CONF-800130--1

STATISTICAL EVALUATION OF DESIGN-ERROR RELATED ACCIDENTS

by

K.O. Ott, Angewandte Systemanalyse, and J.F. Marchaterre

Prepared for

Workshop

on

Organizational and Procedural Aspects

of

Nuclear Reactor Accidents

Laxenburg, Austria

January 28-31, 1980

UofC-AUA-USDOE

ARGONNE NATIONAL LABORATORY, ARGONNE, ILLINOIS

Operated under Contract W-31-109-Eng-38 for the

U. S. DEPARTMENT OF ENERGY

Paper to be presented at the
Workshop on the Organizational and Procedural
Aspects of Management of Nuclear Reactor Accidents
International Institute for Applied Systems Analysis
Laxenburg, Austria
January 28-31, 1980

STATISTICAL EVALUATION OF DESIGN-ERROR RELATED ACCIDENTS

K.O. Ott
School of Nuclear Engineering
Purdue University
and
Angewandte Systemanalyse (ASA)
Arbeitsgruppe in der Arbeitsgemeinschaft
der Grossforschungseinrichtungen (AGF)
and
J.F. Marchaterre
Reactor Analysis and Safety Division
Argonne National Laboratory

Abstract

In a recently published paper (Campbell and Ott, 1979), a general methodology was proposed for the statistical evaluation of design-error related accidents. The evaluation aims at an estimate of the combined "residual" frequency of yet unknown types of accidents "lurking" in a certain technological system. Here, the original methodology is extended, as to apply to a variety of systems that evolves during the development of large-scale technologies. A special categorization of incidents and accidents is introduced to define the events that should be jointly analysed. The resulting formalism is applied to the development of the nuclear power reactor technology, considering serious accidents that involve in the accident-progression a particular design inadequacy.

I. Introduction

The development of large-scale technological systems normally proceeds through three phases. Ideally, proceeding into the next phase would be contingent on the successful completion of the previous phase, practically however, neighboring phases overlap to a certain extent.

The first phase consists of establishing the scientific feasibility and of exploring the prospects of an economically viable technical feasibility, the Exploratory and Experimental Phase. The information derived during this phase shows the general benefits and possible risks of the new technology. The end of this phase is reached, when the basic conceptual design decisions are made and the corresponding economic and risk prospects are evaluated. Achieving these goals normally requires the construction and operation of several small and larger-scale experimental facilities.

If this information is satisfactory, the second phase can be started, during which full-scale models are actually being developed, the Main Developmental Phase. As part of this development, detailed information about the design, its economics and risk characteristics is obtained. Most of this information could not have been obtained during the first phase, since many specifications were not yet developed or defined at that time. Achieving the goals of this phase normally requires the construction and operation of several prototype and demonstration facilities.

After the developmental phase is successfully completed, after the need for deployment is established and after the health and accident risk has been reduced to a sufficiently small magnitude, the new technology can be practically introduced. The experiences during this phase normally lead to additional refinements and thus to a maturing of that technology (Maturation Phase).

Throughout these developmental phases, the engineers learn -- through r&d, but also through (often unexpected) failures, incidents and accidents -- how to accomplish the task at hand more safely, reliably and economically, and with less environmental impact. After the development of a large-scale technology has been essentially completed, there still exists the possibility of accidents, and this series of IIASA Workshops deals with the respective consequences. It should be noted also, that large-scale technologies are always subject to changes, either through technological break-throughs or due to changes in boundary conditions, e.g., resource availability or effluent restrictions; the development towards maturation is normally inhibited by all these changes.

After all the developed safety measures have been applied there still remains a risk, the "residual risk," to which the populace is subjected in the large-scale application of a particular technology, e.g., nuclear power reactors. It is necessary to quantitatively evaluate the residual risk, not only to estimate the residual risk to the populace, but even more, to improve the safety of existing plants.

The risk assessment approach, as developed by Rasmussen (RSS 1975), is based upon a priori systematic search for accident initiating pos-

sibilities, for which -- partly through the fault-tree technique -- expected occurrence frequencies are established. The main characteristics of these accidents are, that they have a priori identified failure modes, although the identification may not be specific in details. We call them here "IFM-accidents," i.e., accidents with "identified failure modes."

However, during the development of large-scale technological systems as well as during their application accidents of a kind which have not been conjectured a priori also occur. Mostly, they are due to design inadequacies, or design errors, that have gone unnoticed and become evident through incidents or accidents. The development of nuclear power is no exception to that rule. The main characteristics of this second class of accidents is that they have not been identified a priori; we call them here "UFM-accidents," i.e., accidents with a priori" unidentified failure modes."

All modern large-scale technologies evolved out of a rigorous design process which includes the development and application of appropriate standards, detailed quality assurance programs, and careful internal and independent reviews. This thorough and comprehensive approach has been devised in order to avoid design deficiencies as far as humanly possible. Still, there appears to be a residue of deficiencies that slips through this careful process, and some of them may lead to UFM-accidents. These UFM-accidents then are not included in an explicit manner in the residual risk as evaluated by the Rasmussen-type approach.

A general methodology to evaluate UFM-accidents has been recently developed (Campbell and Ott, 1979). Since UFM-accidents are normally related to "human errors (HE)," e.g., in the form of design inadequacies,

they were called HE-events by Campbell and Ott (1979). The evaluation procedure aims at estimating a statistical upper limit of the combined residual frequency (L) of yet unknown UFM-accident possibilities still "lurking" in a certain technological system. This residual frequency, L, is apparently the most important quantity of interest for the characterization of these accidents. The residual frequency, L, can be estimated despite the fact the UFM-accident possibilities can neither be described a priori in any specific way, nor can one get any information about their possible number.

As the methodology of Campbell and Ott (1979) addresses only the residual frequency, no information on the contribution to the risk can be obtained, since accident consequences are not considered explicitly. Here, the question of the risk-contribution of UFM-accidents is addressed. Furthermore, the differences of various plants as they appear in the course of the development of a large-scale technology are introduced explicitly. This then allows the application of the evaluation procedure to the combined historical record of the existing variety of system.

