NUREG-XXXXX

A Framework for Integrating Risk and Safety Margins

Mirela Gavrilas, RES/DSARP/NRCA

CONTENTS

LIST OF FIGURES

LIST OF TABLES

GLOSSARY

The **safety variable** indicates the onset of damage for a component or a system. Examples are peak clad temperature, maximum cladding oxidation, and containment pressure.

The **load** (also called challenge) is the probability density function that describes the values of the safety variables experienced by the system or component during duty. An example is the probability density function of peak clad temperature as generated by a best estimate plus uncertainty deterministic calculation for a specific event sequence.

The **capacity** (also called resistance) describes how likely the component/system is to fail as a function of the safety variable value. There is an assumption that the load probability density function is obtained by testing a sufficient number of components/systems to failure. An example of a capacity probability distribution function is the containment fragility curve.

Aleatory uncertainty (also known as stochastic uncertainty) is inherent in a physical process that causes outcomes to be randomly distributed. For example, the temperature of a water storage tank that is exposed to ambient conditions has an aleatory uncertainty associated with it. Aleatory uncertainty cannot be reduced through further experimentation

Epistemic uncertainty is associated with limited knowledge. Epistemic uncertainty includes "unknown unknowns" as well as uncertainties due to incomplete information. An example of an epistemic uncertainty is that associated with a heat transfer coefficient correlation that was determined with limited experimental data. Epistemic uncertainty can be reduced with further experimentation.

The **safety margin** is the distance between the bounding prediction of the load and the point at which failure becomes non-negligible on the capacity probability density function. In the nuclear industry, safety margins are ensured by setting conservative safety limits and keeping operating conditions below the safety limit by an amount that is at least commensurate with the uncertainty of the load.

The **safety limit** is imposed on a safety variable to ensure adequate safety margin. The safety limit is set conservatively under the onset of damage on the capacity curve. When the load probability distribution function stays under the safety limit by all but a negligible amount, adequate safety margin exists.

The **conditional probability of loss of function** is equivalent to the probability that the load will exceed the safety limit given that a specific event sequence takes place. This definition allows the consideration of barrier bypass.

The **risk space** is the set of event sequences that are impacted by a specific plant modification. Impact can be either to the frequency of occurrence of the event sequence, or to the safety margin in that event sequence.

The **rest of time** is a pseudo event sequence that stands for all plant states, including normal operation, that are not explicitly impacted by a given modification.

The **unconditional probability of loss of function** is the product between the conditional probability of loss of function and the probability of occurrence of the event sequence.

Consequences are the amount of radioisotopes released and/or the radiological effects of an event sequence. Consequences can be evaluated within physical barriers as well as to the public and the environment.

Risk acceptance criteria have been set to evaluate the adequacy of a modification in risk-informed regulatory decision-making. The core damage frequency and large early release frequency criteria of Regulatory Guide 1.174, "An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis" [7], and the Commission's quantitative health objectives are risk acceptance criteria.

ABBREVIATIONS

C	Celsius
CD	core damage
CDF	core damage frequency
ΔCDF	change in CDF
CS	containment spray
CSAU	code scaling, applicability, and uncertainty evaluation methodology
CSNI	Committee on the Safety of Nuclear Installations
DM	damage mechanism
ECCS	emergency core cooling system
ES	event sequence
F	Fahrenheit
ft	foot/feet
ft^2	square foot/feet
ft^3	cubic foot/feet
gpm	gallon(s) per minute
GSI	generic safety issue
IE	initiating event
lbm	pound mass
LERF	large early release frequency
ΔLERF	change in LERF
LLOCA	large-break loss-of-coolant accident
LOCA	loss-of-coolant accident
LR	loading roughness
LWR	light-water reactor
MLOCA	medium-break loss-of-coolant accident
MS	mitigation system
NPSH	net positive suction head
NRC	U.S. Nuclear Regulatory Commission
NRR	Office of Nuclear Reactor Regulation
OECD	Organization for Economic Cooperation and Development
PCT	peak clad temperature
PIRT	phenomena identification and ranking table
PRA	probabilistic risk assessment
psi	pound(s) per square inch
psig	pound(s) per square inch gauge
PTS	pressurized thermal shock
PWR	pressurized-water reactor
QHO	quantitative health objective
RES	Office of Nuclear Regulatory Research
ROT	rest of the time
SLOCA	small-break loss-of-coolant accident
SM	safety margin
SMAP	Safety Margins Action Plan
SPAR	Standardized Plant Analysis Risk
SUSA	System for Uncertainty and Sensitivity Analysis
SV	safety variable

NOMENCLATURE

Bn	barrier n
C_{FP}	concentration of fission products within a physical barrier, and
C_{0n}	initial concentration of fission products within barrier n
$C_{P\&E}$	consequences to public and environment
ES_i	event sequence i
K-loss	pressure drop due to minor losses
L	load
$\overline{L}$	mean load
$NPSH_a$	net positive suction head available
$NPSH_r$	net positive suction head required
p_{atm}	pressure head (containment pressure),
$p(f\_Bn \mid ES_i)$	conditional probability that barrier n will loose its function during event sequence i
$p(ES_i)$	probability of occurrence of event sequence i
$p(ES_i \cap f\_Bn)$	probability of occurrence of event sequence i and barrier n loss of function
p_{loss}	friction and K-loss head in the suction side, including losses at the screen
p(S>L)	reliability
p_{stat}	static suction head (sump level)
p_{vap}	vapor pressure (at maximum pumping temperature)
r_i	risk in event sequence i
S	strength
$\overline{S}$	mean strength
σ_L	load standard deviation
σ_S	strength standard deviation
t	transmission factor of radioisotopes through a breached barrier

EXECUTIVE SUMMARY

The U.S. Nuclear Regulatory Commission (NRC), Office of Nuclear Regulatory Research, has developed a framework to integrate deterministic and probabilistic information. This report details the framework for integrating risk and safety margins, and discusses the implicit formalism that results from using probabilities and consequences to compute risk metrics for decision-making.

To date, deterministic analyses and probabilistic risk assessments (PRAs) are used in a complementary manner; the analyses are distinct and separate. In the past, regulators have occasionally discussed the need to integrate risk and safety margins, and the NRC's Office of Nuclear Reactor Regulation recently articulated it in the 2005 draft Office Instruction LIC-504, "Integrated Risk-Informed Decision-Making Process for Emergent Issues." [8] Specifically, Step 3, "Assessment of impact on safety margins," in the "Template for Documenting Risk-Informed Decisions—Information Gathering and Technical Analysis," provided as Enclosure 2 to Office Instruction LIC-504, asks in Question 3.6, "Can the loss of margin be quantified in such a way as to provide input to a PRA evaluation?" This report affirmatively answers this question and provides the framework to carry out this analysis. Because the integrated risk/safety margins framework builds upon existing, established evaluation methods and techniques, its addition to the tools used for regulatory decision-making is seamless.

The framework is intended for use in quantifying changes in safety margins that result from modifications in plant design parameters and operational conditions. Examples of such plant modifications include power uprates, life extensions, mixed oxide fuels, different cladding materials, and changes to technical specifications. A cursory look at this partial list shows that some of these modifications impact safety margins in deterministic analyses, while others impact the reliability of systems and components, and yet others impact safety margins and reliability simultaneously.

Consequently, the objective of this research is to integrate risk and safety margins for the gamut of foreseeable plant modifications or combinations thereof. In very general terms, the safety margin calculated for a specific accident sequence can be translated into a probability of loss of function. This requires safety margin to be defined as having two prongs. The first prong is setting safety limits such that the probability of failure is negligible as long as operating

conditions stay within imposed limits. As is current practice, due consideration is given to unknown sources of uncertainty in setting safety limits. Deterministic regulatory practice treats exceeding the safety limit as analogous to loosing function. The second prong is ensuring that operating conditions stay within the safety limits. Known variables and uncertainties are considered by calculating operating conditions as probability density functions. The two-prong definition of safety margin is consistent with both the intent and current use of the term, as shown in Chapter 2.

For a given accident scenario, the safety margin, as defined in the glossary, gives the probability of loss of function. This conditional probability, multiplied by the probability of occurrence of the accident sequence, is the unconditional probability that the system or component will fail. The metric is obtained by aggregating unconditional probabilities of failure over all relevant accident sequences. Properly expressed, this metric—before and after a modification—is comparable to existing risk acceptance criteria, such as the change in core damage frequencies and change in large early release frequencies of Regulatory Guide 1.174, "An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis." [7] By further considering the consequences of each accident sequence, a higher level metric is obtained that is comparable to the Commission's quantitative health objectives. Chapter 3 details the framework.

Three conditions must be met simultaneously to warrant the application of the framework. The first is that the sufficiency of a safety margin becomes questionable. The second is that the subject safety margin can be reasonably tied to a loss of function. Third, an exemption is needed to justify continued plant operation. In such a case, the integrated risk/safety margins framework can be deployed to complement existing decision-making tools. Two types of safety inquiries would benefit from the application of the framework:

(1) A new phenomenon or process is identified that contributes to risk such that prior conservative deterministic analyses may not show compliance with regulations. An example of such an occurrence is the identification of the potential sump screen blockage and the subsequent loss of net positive suction head for emergency core cooling and containment spray pumps, as identified in Generic Safety Issue (GSI)-191. Chapter 4 includes a proof-of-concept demonstration of integrating risk and safety margins.

(2) A plant modification is proposed for which risk benefits are not evident because probabilities of damage and consequences are affected differently—one increases and the other decreases. An example can also be identified from among GSI-191 issues. The removal of trisodium phosphate from certain pressurized-water reactor containments to eliminate chemical effects reduces the probability of core damage but potentially increases health consequences in all accidents that have offsite releases. The integrated risk/safety margins framework provides the means of evaluating the overall safety impact of such a modification to establish the risk reduction, if any.

Integrating risk and safety margins is accomplished with tools and techniques that have reached substantial levels of maturity in the industry. For example, the process of conducting best estimate plus uncertainty analyses is fundamental to the framework. Also fundamental is the separate treatment of aleatory and epistemic uncertainty for cases in which the decision-maker requires these uncertainties to be propagated independently. This report only mentions such subsidiary tools and techniques to illustrate their role in obtaining the risk metrics required by the integrated framework. The details of individual techniques and the selection process among alternative methods are only tangentially mentioned here. This is in large part because the scope of these subsidiary techniques is quite large and in some cases still evolving; thus, the scope of this report would become impracticable if it attempted to be exhaustive. Instead, the Safety Margins Action Plan (SMAP) Task Group has prepared pertinent survey references concomitantly with the framework for integrated risk and safety margins.[1] The Committee on the Safety of Nuclear Installations of the Organization for Economic Cooperation and Development tasked the SMAP Task Group in 2002 to survey tools and techniques that can be used to quantify global plant safety margins and to provide guidance on the quantification of global margins.

The work on integrating risk and safety margins has met its objective. The framework can capture the safety relevance of any conceivable plant modification as well as any synergistic effects that may occur when several plant modifications act in concert. From a practical

[1] Note that these are the task group's working papers and thus contain views that are occasionally inconsistent or irrelevant to the integrated risk-safety margins framework.

perspective, integrating risk with safety margins offers an additional decision-making tool for instances in which uncertainty plays a dominant role.

1 Introduction and Background

The objective of this research was to develop a framework that can be used to quantify the change in plant safety margins following a broad range of plant modifications. These modifications include power uprates, changes to technical specifications, license extensions, use of mixed oxide fuel, etc. The safety margin measure devised in the course of this research also had to capture any synergistic effects that occur as two or more of these modifications are implemented. Two major constraints were specified to ensure the practical applicability of this research product. First, the framework had to build on existing, tested tools and techniques. Second, the framework had to be applicable to a problem of practical interest to the U.S. Nuclear Regulatory Commission (NRC). These constraints shaped the scope and direction of the research. However, a general approach was taken in the development of the framework to allow its future extension to new areas, such as the licensing of radically different reactors.

