LA-UR- 96-1422

CONF-960912--22

Title: BAYESIAN RISK-BASED ACCEPTANCE CRITERIA

Author(s): Harry F. Martz
Lee R. Abramson
James W. Johnson

Submitted to: Probabilistic Safety Assessment '96

DISCLAIMER

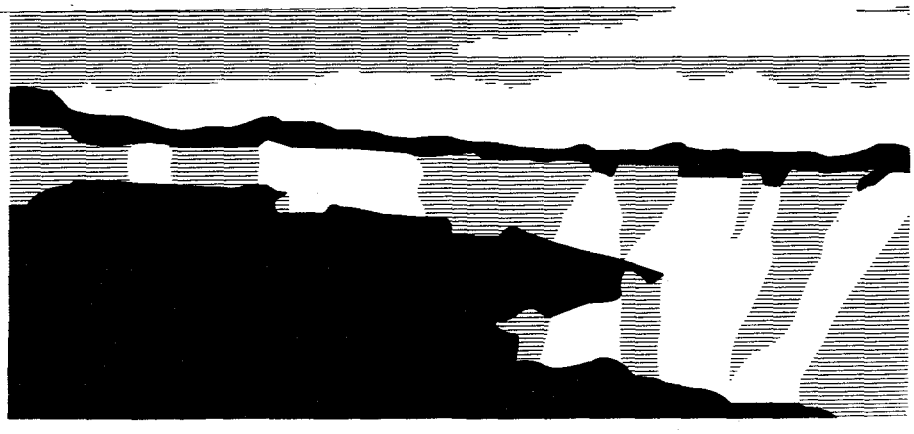

Los Alamos
NATIONAL LABORATORY

Form No. 836 R5
ST 2629 10/91

LA-UR-96-

Bayesian Risk-Based Acceptance Criteria

by

Harry F. Martz
Los Alamos National Laboratory

and

Lee R. Abramson and James W. Johnson
U.S. Nuclear Regulatory Commission

Prepared for:

US Nuclear Regulatory Commission
Office of Nuclear Regulatory Research
Washington, DC 20555

Under:

NRC Contract Job Code W6505
Los Alamos National Laboratory
Harry F. Martz, Principal Investigator

Dr. Lee R. Abramson
NRC Technical Monitor

April 1996

1 INTRODUCTION

Mechanistic (or deterministic) analysis is traditionally performed in the process of designing a new nuclear reactor or reactor core and also as part of the safety analysis of existing reactors or reload cores. *Mechanistic accident analysis* is characterized by the specification of an initial operating condition, an initiating event, and subsequent system faults. These subsequent faults are often chosen, through such mechanisms as the worst single failure criterion, so as to maximize the consequences of the accident. Thus, the behavior of all reactor system hardware is prescribed before this analysis begins. Mechanistic analysis then attempts to predict the dynamic response of the system and usually involves detailed reactor physics and thermal-hydraulic predictions of the system behavior, including such parameters as power distributions in the reactor core, coolant temperatures and flow rates, and fuel clad temperature distributions. The objective of this analysis has typically been to establish and verify reactor operating limits and technical specifications, so that severe core damage (SCD) is prevented for a wide variety of reactor accidents.

Traditional mechanistic reactor safety analysis used a prescriptive, conservative approach in its treatment of reactor accidents. The best example of this is furnished by Appendix K to 10CRF50,[1] which gives a detailed prescription of calculational methods for predicting the peak clad temperature following a large break loss of coolant accident (LOCA). Appendix K also furnishes a limit on this peak clad temperature. A second example is the specification of a minimum departure from nucleate boiling ratio (DNBR) for pressurized water reactor transients. It is assumed that this prescriptive process yields a conservative limit on reactor power and other reactor operating conditions, so that there is a high degree of confidence that no SCD would result from such an accident. However,

this traditional methodology yields neither a realistic estimate of reactor performance nor a complete estimate of the uncertainties in the analysis.

The modern approach to mechanistic analysis is to perform so-called *best estimate* (BE) analysis, in which an attempt is made to predict reactor behavior realistically with explicit consideration of and accounting for uncertainties that are introduced at every stage of the calculation. In 1988, the NRC approved a rule change in the acceptance criteria for emergency core cooling systems (ECCSs) in light water power reactors. The revised rule states that an ECCS evaluation is acceptable if it describes realistically the reactor behavior and provides quantified uncertainties. Using the BE approach, however, system engineers need to demonstrate to a "high level of probability" that the safety parameters stated in 10CFR50.46[2] will not exceed the prescribed values or criteria. The NRC does not specify how high this probability must be, although the Standard Review Plan[3] repeatedly calls for 95% probability at a 95% confidence level that certain limits should not be exceeded.

Modern mechanistic BE analysis of reactor performance under various accident conditions is characterized by two distinguishing features. First, realistic predictions are made of parameters that describe the performance of the reactor during the accident, which we henceforth refer to as *operating parameters*. For example, in the analysis of light water reactor core performance, the principal operating parameter of interest is the time-dependent behavior of the maximum fuel cladding temperature from the onset of the accident until adequate long-term coolability is established. Second, estimates of the uncertainties associated with these operating parameters must be produced. These estimates are obtained by considering uncertainty in the state of knowledge (SOK) regarding the values of the input parameters to the model (or even the model itself) for the accident under consideration. These uncertainties in the input parameters are then

propagated through the mechanistic model by means of an appropriately chosen statistical procedure, such as Monte Carlo simulation or Latin Hypercube Sampling[4,5] (or some other method). The result is a probability distribution on the operating parameter of interest; for example, a probability distribution of the maximum fuel cladding temperature under the given accident conditions.