These effects have been incorporated here into the previous methodology, which leads to generalized description of UFM-accident frequencies. The resulting improved formulation makes the learning process more transparent and should help to accelerate the development of safer systems.

II. Categorization of Accident Consequences

In the Reactor Safety Study (1975) a fairly detailed categorization of possible accident consequences had been introduced. Here, a more coarse consequence-categorization is appropriate. Table 1 lists the categories introduced here, where column 1 gives the category number. The action taken after the accident is an important aspect in this categorization.

Column 2 categorizes the seriousness of the events; column 3 specifies the "actions taken," and column 4 indicates the acceptable chance of recurrence of a particular event in a certain category. In this context, it appears appropriate to distinguish four categories:

1. Very serious accidents, with substantial off-site consequences and a number of casualties; no accident of this kind ever happened with nuclear reactors.
2. Serious accidents, but with relatively small off-site consequences. This category includes accidents in which the core is severely damaged; but it may also include accidents without core envolvement. What these accidents have in common is that they revealed a major design inadequacy, which is subsequently eliminated in existing and later plants such that a recurrence is precluded. We would place five accidents with U.S. nuclear reactors in this category (see Sec. VI).

TABLE 1

Categorization of Accidents and Incidents Involving Design Deficiencies

Category No.	Seriousness of the event	Action Taken	Acceptable Recurrence	Number of events in U.S.A.
1	Very Serious (substantial off-site consequences)	Elimination of the responsible design-error through retrofitting in all systems	none, recurrence precluded by actions taken	0
2	Serious (but insignificant off-site consequences)	same as in 1.	same as in 1.	5
3	Less Serious (no or practically no off-site consequences)	more-than-repair action taken	Probability for recurrence significantly reduced	dozens
4	Not Serious (no off-site consequences)	only repair action taken	probability not altered significantly	hundreds

3. The third category comprises "less serious" events; they have no (or practically no) off-site consequences. These incidents are considered serious enough that they warrent "more than just the repair" of the failed component, i.e., actions in other plants which use the same component appear to be called for. This category also includes the mere discovery of a design flaw, prior to any failure.
4. The fourth category comprises the longer list of abnormal events which reveal a minor design flaw, the character of which is such that merely repair action is needed; a recurrence is deemed tolerable. Subsequent designs will probably improve upon that.

The "numbers of events" for the last two categories are only listed to additionally characterize the particular types of events by giving an order of magnitude of their probable previous occurrence.

As all accidents start with some sort of a failure in the hardware or the operation of the system, this categorization represents also a basis for a description of accident progressions. In most cases one would expect an initial failure to not proceed beyond categories 4 or 3. Only if amplified by compounding failures or errors would the initiating event develop into an accident of category 2 or 1. In this sense, the partitions between these categories can be thought of and utilized as "lines of defense."

III. Umbrella Accident Considerations

Of all the large-scale technologies, that have a significant risk-potential, the nuclear power reactor development is probably the only one which -- during their long years of development and operation -- did not have accidents causing, all in all, a large number of casualties. The fact is that there has never been a "very serious" accident and of the serious ones there were only a very few, without any casualties from civilian-reactor accidents. This is probably neither incidential nor just fortunate.

Safety has always been foremost in the mind of nuclear engineers. A much larger fraction of effort has been devoted to safety than for probably any other technology. The result was a conceptual framework called "defense-in-depth," in which three or four lines of defense against the release of major amounts of fission products have been established.

Within this conceptual framework, the possibility of unforseen and human-error related accidents was fully recognized. Based on this recognition, the defense-in-depth concept has been augmented by detailed considerations of so called "umbrella accidents." Although accidents can differ in many details along the paths of the accident progression, the investigation of certain types of major events, and the protection against these, can provide an "umbrella," which also protects against other major accidents. The umbrella-

type protection tends to prevent the progression of category-2 to category-1 type accidents. It is therefore particularly important for UFM-accidents as pointed out recently by L.B. Luck (1979).

IV. The Single-System Evaluation Approach

A "model" in terms of several model-assumptions needs to be established as basis for the statistical evaluation (Campbell and Ott, 1979):

The types of accidents considered here are caused by human error, particularly in the design of the system. At some time during the operation such a design error may get involved in a chain of events and -- through its presence -- may cause an accident.

Basically, accidents of these categories, "lurking" in a system, can be described by occurrence rates or frequencies λ_k even though these λ_k cannot be determined a priori. Since the design deficiency is present in the system from the beginning of its operation on, one can, in a first approximation, assume that the λ_k's are constant in time. This is the first model assumption. (Note, fatigue related errors, for which the λ's would increase with time, are not considered here.)

The second model assumption is that any λ_k which manifests itself in an accident is subsequently eliminated. Thus, the λ_k's are step-functions, being constant until the time of the accident, and zero thereafter. This elimination holds for the first two categories listed in Table 1.

The first key to the evaluation procedure is the realization that the (average) rate, L, for any one of these accidents to occur at a given time is just the sum of all λ_k's (note that the total number of terms is unknown which is indicated by "..." as "upper limit" of the k-values):

Average occurrence rate: $$L = \sum_{k}^{\cdots} \lambda_k \quad . \tag{1}$$

Prior to the first accident, the sum in Eq. (1) begins with k = 1, prior to the second with k = 2, etc.:

$$L_1 = \sum_{k=1}^{\cdots} \lambda_k \quad ; \qquad L_2 = \sum_{k=2}^{\cdots} \lambda_k \cdots \quad . \tag{2}$$

If several accidents have already occurred (say K), one wants to know the sum of all residual frequencies at that time, i.e.,

$$L_{rest} = \sum_{k\text{-rest}}^{\cdots} \lambda_k = L_K = \sum_{k=K}^{\cdots} \lambda_k \quad . \tag{3}$$

The second key to the evaluation procedure is the realization that any inverse L gives the average time between the previous and the subsequent event. Thus, $1/L_1$ is the average time for the first event to occur, $1/L_2$ the average time <u>interval</u> between the first and the second event, etc.