Most regulatory decisions are based on design-basis analyses. The rules and criteria that govern design-basis analyses were largely established during the 1970s. In recognition of the fact that some of the models involved in design-basis analyses were rather crude, ample conservatism was build into all stages and aspects of the analysis. Appendix K, "ECCS Evaluation Model," to Title 10, Part 50, "Domestic Licensing of Production and Utilization Facilities," of the *Code of Federal Regulations* (10 CFR Part 50), for example, prescribes these conservatisms. [1] The concept of safety margin, as a means of coping with uncertainty, was introduced at this point as fundamental to conducting design-basis analyses.

Between the mid-1970s and the 1990s, the industry embarked on extensive research programs to refine fundamental analytical methods. These research products were used to justify relaxing some of the most stringent requirements of Appendix K and opening the way for best estimate plus uncertainty analyses. Researchers developed proper methods for determining the best-estimate value of a safety variable and its associated uncertainty band, starting with the code scaling, applicability, and uncertainty (CSAU) evaluation methodology of NUREG/CR-5249, "Quantifying Safety Margins: Application of Code Scaling, Applicability, and Uncertainty Evaluation Methodology to a Large-Break Loss-of-Coolant Accident." [2]

Risk assessments are a long-standing practice in the nuclear industry. In the 1960s, General Electric would conduct technical risk evaluations for major products, which included nuclear

power reactors. [3] The first substantial risk study was the 1975 Reactor Safety Study, WASH-1400, which was initiated to support Congress in a decision regarding the Price Anderson Act. [4] With the Three Mile Island Unit 2 accident, risk assessments started to assume a more prominent role. The industry expanded the role of safety analyses to include different types of transients, operating procedures, and severe accidents. Risk analyses evolved as an acceptable basis for regulatory decisions, and Commission's Safety Goal Policy Statement formulated acceptance criteria. [5]

Risk analyses also began using the term "safety margin" as a general qualifier. For example, phrases like "sufficient safety margin" and "increased/decreased safety margins" are often used in relation to accidents that are not part of the design basis. "Safety margin" became a qualitative descriptor of plant safety that could be used without the burden of quantification (e.g., "release margins" or "recover margins"). Furthermore, different people have different interpretations of the term "safety margin."

The Safety Margins Action Plan (SMAP) Task Group assembled much of the history excerpted above for the Committee on the Safety of Nuclear Installations (CSNI) within the Organization for Economic Cooperation and Development (OECD). CSNI tasked the SMAP Task Group to develop consensus on a methodology that can be used to ascertain changes in safety margins induced by one or more plant modifications. The group expended a substantial amount of effort surveying existing tools and techniques to identify those most suitable for quantifying global safety margins. The framework for integrating risk and safety margins contained in this report has evolved with due consideration to the state-of-the-art practices in the area of regulatory decision-making.

To integrate risk and safety margins, Chapter 2 of this report adopts a unique, well-specified definition of safety margin. This definition is consistent with the intended use of this important safety concept as well as with the common usage of the term in current regulatory practice. As defined, the concept of safety margin links to conditional probability of loss of function. This conditional probability becomes part of a risk metric constructed using the probability of occurrence of an event sequence and, if necessary, its consequences. This risk metric is aggregated over a relevant set of event sequences. Chapter 3 covers the process of building the risk metrics. The framework is developed to eliminate conservatisms wherever the state of

knowledge permits it. This is accomplished by the explicit treatment of uncertainty, which is propagated into the final metric.

The integrated risk/safety margins framework can complement existing regulatory decision-making tools in two types of applications, which both involve cases in which uncertainty plays a significant role. The first type of application involves a scenario in which design-basis analyses cannot demonstrate sufficient safety margin. Such applications arise when previously unknown phenomena are identified. The proof-of-concept example of Chapter 4 is one such application. It involves the potential loss of net positive suction head (NPSH) in accidents that generate debris inside a pressurized-water reactor (PWR) containment. The example included in this report is developed with a simplified model that uses generic data and, thus, adds no quantitative insight into the resolution of generic safety issue (GSI)-191, Assessment of Debris Accumulation on PWR Sump Performance," as discussed in NRC Bulletin 2003-1, "Potential Impact of Debris Blockage on Emergency Sump Recirculation at Pressurized-Water Reactors." [6]

The second type of application involves a tradeoff. Two such examples can be conceptualized from processes associated with GSI-191. For instance, removing a buffering agent from the containment sump can reduce the probability of core damage due to chemical effects but can increase the radiological consequences of accidents that cause iodine releases. Another tradeoff example is enlarging the debris screen so much that the reduction in core damage frequency (CDF) due to loss of NPSH is offset by the increase due to downstream effects. In both these cases, the net change in risk caused by the modification can only be assessed with realistic calculations of individual contributors. The integrated risk/safety margins framework makes it possible to deal with each contributor to risk in a best estimate plus uncertainty manner.

Metrics used to ascertain safety margin sufficiency throughout the risk space lend themselves to using existing risk criteria for the acceptability of a margin change as a result of a plant modification. In particular, for safety margins that are subsidiary to CDF or large early release frequency (LERF), one can screen using the risk acceptance criteria set forth in Regulatory Guide 1.174, "An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis." [7] This makes the framework applicable to Step 3, "Assessment of impact on safety margins," in the "Template for Documenting Risk-

Informed Decisions—Information Gathering and Technical Analysis," provided as Enclosure 2 to Office Instruction LIC-504, "Integrated Risk-Informed Decision-Making Process for Emergent Issues," prepared by the NRC Office of Nuclear Reactor Regulation. [8] In particular, the framework applies to Question 3.6 of Step 3, which asks, "Can the loss of margin be quantified in such a way as to provide input to a probabilistic risk assessment (PRA) evaluation?"

2 Definition of Safety Margin

Although the term "safety margin" is fundamental to the nuclear regulatory framework, no universal definition exists. This does not imply ambiguity in the current use of the term, which is captured in "adequate safety margins" during normal operations and design-basis events. Nonetheless, to derive a more formal definition for the term, one must understand what the "safety margin" was intended to accomplish, and what is meant by "adequate safety margin" in its current, established usage. This chapter defines "safety margin" as it applies to a single event sequence. This event sequence can be a design-basis accident or a sequence from an event tree.

In general, the safety margin has been devised to cope with uncertainty. The challenge in quantifying margin lies in the fact that, in the nuclear industry, uncertainty must include both aleatory uncertainties (those attributable to quantities that are inherently random or stochastic) and epistemic uncertainties (those attributable to lack of knowledge, which can be reduced with the acquisition of additional data). [9][2] To further complicate matters, epistemic uncertainty may reflect limited lack of knowledge as well as complete ignorance, the "unknown-unknowns".

2.1 The Two Prongs of Safety Margin

"Adequate safety margins" are inextricably linked to safety limits—limiting values imposed on safety variables (e.g., peak clad temperature (PCT) and containment pressure). Thus, when operating conditions stay within safety limits, the barrier or system has a negligible probability of loss of function, and an adequate safety margin exists. Therefore, the first prong of ensuring adequate safety margin is to set safety limits such that the probability of loss of function is negligible, so long as operating conditions stay within those criteria. Figure 1 illustrates this

concept. The strength, S, is sometimes called capacity or resistance, and represents the probability density function obtained when the barrier is tested to failure a sufficiently large number of times.

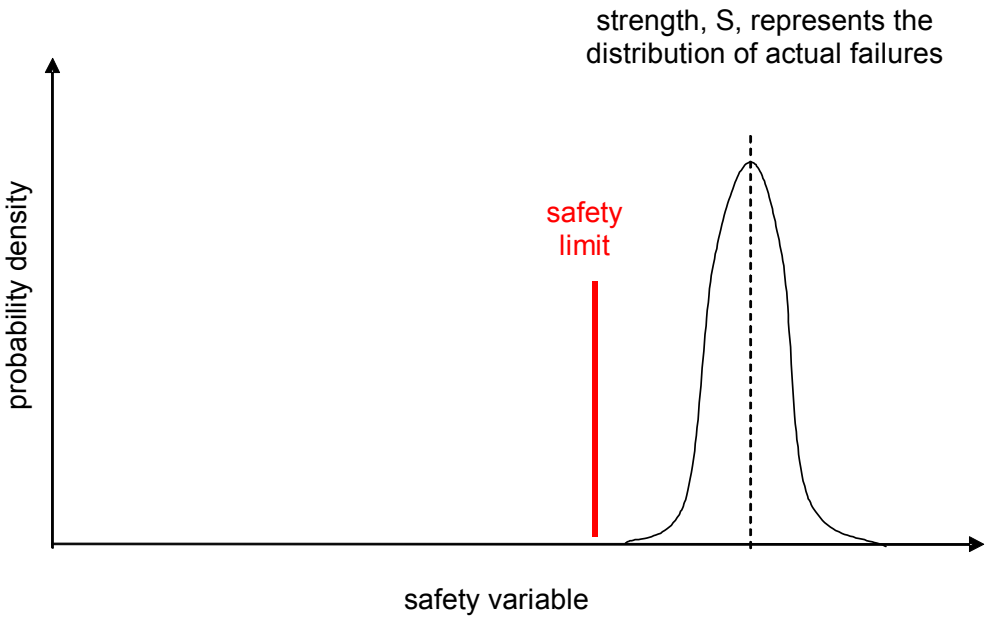

Figure 1 Setting the safety limit for a specific safety variable

One or more safety variables characterize operating conditions. For example, for the fuel barrier of a nuclear reactor, PCT and total clad oxidation are safety variables. These safety variables depend on the physical characteristics of the barrier or system being analyzed. In the case of the fuel, both PCT and clad oxidation can be measures of the embrittlement damage mechanism. In setting conservative safety limits for safety variables, the industry builds in margin for lack-of-knowledge uncertainties. The intent is to allow margin for phenomena and processes that are inadequately considered in generating models to simulate the behavior of the given system or physical barrier. Epistemic uncertainty is reflected, for example, in setting the safety limit for maximum PCT in a light-water reactor (LWR). That safety limit, 1204 °C (2200 °F), lies below the onset temperature for autocatalytic oxidation of zirconium, which, in turn, is below the point at which significant radioactive releases are expected from the fuel. Therefore, adequate

---

[2] Please refer to Reference 9 for a formal discussion of uncertainty that is consistent with the use adopted in this report.

safety margin exists if operating conditions are such that PCT remains under 1204 °C (2200 °F)[3].