It is also characteristic of BE methods that a statistical assessment be made of the damage threshold for each operating parameter, where a *damage threshold* is a limit on an operating parameter that is imposed to prevent SCD. Thus, instead of prescribing a fixed, conservative value at which the cladding will fail, the experimental evidence of clad failure is used to obtain a probability distribution function for the temperature at which the cladding fails. The mean of this distribution is usually considerably higher than the old conservative limit. It is typically assumed in this BE analysis that SCD occurs if and only if the operating parameter exceeds its damage threshold.

Martz et al[6] present a formal method for coupling mechanistic analyses and probabilistic risk assessment (PRA). We briefly describe the method. First, a particular initiating event is selected. The frequency of occurrence of this event is determined using standard PRA techniques. Next, a so-called *system-mechanistic analysis event tree* (SMAET) is constructed. Each of the event sequences is then analyzed by means of modern mechanistic methods, and a conditional probability of SCD is obtained. (It is at this point that a significant departure from traditional methods occurs. Traditional techniques would have assigned a value of 0 or 1 to the conditional probability of SCD for each individual event sequence.) This probability is then introduced into the event tree as a final split fraction. The contribution from this single event sequence to the overall SCD frequency is then calculated as the frequency of the initiating event times the product of the appropriate success or failure probabilities for each node of this

branch of the tree times the conditional probability of SCD obtained from the mechanistic analysis for this sequence. In this method, every event sequence, including the top-most sequence in the event tree, where all systems function as designed, yields some non-zero contribution to the probability of SCD. This occurs because best-estimate analysis always yields some finite probability of SCD, no matter how benign the reactor transient.

The effect of this new methodology is to reduce the contribution to SCD of event sequences that were traditionally assumed to lead to SCD (the frequencies of these sequences are multiplied by probabilities less than one) and to increase the contribution of event sequences that were assumed to lead to no SCD (the probability of SCD is now a small but non-zero number).

Several important benefits are obtained from this methodology. First, a substantially more accurate estimate of SCD frequency is obtained. Second, this framework permits mechanistic acceptance criteria to be derived based on safety goals that establish overall plant risk. To meet safety goals, we must establish acceptance criteria so that accident sequences with relatively high frequencies of occurrence will have relatively low probabilities of SCD. Accidents with relatively low frequencies of occurrence will require less stringent criteria. Third, the impact of the uncertainties in the mechanistic analysis on SCD frequency can be evaluated, and research can be focused on those uncertainties with the largest impact.

A fundamental problem for reactor safety analysis has been the derivation of acceptance criteria for specific accidents. Heretofore, there has been no rational basis for requiring a particular margin between an operating parameter, such as peak clad temperature, and its damage threshold, such as 2200°F. It has been recognized for a long time that there are uncertainties in calculating the parameter itself and uncertainties in its damage threshold. Adoption of the BE

methodology has yielded a quantification of these uncertainties, but it provides no guidance about how closely the calculated BE value of the parameter should be allowed to approach the BE of the damage threshold. Because the methodology yields distribution functions for both of these quantities, the question becomes a statistical one relating to the overlap (or convolution) of these two distributions. The mean value of the operating parameter must be kept sufficiently below the mean value of the damage threshold so that the probability of the operating parameter exceeding the damage threshold is sufficiently small. We require a precise limit for the probability that the damage threshold is exceeded at a specified confidence level.

The NRC, in attempting to answer the question "How safe is safe enough?" has established safety goals. For example, a typical goal is 10^{-5} per reactor-year on the overall frequency of SCD. Clearly it is impossible to obtain from this one safety goal *unique* acceptance criteria for all of the very large number of events that can be postulated to occur in a reactor system. What we have is a limit on the sum of a very large number of terms, each term representing a positive contribution from one particular event sequence to the overall SCD frequency. This sum can be kept below the limit only by keeping each term in the sum substantially below the limit. This can be done only through an allocation of a small portion of the total risk of SCD to each and every postulated initiating event and, further, to every event sequence following from this initiating event. If it proves difficult to design to meet the allocated risk limit for LOCAs or earthquakes, then the designer can allocate a larger portion of the risk to such accidents, provided that the risk from other accidents is sufficiently reduced so as to meet the overall safety goal.

This paper extends the work of Martz et al[6] and describes a two-step procedure that can be used to determine Bayesian risk-based acceptance criteria. The two-steps are:

Step 1. the development of a new method based on Bayes theorem that can be used to allocate a probabilistic safety goal to the mechanistic probabilities that the corresponding operating parameters will exceed their respective thresholds; and

Step 2. the development of corresponding Bayesian risk-based acceptance criteria that can be used to set appropriate conditions on the mechanistic BE analysis so that the probabilities in (1) are not exceeded at a specified credibility (confidence) level.

These two steps are described in Sections 2 and 3, respectively, and an example is presented in Section 4.

2 STEP 1: SAFETY GOAL ALLOCATION

Traditionally, acceptance criteria are bounding or limiting values imposed on reactor operating parameters for the purpose of ensuring that minimum acceptable safety margins are maintained. Here the term "safety margin" is used to mean the difference in magnitude between some point estimate of an operating parameter and its associated damage threshold, referred to as X and Y, respectively. Examples of traditional acceptance criteria are heat flux limits that provide assurance that critical heat flux is not exceeded and fuel cladding temperature limits that prevent structural damage to the fuel. In the traditional design process, development of appropriate acceptance criteria involved an identification of those operating parameters that should be maintained within specified limits to ensure that safety objectives were met. For each operating parameter identified, a minimum acceptable safety margin was established.