Then, aside from statistical fluctuations, the inverse of any period between two UFM-accidents (say Δt_K) gives an estimate of L_{rest}, i.e., of <u>all the residual</u> λ_k's, prior to the K'th event:

$$\tilde{L}_K = \frac{1}{\Delta t_K} \quad . \tag{4}$$

There are however, statistical fluctuations in the length of the time intervals between accidents. If incidentally, there should be a long period (Δt_{long}) without an accident, followed by a short one (Δt_{short}), the inverse of Δt_{short} would give an L_{rest}-estimate that is much larger than the previous estimates, $1/\Delta t_{long}$. In actuality however,

the L_K-values decrease, as another term (λ_k) is removed from the sum, Eq. (3). Therefore, an improved estimation procedure is needed that yields a monotonically decreasing set of L_{rest}-estimates, reflecting the actual monotonic decrease in the L_K-values. The procedure of <u>isotonic regression</u> has the desired characteristics and has therefore been applied by Campbell and Ott for the L_{rest}-estimation. The resulting estimates are denoted by $\hat{L}_K$:

$$\hat{L}_K = \left(\frac{1}{\Delta t_K}\right)_{IR} , \qquad (5)$$

where the subscript IR indicates the modification of Δt_K by means of the isotonic regression procedure as presented by Campbell and Ott.

In the application presented below, the Δt_K values are such that the $\tilde{L}_K$-estimates automatically appear in a monotonically descending order. Therefore, just forming the inverse of Δt_K already yields the desired estimates; i.e., the $\tilde{L}_K$ are then equal to the $\hat{L}_K$.

Before its application to nuclear power reactors, the simple evaluation approach is illustrated in two trivial idealized examples, in which all λ_k's are assumed to be known. For simplicity, the λ's are assumed to be the same; let λ be equal to 1/12 years in case A:

$$\lambda_k = \lambda = \frac{1}{12} \text{ per year, for all k (case A).} \qquad (6)$$

Suppose also the number of possible events is given. Then, the corresponding L_K values are known. The statistical fluctuations of the occurrence times t_K are neglected in these idealized example. The corresponding Δt_K-values are then equal to $1/L_K$:

Since statistical fluctuations are disregarded, the first event occurs at time t_1, with

$$t_1^A = \frac{1}{4\lambda} = 3 \text{ years } \left(= \frac{1}{L_1}\right) \quad . \tag{7a}$$

After the first occurring event has been eliminated, there are three possibilities left. The respective second event occurs then after the time interval Δt_2:

$$\Delta t_2^A = \frac{1}{3\lambda} = 4 \text{ years } \left(= \frac{1}{L_2}\right) \quad . \tag{7b}$$

Note that the time interval is stretched because there are fewer possibilities left than before. This is even more so in the subsequent events:

$$\Delta t_3^A = 6 \text{ years}; \quad \Delta t_4^A = 12 \text{ years} \quad . \tag{7c}$$

The inverse time intervals of Eqs. (7) yield the estimates of the residual L as given by Eq. (4). The learning process becomes evident through the strong decrease in the $\tilde{L}_K$-values as illustrated by the solid lines in Fig. 1.

In a second example, case B, λ is assumed to be ten times smaller, and the number of possibilities ten times larger. The time t_1^B is the same as t_1^A, but by eliminating a very few accident possibilities, the number of "lurking" accident-possibilities is reduced only marginally (from 40 to 39 to 38 etc.). In this case then the learning from a smaller number of accidents would not improve the safety in a very significant way.

