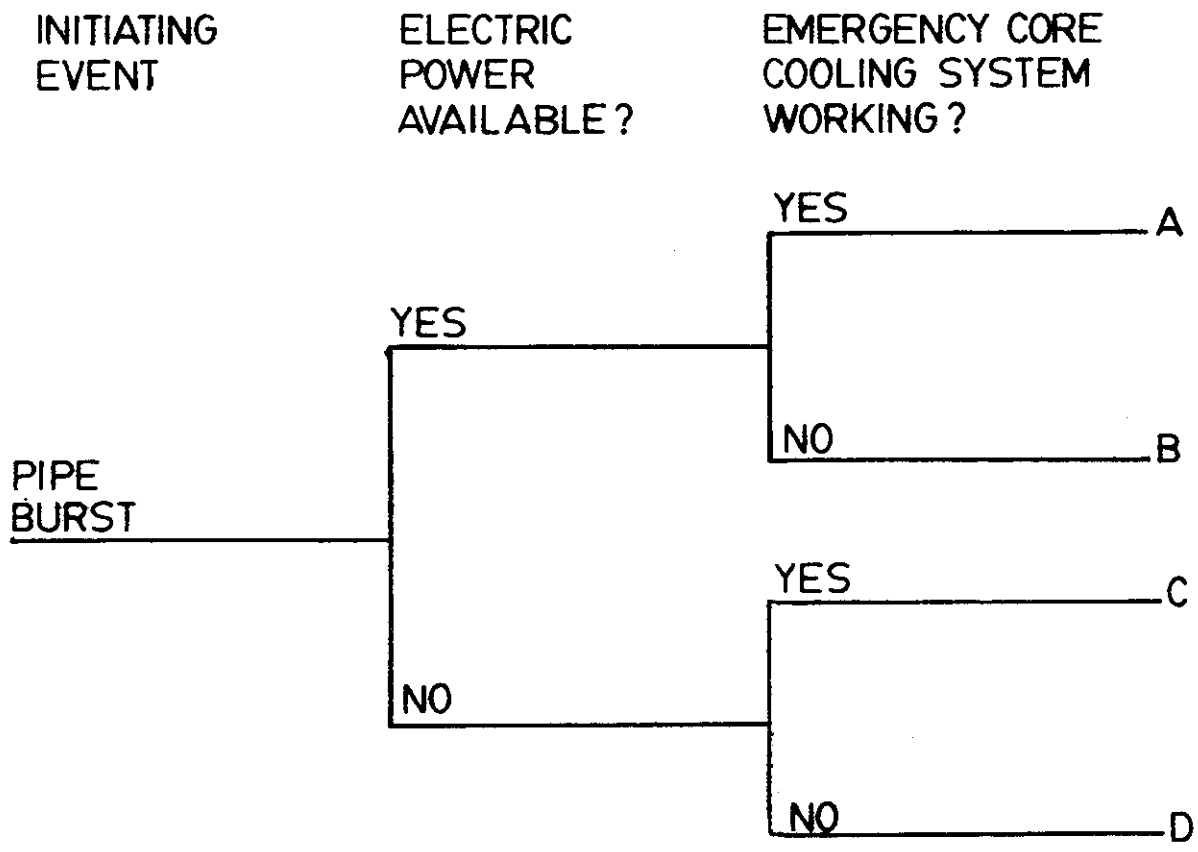

FIGURE 2. A SIMPLE EVENT TREE LEADING TO FOUR ACCIDENT SEQUENCES

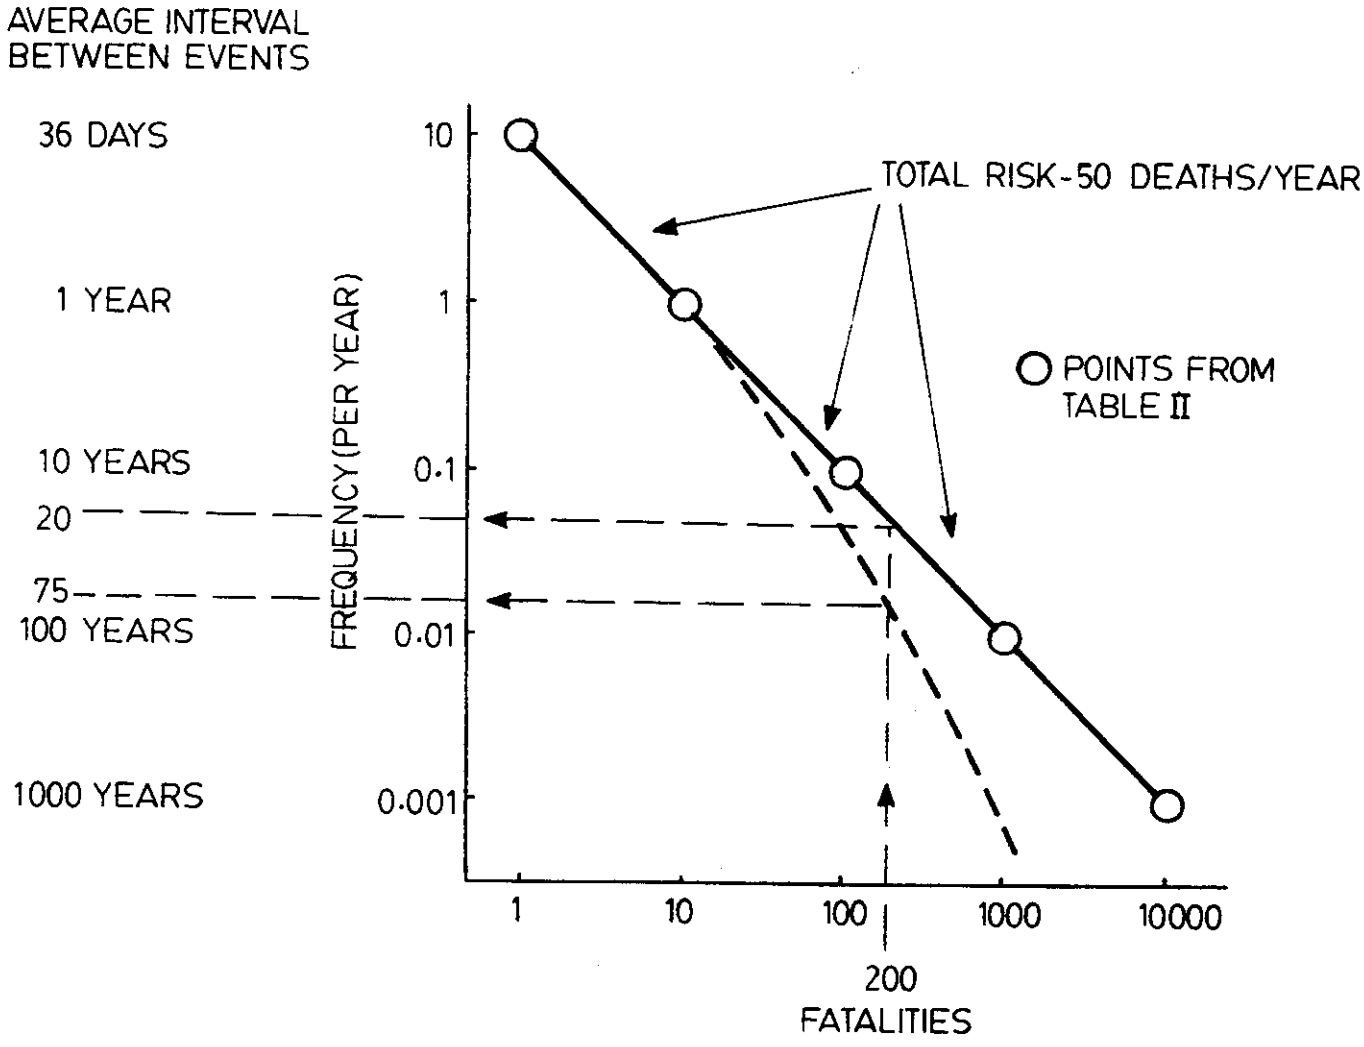

FIGURE 3. FREQUENCY-CONSEQUENCE CURVE FOR THE EXAMPLE OF TABLE 2 (FULL LINE)

For the purposes of this illustration the items in the table are used as the components of one single risk, not five equal alternatives. The dotted line is of a more acceptable shape.