As discussed in Section 1, there are two generally accepted methodologies for mechanistic analysis of core and system performance: the conservative approach and the BE approach. In light water reactor applications, the first method is also called the *evaluation model approach* and is described in 10CFR50.46[2] and in Appendix K to 10CFR50[1]. The intent of this method is to calculate (in a conservative manner) point estimates for each operating parameter and its associated damage threshold for each event sequence (or for a limiting event sequence that would represent the worst consequences within a group of event sequences). The degree of conservatism in the calculation of the operating parameter and its damage threshold is not usually quantified. This methodology enjoyed wide application in the past and led to the use of traditional acceptance criteria. For a purely conservative approach, the probability of SCD, P, becomes a simple step function that is 0 for X<Y and 1 for X>Y.

In Section 1 we discussed the evolution of conservative methods to more modern BE techniques. In the general case, when the BE approach is used for a particular event sequence, uncertainties in all parameters are propagated through the model. This method yields probability distributions, rather than just point estimates, for all operating parameters and their damage thresholds. More restrictive applications of this method are currently in use. Regulatory Guide 1.157[7] considers only uncertainties associated with the operating parameters (e.g., peak clad fuel temperature) for the specific case of ECCS analysis of light water reactors. NUREG-0800[3] considers only uncertainties associated with the damage thresholds (e.g., minimum DNBR) for the specific case of analysis of Anticipated Operational Occurrences (AOOs) for light water reactors.

In this paper we employ a more general BE approach that considers uncertainties in both the operating parameter as well as its associated damage threshold. Of course, the case in which there is no uncertainty in the threshold is

a special case of the more general method we consider. For a BE approach, the probability of SCD can be represented as a family of curves dependent not only upon the mean point estimates of the operating parameter and its damage threshold but also on the statistical variation of these parameters about their means.

In this section, we propose a method for allocating a prespecified overall probabilistic safety goal to the conditional mechanistic probability of SCD [the overlap probability; i.e., P = Prob(X > Y)] for each accident sequence. To keep things simple, consider a specific class of initiating events, such as small break LOCAs, loss of offsite power (LOOP) events, or nonspecific reactor trips (general transients). Let C_i, i = 1, ..., m, denote the various initiating event classes (or event trees) where m is the total number of classes of initiating events considered. For example, class C_1 might denote the class of LOOP-initiated accidents, C_2 the class of small-break LOCAs, and so forth. Also, let S_{ij}, j = 1, ..., n_i, denote the jth accident sequence within class C_i, where n_i denotes the total number of accident sequences considered in initiating event class C_i.

Martz et al[6] suggest the definition and use of a fault tree logic model to relate the PRA-based accident sequence frequencies (based solely on safety and support system response to the initiating event) and overlap probabilities given the system state (or system conditions) defined by each accident sequence. For initiating event class C_i, suppose that there are n_i accident sequences of interest for which mechanistic analysis (and the associated acceptance criteria) is under consideration. Further suppose that there are k_{ij} mechanistic operating parameters of interest and which are appropriate for accident sequence S_{ij}, j = 1, ..., n_i; i = 1, ..., m. Alternatively, k_{ij} may represent different competing models or computer codes used to calculate fewer than k_{ij} operating parameters associated with sequence S_{ij}. The fault tree which relates

the PRA-based accident frequencies and mechanistically-produced overlap probabilities is a direct extension of the SMAET considered by Martz et al[6]; thus, we denote this tree as a *system-mechanistic analysis fault tree* (SMAFT).

Figure 1 illustrates the SMAFT, where "AC" denotes *acceptance criterion.* Note that the top event is a mutually exclusive "OR" gate, labeled "M". Note also that different sequences may have a different number of acceptance criteria. Also, note that for those sequences which are either deemed risk-insignificant or for which the probability of SCD is either 0 or 1 with certainty, k_{ij} will be 0 (i.e., for such sequences there is no corresponding mechanistic BE predicted SCD node in the SMAFT). The "basic events" labeled "S_{ij} occurs" in Figure 1 simply denote the occurrence of accident sequence S_{ij}, as delineated in the SMAET. These "basic events" have frequencies which have been calculated using standard PRA methods for event tree sequences. In the case where k_{ij} represents different models or computer codes used to calculate fewer than k_{ij} operating parameters associated with sequence S_{ij}, the corresponding lower-most gates in Figure 1 may each represent q_{ij}-out-of-k_{ij} voting logic; i.e., $1 \leq q_{ij} \leq k_{ij}$ of the models must predict SCD in order that mechanistic BE predicted SCD given S_{ij} occurs.

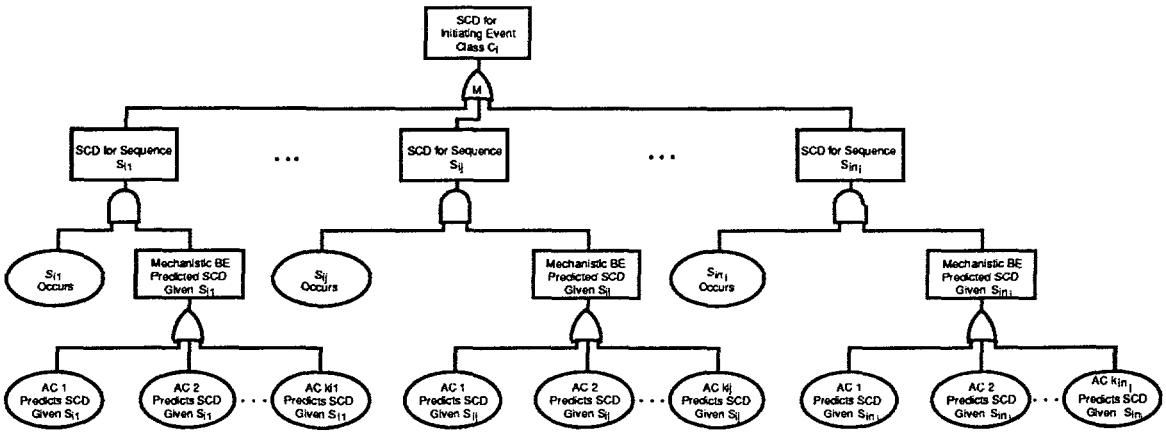

Figure 1. System-Mechanistic Analysis Fault Tree (SMAFT)

Now suppose that an overall safety goal has been allocated to each initiating event class and consider the top-down apportionment of this goal for event class C_i to the individual acceptance criteria at the bottom of Figure 1. Martz and Almond[8] present a top-down method based on the use of Bayes' theorem for allocating higher-level data in a fault tree to the basic events, which they refer to as *back-estimation*. We use this same method here for apportioning the goal for event class C_i to the acceptance criteria.

