

Advances in Probabilistic Safety Assessment

George Apostolakis

University of California, Los Angeles, CA USA

Peter Kafka

Gesellschaft für Reaktorsicherheit (GRS) mbH, Garching, FRG

ABSTRACT

This paper reviews the evolution of Probabilistic Safety Assessment in the last twenty years, the basic methodology, lessons learned, as well as several current issues. These include the formal use of expert opinions, model uncertainties, the notion of a living PSA, and source-term uncertainties.

1. INTRODUCTION AND PSA OVERVIEW

The complexity of large technological systems like nuclear power plants and the requirements for efficient and safe operations have created the need for the development of models that accurately represent these systems. The occurrence of major accidents (e.g., Three Mile Island and Chernobyl) has focused the attention of the public on the safety of these facilities and has accelerated the development and use of these models. It has also made it very clear that the events of interest are rare and any decision-making process that involves them must include the large uncertainties that are associated with their occurrence.

The methodology that consists of these models is known as Probabilistic Safety Assessment (PSA), although essentially the same methodology can be employed for reliability assessment. It essentially consists of two parts:

- i. The identification of scenarios (accident sequences) that lead to the consequence of interest, e.g., system unavailability, the release of radionuclides, and so forth; and
- ii. The quantification of the uncertainty associated with the occurrence of these scenarios.

The identification of the scenarios follows the logic depicted in Figure 1 [1]. The process begins with the selection of a set of Initiating Events (IE), i.e., events that have the potential to start an accident scenario, e.g., loss of coolant accidents (LOCA) of various sizes (for a reasonable list of IES see [2]).

The possible responses of the plant to the occurrence of the IE are modeled next employing event trees (ETs). Each branch point of the ET represents the possibility that the corresponding safety system (or function) named above this point is available or not. The paths through the ETs define the accident sequences. These sequences lead to a "plant state," that is, to a particular state of the reactor. PSAs that end at this point are called 'Level 1' PSAs [2].