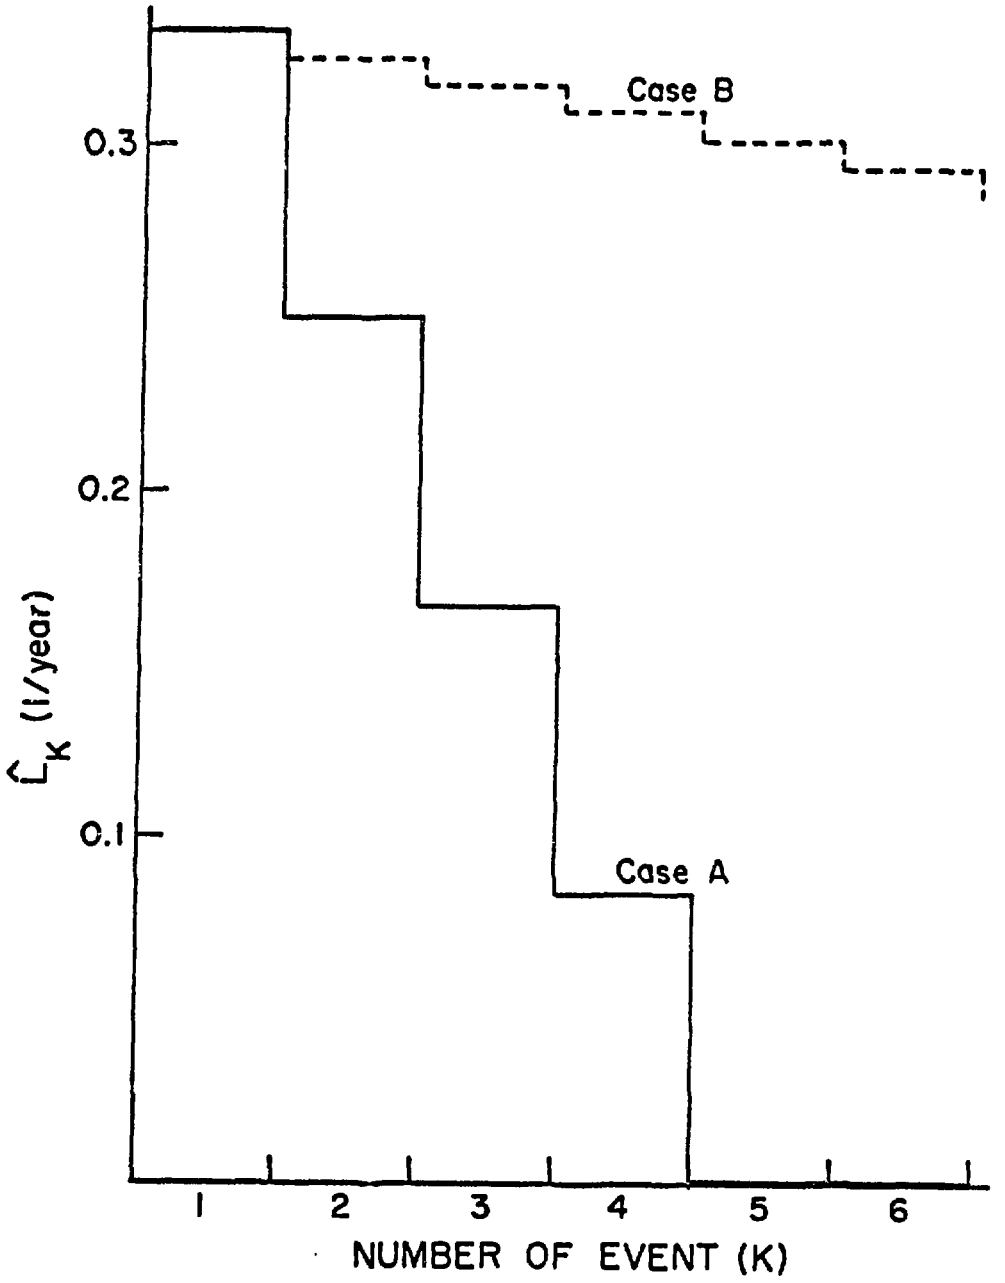

Figure 1: Illustration of the statistical evaluation of the residual frequencies, $\hat{L}$. Case A: four event-possibilities of equal frequency (solid line). Case B: 40 event-posibilities with 1/10 of the frequency as in case A (dashed line).

V. The Generalized Evaluation Model

The model developed by Campbell and Ott (1979) formally considers a single type of a system, with the individual members distinguished only by their different operational schedules. It was shown that the operation times of all systems between two accidents have to be combined in order to obtain the proper temporal spacing for the L_K-estimates.

Here, this model is generalized by introducing the λ_k's of individual reactors rather than of an entire class. Then, the sequence of λ_k-values is to be replaced by a λ-matrix, in which the components λ_{ik} give the frequency of accident-type k for the individual system i. The combined occurrence rate, L, is then given by

$$L = \sum_k \sum_i \lambda_{ik} \quad . \tag{8}$$

If one sums over i and defines

$$\lambda_k = \sum_i \lambda_{ik} \quad , \tag{9}$$

one has formally the same expression as in Campbell and Ott (1979) but the individual terms in the sum of Eq. (9) may be different. This then allows one to account for differences in systems as they certainly occur during the developmental evolution of a technology.

The λ's for a particular system contribute to L only while there is in principle the possibility of an accident. These periods of times include the operation of the system, (at any power level)

as well as, periods such as refueling or control rod drive maintenance. Let $c_i(t)$ describe the time schedule when the reactor is in one of the states delineated above, which are abbreviated as operation in quotation marks, "operation:"

$$c_i(t) = \begin{cases} 1 \text{ if the reactor is in "operation."} \\ 0 \text{ if the reactor is not in "operation."} \end{cases} \qquad (10)$$

Inserting $c_i(t)$ in Eq. (8) gives the generalized expression

$$L(t) = \sum_k^{\cdots} \sum_i c_i(t)\lambda_{ik} \quad . \qquad (11)$$

Equation (11) reflects the time dependencies as they result from the various operational schedules. In the range of accuracies, that can be achieved in an evaluation like this, the description of this kind of detail is not needed. We therefore replace in Eq. (11) the $c_i(t)$ by corresponding average "load"-factors, $\bar{c}_i$:

$$L = \sum_k^{\cdots} \sum_i \bar{c}_i \lambda_{ik} \quad , \qquad (12)$$

or even simpler

$$L = \bar{c} \sum_k^{\cdots} \sum_i \lambda_{ik} \quad . \qquad (13)$$

Equation (13) accounts the reduction of the occurrence rate (per calender year) due to off-operation periods simply by a typical average "load" factor, $\bar{c}$.

Equations (11) to (13) are generalizations of the previously used sum, Eq. (1). The main additional aspect as compared to Eq. (1) is that -- through the consideration of individual systems -- the requirements on commonality of essential design features to all systems has been eased if not removed.

Still, there is a substantial commonality between various nuclear reactors and reactor-types. Therefore, an accident in one system can provide quite important information for other systems although they may be different in many respects from the one which had the accident. Thus, the learning from a particular accident can lead to a removal of this type of accident-possibility in all systems.

The previous consideration addresses only the learning from accidents. There is however, substantial learning through r & d, or as the result of minor incidents as comprised in category 3 of Table 1. This additional learning can be represented by a factor $d_{ik}(t)$ which describes the corresponding reduction of the original λ_{ik}; i.e., $d_{ik} = 1$ prior to the implementation of that learning event, and $d_{ik}(t) < 1$ after that point in time. The resulting expression for L is given by

$$L = \sum_{k}^{\cdots} \sum_{i} \bar{c}_i d_{ik}(t) \lambda_{ik} \quad . \tag{14}$$

The L-estimation procedure is the same as described above: the inverse time intervals between events give the respective previous L_K-estimate per Eq. (4) or Eq. (5). In practically determining the Δt_K for a variety of systems one can in a first approximation just combine the calendar years in which these systems were in "operation." The inverse of the so determined Δt_K-values automatically contains the reduction through an average "load" factor as required by Eqs. (12) or (13).

VI. *Category-2 Accidents*

We have limited our examination of Category 2 accidents to power reactors but we have included small experimental reactors in this category if they were developmental precursors or prototypes of power reactors. We have not included production, research or naval reactors, and have limited ourselves to U.S. experience. The following are the reactors that had accidents so categorized; the dates of the accidents are also given:

(1)	EBR-I	November, 1955
(2)	SL-1	January 3, 1961
(3)	Enrico Fermi	October 5, 1966
(4)	Browns Ferry	March 22, 1975
(5)	Three Mile Island	March 28, 1979

What these accidents have in common is that they are on the one hand human-error related, on the other hand they led to a widely recognized learning experience that significantly influenced the engineering development of nuclear power reactors. The fact that two of these reactors had sodium as coolant (EBR-I and Enrico Fermi) and that one of them was a military facility (SL-1) is of a much lesser importance than the wide application of the respective learning experience. In addition, the differences of systems as they evolve during the development of a large-scale technology are accounted for by the general procedure presented above.