For simplicity, denote the basic events in the SMAFT as Bj; thus, j is a basic event index. Further, denote the top event in the SMAFT as Si; thus, as before, i is an initiating event class index. Let $\overline{Si}$ denote the event that Si does not occur. The Martz and Almond[8] method requires consideration of the probability of Si rather than its frequency. We can consider probability rather than frequency by conditioning on an occurrence of the initiating event and considering the conditional probability of occurrence of each sequence given that an initiating event has occurred. This effectively removes the initiating event as the first event in the corresponding event tree. Thus, let $p_{Si} = Si(\mathbf{p})$ denote the conditional probability (or approximate probability) of Si (given that an initiating event has occurred) where **p** denotes the vector of basic event probabilities with elements p_j.

In order to assess the influence (or importance) of each basic event to the occurrence of Si, we use the well-known *risk increase ratio* (RIR[9]).). RIR is an indication of how much the probability of Si would increase if the basic event occurred with probability 1.0. In equation form, based on **p**, RIR for basic event Bj is given by

$$RIR_j = P(Si \mid Bj) / P(Si) = Si(\mathbf{p} \mid 1_j) / Si(\mathbf{p}), \qquad (1)$$

where $\mathbf{p} \mid 1_j$ denotes the vector **p** in which the probability of basic event Bj is set to 1.0.

Prior to considering the goal for class C_i, we have $RIR_j = P(Si \mid Bj) / P(Si)$, $P(Si \mid Bj) = RIR_j\, P(Si)$, $P(\overline{Si} \mid Bj) = 1 - RIR_j\, P(Si)$, and $P(\overline{Si}) = 1 - P(Si)$. Thus, $P(Si, Bj) = P(Si \mid Bj)\, P(Bj) = RIR_j\, P(Si)\, P(Bj)$. Now define $I_j \equiv P(Bj \mid \overline{Si}) = RIR_j\, P(Bj)$. Similarly, $P(\overline{Si}, Bj) = [1 - RIR_j\, P(Si)]\, P(Bj)$ from which we define $I_j' \equiv P(Bj \mid \overline{Si}) = [1 - RIR_j\, P(Si)]\, P(Bj) / [1 - P(Si)]$. Using the law of total probability, we now consider I_j and I_j' as weights which we apply to the goal probability of Si, which we denote by p_{Si}^+, and its complement as follows:

$$\begin{aligned} p_j^+ &= I_j\, p_{Si}^+ + I_j' \{1 - p_{Si}^+\} = (I_j - I_j')\, p_{Si}^+ + I_j' \\ &= p_j\, (a_j + b_j\, p_{Si}^+), \end{aligned} \tag{2}$$

where we have defined

$$a_j \equiv (1 - p_{Si}\, RIR_j) / (1 - p_{Si}), \qquad b_j \equiv RIR_j - a_j. \tag{3}$$

In (2), p_j^+ denotes the portion of the overall goal p_{Si}^+ which is allocated to the basic event Bj on the basis of its importance to Si. Let $\mathbf{p}^+$ denote the vector of apportioned basic event probabilities whose elements are given by (2).

It is interesting and informative to examine the performance of (2) for three bounding cases of particular interest. First, if basic event Bj does not appear in any minimal cut set for Si, then $RIR_j = 1.0$. In this case, we see from (2) and (3) that $p_j^+ = p_j$; that is, p_j is unaffected by the goal. Second, if Bj is a minimal cut set for Si, then $RIR_j = 1.0/p_{Si}$ and, again using (2) and (3), we have $p_j^+ = p_j\, p_{Si}^+ / p_{Si}$. In this case, p_j is modified proportional to the ratio of the goal probability of Si to the

approximate probability of Si predicted using **p**. Third, if $p_{Si}^{+} = p_{Si}$, then we again see that $p_j^{+} = p_j$; that is, there is no change in p_j.

The goal may not be apportioned exactly right when using (2) and (3) in the sense that $Si(\mathbf{p}^{+})$ may not be equal p_{Si}^{+}. The reason for this is that the importances RIR_j calculated from the initial vector **p** may not agree with those calculated using $\mathbf{p}^{+}$. However, this shortcoming can be rectified by iterating the above procedure until convergence occurs. For example, a second iteration of the above procedure would begin by using $\mathbf{p}^{+(1)}$ ($\mathbf{p}^{+}$ from the first iteration) to calculate "updated" RIR_j values using (3). These updated RIR_j values are then used to recalculate $\mathbf{p}^{+}$ again using (2). Denote the value of $\mathbf{p}^{+}$ on the second iteration as $\mathbf{p}^{+(2)}$; in general, denote the nth iterated value of $\mathbf{p}^{+}$ as $\mathbf{p}^{+(n)}$. We continue this iterative process until $Si(\mathbf{p}^{+(n)})$ is sufficiently close to p_{Si}^{+}; that is, until convergence is achieved. We will illustrate the performance of this iterative procedure in Section 4.