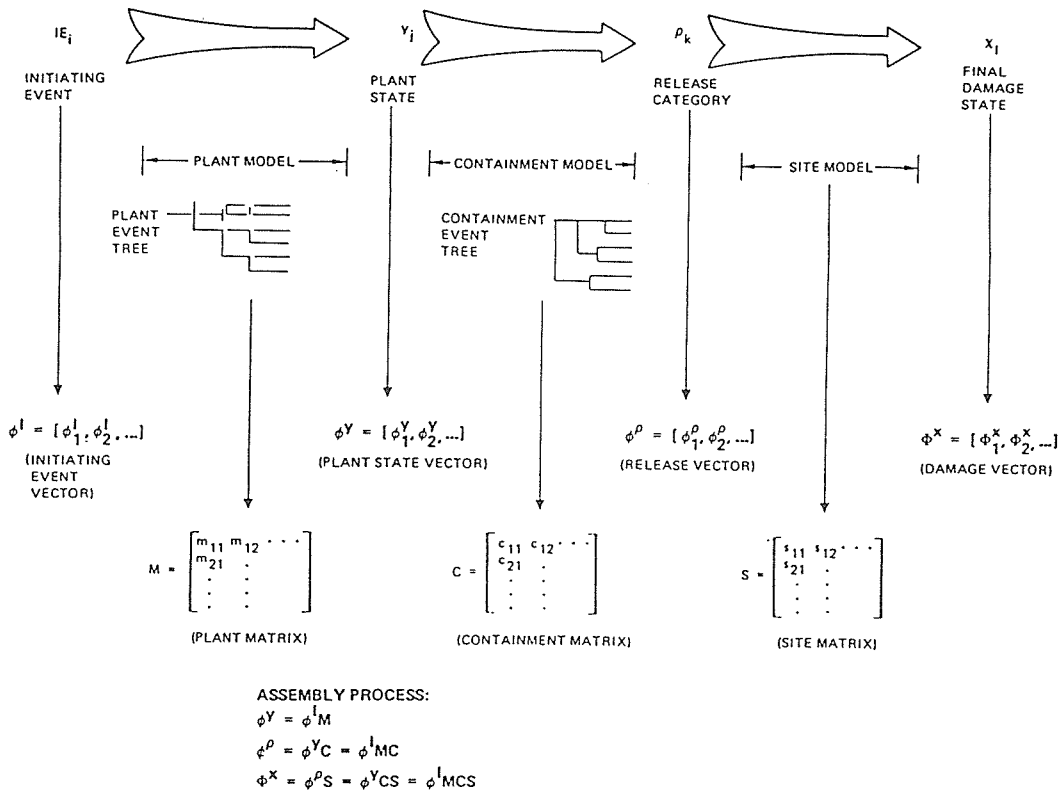


Fig. 1. Overview of the assembly process for nuclear plant risk analysis, showing relationship of pinch points, frequency vectors, event trees, and transition matrices [1].

Given damage to the reactor, as specified by the plant states, we continue to develop the accident scenarios by following the progression of physical phenomena that are possible after the core is damaged. The end points of the accident scenarios are now the various releases of radioactivity to the environment. A 'Level 2' PSA would stop at this point of the release categories.

Finally, the accident sequences are completed by including site characteristics that may influence health effects ('Level 3').

The question that we ask in the preceding development is "if event X (e.g., an IE) occurs, how can state y (e.g., a plant damage state) be reached?" Another type of question that we often ask is "how can this event, e.g., system failure, happen?" We usually answer this question by employing fault tree analysis (FTA). Typically, FTA produces more scenarios that lead to the specified undesirable consequences.

The major problems that usually arise in connection with the development of the scenarios are related to the question of the completeness of the analysis. The possible omission of important events or processes from the accident scenarios is a perennial problem.

After the scenarios have been identified there remains the issue of quantifying our uncertainty regarding their occurrence. While various methods are employed to varying degrees of sophistication, the rigorous and formal methods of Bayesian or Subjectivistic Probability Theory [3,4], are gaining increasing acceptance. Probability is interpreted as a measure of degree of belief in

accordance with de Finetti [5]. Admissible evidence is not just of the conventional statistical type, but also expert judgment stemming from knowledge of the process and operational experience. The use of expert opinions is necessitated by the fact that the events of interest are rare and statistical evidence is either very weak or nonexistent. This extensive use of judgment often creates conflicts between industry and regulatory agencies. While the formal methods that we cited above help in the understanding of the differences, they are unable to eliminate all of them.

To make the discussion concrete, and without loss of generality, we consider the failures of a type of pump. The random variable is T , the failure time. Then, the distribution of this time is usually taken to be exponential, i.e.,

$$F(t/\lambda) = 1 - \exp(-\lambda t) \quad (1)$$

This is the probability that T is smaller than t , i.e., that the pump fails before t .

The parameter λ , the failure rate, of Eq. (1) specifies $F(t)$. Its value depends on the kinds of pumps that we have included in our class of pumps and on the conditions of their use. Thus, the value of λ depends on what we include in our model. It is important to realize that the pumps and conditions of operation that are included in the model are assumed to be completely equivalent (as far as the behavior of T is concerned). That is, if there is no distinction between two different plants, we assume that two pumps, of the type of interest, at these two plants are not distinguishable.

After we define the model as above we realize that the numerical value of λ is not known precisely. Thus, we use a subjective probability density function, pdf, $\pi(\lambda)$ to express this uncertainty (it is common practice to use the lognormal distribution for $\pi(\lambda)$.) Then, $F(t/\lambda)$ of Eq. (1) can be viewed as a conditional probability, i.e., it assumes knowledge of the numerical value of λ . The unconditional probability of failure before t is

$$\Pr[T \leq t] = \int_0^{\infty} F(t/\lambda) \pi(\lambda) d\lambda \quad (2)$$

Some authors (e.g., Reference 4) wish to make very clear the distinction between $F(t/\lambda)$ and $\pi(\lambda)$ and call the former "frequency" and the latter "probability" (hence, the probability-of-frequency formulation). We may also say that $F(t/\lambda)$ represents statistical or stochastic uncertainties (in the terminology of decision theory, it is our model of the world), while $\pi(\lambda)$ represents state-of-knowledge uncertainties.