A brief description of each accident is presented below:

(1) The EBR-I Core Melting Accident

While EBR-I was a power producer (the first reactor to produce electrical power) it was designed primarily to study the physics and behavior of fast reactors. In the course of the experimental program it was found that the power coefficient of reactivity had a relatively large positive component, especially under reduced flow conditions (Thompson, 1964). Since there was some doubt as to the origin of this reactivity component, it was decided to repeat an experiment done earlier. It was planned to place the reactor on a long positive period and to end the experiment on a short negative period. The experiment was conducted with the coolant flow shut off. In previous experiments it had been possible to end the excursion by use of slow response motor-driven control rods. In the shutdown procedure laid out for this experiment, the technician at the panel was expected to use the fast acting shut off rods upon receipt of instructions from the scientist in charge. Upon receiving the instruction the technician repeated the use of the slower control rods. The staff scientist, as soon as he realized the situation reached over and pressed the rapid shut off button and simultaneously the automatic power-level trips activated the shut off rods. When it was recognized that reactor power was still increasing, the reactor scram button was pushed and at the same time an instrument scram was indicated. (The reactor scram caused the blanket to drop away from the reactor.) Subsequent examination showed that 40% to 50% of the (small) core had melted because of the delayed shut down.

After the accident, the positive reactivity coefficient was recognized as being caused by fuel element bowing, and different criteria have been established especially for fast reactor core

designs as a result. Specifically, the methods of core support have been changed to insure that dimensional changes as a result of thermal expansion are such that this component of reactivity is negative.

(2) The SL-1 Accident

An accident with the reactor occurred while three reactor operators were carrying out supposedly routine maintenance while the reactor was shutdown

The accident was apparently caused by a rapid withdrawal of the central rod beyond the point which it should have been moved. This caused a severe nuclear excursion. The reactor was completely destroyed; the three operators were fatally injured.

It is worth noting the conclusion regarding the accident by Thompson (1964). "As has been pointed out in the introductory chapter, most accidents involve design errors, instrumentation errors, and operator or supervisor errors. The SL-1 accident is an object lession in all of these. There has been much discussion of this accident, its causes, and its lessons, but little attention has been paid to the human aspects of its causes."

As a result of the SL-1 accident, it has become a general criteria in reactor design that control rod systems be designed such that the reactor cannot be made prompt critical by the withdrawal of single control rod.

(3) The Fuel Melting Accident in the Enrico Fermi Atomic Power Plant

The fuel melting accident in the Enrico Fermi reactor occurred on October 5, 1966 during a slow power increase (APDA, 1968). The reactor had been operated without incident since August 23, 1963. The reactor

was shutdown after it was observed that several subassemblies had abnormally high outlet temperatures and after radiation alarms sounded due to the leakage of a small amount of fission products into the reactor building.

Inspection of the core after the event revealed that melting had occurred in two subassemblies. Inspection of the inlet plenum revealed two loose objects which were identified as two of six zirconium segments which were originally installed on the conical flow guide as part of a system to mitigate a core meltdown. From the analyses and inspections it is believed that the fuel melting was caused by partial blockage of the inlet nozzles of four adjacent subassemblies. This blockage caused fuel in two of the subassemblies to melt and the fuel in the other two to overheat.

Since the accident with the Enrico Fermi reactor, it has become a clear criteria in reactor design that the reactor be designed such that loose objects in the primary system cannot cause a significant blockage of flow to the core.

(4) The Browns Ferry Accident

The Browns Ferry accident was the result of a fire which started "when workmen used a candle to check for air leaks between the cable room and the Unit 1 reactor building, which was kept at a negative air pressure. The urethane seal used where cables penetrated the wall caught fire and destroyed more than a thousand cables, knocking out inplant operating and safety systems, including emergency core-cooling systems. Despite the loss of some instrumentation, though, indicators for reactor water level, temperature, and pressure continued to work and both Units 1 and 2 were safely shut down from 100 percent power" (Nuclear News 1976).

As a result of the Browns Ferry accident corrective actions were taken in fire prevention in nuclear plants and strict criteria for cable separation were implemented.

(5) The Three Mile Island Accident

As other papers at this meeting will discuss the Three Mile Island Accident experience in some detail, and further examinations will be necessary to completely assess the core damage, a detailed discussion of the accident will not be included here.

The accident occurred at 4:00 a.m. on March 28, 1979; the President's Commission (1979) concluded that "the accident was initiated by mechanical malfunctions in the plant and made much worse by a combination of human errors in responding to it." While it is too early to assess all the corrective actions that will be taken as a result of the Three Mile Island accident substantial actions have already taken place. Among these are some design modifications, operator training to deal with such emergencies, and an in-depth probabilistic safety analysis of existing plants.

The TMI-accident was not as clearly caused by a single design deficiency in the hardware as the other accidents included in Category 2. One may even raise the question, whether the TMI-accident was an IFM- or UFM-accident. In our assessment, the TMI-event clearly was a UFM-accident since it assumed accident proportions due to inadequacies in the "design" of the system, with "design" understood more broadly than just pertaining to the hardware. In this more general sense, "design" should include also procedures and instrument lay-out such as:

-- probabilistic risk evaluation for individual plants,
(as part of the design/review process)
-- expedient learning from precursor incidents,
-- clarity of the layout of the control panel and
information devices, especially for the diagnosis
of off-normal operation,
-- operator training for handling off-normal events,
-- adequate emergency preparedness.