3 STEP 2: BAYESIAN ACCEPTANCE CRITERIA

The recent development of BE methods requires a more general approach to the specification of acceptance criteria. Because we now have distribution functions for both operating parameters and damage thresholds, it is not sufficient to require simply that an operating parameter be less than its damage threshold. To develop a general concept of acceptance criteria that is consistent with the use of BE analysis, we reexamine the basic definition of these criteria. As discussed earlier, acceptance criteria have traditionally placed limits on the operating parameters so as to preserve, in some imprecise sense, an adequate margin of reactor safety. In the context of BE methods, we can make this concept more precise. For illustrative purposes, let us assume (as discussed in Section 1)

that an overall SCD frequency (or probability) goal has been established and that this goal has been allocated among all classes of events and event sequences. Using the allocation procedure described in Section 2, we can obtain the maximum allowable conditional probability of SCD for any particular sequence of interest consistent with the goal for this initiating event class. We must therefore require that the mechanistic BE analysis produce a conditional probability of SCD that is no greater than this allowable probability. This limit on the mechanistic BE analysis restricts the amount of overlap between the distribution functions of the operating parameter and its damage threshold. Thus, this maximum allowable conditional probability of SCD serves precisely the same role in BE calculations that traditional acceptance criteria serve for conservative methods. For this reason, we define the acceptance criterion for BE analysis of a particular event sequence as the maximum allowable probability of SCD obtained by an allocation of the overall SCD probability goal. This concept of acceptance criteria as probabilities, rather than as absolute (dimensional) limits on operating parameters, is a natural extension of this concept to the BE methodology.

We note that this new concept of acceptance criteria also gives rise to a new and more precise definition of safety margin. Since the acceptance criterion is now the maximum allowed conditional probability of SCD, its complement is in some sense a measure of the safety of the system. We will define the complement of an acceptance criterion as a *probabilistic safety margin* for a particular event sequence, and later we will see how a dimensional safety margin can be obtained once an acceptance criterion is stated. For all event sequences analyzed by the BE approach, second-level uncertainties will typically exist when estimates are computed for their conditional probabilities of SCD. Therefore, confidence levels must also be given in order to fully specify these acceptance criteria (or safety margins).

Following Martz et al[6] we assume that both the operating parameter X and associated threshold Y follow a Gaussian distribution with true (but unknown) means μ_x and μ_y and variances σ_x^2 and σ_y^2, respectively. Also, we assume that the random variables X and Y are statistically independent. Thus, as discussed in Martz et al[6], we find that

$$P = \text{Prob}(X>Y) = 1 - \Phi\left(\frac{\mu_y - \mu_x}{\sqrt{\sigma_y^2 + \sigma_x^2}}\right), \tag{4}$$

where $\Phi(\bullet)$ denotes the usual standard Gaussian cumulative distribution function (cdf).

In practice, the means and variances in (4) are unknown and estimates are required, and these estimates produce uncertainty about P. Here we assume that the mechanistic BE analysis produces a sample mean (or average) $\bar{x}$ and sample standard deviation s_x of the appropriate reactor operating parameter based on n observations (runs) of the appropriate computer code used in the mechanistic BE analysis. For example, the extreme value of the peak clad temperature might be one such parameter where the peak value corresponds to some appropriately defined location or regional average within the neutronic core. In addition, we assume that the corresponding damage threshold analysis for SCD also produces a sample mean $\bar{y}$ and sample standard deviation s_y based on m observations. For the example of peak clad temperature, the extreme value of interest occurs at the point in time where $(\bar{y}-\bar{x})$ is at its minimum. It is also assumed here that the sample variances s_x^2 and s_y^2 have been computed using a factor (n-1) and (m-1) in the denominators, respectively. This results in unbiased estimates of the corresponding population variances.

If the Gaussian distributions of X and Y were precisely known (i.e., if μ_x, μ_y, σ_x, and σ_y were known with certainty), then we could simply compute P using (4), and there would be no second level of uncertainty regarding the probability of SCD. However, we do not know the true values of μ_x, μ_y, σ_x, and σ_y. Because of uncertainty regarding the true values of μ_x, μ_y, σ_x, and σ_y, the true overlap probability P in (4) is thus also uncertain (the second level of uncertainty). We desire the distribution of P and use a parametric Bayesian bootstrap approach for determining this distribution.

First, consider the operating parameter X. Suppose that we consider a joint noninformative Jeffreys prior distribution for (μ_x, σ_x); namely, $\pi(\mu_x, \sigma_x) \propto \sigma_x^{-1}$. Box and Tiao[10] show that the corresponding marginal posterior distribution of σ_x is an inverted-gamma 2 distribution with density function given by

$$\pi(\sigma_x \mid s_x) = \left[0.5\,\Gamma(\nu/2)\right]^{-1}\left(\nu s_x^2/2\right)^{\nu/2}\sigma_x^{-(\nu+1)}\exp\left(-\nu s_x^2/2\sigma_x^2\right), \tag{5}$$

where $\nu = n-1$. The corresponding marginal posterior distribution of μ_x has density function given by

$$\pi(\mu_x \mid \bar{x}, s_x) = \frac{\Gamma\left(\frac{\nu}{2}+0.5\right)(s_x/\sqrt{n})^{-1}}{\Gamma\left(\frac{\nu}{2}\right)(\pi\nu)^{0.5}}\left[1+\frac{n(\mu_x-\bar{x})^2}{\nu s_x^2}\right]^{-(\nu+1)/2}. \tag{6}$$

Using (5), it can be shown that $w = \nu s_x^2/(2\sigma_x^2)$ has a gamma distribution with shape parameter $\nu/2$ and scale parameter 1. Thus, to simulate a value of σ_x^2, we first simulate a value of w, after which we then calculate $\sigma_x^2 = \nu s_x^2/(2w)$. From (6), it can also be shown that $t = \sqrt{n}(\mu_x - \bar{x})/s_x$ has a t-distribution with ν degrees of

freedom. To simulate a value of μ_x, we first simulate a value of t, after which we then compute $\mu_x = \bar{x} + s_x t/\sqrt{n}$.