The generalization of Eq. (2) is

$$\Pr[T \leq t] = \int d\phi \pi(\phi) F[T \leq t/\phi] \quad (3)$$

Where $F(T \leq t/\phi)$ is the stochastic model for the random variable T and ϕ is the vector of its parameters. The state-of-knowledge pdf $\pi(\phi)$ represents what we know about ϕ .

When new, usually statistical, evidence becomes available, the state-of-knowledge distribution is updated using Bayes' Theorem, i.e.,

$$\pi(\phi|E) = k^{-1} L(E|\phi) \pi_0(\phi) \quad (4)$$

where:

$\pi_0(\underline{\phi})$: our prior state of knowledge about the unknown vector $\underline{\phi}$ (prior to receiving the evidence)

E: the evidence

$L(E|\underline{\phi})$: the likelihood function, i.e., the likelihood of the evidence given that the true value of the vector is $\underline{\phi}$

$\pi(\underline{\phi}|E)$: our posterior state of knowledge about the unknown vector $\underline{\phi}$ given that we have received the evidence

and k^{-1} is a normalization factor that makes $\pi(\underline{\phi}|E)$ a probability distribution function.

2. HISTORICAL PERSPECTIVE

The first major study to address reactor safety issues was WASH-740 more than thirty years ago [6]. This study was done with very limited design information and resorted to the estimation of rough upper bounds of the consequences from a large release of radioactivity. While the authors of that report cautioned its readers that their conservative results could not be appraised independently of their probabilities, they also admitted that estimating those low probabilities would present great problems.

Farmer's paper in 1967 [7] described a "new approach" that was based on the premise that "a measure of risk can be obtained by estimating the probability of the failure and assessing its consequences." He proposed his now-famous line that defined "acceptable" from "unacceptable" releases. He also warned that "it is not possible by extrapolation from past experience to establish the validity of a failure probability as low as 1 in 10^5 to 1 in 10^6 in the lifetime of a single plant."

The Reactor Safety Study (RSS, [8]) followed and basically established the PSA methodology that is used today. Its results and methodology became highly controversial. Its Executive Summary contained figures that conveyed the message that the risks from nuclear reactors are very small compared with other risks (natural and manmade) that already exist. State-of-knowledge uncertainties were not displayed and this omission was criticized a great deal. The Risk Assessment Review Group [9] undertook the task of seriously reviewing the RSS. Its report suggested further use of the methodology in the regulatory process and identified several problems that required re-examination.

The German Risk Study [10] followed in the late seventies. It essentially applied the RSS methodology to the reference nuclear power plant Biblis B sited in the Federal Republic of Germany. At the same time the utilities in the USA started sponsoring their own PSAs. The first two such studies that were released after the accident at Three Mile Island were the Zion and Indian Point PSAs [11-12]. Other studies include that for Sizewell B [13] and those for Swedish plants [14]. The latest major effort sponsored by the USNRC is NUREG-1150 [15].

3. LESSONS LEARNED

Several general conclusions that can be reached from the completed PSAs are as follows [16-19]:

i. Core Melt Frequencies

Core melt frequencies range from 10^{-6} to 10^{-3} per reactor year (including plant-to-plant variability). A good part of this wide range is due to the different methods and scope of the various studies. Nevertheless, these numbers are, in general, greater than had been thought prior to the PSAs.

Fig. 2 [17] presents a picture of the relative contributions of accident scenarios to core melt frequency. Some of the studies surveyed include "external" events (earthquakes, fires and others). Figs. 3-4 [19] provide an indication of the contribution of external events.