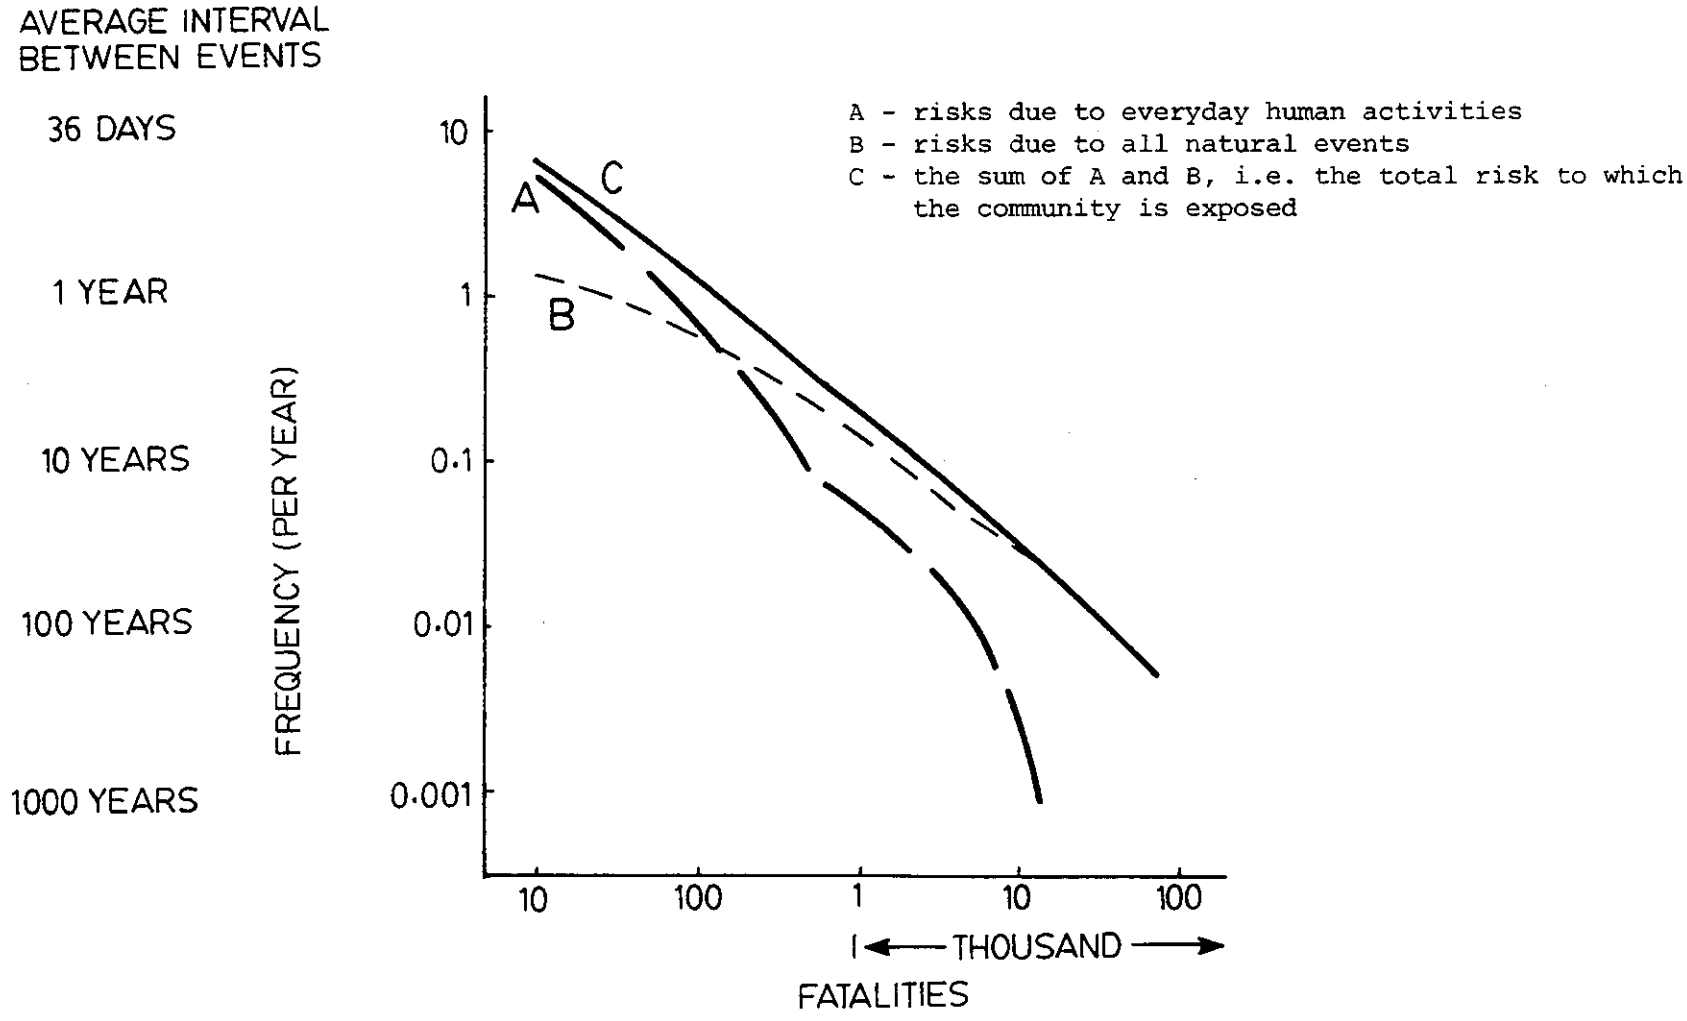

FIGURE 4. FREQUENCY-CONSEQUENCE CURVES FOR DEATH BY ACCIDENT IN THE USA

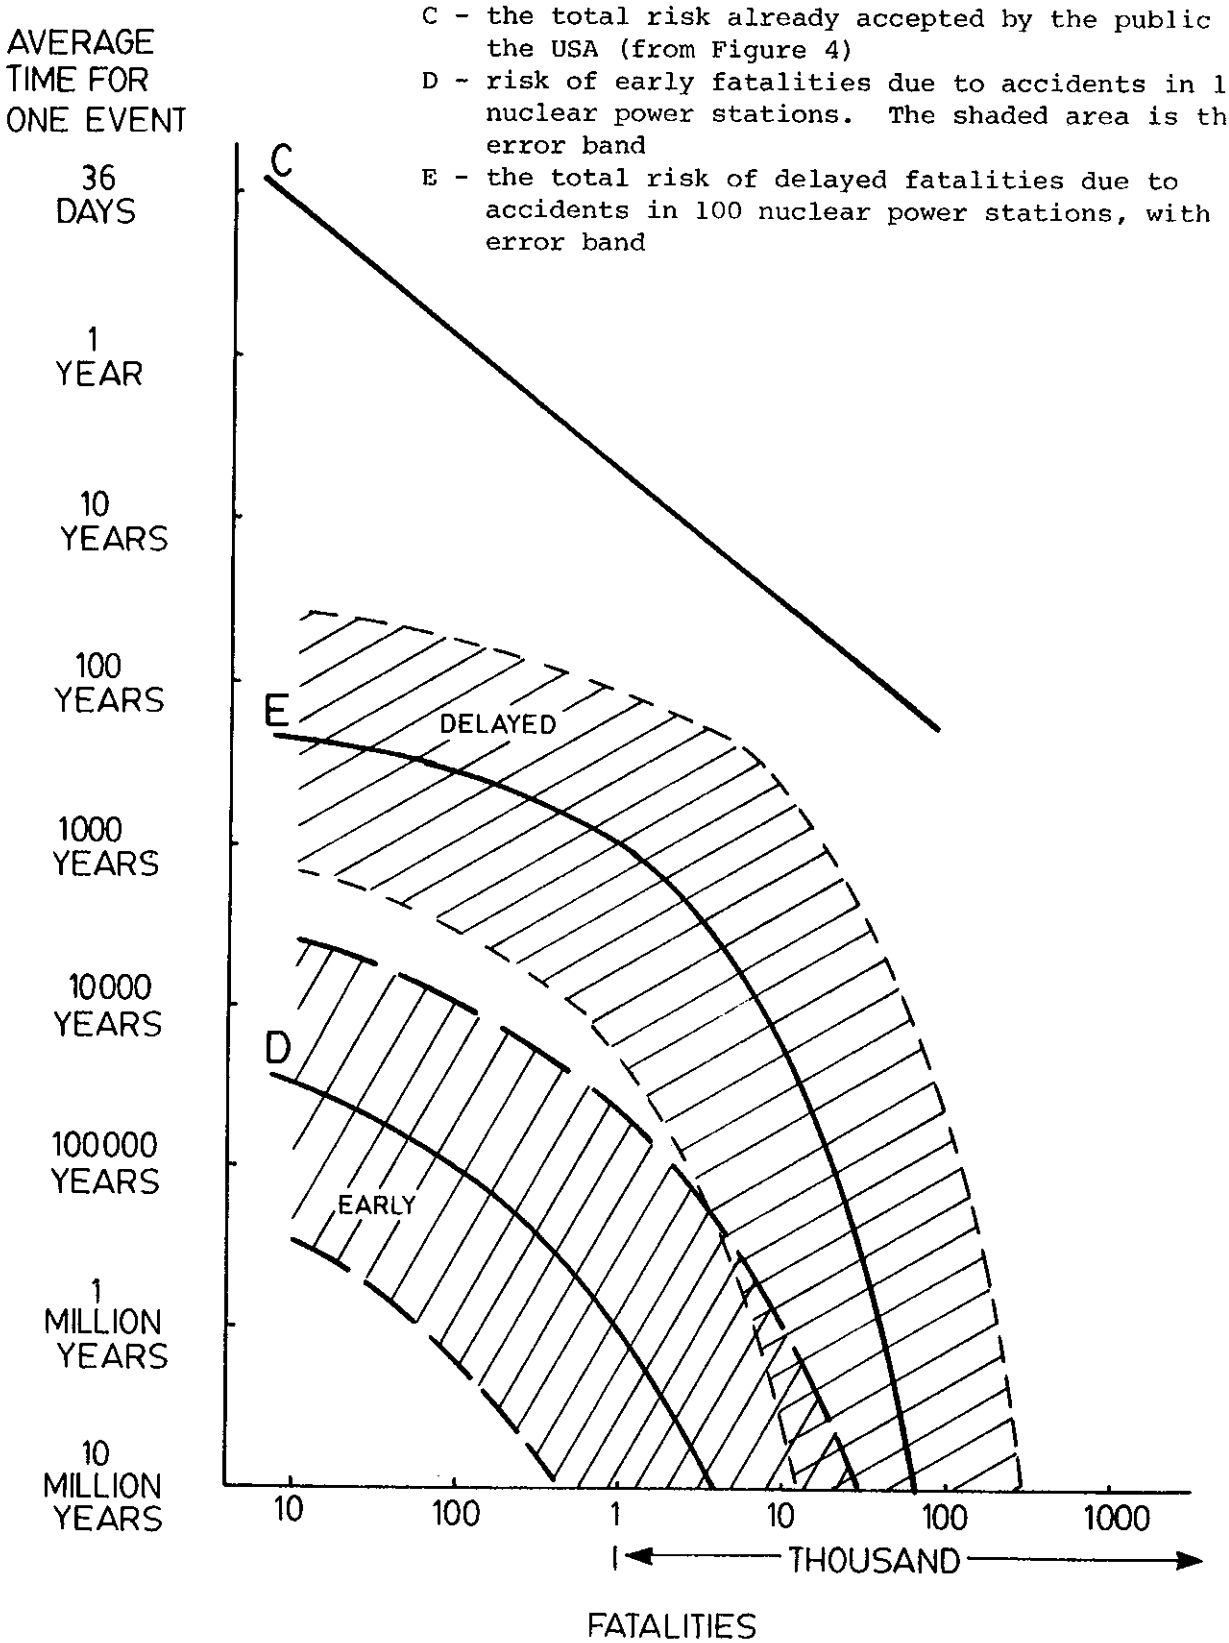

FIGURE 5. RSS RESULTS, EXPRESSED AS FREQUENCY-CONSEQUENCE CURVES

APPENDIX A

RISK ASSESSMENT REVIEW GROUP REPORT TO THE US NUCLEAR REGULATORY COMMISSION (NUREG/CR-0400)

This report [Ref. 12] became available too late to be considered by the present review. The report is by far the most authoritative and up-to-date study of RSS yet to appear, and considered all previous criticisms. The Summary is reproduced below.

Compared with the AAEC investigation this review group had more effort available and a much wider range of expertise, as well as direct access to both the authors and the critics of RSS, and to the manufacturers and operators of nuclear power stations. It is gratifying to find that the report is in substantial agreement with the opinions expressed in the present review.

SUMMARY

The Risk Assessment Review Group was organized by the U.S. Nuclear Regulatory Commission on July 1, 1977, with four elements to its charter:

(1) Clarify the achievements and limitations of WASH-1400, the "Rasmussen Report."*

(2) Assess the peer comments thereon, and responses to those comments.

(3) Study the present state of such risk assessment methodology.

(4) Recommend to the Commission how (and whether) such methodology can be used in the regulatory and licensing process.

The group was formed to represent a wide spectrum of views about nuclear safety, though each member was chosen for his technical expertise. We have profited from a year of study and testimony, and wish to acknowledge the outstanding cooperation we have received from the staff of the Nuclear Regulatory Commission, the nuclear industry, and concerned scientists and citizens.