Provided the random number generator permits conditional simulated values, there is an even simpler way to simulate a value of μ_x. We recognize that the conditional posterior distribution of μ_x given σ_x is a Gaussian distribution with mean $\bar{x}$ and variance σ_x^2/n. Thus, after first simulating a value of σ_x^2, we then simulate a value of μ_x from this Gaussian distribution conditional on the given simulated value of σ_x^2. In a completely analogous manner, the marginal posterior distributions of σ_y and μ_y are also given by (5) and (6), except that now m replaces n and y replaces x in (6) and (7); thus, v = m-1.

Given $\bar{x}$, s_x, $\bar{y}$, and s_y, simulated values of μ_x, μ_y, σ_x^2, and σ_y^2 are then successively used in a Monte Carlo fashion to determine the simulated (approximate) distribution of P in (4). Upper percentiles of this distribution, such as the 90th and 95th percentiles, provide the required Bayesian credibility (or confidence) that P will not exceed these quantiles at the stated confidence level. We illustrate how these upper percentiles can be used to determine corresponding dimensional limits on the operating parameter X (in the form of limits on $\bar{x}$) in the next section.

To illustrate this process, consider the following example. Suppose that the peak fuel clad temperature of a certain light water reactor is of interest. Further suppose that for a large pipe break LOCA the reactor responds as designed. For this event sequence (i.e., for a specified initiating event and plant safety system response), suppose further that the corresponding mechanistic T-H analysis yields $\bar{x}$ = 1500°F and s_x = 75°F based on n = 10 Latin hypercube sampling (LHS) runs. Further, suppose that the mean fuel damage threshold is calculated to be $\bar{y}$ = 2200°F with s_y = 50°F based on m = 5 observations. Implementing the above Monte Carlo procedure in Excel[11] using Crystal Ball[12],

the mean and median of P in (4) are computed (based on 10,000 Monte Carlo replications) to be 7.6E-5 and 1.9E-4, respectively. The standard deviation is also calculated to be 1.8E-3, while the upper 90th and 95th percentiles of the distribution of P are 3.4E-4 and 3.6E-4, respectively. Thus, for example, we can claim that P is less than 3.6E-4 with 95% probability (or Bayesian credibility).

4 EXAMPLE

Consider the proposed nuclear reactor facility example first presented by Martz et al[6], and consider the class of general transient events, C_4. The system portion of the SMAET transient event tree has 10 systems and 12 accident sequences, labeled here as S_{41} through $S_{4,12}$. For simplicity, only a single mechanistic criterion is considered in this example. Based upon neutronic core similarities between the proposed reactor facility and an existing facility, the calculated operating heat flux (OHF), Φ_{OHF}, at its spatially peaked location was chosen as the operating parameter X. The corresponding damage threshold Y is the onset of (ledinegg) flow instability (OFI), which we express in terms of a heat flux limit (or damage threshold), Φ_{OFI}.

First consider step 1: safety goal allocation. Based on similar reactor designs and corresponding PRA analysis, Table 1 gives the estimated conditional probabilities for each of the 12 transient accident sequences, given a general transient initiating event.

Table 1. Transient Accident Sequence Conditional Probabilities of Occurrence

Sequence	Probability	Sequence	Probability
S_{41}	9.99E-1	S_{47}	3.26E-8
S_{42}	5.36E-4	S_{48}	1.79E-11
S_{43}	1.39E-6	S_{49}	9.87E-15
S_{44}	1.25E-5	$S_{4,10}$	1.90E-9
S_{45}	6.88E-9	$S_{4,11}$	7.71E-8
S_{46}	6.89E-9	$S_{4,12}$	5.05E-9

To begin the allocation process, suppose that an overall safety goal on the total frequency of internally initiated events of 1.0E-6 per reactor-year (R-Y) is to be equally allocated to four classes of events: large LOCAs, small LOCAs, station blackout, and general transients. Thus, the portion of the overall frequency goal to be allocated to transient events is 2.5E-7 per R-Y. The expected frequency of general transients at the proposed facility is estimated to be 3 events per R-Y. Thus, the overall expected conditional probability goal given a transient event is 2.5E-7/3 = 8.33E-8, and it is this goal which we wish to allocate to the corresponding 12 mechanistic basic events (1 for each accident sequence) in Figure 1.

The Bayesian allocation procedure in Section 2 requires initial estimates of the expected value of P for each mechanistic BE acceptance criterion (i.e., an estimate of the expected value of P for each accident sequence). Based on transient accident analysis of accident sequences similar to S_{41} and S_{42} performed on an existing reactor facility[13], means and standard deviations for these sequences were computed to be $\bar{x}$ = 100, s_x = 10 ,n = 8, and $\bar{y}$ = 300, s_y=48, and m = 30. Using these values in the Monte Carlo simulation of Section 3

produces mean and median values of P of 1.9E-4 and 4.3E-5, respectively. The standard deviation of P is also found to be 5.6E-4, while the 90th and 95th percentiles of P are 4.5E-4 and 8.5E-4. Thus, we initially consider the expected value of P for each of the 12 sequences to be 1.0E-4 and, using this data in conjunction with the data in Table 1, Figure 2 gives the initial SMAFT for this example. Note that, for legibility, gates corresponding to only 4 of the 12 sequences are shown. The minimal cut set form of the Boolean equation for the top event in Figure 2 is $\Psi(\mathbf{x}) = x_1x_2 + x_3x_4 + \ldots + x_{23}x_{24}$, where x_i is a Boolean variable corresponding to each successive basic event in Figure 2. Because the uppermost gate is a mutually exclusive "OR" gate, the probability of the top event is simply $p_{S4} = S4(\mathbf{p}) = p_1p_2 + p_3p_4 + \ldots + p_{23}p_{24}$.