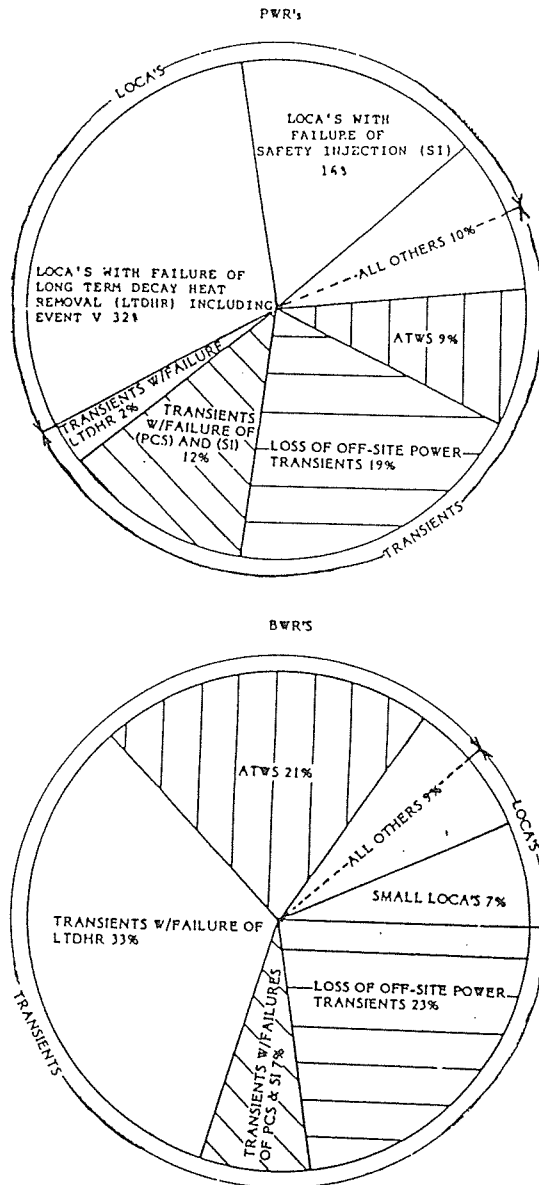


Fig. 2. Composite of major contributors to core melt frequency [17].

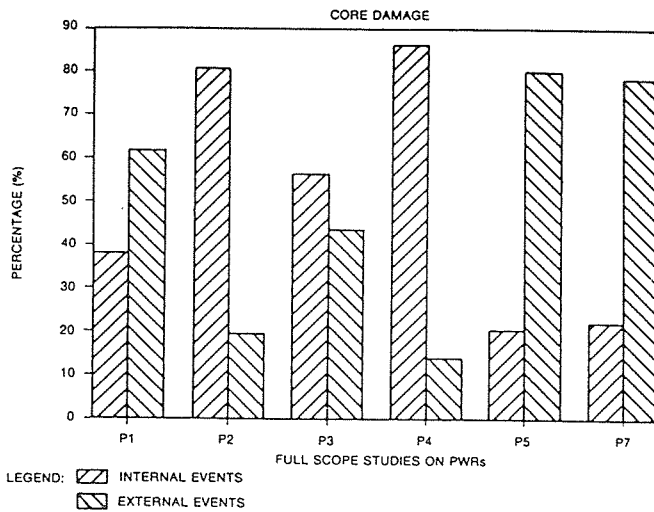


Fig. 3. Internal versus External Contributors to Core Damage [19].

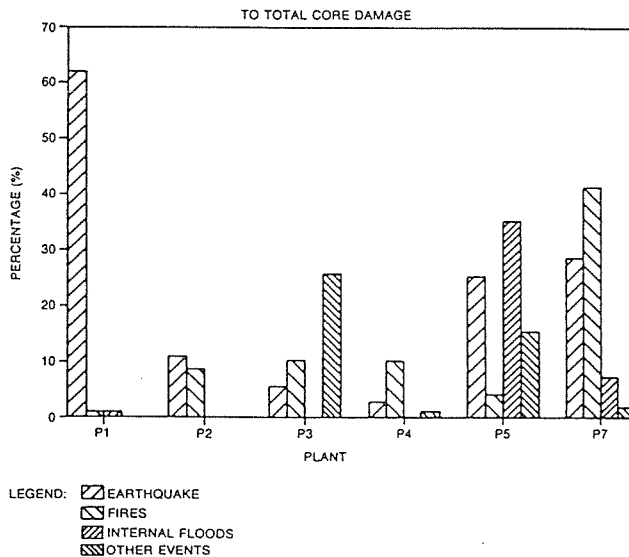


Fig. 4. Contribution of External Events [19].

ii. Radionuclide Releases

Fig. 5 [17] displays the results of several studies for iodine releases. The dominant accident sequences are those that lead to early containment failure, or to containment bypass.