Apparently, the TMI-accident revealed "design-deficiences" in several of the areas listed above, and the Kemeny commision (1979) concluded "we are convinced that an accident like Three Mile Island was eventually inevitable." Thus, in the Kemeny commission's judgment, the TMI-accident satisfied the first "model assumption" described above, i.e., it had a preexisting occurrence rate λ_k associated with it

The second "model assumption," that a recurrence of this type of an accident be made (practically) impossible, is apparently also satisfied, judging from the changes that have already been made and that are being implemented or are under consideration. The expedient learning from precursor incidents alone could have prevented the TMI-accident.

Thus, the TMI-accident was, in this broader sense, design-error related, and as such, it satisfies both model-assumptions needed for the presented analysis of UFM-accidents.

VII. Statistical Evaluation of Category-2 Accidents

The category-2 accidents on U.S. power reactors or on their developmental precursors have been described in the previous section. The operational periods of all U.S. reactors of these types between the listed accident dates allows an estimation of the residual frequencies, $\hat{L}_K$, "lurking" in all these reactors. The approximate operational periods, up to and between the accident dates have been extracted from the World List of Nuclear Power Plants (Nuclear News, 1979) and from Nuclear Engineering International (1979).

The results are given in Table 2. The first two columns give the reactor names and accident dates, the next two the accumulated and incremental operational periods in calendar years. The listed periods are not reduced by "load" factors so that the inverse periods can give the $\hat{L}$-estimates directly in calendar- and not in operation-years [see Eqs. (12) and (13)]. Since the incremental periods increase monotonically, the L-estimates per Eqs. (4) and (5) agree, i.e., the isotonic regression modification is not needed (see Campbell and Ott, 1979). The last column in Table 1 gives the respective results, which are also depicted in Fig. 2. The L-estimates strongly decrease from event to event. The latest value of $\hat{L}$, (i.e., $\hat{L}_5$) is determined by the TMI-accident to be 0.0046, i.e., a TMI-type accident was "lurking" in the present reactors corresponding to an L of about one per 220 years. The tremendous learning experience of this accident should result in a substantial reduction of the residual frequencies after March 1979. The last line in Table 1

Table 2

Statistical Evaluation of Category-2 Accidents

Month-Year	Reactor	Reactor Years (t_K)	Incremental Reactor Years (Δt_K)	$\hat{L}_K = \tilde{L}_K$ (1/year)
11-1955	EBR-I	4	4	0.25
1-1961	SL-I	15	11	0.091
10-1966	Enrico Fermi	75	60	0.017
3-1975	Browns Ferry	164	89	0.011
3-1979	TMI-2	380	216	0.0046
1-1980	--	(430)	(50)	