The second prong of ensuring adequate safety margin is to keep operating conditions within safety limits. Figure 2 illustrates this concept. The load, L, is the probability density function obtained for the safety variable by propagating contributing uncertainties. In the computation of PCT in a specific large-break loss-of-coolant accident (LOCA) scenario, for example, uncertainties associated with boundary and initial conditions, heat transfer coefficients, and other modeling assumptions, should all be captured in the load. The 1989 CSAU method of NUREG/CR-5249 has laid the foundation for generating the probability density function associated with the load. [2] The fundamental process of identifying key phenomena and variables introduce by CSAU is essential to integrating risk and safety margins as presented in this report. Several advances have been introduced in more recent best estimate plus uncertainty methods, most notable the extension of the Gesellschaft fuer Anlagen- und Reaktorsicherheit System for Uncertainty and Sensitivity Analysis (SUSA) methodology into SUSA-AB to deal separately with epistemic and aleatory uncertainties; see also the discussion of Section 2.5. [10]

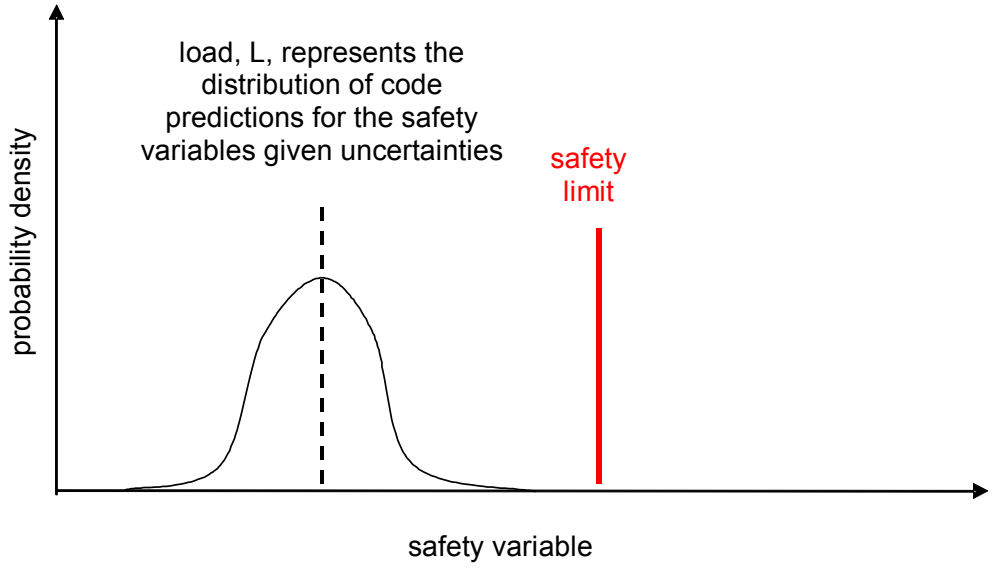

Figure 2 Keeping operating values of a specific safety variable under the safety limit

[3] PCT is one of the two safety variables used to ensure that fuel cladding does not become embrittled. The other safety variable is total clad oxidation, which has an acceptance limit of 17 percent.

The two-prong approach to safety margin discussed above is consistent with both the original intent of the framework and the more general definition of safety margin. The more general definition of safety margin was cast for structural-mechanics analyses, recognizing the fact that both load, L, and strength, S, are distributed parameters (see, for example, Reference 11). Figure 3 shows probability densities for load and strength, which form the bases for the more general definition of safety margin.

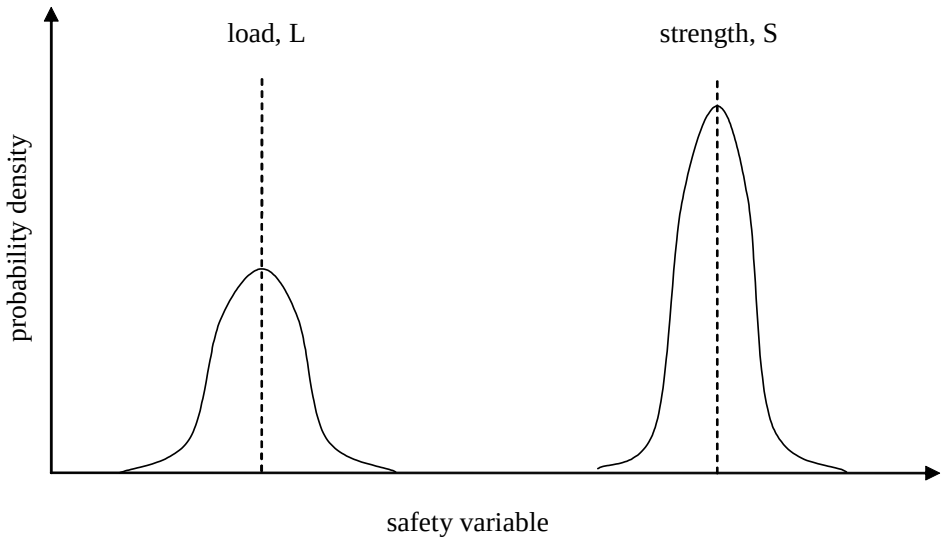

Figure 3 Probability densities for load and strength

Two quantities—namely, safety margin, SM, and loading roughness, LR—describe the reliability of a barrier or system in light of load-strength considerations. These quantities are computed from the following equations:

$$SM = \frac{\overline{S} - \overline{L}}{\sqrt{\sigma_S^{\,2} + \sigma_L^{\,2}}}, \text{ and} \qquad \textbf{Eq. 1}$$

$$LR = \frac{\sigma_L}{\sqrt{\sigma_S^{\,2} + \sigma_L^{\,2}}}, \text{ respectively,} \qquad \textbf{Eq. 2}$$

where $\overline{S}$ is the mean strength, $\overline{L}$ is the mean load, σ_S is the strength standard deviation, and σ_L is the load standard deviation. Thus, the safety margin and loading roughness are indirect measures of the overlap in the probability density functions and can be used to estimate the probability that the load does not exceed the strength (i.e., the reliability):

$$p(S > L) = \int_0^\infty f_L(L)\left[\int_L^\infty f_S(S)dS\right]dL\,, \qquad \textbf{Eq. 3}$$

where $f_S(S)$ and $f_L(L)$ are the probability density functions for strength and load, respectively.

For normally distributed strengths and loads, the probability of failure (i.e., 1-reliability) can be expressed as a function of safety margin (Equation 1) alone. This is one reason why safety margin emerged as the sole proxy for reliability in many applications. The other reason is that the design goal (in the nuclear industry, as well as other fields such as civil engineering or pressure vessel construction) is to build components and systems that have negligible failure probabilities. This can be attained by having sufficient safety margin (i.e., a large separation between mean strength and load relative to their combined standard deviations). This solidified the generalization that having adequate safety margin is a sufficient condition for high reliability. Thus, a highly reliable system (i.e., one in which the probability of failure is negligible) looks like Figure 3, with practically no overlap between the probability densities of strength and load.

Figure 4 is a schematic representation of the probability of failure. Given sufficient information with regard to load, strength, and their standard deviations, reliability can be precisely computed. However, such information is often beyond the current state of the art. In the nuclear industry, for example, probability functions for strengths of fuel or containments are prohibitively expensive to obtain.

2.2 Calculating the Conditional Probability of Loss of Function in an Event Sequence

Figure 5 shows the approach taken to ensure margin sufficiency in the nuclear industry. The probability density functions in the figure are used for illustrative purposes. The safety limit is conservatively set below the strength probability density function. Simultaneously, the code predicts values used to assess acceptability under conservative assumptions set forth in the emergency core cooling system (ECCS) evaluation models discussed in Appendix K to 10 CFR Part 50 [1] or more realistic alternatives. Although it cannot be strictly proven, a conservative calculation of the type imposed by Appendix K is expected to be sufficiently conservative to be more restrictive than one obtained from a more realistic approach (e.g., one that computes the bounding 95th percentile of the safety variable value with 95 percent confidence). Thus, a conservative Appendix K calculation would leave even more room for epistemic uncertainty than simply setting a conservative safety limit.

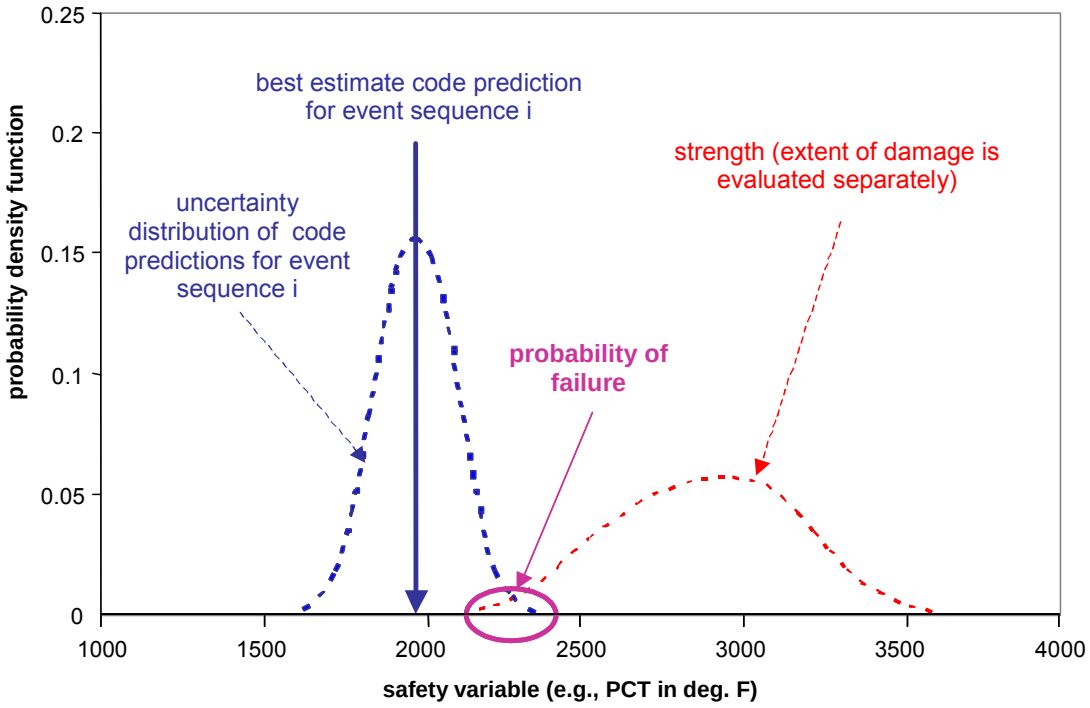

Figure 4 The probability of failure in an event sequence

In general, approaches used to compute the limiting value of a safety variable are classified as very conservative (Appendix K), bounding best estimate, realistic conservative, and best estimate plus uncertainties. The latter is ideally suited for integrating risk and safety margins. However, where sufficient margin exists, simpler, more conservative approaches can be used, which effectively reduces to current regulatory practice. Best estimate plus uncertainty methods have evolved substantially over the years and include Monte Carlo analyses, response surface methods, tolerance limit methods, internal assessment of uncertainty and other approaches practiced in other technical fields. The SMAP Task Group prepared a thorough survey of these methodologies, which can be used to select the most appropriate approach given a particular application. [12]

The use of safety limits instead of the onset of damage is more suitable for integrating risk and safety margins for two reasons—convenience and consideration of the unexpected. Obtaining the strength probability functions for physical barriers (e.g., fuel, reactor coolant boundary system, and containment) for each damage mechanism will continue to be prohibitively expensive. Therefore, it is convenient to set the safety limit below the onset of damage, by an amount that is commensurate with the lack of data and the importance of the subject safety variables. This gives the requisite confidence that, if operating conditions remain within safety

limits, the probability of failure will be negligible and some additional margin will be available for unknown events and phenomena.

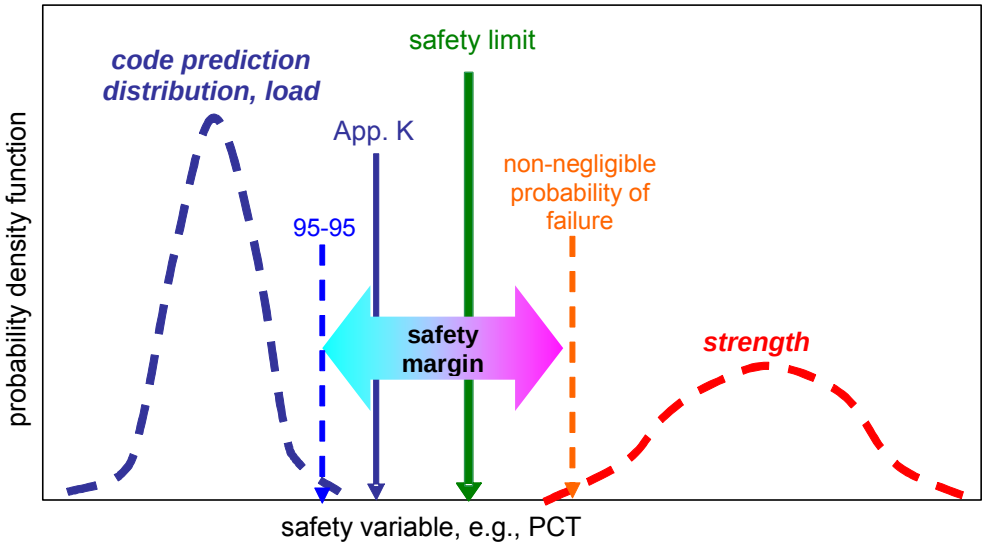

Figure 5 Ensuring adequate safety margins by setting a conservative safety limit and using bounding code prediction values to assess acceptability

It is important to note that it may be necessary to rethink the appropriate value of a safety limit for risk calculations. For example, the design-basis limit for containment pressure may be justifiably considered overly-conservative in light of the epistemic uncertainty associated with the containment fragility curve. In this case, the value of the limit used to determine the existence of sufficient margin when integrating risk and safety margins may differ from a design-basis safety limit.