Fig. 5. Iodine release fractions from listed PRA studies [17].

iii. Public Risk

Off-site consequences are generally found to be very low. The estimated public risk is within the safety goals that have been issued by the USNRC.

4. CURRENT ISSUES

4.1. Expert Opinions

NUREG-1150 has attracted attention, once more, to the problems associated with the use of expert judgment in PSAs. This is a subject of continuing interest and debate [20-21].

The use of expert opinions, when the research workers in the field believe that physical insights can be gained by models and experiments, is perceived as "soft science," especially by scientists and engineers. On the other hand, the decisionmaking responsibilities of regulatory authorities and the resultant need for timely assessments of the current state of knowledge are not always fully appreciated by these groups. The perception, at this time, is that NUREG-1150 has gone too far in the "soft" science path and that its choice of the experts has been biased. A positive outcome of this debate is the fact that the elicitation and use of expert opinions (which, incidentally, are not new to reactor safety assessments and were used under the more acceptable name of "engineering judgment" long before PSAs became popular) have become visible, thus making the PSA community at large aware of the fact that related issues have been studied by psychologists, decision scientists and others.

After the expert opinions have been elicited, they must be used in some way to produce a final result. In NUREG-1150 simple arithmetic averages are taken. The full implications of this practice have not been explored. For example, the between-expert variability tends to be suppressed this way. More sophisticated methods are discussed in [22], and the Post-SMiRT 8 Seminar on the role of data and judgment in PSA [23].

4.2. Model Uncertainties

While most of the statistical literature deals with the assessment of $\pi(\phi)$, (Eq. (4)), in the light of some evidence (usually in the form of a random sample), it is becoming increasingly evident in PSA that a major source of uncertainty could be the model itself. In some cases the choice of a model is fairly straightforward. For example, it is common practice now to include in a system unavailability analysis the dependencies between various failures. This was not the case before the publication of the Reactor Safety Study in 1975. Many system unavailabilities were calculated without CCFs (common cause failures) and were much smaller than the values that are typically calculated today (although there is no universally accepted CCF model yet).

Two problems for which several models are available in the literature are the analysis of common-cause failures and post-initiator (dynamic) human reliability analysis (HRA). A widely used model for CCF analysis is the beta-factor model (and its extension, the multiple greek letter (MGL) model [24]. Both the definitions and the statistical analysis of the parameters of these models have been questioned in [25].

Several models are also available for HRA, e.g., THERP [26], SLIM-MAUD [27] and HCR [28]. The PSA models for human reliability were reviewed and discussed at length at the Post-SMiRT 9 Seminar on Accident Sequence Modeling [29]. The high degree of skepticism on the part of the psychologists concerning the basic assumptions of the engineering models makes the issue of model uncertainties a principal one in HRA.

4.3 Living PSA

The objectives of a "living" PSA are to provide an analysis tool that:

- allows time-dependent risk monitoring of the plant
- allows evaluations to be made of the risk impacts of design and operating changes of the plant
- supports operating, maintenance and repair decisions
- is an aid in managing accidents
- evaluates the risk impact of regulatory issues and assesses the risk based prioritization of regulatory activities.

All the available systems (PRISM [30], RAPID [31]) and the concepts under development (ORACLE [32], EXPRESS [33]) in this context consist of a large data base, a computerized PSA model, and a user-friendly man-computer interface for updating the PSA model with new plant-specific data or plant-specific system/component status. The important idea of all these systems is not only to design and construct the plant with low risk, but also to operate the plant under various conditions, while maintaining that low risk level.