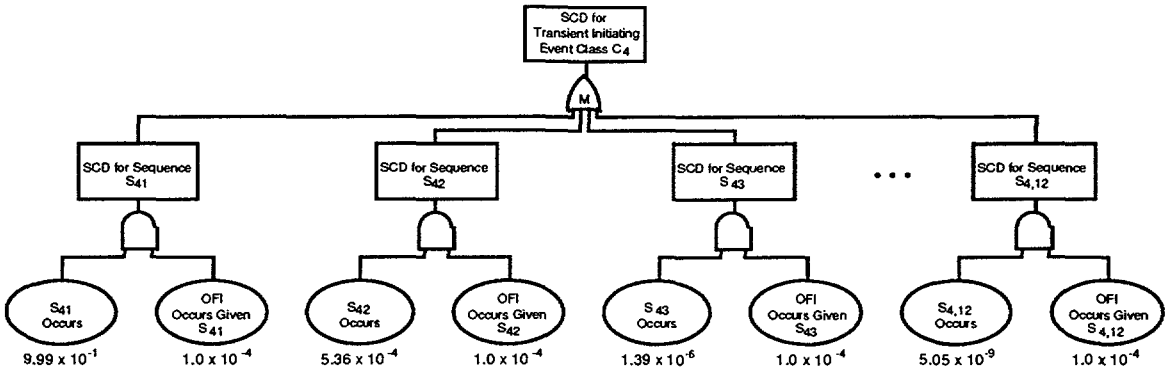

Figure 2. Initial SMAFT for the Transient Initiating Event Class C_4

We now apply the iterative allocation scheme described in Section 2 to determine the allocation of the conditional probability goal 8.33E-8 to the mechanistic basic events in Figure 2. The procedure converges to 8.33E-8 after 17 iterations, and the iterative results are shown in Figure 3.

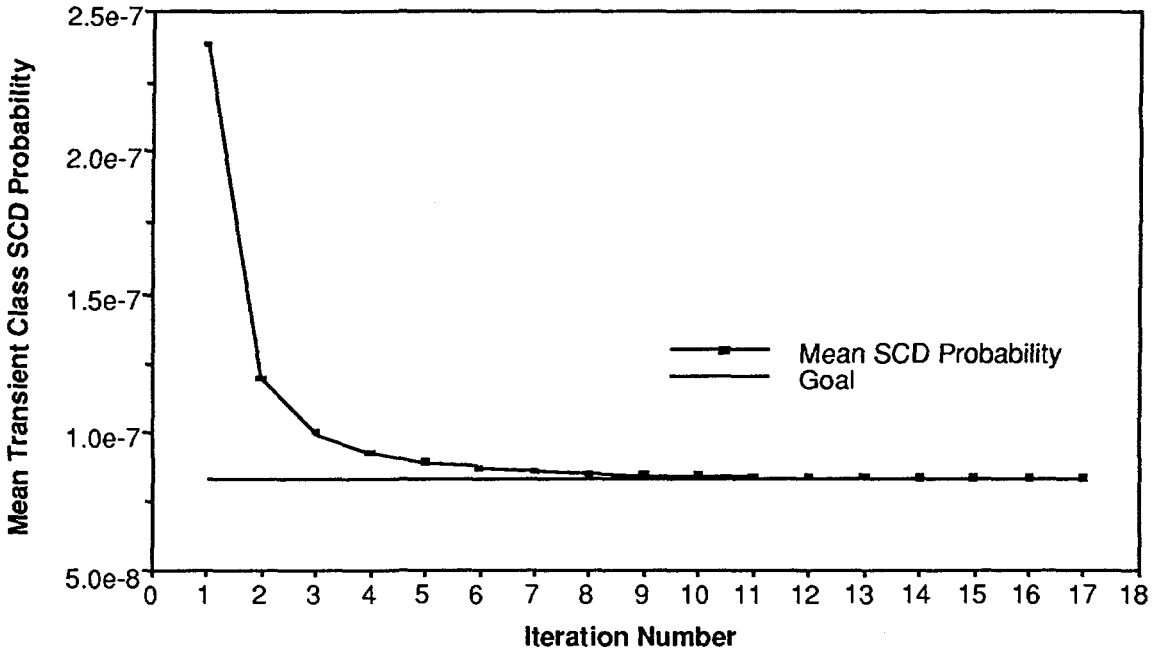

Figure 3. Convergence of the Iterative Safety Goal Allocation Procedure

The final allocation (after 17 iterations) of the 8.33E-8 safety goal to each of the mechanistic basic events is given in Table 2.

Table 2. Safety Goal Allocation to the Mechanistic BE Analysis Associated with Each Transient Event Sequence

Sequence	Allocation	Sequence	Allocation
S_{41}	2.84E-8	S_{47}	1.00E-4
S_{42}	9.99E-5	S_{48}	1.00E-4
S_{43}	9.99E-5	S_{49}	1.00E-4
S_{44}	9.99E-5	$S_{4,10}$	1.00E-4
S_{45}	1.00E-4	$S_{4,11}$	1.00E-4
S_{46}	1.00E-4	$S_{4,12}$	1.00E-4

By comparing Tables 1 and 2, we immediately see the tradeoffs between the availability/unavailability of the safety systems and corresponding conditional

mechanistic BE probabilities of SCD. For example, for sequence S_{41} the mean probability that the sequence occurs is 0.999; thus, the corresponding mechanistic BE probability allocation for this sequence must not exceed 2.84E-8. On the other hand, the mean conditional probability that sequence S_{49} occurs is already extremely small at 9.87E-15; thus, the mean mechanistic BE probability of SCD for this sequence is unchanged from its initial value of 1.0E-4.