2.3 Caveats in Adopting This Definition of Safety Margin for Risk Investigations

There are three caveats with regard to the definition of safety margin as presented above. The first involves setting the safety limit confidently below the onset of damage, which can be achieved for most physical barriers. One can imagine that for certain damage mechanisms and certain barriers, the uncertainty associated with the onset of damage could be so large as to preclude the ability to set a safety limit such that operations stay below it for certain accidents. This, however, is not the case with barriers of existing LWRs, so discussions on dealing with large uncertainty in the capacity density function will be deferred.

The second, and somewhat related, caveat is that one can make definitive statements with regard to keeping operating conditions below safety limits for design-basis events. However, the same is clearly not true in the risk space. For example, in a classic large-break LOCA event tree, many event sequences end in core damage. That presumes that the safety limit of 1204°C (2200 °F) was exceeded. Thus, there is an additional consideration of frequency of exceedance that should be associated with a given safety limit. If the frequency of exceedance is linked to a high-level risk acceptance criterion (e.g., the Commission's safety goal [5]) for a given plant, the threshold safety limit and exceedance frequency form a unique point on the frequency-consequence curve. Thus it is possible to overspecify the problem. The resolution of this issue is beyond the immediate scope of this research, but a promising solution can be envisioned using the frequency-consequence approach.

The third caveat is that the change in safety margins captured in this report pertains only to cases where a significant fraction of the load probability density function exceeds the safety limit. This is insufficient for those researchers who believe that any change in operating conditions that moves the plant closer to the safety limit is an effective loss of safety margin, whether the safety limit is exceeded or not. Earlier work on this framework suggested that it is possible to quantify loss of margin that occurs far away from the safety limit, where far is determined by the standard deviation in the load probability density function. A synopsis of this approach is included in Appendix A. This approach was not pursued further because it failed to meet one of the constraints imposed on this work: demonstrate the methodology on a problem of practical interest to the Agency. The approach could not have been demonstrated because no acceptance criteria exist for modifications that reduce margins but do not impinge on the safety limit.

2.4 The Conditional Probability of Loss of Function

The probability of exceedance is a well-established concept in PRAs. When a safety limit exists, the cumulative probability of the load curve that exceeds the safety limit is the probability of exceedance as shown in Figure 6. The load probability density function is generated through well-established methodologies such as CSAU or SUSA. [2 and 10] Simple approximations for exceedance probability can be devised (see, for example, Reference 13). The exceedance probability is conditioned on the occurrence of the event sequence that was simulated to generate the probability density function for strength. Thus, the proper term is conditional

probability of exceedance. To integrate risk and safety margins, the assumption is made that function is lost when the safety limit is exceeded. Therefore, the conditional probability of exceedance equals the conditional probability of loss of function in this approach.

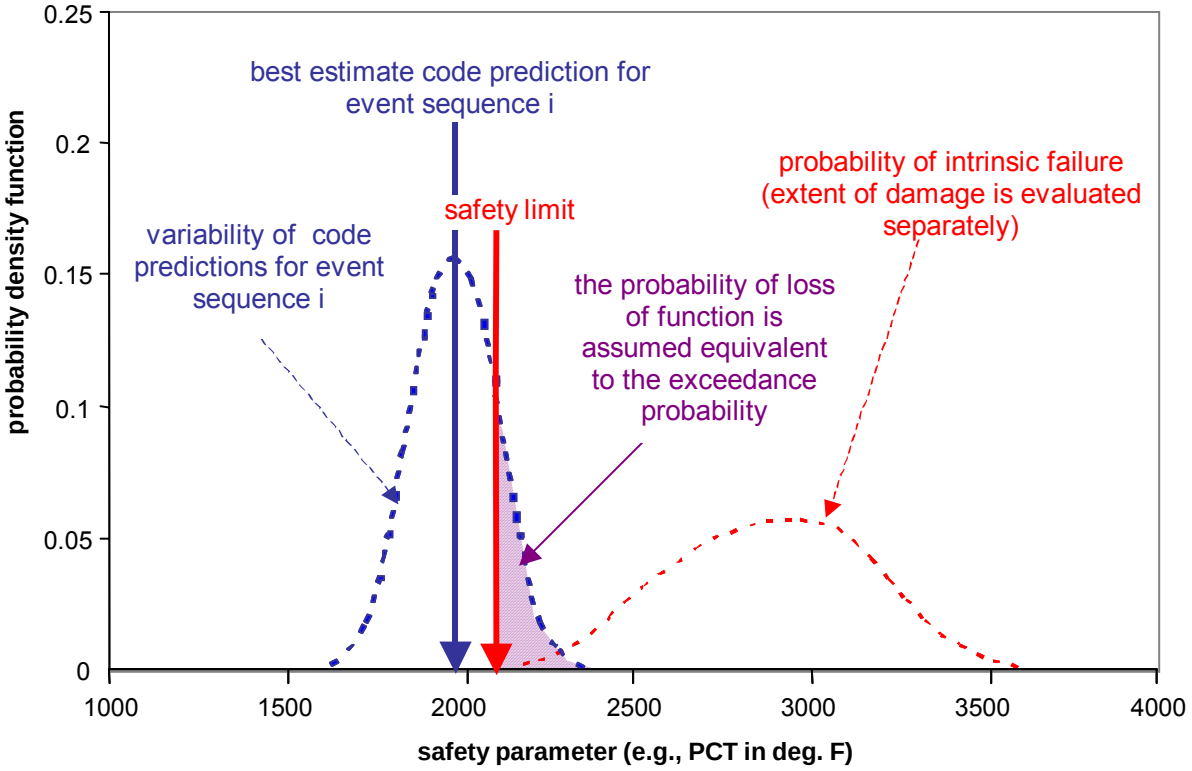

Figure 6 Calculating the probability of failure in an event sequence

This is, potentially, the most contentious step in integrating risk with safety margins, but it grounds this framework in existing regulatory practice. It is fully consistent with the assumption made in deterministic regulatory analyses that function is lost when the safety limit is reached. This assumption is fully justified if one remembers that an important driver in setting the safety limit below the onset of damage is to cope with "unknown unknowns." Because safety limits are set commensurate with the lack of knowledge and the importance of the subject variable, and because both these considerations are equally applicable in risk assessments, it is wise to extend this assumption to PRA analyses.

Once the conditional probability of loss of function is defined, the meaning of the term "safety margin" becomes unambiguous. The phrases "sufficient margin" and "loss of margin" also become clear. Sufficient margin exists if the probability of loss of function is negligible. Margin

is lost if and only if a change occurs in the probability of loss of function. In other words, the framework that integrates risk and safety margins is insensitive to changes that move the entire load probability density function through the space below the acceptance limit. For example, if, following a power uprate, the PCT in a transient changes from 800 °C (1472 °F) to 850 °C (1562 °F), the change is imperceptible to the risk metric calculated by integrating risk and safety margins. To some this change represents an erosion of margin and should be captured. One can devise means of capturing such changes (see, for example, Reference 13), but judging the acceptability of such an increase requires setting new acceptance criteria, which is beyond the scope of this work.

2.5 Treatment of Aleatory and Epistemic Uncertainties in Margin Calculations for an Event Sequence

The safety limit is set with due consideration to both epistemic and aleatory uncertainties associated with the probability density function that describes the strength of the barrier. In theory, the uncertainty that is considered in both the load and the strength density functions of Figure 6 must be of the same nature (i.e., either aleatory or epistemic). Krzykacz-Hausmann makes a thorough argument, which includes illustrative examples, for separating epistemic and aleatory uncertainty. [14] At the same time, many argue that there is no fundamental distinction between aleatory and epistemic uncertainty and that the price of this separation is not always warranted. [15] Also, one has to recognize that separating epistemic and aleatory uncertainty is often problematic and sometimes impossible. [16]

Even when such a separation is not warranted, it is helpful to understand why and how epistemic and aleatory uncertainties are treated differently in generating the load curve of Figure 6. One acceptable approach was employed in setting the pressurized thermal shock (PTS) screening criterion. [17] In simple terms, for a single value of each epistemic contributor, each aleatory contributor is sampled in its entirety. This is consistent with physical insight, yet a single probability density function is created that contains both aleatory and epistemic uncertainty information. Section 3.4 further discusses the treatment of aleatory and epistemic uncertainties for the complete integrated risk/safety margins framework.

In adopting the safety margin definition to determine the probability of loss of function, the strength probability density function is replaced with a distinct value—the safety limit. The

safety limit carries no uncertainty, neither epistemic nor aleatory, and poses no limitation on the nature of uncertainty contained in the strength probability density function.

3 Developing the Risk Metrics

Recognizing the fact that the integration of risk and safety margins described in this report is suitable to currently operating reactors as well as radically different reactor designs, two fundamental premises cast the framework in a technology-neutral context:

(1) Any foreseeable nuclear power plant can be summarily described as a volume that contains the fuel and fission products surrounded by one or more physical barriers.

(2) For any physical barrier, safety variables can be identified to demarcate the transition from “intact” to “lost function.”

The first premise is self-evident. The role of the regulator is, and will continue to be, to protect the public and the environment from inadvertent releases of radionuclides from the barrier(s) that contain the loci of fission. The second premise is based on inherent properties of physical barriers. The integrity of physical barriers (i.e., those made of materials, as well as different confinement systems, such as electromagnetic confinements) is subject to operation within acceptable ranges of dominant safety variables. Examples of such variables for the physical barriers of the existing LWR fleet are pressure, temperature, and strain. To determine barrier integrity, these variables must be directly or indirectly measurable, and their values must be predictable for plant conditions during normal and emergency operation.

Furthermore, the ranges over which barrier integrity is maintained must be determined analytically or experimentally. If necessary, the proper function of a physical barrier is ensured by systems and components that maintain safety variables within the range in which the barrier retains its function.

Safety variables that determine barrier integrity are suitable for use in establishing safety limits and quantifying the conditional probability of loss of function, as discussed in Chapter 2. The more generic term, “probability of loss of function,” can encompass failure as well as bypass of a physical barrier.

The process of quantifying the probability of loss of function begins with individual event sequences, which can be either design-basis accident sequences or all sequences that comprise the plant's risk space. In this context, the risk space includes all plausible event sequences of nonnegligible frequency of occurrence, regardless of the associated consequences. The risk space includes success paths, such as normal operations.

3.1 Probabilities

In PRAs, the frequency of occurrence is the predominant measure of the likelihood of failure. This is primarily a matter of convenience. If sufficient data exist, both the frequency of failure and the failure rate can be determined. Similarly, when expert opinion is solicited on the likelihood of an initiating event occurring, it can be sought either in terms of frequency or in terms of rate.

This section uses probabilities to construct the framework for integrating risk and safety margins because many are more comfortable with probabilities than with frequencies. Moreover, in integrating risk and safety margins, it is useful to construct the framework using probabilities because they impart a formalism that is not readily evident with frequencies. However, because frequencies are prevalent in PRA and, more importantly, risk acceptance criteria are based on frequency, this report often uses frequencies as examples. A simple derivation in Appendix B to this report shows that, for very rare events, the failure rate and the failure frequency assume identical numerical values.