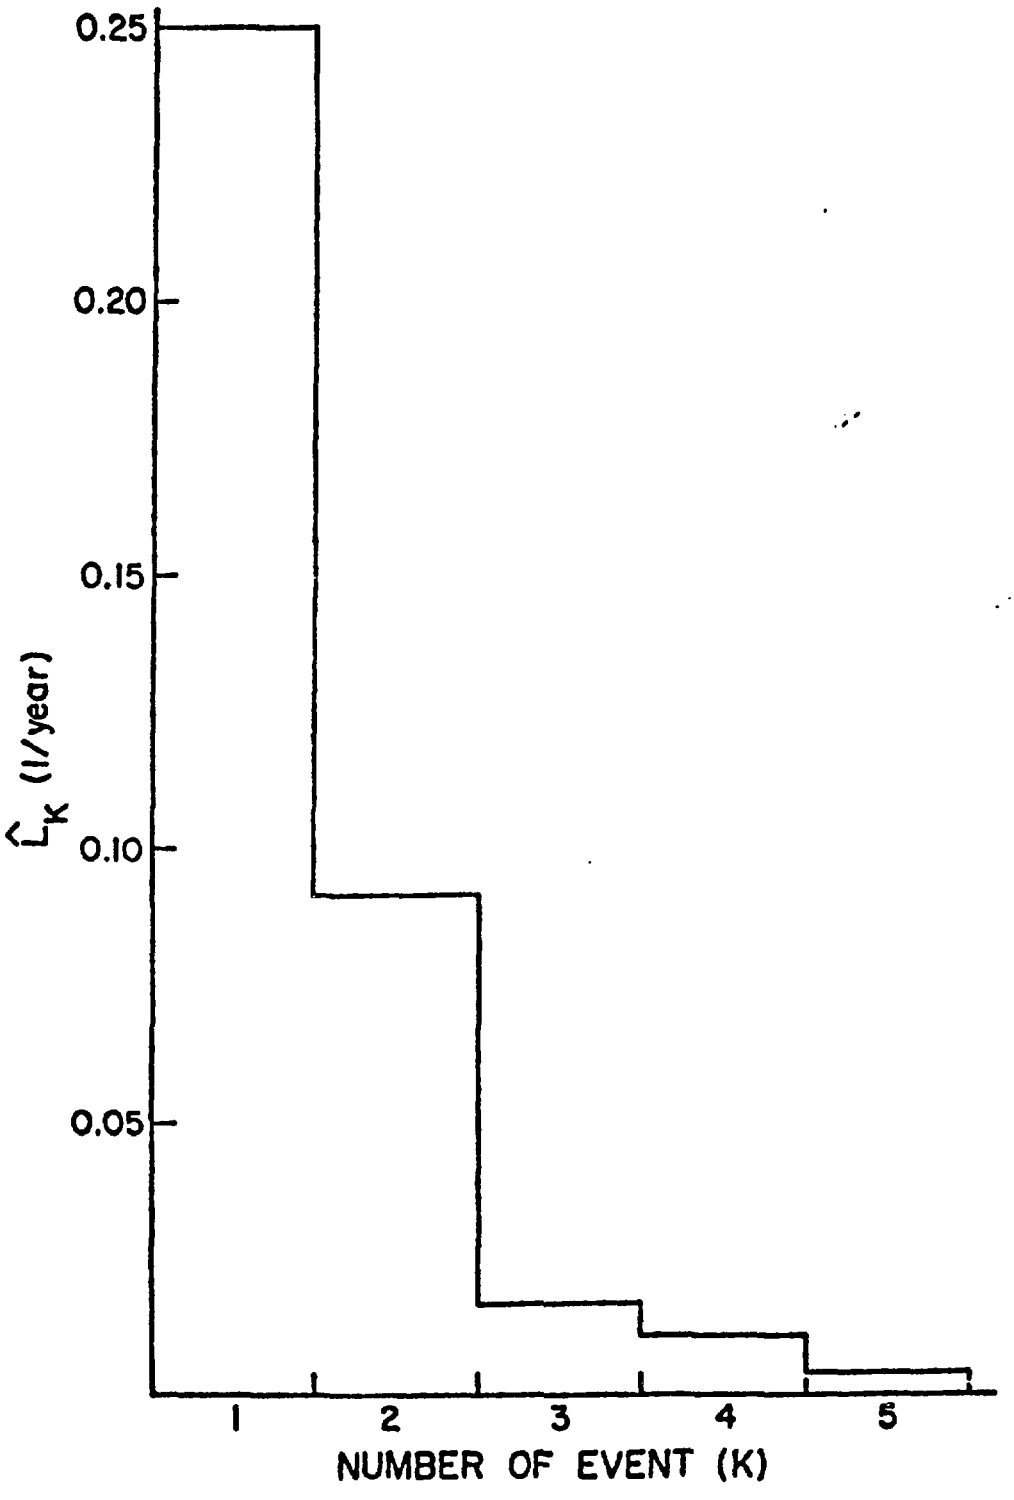

Figure 2: Estimated residual frequencies ($\hat{L}_K$) of category-2 accidents for U.S. power reactors and developmental precursors: $\hat{L}_K$ applies to the time prior to the K'th event.

shows the approximate subsequent operational period. It is still too small to give an indication of a possible reduction of $\hat{L}_{K+1}$.

VIII. Time-Trends of UFM- and IFM-Accident Frequencies

In general, one would expect a different trend of UFM- and IFM-accident frequencies with time:

Design-deficiency related (UFM-accidents) are more likely to happen during the developmental evaluation and the early years of operation. As the technology matures, the UFM-accidents should become quite unlikely. Thus, the residual frequency of UFM-accidents should decrease strongly, as it is indicated in Fig. 2 for nuclear reactors.

Alternately, the IFM-accident possibilities represent more of a "background" to which the overall accident frequency eventually recedes. A reduction of this IFM-background should also be possible, but one would expect it to occur at a slower pace than for UFM-accident possibilities. The IFM-background reduction relies primarily on r & d including the further improvement of the overall design and review process. It is formally described by a reduction of factors d_{ik} as included in Eq. (14).

Both time trends are qualitatively illustrated in Fig. 3. The category-2 accidents of nuclear power plants do not give an indication yet of having reached an IFM-background. It would be of great interest to investigate how technologies that are older than nuclear power behave in that respect. Of particular importance would be to investigate