We find that WASH-1400 was a conscientious and honest effort to apply the methods of fault-tree/event-tree analysis to an extremely complex system, a nuclear reactor, in order to determine the overall probability and consequences of an accident. We have reviewed the methodology, the data base, the statistical procedures, and the results.

*U.S. Nuclear Regulatory Commission, Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants, WASH-1400 (NUREG-75/014), October 1975. Available from National Technical Information Service, Springfield, VA 22161.

We have found a number of sources of both conservatism and nonconservatism in the probability calculations in WASH-1400, which are very difficult to balance. Among the former are inability to quantify human adaptability during the course of an accident, and a pervasive regulatory influence in the choice of uncertain parameters, while among the latter are nagging issues about completeness, and an inadequate treatment of common cause failure. We are unable to define whether the overall probability of a core melt given in WASH-1400 is high or low, but we are certain that the error bands are understated. We cannot say by how much. Reasons for this include an inadequate data base, a poor statistical treatment, an inconsistent propagation of uncertainties throughout the calculation, etc.

Also, both the dispersion model for radioactive material and the biological effects model should be improved and updated before they are applied in the regulatory and licensing process.

We do find that the methodology, which was an important advance over earlier methodologies applied to reactor risks, is sound, and should be developed and used more widely under circumstances in which there is an adequate data base or sufficient technical expertise to insert credible subjective probabilities into the calculations. Even when only bounds for certain parameters can be obtained, the method is still useful if the results are properly stated. Proper application of the methodology can therefore provide a tool for the NRC to make the licensing and regulatory process more rational, in more properly matching its resources (research, quality assurance, inspection, licensing regulations) to the risks provided by the proper application of the methodology. NRC has moved somewhat in this direction, and we recommend a faster pace.

Among our other findings are the well-known one that WASH-1400 is inscrutable, and that it is very difficult to follow the detailed thread of any calculation through the report. This has made peer review very difficult, yet peer review is the best method of assuring the technical credibility of such a complex undertaking. In particular, we find that the Executive Summary is a poor description of the contents of the report, should not be portrayed as such, and has lent itself to misuse in the discussion of reactor risks.

In summary we find that the fault-tree/event-tree methodology is sound,* and both can and should be more widely used by NRC. The implementation of this methodology in WASH-1400 was a pioneering step, but leaves much to be desired.

---

*One of us (F.v.H) is doubtful that the methodology can be implemented so as to give a high level of confidence that the probability of core melt is well below the limit set by experience.

GLOSSARY

The following definitions will assist the reader not familiar with some of the engineering and other terms used in this paper.

boiling water reactor (BWR) A reactor in which the cooling water is permitted to boil in the core, the resulting steam being used to drive the turbo-generators (Figure 1).

data-base An extensive collection of failure-rate data for a range of components such as valves, pumps, relays, etc. These rates were derived from failures reported by industry (including the nuclear industry) over a number of years.

dominant accident sequences Those accident sequences contributing the larger probabilities within a given release category (q.v.): i.e. the accidents responsible for most of the total risk.

emergency core cooling systems (ECCS) A group of water supply systems working at a range of pressures which are designed to cool the reactor core if water is lost from the primary circuit, for any conceivable leak (apart from a pressure vessel rupture).

engineered safety features Equipment which is intended solely to mitigate the consequences of a range of postulated accidents, and which plays no part in the normal reactor operation. One example is the ECCS and another is the containment water sprays intended to wash out fission products.

event trees A simple graphical method of relating a number of events so that no possible combination of them is omitted from consideration (Figure 2). The tree is quantified by evaluating the probability of each event occurring or failing to occur, depending upon prior events in the tree.

fault tree A graphical method of displaying the logical relationships between the failure of the components of a system and the failure of the whole system.

initiating event The 'accident' itself, such as a burst pipe, which initiates a series of accident sequences.

low pressure water injection system The component of the ECCS responsible for supplying a large quantity of water at low pressure.

pressurised water reactor (PWR) A reactor in which no boiling is permitted in the core: a heat exchanger provides the steam for the turbo-generators within a separate piping system (Figure 1).

release category Considering all possible accident sequences and all possible toxic materials, a very wide range of releases is obtained. To reduce this complexity a limited number of release categories is defined so that each sequence can be allocated to a category having a release magnitude which is not smaller than the actual one for that sequence. There are nine categories for the PWR and five for the BWR.