The first step in computing the risk metrics is to obtain the unconditional probabilities of loss of function for each event sequence. The probability of loss of function is calculated based on "distance" to the safety limit described in Chapter 2. Deterministic calculations that assume a specific progression of events are used to generate the load probability density function in Figure 4. For example, a deterministic calculation is carried out using a thermal-hydraulic code for event sequence number 7 in the large-break LOCA tree (LLOCA 07) obtained from a Standardized Plant Analysis Risk (SPAR) model. Thus, the strength curve obtained from these runs will yield the conditional probability that the fuel barrier will lose its function because the 1204 °C (2200 °F) safety limit is exceeded. In other words, the predicted loss of function probability is conditioned upon the occurrence of the sequence of events simulated through the

thermal-hydraulic calculation. To obtain the unconditional probability of loss of function for event sequence LLOCA 07, the conditional probability of loss of function must be multiplied by the frequency of occurrence of LLOCA 07.

More formally, when the strength and load in Figure 4 pertain to a safety variable that governs the loss of function of barrier n, Bn, then $p(f\_Bn \mid ES_i)$ is the conditional probability of loss of function for barrier n during event sequence i (ES_i):

$$1 - p(S>L) = p(f\_Bn \mid ES_i). \qquad \textbf{Eq. 4}$$

In its most general form—one that ignores the need to afford due consideration to epistemic uncertainty—p(S>L) is the reliability from Equation 3. In the approach taken here, p(S>L) relates to the exceedance probability of the safety limit as defined in Chapter 2.

Two things must happen in order for fission fragments to be released beyond barrier B1—first, ES_i has to occur, and, second, the barrier B1 must loose its function. This is expressed as follows:

$$p_i^{B1} = p(ES_i \cap f\_B1) = p(ES_i) \cdot p(f\_B1 \mid ES_i), \qquad \textbf{Eq. 5}$$

where:

$p(f\_B1 \mid ES_i)$	is the conditional probability that barrier 1 will lose its function given ES_i,
$p(ES_i)$	is the probability of occurrence of event sequence i, and
$p(ES_i \cap f\_B1)$	is the probability of occurrence of ES_i and barrier 1 loss of function.

Equation 5 can be generalized to any subsequent barrier, Bn. This is a natural conclusion of the fact that deterministic computations to calculate the values of safety variables for barrier n simulate the barrier's response given the initiating event and the breach of previous barriers. Thus, the probability of loss of function for barrier n is conditioned on the occurrence of event sequence i, as well as the conditional loss of function of preceding barriers and can be expressed as follows:

$$p^{Bn}_{i} = p(ES_i \cap f\_B1 \cap f\_B2 \cap ...) = \\ = p(ES_i) \cdot p(f\_B1 | ES_i) \cdot p(f\_B2 | ES_i \cap f\_B1)... \quad \textbf{Eq. 6}$$

where:

$p(ES_i)$	is the probability of occurrence of ES_i,
$p(f\_B1 \mid ES_i)$	is the conditional probability that barrier 1 will lose function given ES_i, and
$p(f\_B2 \mid ES_i \cap f\_B1)$	is the conditional probability that barrier 2 will lose function given ES_i and the loss of function of barrier 1.

If n is the ultimate barrier, the following equation can approximate the probability of exposing the public and environment to fission products because of event sequence i:

$$p_i \approx p(ES_i) \cdot p(f\_B1 | ES_i) \cdot p(f\_B2 | ES_i \cap f\_B1) \cdot ... \cdot p(f\_Bn | ES_i \cap f\_B1 \cap ... \cap f\_Bn-1). \quad \textbf{Eq. 7}$$

3.1.1 Evaluating Acceptability Given a Core Damage Frequency Criterion

As discussed above and as shown in Appendix B to this report, the probability of occurrence of ES_i is numerically the same as the frequency of occurrence of ES_i, if the subject event is very rare. Thus, the relationships derived above can be used to calculate the unconditional occurrence frequencies of any barrier. In Equations 5, 6, or 7, one replaces the probability of occurrence of ES_i with the frequency of occurrence of ES_i. In fact, the derivation followed above can be repeated using frequency instead of the probability of occurrence of ES_i, and the rare event limitation can be eliminated. However, working with probabilities ensures that the analysis will include all event sequences that must be considered, and that the end state is consistent throughout the risk space.

To evaluate the acceptability of a modification, the first step is to generate the risk space. The risk space is the set of all event sequences that the modification impacts, either in terms of the distance to the safety limit or the frequency of occurrence of the event sequence. Using the conditional probabilities of loss of function before and after the modification and the associated occurrence frequency of each event sequence, one can generate the expected probability of occurrence before and after the modifications. As outlined in the CSNI/SMAP technical note for Task 2 [18], the questions of 10 CFR 50.59, “Changes, Tests and Experiments,” [19] can be

adapted into a rigorous process for determining the changes that are needed to the PRA model to capture a given modification.

The framework to integrate risk and safety margins makes it possible to evaluate the available margin for a specific function that comes into question at any given time. The approach is best demonstrated using a highly abstracted example. Consider a reactor that has the risk space depicted in Figure 7. For the first barrier (i.e., the core), two known initiating events (IEs) can lead to damage—IE1 and IE2 (e.g., a LOCA and a reactivity insertion accident). The rest of the time (ROT), the reactor is operating without incident or is shut down. There are two mitigation systems (MSs)—MS1 mitigates IE1, and MS2 mitigates IE2. For example, MS1 is a makeup system for the LOCA, and MS2 is a neutron-poison injection system for the reactivity insertion accident.

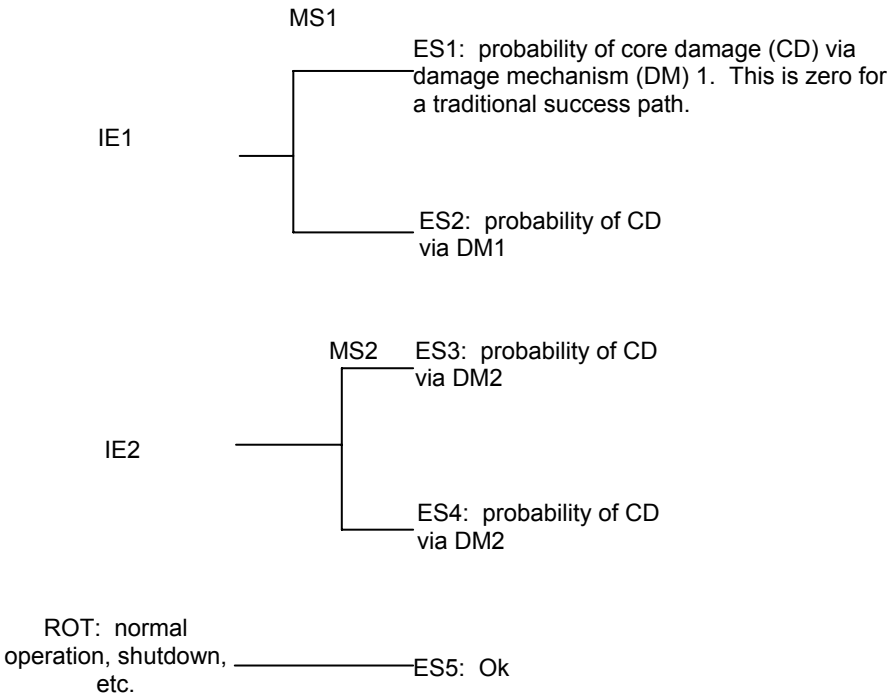

Figure 7 Risk space of a representative reactor

To better illustrate the applicability of integrating risk and safety margins, the end states are identified for all possible damage mechanisms. Core damage can occur through one of two damage mechanisms (DMs), DM1 (e.g., embrittlement of the first barrier) or DM2 (e.g., cracking), which can lead to the release of fission products from the core.

The embrittlement damage mechanism occurs as a consequence of an increase in the safety variable (SV), SV1 (e.g., PCT). Similarly, SV2 (e.g., enthalpy deposition rate) governs the initiation of cracking, DM2. It is important to refine the event trees to sufficient detail, such that only one possible independent damage mechanism is present at the end of an event sequence.[4] This ensures the integrity of the conditional probability of loss of function in the computation of the risk metric.

Table 1 specifies the probabilities of occurrence of event sequences and the conditional loss of function probability for each event sequence in Figure 7. For example, IE1 triggers ES2; MS1 does not work. The probability of occurrence of this event sequence, given the expected probability of occurrence of IE1 and reliability of MS1, is 5×10^{-5}. The conditional probability of loss of function is calculated from the distribution of deterministic code predictions for SV1 given the known input/model variabilities and the safety limit, as discussed in Section 2.2. For ES2, the conditional probability of loss of function is 50 percent; multiplying this by the frequency of occurrence, the unconditional frequency of loss of function for ES2 is 2.5×10^{-5}.

The expectation value for CDF is determined by computing the unconditional frequency of loss of function for each event sequence, and then adding them for all sequences that comprise the risk space.

Table 1 Probabilities of Loss of Function for the Representative Reactor

Event Sequence	Probability of Occurrence of the Event Sequence	Conditional Probability of Loss of Function	Unconditional Probability of Loss of Function
1	1.00E-04	0.00	0.00E+00
2	5.00E-05	0.50	2.50E-05
3	3.00E-03	0.20	6.00E-04
4	2.00E-07	0.90	1.80E-07
5	9.97E-01	0.00	0.00E+00
Expectation value for core damage due to DM1 and DM2			**6.25E-04**

[4] However, it is acceptable to have several safety variables related to a single damage mechanism (e.g., both PCT and total clad oxidation can be tracked if the subject damage mechanism is embrittlement).

For economic reasons, the licensee operating the representative reactor proposes two modifications, including a reduced testing schedule and a reduced injection capacity for MS1. Both modifications impact only the embrittlement damage mechanism, DM1, and have no bearing on core cracking, DM2. Assume that guidelines exist (similar to those of Regulatory Guide 1.174 [2]) with regard to the maximum increase in CDF allowable for the representative reactor.

Given the proposed modifications, the reduced injection capacity challenges the PCT safety limit, SL1, but not the cracking safety limit, SL2. Thus, the risk space for the inquiry can be reduced as shown in Figure 8. Because the modification does not impact the core damage triggered by cracking, the change in the expectation value for CDF is given by the change in the expectation value for DM1 frequency.

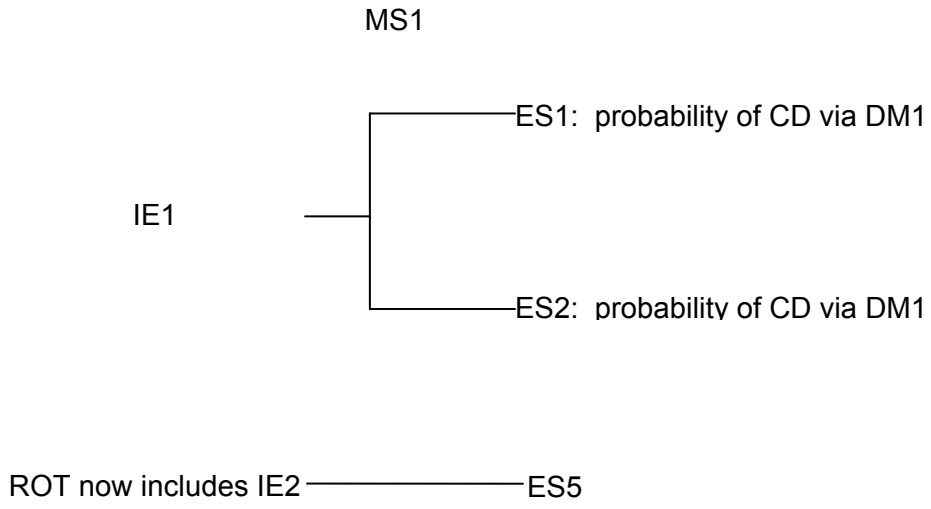

Figure 8 Reduced risk space for the example safety inquiry

For the representative reactor, the unconditional frequency of core damage via embrittlement, DM1, is calculated from the values shown in Table 2. The reduced testing schedule lowers the reliability of MS1 by 5×10^{-5}. The reduced injection capacity increases the best-estimate maximum value and alters the probability density function of PCT, SV1, such that the conditional probability of loss of function after the modification increases in ES2 from 50 percent to 75 percent. The change in CDF due to DM1 is 2.00E-05. This value can be compared to the permissible change in CDF to determine the acceptability of the proposed modifications.