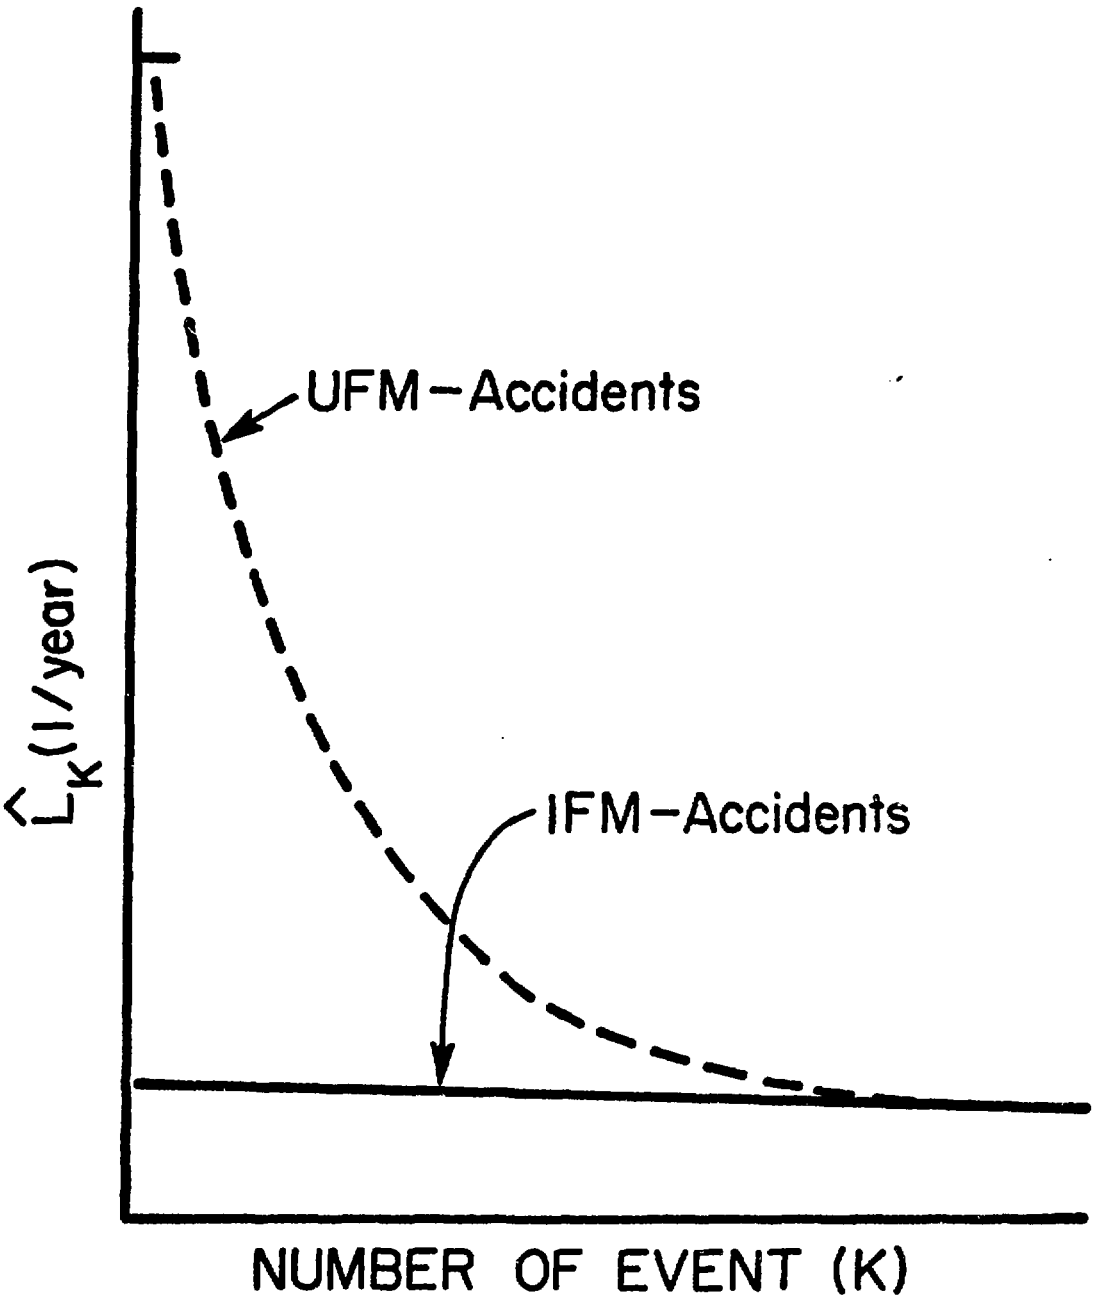

Figure 3: Qualitative comparative time dependencies of UFM- and IFM-accident rates.

for "older" technologies (that also have a substantial risk-potential) the following questions:

- How (i.e., through what particular actions) has the process of "learning from accidents" reduced the rate of UFM-accident occurrence?
- Is a transition from UFM- to IFM-accidents clearly discernable?
- How does a IFM-background depend on aspects such as
  -- technical matters (e.g., design practices, standards, review processes),
  -- human factors,
  -- man-machine interfaces,
  -- costs,
  -- risk aversion considerations,
  -- public acceptance?

The possibility of a broader study of the general comparison of UFM- and IFM-accidents for a wide range of large-scale technologies is being currently considered by the authors of this paper. It is hoped that the comparison of several "older" technologies will reveal information on the maturing process of these technologies. It may then be possible to draw more general lessons that help to accelerate the learning process in newer technologies. Furthermore, a better understanding of the entire maturing process of large-scale technologies could help to create a long-term harmonious relation of the society and its technological "life-support system."

Summary and Conclusions

Two broad groups of accidents and incidents are distinguished in this paper: Events with a priori Identified Failure Modes (IFM), and with a priori Unidentified Failure Modes (UFM). IFM-accidents are considered in a typical safety and risk analysis (e.g., the U.S. Reactor Safety Study). UFM-accidents on the other hand happen due to generally unexpected failures or combination of failures, mostly involving human error in the "design" of the systems hardware or "software" (e.g., operation or review procedures).

Occurrence rates of UFM- and IFM-accidents have different trends with time. The UFM-accident possibilities can be expected to become largely eliminated during the maturing process of a technology. The overall accident rate should then recede to an IFM-accident-rate background.

This paper is primarily concerned with a statistical analysis of the actuarial UFM-accident experience. A categorization of UFM-accidents is introduced based both on the seriousness of the events and on the actions taken and lessons drawn afterwards.

Category 1 contains the "very serious" accidents with substantial off-site consequences and Category 2 comprises the "serious" accidents which have only insignificant off-site consequences. For both of these categories the lessons drawn result in elimination of the responsible design-error through retrofitting in all systems, which should preclude recurrence of these types of accidents.

Two additional categories are introduced, with Category 3 comprising incidents for which more-than-repair action is taken; i.e.,

the Category 3 events contribute significantly to the learning experience and thus to the maturing process of the technology.

In a sense, the categorization of UFM-accident possibilities described above is reflected in the safety design of nuclear reactors: The progression of UFM-accidents from category 2 to the ("very serious") category 1 is made very unlikely by the provisions resulting from umbrella-accident evaluations. In the same way, learning from accidents, such as TMI, consists largely of making the progression of category-3 type incidents to category-2 type accidents very unlikely.

The analysis of UFM-accidents developed here (and in an earlier paper) pertains to events for which a recurrence is precluded by proper actions (Categories 1 and 2). The evaluation procedure aims at the estimation of the residual occurrence rate of UFM-accident possibilities "lurking" in the system. No category-1 accident has happened on a nuclear reactor. The application to category-2 accidents shows a decrease of about a factor of 50 up to the time of the TMI-accident. This is in stark contrast to an illustration presented in Fig. 1 (case B) of a technology with many low-frequency UFM-accident possibilities, for which the learning experience from a single accident would be marginal. This analysis clearly indicates that the more likely accident possibilities have already been removed from the system and that now, in all likelihood, large numbers of accidents with a high expectancy cannot be "lurking" in the system.

The learning experience of the TMI-accident itself is quite broad. Sweeping reforms and improvements have already been and are being implemented or defined. One therefore would expect a further substantial reduction of the possible future occurrence rate of category-2 accidents.

References

1. Reactor Safety Study (RSS) WASH-1400 (1975).

2. Campbell, G. and Ott, K.O. (1979), Statistical Evaluation of Major Human Errors During the Development of New Technological Systems, Nucl. Sci. Eng. 71: 267-279.

3. Luck, L.B. (1979) private communication.

4. Thompson, T.J. Technology of Nuclear Reactor Safety, Vol. 1, Reactor Physics and Control, MIT Press, Cambridge, Mass. (1964).

5. Report on Fuel Melting Incident in the Enrico Fermi Atomic Power Plant on October 5, 1966 APDA 233 December 15, 1968.

6. Browns Ferry - Nuclear News, April 1976.

7. Three Mile Island - Report of the President's Commission on the Accident at Three Mile Island, October 1979, Washington, D.C.

8. World List of Nuclear Power Plants, Nuclear News, p. 81, August 1979.

9. Nuclear Engineering International, p. 22, July Supplement, 1979.