Now consider step 2: Bayesian acceptance criteria. For the case in which n = 8, s_x is equal 10% of $\bar{x}$, $\bar{y}$ = 300, s_y = 48, and m = 30, we estimate the distribution of P in (4) as a function of $\bar{x}$ using Monte Carlo simulation with 10,000 replications. Figure 4 gives the median mean, 75th, 90th, 95th, and 97.5th percentiles of the distribution of P as a function of $\bar{x}$. In addition, Figure 4 also contains the maximum observed value of P in 10,000 replications for each value of $\bar{x}$ considered in the simulation (the curve labeled 100%). Note that, because of poor reproduction of this figure, the median curve is at the bottom, the 75th percentile is the second curve up from the bottom, and so on. Note also that the mean, which crosses the 75th percentile curve, is dashed.

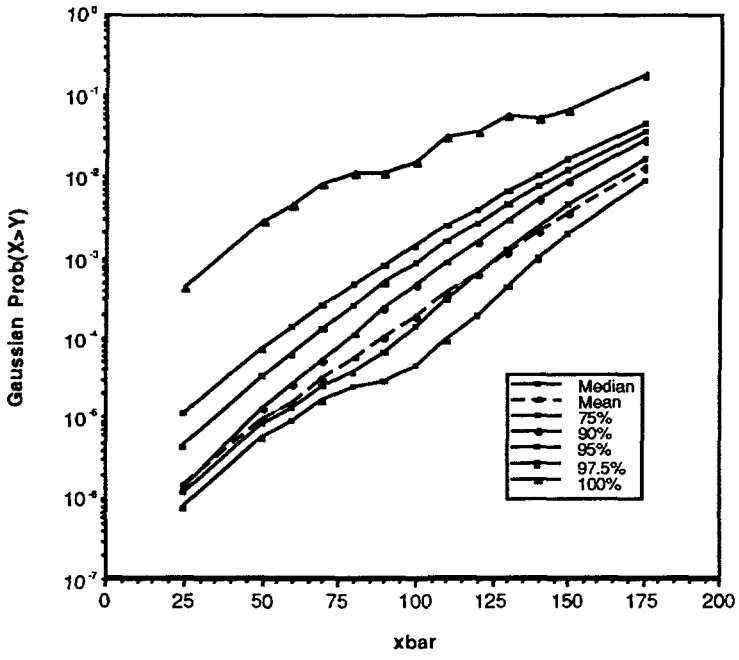

Figure 4. Mean and Selected Percentiles of the Distribution of P

Figure 4 can be used to determine dimensional acceptance criteria on Φ_{OHF} for the probability P at any desired confidence level. For example, for transient accident sequence S_{42}, from Table 2 we see that the portion of the probabilistic safety goal that has been allocated to the corresponding mechanistic BE analysis for this sequence is 9.99E-5. Suppose we forego our interpretation of this as a mean value and desire a corresponding acceptance criterion for this value of P with at least 95% confidence. In this case, we find from Figure 4 that $\bar{x}$ must not exceed 66 for the mechanistic BE analysis for this sequence. If this turns out to be the case, then we can claim that P = 9.99E-5 has been accepted with at least 95% confidence (Bayesian credibility).

ACKNOWLEDGMENTS

This work was supported by the U.S. Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, under Job Code Number W6505 (Dr. Lee R. Abramson, NRC Technical Monitor) at Los Alamos National Laboratory. We are grateful for this support.

REFERENCES

1. Title 10, Part 50, Appendix K, "Emergency Core Cooling Systems Evaluation Models," U.S. Code of Federal Regulations. Washington, DC, 1990.

2. Title 10, Part 50.46, "Acceptance Criteria for Emergency Core Cooling Systems for Light Water Nuclear Power Reactors," U.S. Code of Federal Regulations. Washington, DC, 1990.

3. U.S. Nuclear Regulatory Commission, Standard Review Plan for the Review of Safety Analysis Reports for Nuclear Power Plants, NUREG-0800. Washington, DC, July 1981.

4. McKay, M., Conover, W. & Beckman, R., A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, **21** (1979), 239-245.

5. Iman, R. & Helton, J., A comparison of uncertainty and sensitivity analysis techniques for computer models. NUREG/CR-3904 (SAND84-1461), Sandia National Laboratories, March 1985.

6. Martz, H. F., Hamm, L. L., Reed, W. H., & Pan, P. Y., Combining mechanistic best-estimate analysis and Level 1 probabilistic risk assessment. *Reliability Engineering and System Safety*, **39** (1993), 89-108.

7. U.S. Nuclear Regulatory Commission, Best-Estimate Calculations of Emergency Core Cooling System Performance, Regulatory Guide 1.157. Washington, DC, May 1989.

8. Martz, H. F. & Almond, R. G., Using higher-level failure data in fault tree quantification, LA-UR-96-1100, Los Alamos National Laboratory, March 1996.

9. Russell, K. D., Atwood, C. L., Galyean, W. J., Sattison, M. B. & Rasmuson, D. M., Systems Analysis Programs for Hands-on Integrated Reliability Evaluations (SAPHIRE) version 5.0, vol. 1 - technical reference manual.

EGG-2716 (NUREG/CR-6116), Idaho National Engineering Laboratory, July 1994.

10. Box, G. E. P. & Tiao, G. C., *Bayesian Inference in Statistical Analysis.* Addison-Wesley, Reading, Massachusetts, 1973.

11. *Microsoft Excel, Version 5.0.* Microsoft Corporation, 1994.

12. *Crystal Ball, Version 3.0.3.* Decisioneering, Inc., 1993.

13. Hamm, L. L., Aleman, S. E., and Laurinat, J. E., Effects of heated wall, wall voidage, and dissolved gases on flow instability limits. WSRC-RP-89-406, Westinghouse Savannah River Company, October 1989.