Table 2 Data Used to Calculate the Change in Expected Unconditional Probability of Core Damage Before and After the Modifications Proposed for the Representative Reactor

Probability of Occurrence of the Event Sequence	Conditional Probability of Loss of Function of the Safety Limit	Unconditional Probability of Loss of Function
	Before Modifications	
1.00E-04	0.00	0.00E+00
5.00E-05	0.50	2.50E-05
1.00E+00	0.00	0.00E+00
Expectation value for core damage due to DM1 before modification		***2.50E-05***
	After Modifications	
9.00E-05	0.00	0.00E+00
6.00E-05	0.75	4.50E-05
1.00E+00	0.00	0.00E+00
Expectation value for core damage due to DM1 after modification		***4.50E-05***

This abstraction shows how the probability of loss of function can be integrated within PRA results and used directly when subsidiary risk acceptance criteria (e.g., for ΔCDF or ΔLERF) exist. Applying the ΔLERF limit is largely similar, but it involves the introduction of another conditional probability—the probability that the time between the loss of the core and loss of containment function is shorter than a prespecified interval.

3.2 Consequences

When a modification impacts the consequences of accidents, not just the frequencies of their occurrence, then it is necessary to include a measure of consequence in the risk metric. For example, consider the case of a power uprate achieved by flattening the axial profile. In a reactor with a flat power profile, a perturbation that leads to exceeding the safety limit affects more fuel bundles than in a reactor with a higher peaking factor. Another example is a modification that affects both CDF and consequences, such as the proposal to remove trisodium phosphate from the containment of certain PWRs in response to GSI-191. This modification lowers the probability of chemical effects and thus the CDF. Simultaneously, offsite and personnel doses are expected to increase for all accidents that involve the release of iodine. A proper evaluation of the risk implications of such a modification can only be done if consequences are considered.

Consequences can be considered in a generic form that is suitable to existing as well as future reactor concepts. Figure 9 depicts the premise that any power-generating reactor consists of fuel and fission products contained within concentric physical barriers. As in current practice, an initial source term must be computed or assumed. If a concentration of fission products, C_{FP}, is contained within the first barrier at the time the event sequence occurs, the decrease in the concentration of fission products as they pass through successive barriers is a function of many factors, including the following:

- volume confined by each barrier
- extent of damage to the barrier
- scrubbing by sprays and water pools
- time between the breaches of successive barriers

Deterministic calculations using severe accident type codes can calculate a transmission factor through a barrier, t, which reflects dependencies on dilution, extent of damage, and other factors. This practice is common in current severe accident analyses.

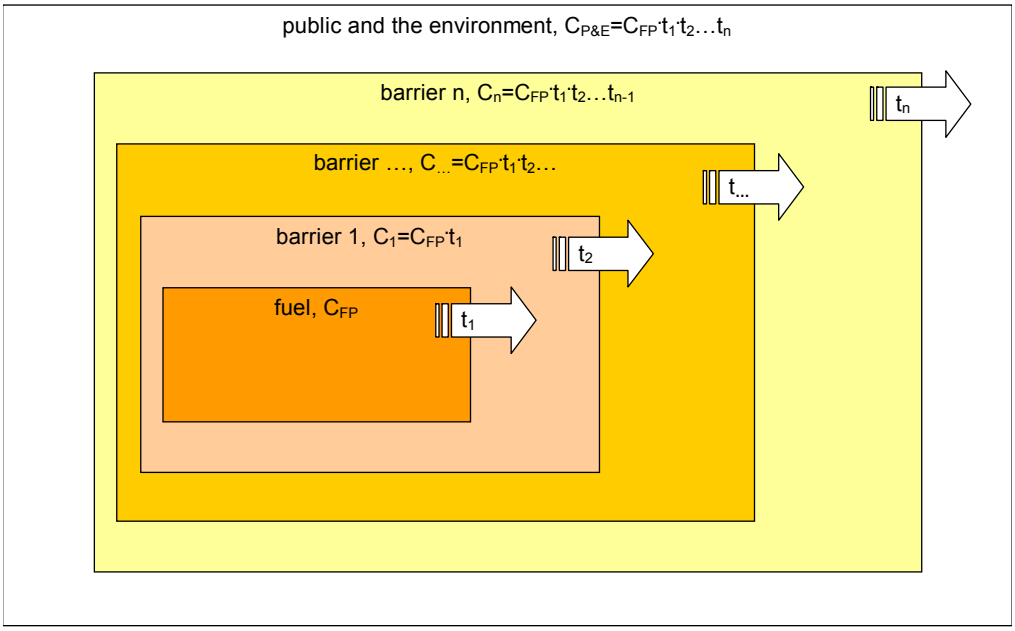

Figure 9 Schematic representation of multiple barriers containing the fuel and fission products

The consequence of an event sequence within a barrier is quantified by the sum of radioisotope concentrations confined by the barrier prior to the initiation of the event and transferred from preceding barriers that have been breached during the event sequence. Thus, for barrier n, the concentration, C_n, can be represented by the concentration of fission products within the confines of that barrier and calculated as follows:

$$C_n \approx C_{0,n} + C_{FP} \cdot t_1 \cdot t_2 \ldots t_{n-1}, \qquad \textbf{Eq. 8}$$

where:

C_{FP} is the concentration of fission products within the primary barrier, and

C_{0n} is the concentration of fission products within barrier n at the initiation of the event sequence.

Note that Equation 8 includes contributions from isotopes that are present in areas outside of the first barrier. This is particularly important if a barrier bypass event sequence is being considered. The formulation of consequences within the confines of barrier n is also useful in calculating risk to personnel. It is not necessary to compute transmission factors for each event sequence; they can be grouped according to barrier, damage mechanism, extent of damage, time lapsed since the breach of the previous barrier, and other factors. Also, conservative transmission factors (e.g., an extreme value of 1) can be used to assess the risk posed by individual event sequences, provided that the plant has sufficient margins to radiological damage limits.

The consequences to the public and the environment are calculated from a generalization of Equation 8 to transport beyond the ultimate barrier. In sequence i, the consequences, $C_{P\&E}$, can be computed from the following equation:

$$C_{P\&E} \approx C_{FP} \cdot t_1 \cdot t_2 \ldots t_n. \qquad \textbf{Eq. 9}$$

A consequence measure related to the one computed above may be better suited for application within existing regulations (e.g., person-rem), but the form above is sufficiently descriptive for the current discussion. The approach described in the preceding paragraphs has already been developed and refined, and is employed in Level 3 PRA calculations.

3.3 Risk

In its most general form, risk is the product between the probability of occurrence of an event and its consequences. The risk to the public and the environment because of event sequence i, r_i, is the product between the probabilities described in Section 3.1 and the consequences discussed in Section 3.2:

$$r_i = p_i \cdot C_{P\&E,i}, \qquad \textbf{Eq. 10}$$

where the probability of release to the public, p_i, is computed from Equation 7, and the consequences of event sequence i, $C_{P\&E,i}$, are computed from Equation 9.

The expected risk for the plant can be calculated assuming that only one event sequence can occur at any given time. In other words, it is fair to assume that at any given time the plant is in one distinct end state. To ensure the validity of this assumption, it is preferable to use probabilities in deriving this framework. The expected risk is the arithmetic sum over all event sequences:

$$\text{expected risk} = \sum_i r_i \,. \qquad \textbf{Eq. 11}$$

With the consequences, $C_{P\&E}$, of Equation 9 cast in the appropriate form, the expected risk calculated using Equation 11 is suitable for comparison with the risk acceptance criteria in the Commission's safety goals. [5]

3.4 Constraints Imposed and Opportunities Afforded by the Integration of Risk and Safety Margins

Integrating risk and safety margins requires care in merging engineering, deterministic, and probabilistic data because the loss of function is conditioned on exceeding the safety limit, and deterministic calculations are conditioned on probabilistic event sequences. The literature has addressed many of the specific constraints that must be imposed to assure the needed consistency. Of particular interest are the general discussions in support of dynamic PRA (see, for example, Reference 20), and some particular constraints raised by the SMAP Task Group

(see Reference 12). Many find this constraint (i.e., requiring coordination between key disciplines that provide data to the decision-maker) to be a very beneficial aspect of integrating risk and safety margins.

The framework described here takes advantage of the state of the art in deterministic and probabilistic analyses and is intended to continue to grow as methods and tools develop. The SMAP Task Group report, "Safety Margins Action Plan Technical Note for Task 3: Safety Margin Evaluation Methods," provides a good summary of methods and tools that are appropriate for use in integrating risk and safety margins. [12] In addition, the framework presented in this report reduces to traditional deterministic and probabilistic analyses if conventional assumptions are made.

Another advantage of integrating risk and safety margins as described in this report is that it is inherently compatible with the separate treatment of epistemic and aleatory uncertainty through traditional probability theory techniques. Much research has been conducted into how, and how much, epistemic and aleatory uncertainties have to be separated. One of the most authoritative references on the subject cautions that given the expense associated with separately propagating epistemic and aleatory uncertainty into the final results, one should duly consider how the additional information will be used, if at all. [15] No criteria or guidelines exist on deciding if separation of epistemic and aleatory uncertainty should be pursued and, if so, how much separation is needed. Such criteria/guidelines could prove very valuable. It is, however, universally recognized that complete separation is neither possible nor necessary for most realistic risk studies.

However, in many cases, the decision-maker needs to know how much uncertainty can be reduced and at what cost. To make this judgment properly requires knowledge of the separate contributions of these different types of uncertainties and knowledge about the sensitivity of results to various sources of epistemic uncertainty. Several techniques have been evaluated for compiling this information (e.g., traditional probability theory, fuzzy set theory, possibility theory, evidence theory); these techniques have various degrees of promise. [21] The most commonly used method of dealing with the two types of uncertainty remains traditional probability theory with strict separation of aleatory and epistemic uncertainty (e.g., as exercised in the PTS investigation described in Reference 17). The SMAP Task Group addresses this topic in Reference 12.

4 Example Application

One of the requirements imposed on the development of this framework has been to identify an example application of current regulatory interest. The example had to demonstrate the value added by integrating risk and safety margins as a complement to existing decision-making tools. Integrating risk and safety margins adds value to decisions in which margin sufficiency is an issue and uncertainty plays a significant role. The framework is most useful when a plant modification has both positive and negative effects on safety. In that case, the framework eliminates unnecessary conservatism to prevent one set of outcomes from overshadowing the other.

Potential candidates for the proof-of-concept demonstration were identified. They include cases of limited margin, which are typically of interest to multiple stakeholders. Because of this, it is especially important to note that the example included in this report has no intrinsic value in drawing safety conclusions. It is strictly a highly simplified, abstracted application to a generic increase in the sump debris screen of a PWR.

The phenomena considered do not constitute a comprehensive list. For example, the increase in screen size only affects the change in minor losses in the suction part of the recirculation pump piping. No consideration is given to changes in downstream-effects that could be induced by increasing the screen size. Furthermore, the values used to illustrate the framework are generic and do not represent any particular plant or grouping of plants. Data used to compute minor form losses due to accumulation of debris on the screen is excerpted from an industry survey. [22] For these reasons, no conclusions can be drawn regarding risk reduction achieved by increasing sump debris screens from the current example.

The case involves the following issues. After a LOCA, debris can travel to the sump screen and potentially cause a loss of NPSH for ECCS and containment spray system pumps as suction headers become blocked. The postulated amount of blockage exceeds that for which the system was designed and, thus, emergency core cooling and containment spray functions are lost. In the absence of emergency core cooling, the core becomes damaged and fission products escape from the first barrier. The consequences of the event can be significant because the loss of containment spray function increases the probability of releases beyond the ultimate barrier.

To evaluate the effect of increasing the debris screen size, one can examine the impact on CDF before and after the modification. This can be done with traditional probabilistic analyses. However, in the case of NPSH margin, substantial uncertainties are associated with parameters that determine whether core damage occurs. Thus, a realistic calculation without regard to uncertainties can be misleadingly optimistic. If the uncertain parameters are treated conservatively, the picture will be overly pessimistic. By integrating risk and safety margins, the uncertainty becomes part of the calculated core damage probability. Furthermore, conservatism is only required when lack of data demands it. Thus, one obtains a realistic picture that is informed by participating uncertainties and in which conservatism is used only where necessary.

The example considers a PWR that has debris screens of 125 square feet (ft^2); this value is representative of current PWR debris screens. The PWR will increase the screen to 1,100 ft^2; this value is close the median proposed new screen size for the 69 PWRs operating in the United States. The proof-of concept example examines the effect of this plant modification on CDF.

Assuming that the proposed modification has no impact other than changing the pressure drop through the debris bed formed on the screen, one can link the change in NPSH margin to CDF. Specifically, the probability of losing emergency cooling because of lost NPSH margin can be calculated before and after a modification. One can reasonably assume that loss of NPSH leads to failure of ECCS recirculation, which in turn leads to the loss of function of the first barrier. Therefore, the probability of loss of NPSH margin is equivalent to the conditional probability of loss of function for the first barrier, and can be used directly to determine the impact on CDF. Specifically, the product of the probability of losing NPSH margin in an event sequence and the frequency of occurrence of that event sequence is the unconditional probability of core damage due to loss of NPSH for the particular event. Because the computed metric is the unconditional CDF, only event sequences in which the margin is inadequate need to be considered before and after the modification.

Integrating risk and safety margins starts with generating (1) the risk space (i.e., all the event sequences that the modification affects) and (2) a phenomena/variables identification table used to compute the conditional probability of loss of function for each event sequence. To generate the risk space, one must consider all initiating events that challenge NPSH margin. In general,

the questions in 10 CFR 50.59 [19] can be modified to systematically determine how event trees change as a result of a plant modification; the SMAP Task Group technical note for Task 2 [18] addresses tailoring the questions in 10 CFR 50.59 to examine changes in event trees. Event sequences must be refined to capture important input variabilities. For NPSH margin, the variabilities would include actuation of containment spray, choosing to start only one makeup injection/core spray train at a time, and others. The process of identifying refinements to event sequences requires knowledge of the phenomena that impact NPSH, as mentioned above. This intrinsic link between PRA and deterministic analyses makes the process of generating the risk space iterative with the process of identifying key safety variables. Refinements to event trees exceed the proof-of-concept scope of this example.

Several practical simplifications can be made that are consistent with current practice in PRAs. For example, one could limit the risk space to medium LOCAs when it can be shown that they dominate the risk. For some plants, it may be reasonable to assume that for small LOCAs, alternative means of making up water can be found to preclude the need for recirculating from the sump[5]. Large LOCAs have relatively low initiating event frequencies so they can be ignored in rough calculations of risk. However, to fully illustrate the framework, this example considers event sequences for small, medium, and large LOCAs. Figure 10 depicts the event tree for the large-break LOCA initiating event.

Another simplification to the risk space is that core damage paths do not need to be considered because paths that lead to core damage prior to NPSH considerations do so by different mechanisms of damage. This proof-of-concept example assumes that increasing the size of the sump screen does not impact these mechanisms of damage, and thus the CDFs along those paths do not change before and after the modification. This type of simplification is possible because of the use of probabilities in developing the integrated risk/safety margins framework, which assures consistency in decision-making metrics. Specifically, it is possible if one ensures that the end states of all event sequences of interest result from distinct damage mechanisms. Given this simplification, for the large-break LOCA event sequences of Figure 10, NPSH margin needs to be determined only for path one, because all other paths lose function due to other mechanisms.

[5] In PWRs, this may be limited by the reactivity insertion that results if deborated water is used.

Furthermore, it is common practice to truncate below a certain frequency threshold. One should ensure that the sum of all truncated event sequences is not of a magnitude that would change the decision. This exercise treats event sequences with frequencies of less than 10E-6 as failed. These will not show up in the ΔCDF, but the baseline CDF includes their sum before and after the modification; thus, their total contribution can be assessed by inspection. A close examination of the scope of each safety inquiry can lead to additional simplifications. For example, in a given plant, one may be able to eliminate an entire range of break sizes that could not generate enough debris to pose blockage problems regardless of break location. No such additional simplifications have been attempted for the proof-of-concept NPSH example.

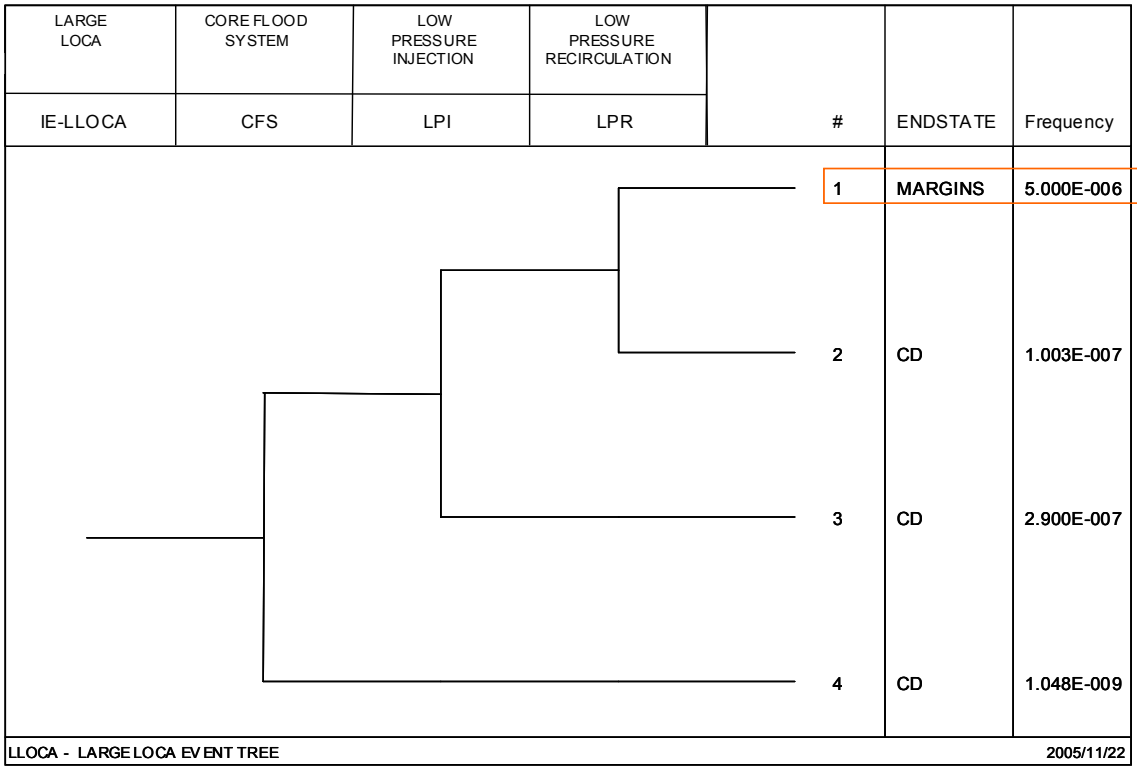

Figure 10 Large LOCA event tree for NPSH margin calculation

The next step is to identify the variables that determine the amount of NPSH margin available in each event sequence. The definition of NPSH is a good starting point for the development of the phenomena/variable list. In most applications, a phenomena identification and ranking table (PIRT) developed by a panel of experts would be available as a starting point. Los Alamos National Laboratories generated some earlier PIRTs for GSI-191, but they are not directly

relevant to the development of this proof-of-concept example. Instead, a list of phenomena and variables was developed from first-principle considerations.

A pump-specific amount of NPSH is necessary to ensure that the pump functions without cavitation in the impeller region. Both the injection capacity and the reliability of a pump are predictable only as long as the required NPSH ($NPSH_r$) is less than the available NPSH ($NPSH_a$). The factors that increase the available NPSH are the containment pressure and the height of the water in the sump. $NPSH_a$ deteriorates with increased pressure drops in suction piping and with increased sump water temperature and can be expressed as follows:

$$NPSH_r \leq NPSH_a = p_{atm} + p_{stat} - p_{vap} - p_{loss} \qquad \textbf{Eq. 12}$$

p_{atm} is the pressure head (containment pressure),

p_{stat} is the static suction head (sump level),

p_{vap} is the vapor pressure (at maximum pumping temperature), and

p_{loss} is the friction and K-loss head in the suction side, including losses at the screen.

Keeping with the notion that having margins requires room for “unknown-unknowns” epistemic uncertainties, it is reasonable to assume that loss of safety function occurs when the $NPSH_a$ is less than the $NPSH_r$ required for the specific pump. No other attempt was made to separate aleatory and epistemic uncertainty in the example calculation. Examining the terms of Equation 12, one can generate the table of phenomena that govern the availability of NPSH margin; see Table 3. The same table lists some of the variables and considerations that are necessary to determine the probability density function of NPSH margin. The table is not intended to be exhaustive but to illustrate the type of information that must be collated to integrate risk and safety margins.

For individual plant cases, the analysis would proceed by running each event sequence with a deterministic code (e.g., RELAP5 or TRACE) to obtain the ranges of values necessary to compute the NPSH margin distribution. Specifically, given variabilities in code models and input/boundary conditions, one would obtain distributions for sump water temperature, sump level, and containment pressure. This was not done for the current example; instead, generic ranges were obtained from industry and NRC documents (e.g., References 22 and 23 as shown in Table 4).

Table 3 Variables that Determine the Available NPSH

Pressure Head	
containment pressure	operator depressurization, evolution of the event sequence
containment leakage	ranges from negligible to that allowed by the technical specifications
containment spray duration/capacity	affected by measures taken to decrease the need for going to recirculation
containment temperature	accident sequence, spray action, initial and boundary conditions
Static Suction Head	
sump level	break size
makeup injection	affected by measures taken to decrease the need for going to recirculation
water hideout	compartment geometry
water density	*sump water temperature*
impurities (solutes and particulates)	debris dissolved or suspended in the sump water
Vapor Pressure	
thermodynamic properties	*sump water temperature*
impurities	debris dissolved or suspended in the sump water
Friction Head	
suction piping	piping configuration
impurities (solutes and particulates)	debris dissolved or suspended in the sump water
viscosity	temperature and impurities
losses due to debris	amount and composition of debris (accident sequence)
screen configuration	vendor
debris distribution	debris source, initiating event, and accident sequence
dispersed obstructions (gloves, reflecting metal)	debris source
presence of sludge	

The type of information contained in Table 4 is similar to that in the PIRTs and is consistent with information developed to identify uncertainty in deterministic calculations. In addition to being a requisite for the integration of risk and safety margins, Table 4 has another important attribute—it lends transparency to the process. The analyst or the regulatory decision-maker can focus on elements such as the completeness of information in the table, the ranges of values, and the adequacy of the source material. For example, uncertainties associated with pool level can be substantial depending on the potential for water-hideout in a particular containment. Similarly, a complicated suction-piping configuration will have a substantial uncertainty associated with minor and major pressure losses. All these sources of uncertainty become important if the licensee is only able to calculate a margin that is less that a foot. Conversely, a licensee who has indeed treated all sources of uncertainty in a conservative manner can easily indicate so in

Table 4. Thus, this table is also useful in deciding when it is cost-beneficial to use accurate ranges as opposed to a conservative value.

Table 4 Variables and Values Used to Generate the NPSH Margin Distributions for the Large-Break LOCA Event Sequence in the Proof-of-Concept Example

Variable	Units	Nominal Value	Minimum (% of nominal)	Maximum (% of nominal)	Source Reference and Comment
Mineral wool volume	ft^3	126	40	100	NUREG/CR-6808 [22]: 10 to 25% range of total (5 to 10% if no CS)
Dirt-dust mass	lbm	170	40	100	NUREG/CR-6808: 10 to 25% range of total (5 to 10% if no CS)
Qualified epoxy mass	lbm	260	40	100	NUREG/CR-6808: 10 to 25% range of total (5 to 10% if no CS)
Paint chips mass	lbm	95	40	100	NUREG/CR-6808: 10 to 25% range of total (5 to 10% if no CS)
Flow rate through strainers	gpm	8700	95	105	representative of 10% controller range
Screen area	ft^2	**125/1100**	80	100	allow for up to 20% obstruction
Water temperature	°F	187	100	130	NUREG/CR-6224 [23]: ranges from 187 °F to 243 °F
Screen losses (nominal)	ft	-32/-0.35			calculated according to NUREG/CR-6224
Containment pressure ($p_{st\text{-}part}$)	psi	14.7–21.7	80	100	NUREG/CR-6224: ranges from 0 to 7 psig; conservative
Pool level above suction	ft	25	90	110	representative pool level
Friction and K losses	ft	-3.00	80	100	account for impurities
Cavitation pressure	ft	20–21.7	100	100	corresponding to sump temperature
$NPSH_r$	ft	-13	90	110	deterioration due to viscosity
$NPSH_a$	ft	20/51			calculated according to Equation 12
Mean NPSH margin	ft	-6.7/12.8			$NPSH_a$-$NPSH_r$

The variables listed in Table 4 were sampled to generate the probability density of NPSH margin. Simple Monte Carlo sampling was used. The probability density functions were generated using 500 samples. The variables were assumed to range uniformly between the maximum and minimum values of Table 4. Figure 11 and Figure 12 show the probability density functions and the integral loss of function probabilities for NPSH margin given a 125-ft^2 and 1100-ft^2 screen, respectively. The probability of losing function because of inadequate

NPSH is more than 80 percent in LLOCA1 if a small debris screen is used. The probability drops to less than 20 percent if the screen is enlarged to 1100 ft^2.

It is important to note that the spread of the distributions in Figure 11 and Figure 12 is highly relevant. The generic variable values and ranges used to generate the plots are representative of actual plant conditions. Therefore, the ±15-foot band that captures most of the trials is not unreasonable, in light of the uncertainties associated with NPSH margin. A sensitivity study showed the impact of sump temperature to be a dominant factor in determining the spread of the NPSH distribution even if different temperature distribution shapes are used. This means that the only acceptable conservative calculation of NPSH is one in which the temperature takes its most limiting value for the time of the computation. This conclusion is important because, if an analyst computes an NPSH margin of 0.4 feet that was calculated with the mean of the temperature range, he/she is effectively reporting a 50 percent probability of failure due to NPSH margin loss if the breadth of the uncertainty range is taken into consideration.

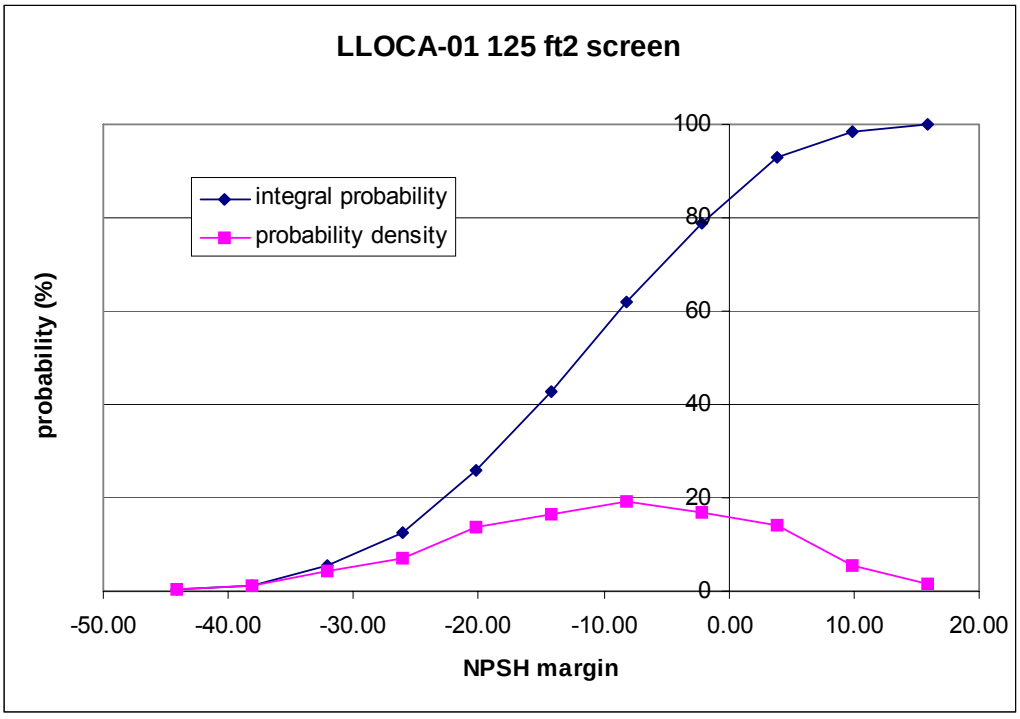

Figure 11 Distributed and cumulative probability of loss of NPSH with a 125-ft^2 debris screen

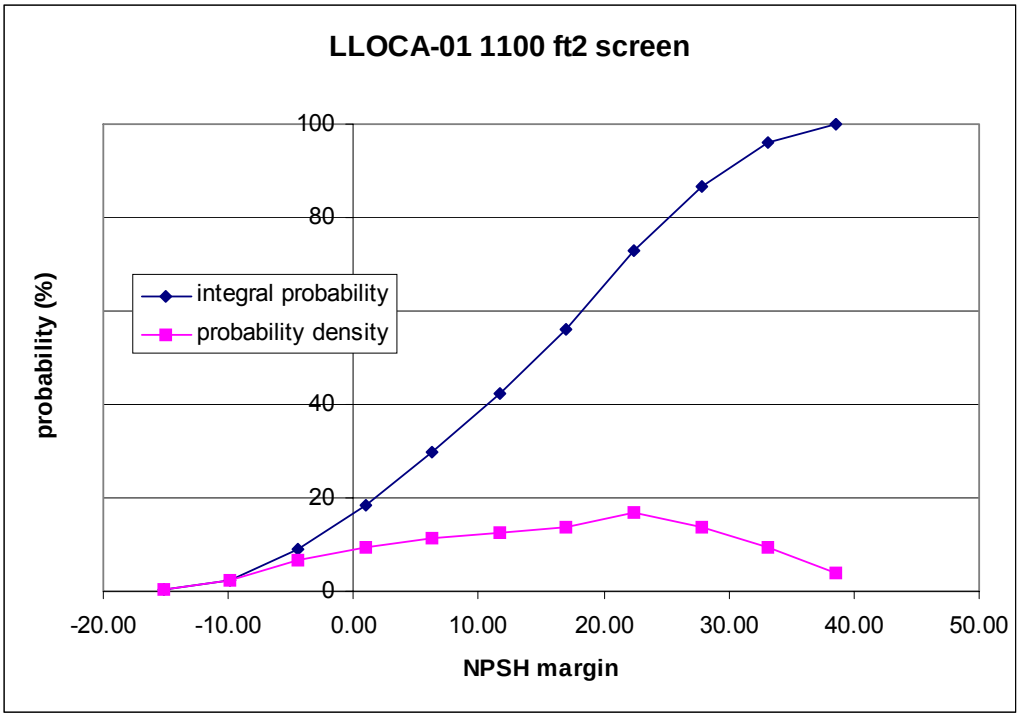

Figure 12 Distributed and cumulative probability of loss of NPSH with a 1100-ft^2 debris screen

The probability density functions for NPSH margin are calculated for all the LOCA event sequences that do not lead to core damage by other mechanisms. Table 5 lists all the LOCA event sequences, as obtained from a SPAR model, and their frequencies of occurrence. For every event sequence that was identified as acceptable before considering NPSH margin, the conditional probability of loss of function was calculated as demonstrated for LLOCA-01 first for a small debris screen and then for a large screen (see Figure 11 and Figure 12). Blank entries under conditional probability of failure in Table 5 indicate that the particular event sequence leads to core damage by other mechanisms.

The unconditional frequency of loss of function due to loss of NPSH margin was computed for each event sequence and each screen size. For every event sequence, the increase in screen size reduced the conditional probability of loss of NPSH margin and, thus, the unconditional

probability of core damage.[6] The last column of Table 5 lists the change in unconditional frequency of core damage due to loss of NPSH margin for every affected event sequence.

Table 5 Calculation of ΔCDF from Conditional Probability of Loss of NPSH and Event Sequence Frequency

Event Sequence Designator from SPAR Model, (ES)	Frequency of Occurrence of the Event Sequence, f(ES)	Small Screen		Large Screen		Change
		Probability of Loss of NPSH Margin in the Event Sequence, p(loss|ES)	Unconditional Frequency of NPSH Loss in the Event Sequence, f(ES)	Probability of Loss of NPSH Margin in the Event Sequence, p(loss|ES)	Unconditional Frequency of NPSH Loss in the Event Sequence, f(ES)	
ROT	1.00E+00					
LLOCA-01	5.00E-06	70%	3.50E-06	18%	9.00E-07	2.60E-06
LLOCA-02	1.00E-07					
LLOCA-03	2.90E-07					
LLOCA-04	1.05E-09					
MLOCA-01	4.00E-05	17%	6.80E-06	0%	0.00E+00	6.80E-06
MLOCA-02	1.90E-07	100%	1.90E-07	100%	1.90E-07	
MLOCA-03	4.60E-10					
MLOCA-04 to 09	2.53E-10					
SLOCA-01	4.00E-04	8%	3.20E-05	0%	0.00E+00	3.20E-05
SLOCA-02	3.31E-06	20%	6.62E-07	0%	0.00E+00	6.62E-07
SLOCA-03	1.03E-06					
SLOCA-04	4.00E-07	100%	4.00E-07	100%	4.00E-07	
SLOCA-05	2.19E-08					
SLOCA-06	4.83E-09					
SLOCA-07	1.48E-09	100%	1.48E-09	100%	1.48E-09	
SLOCA-08	1.20E-11	100%	1.20E-11	100%	1.20E-11	
SLOCA-09	3.24E-12					
SLOCA-10	1.45E-12	100%	1.45E-12	100%	1.45E-12	
SLOCA-11	5.74E-14					
SLOCA-12	1.91E.13					
SLOCA-13	8.00E-07	100%	8.00E-07	100%	8.00E-07	
SLOCA-14	6.64E-09	100%	6.64E-09	100%	6.64E-09	
SLOCA-15	2.05E-09					
SLOCA-16	8.00E-10	100%	8.00E-10	100%	8.00E-10	
SLOCA-17	4.36E-11					
SLOCA-18	1.60E-07	100%	1.60E-07	100%	1.60E-07	
SLOCA-19	7.63E-10					
SLOCA-20	1.60E-08					
SLOCA-21	8.18E-10					
TOTAL			4.43E-05		2.27E-06	4.21E-05

[6] Note again that this is a highly simplified proof-of concept example and that the conclusion of reduced CDF with increased debris screen size is by no means general.

In the last row, the total change in CDF is computed by adding up the changes in frequencies calculated for all the LOCA event sequences. For the simplified model and generic numbers used in the proof-of-concept example, the expected CDF is calculated to decrease by 4E-5 if the debris screen is increased from 125 ft^2 to 1100 ft^2. Again, this value has no significance in the context of GSI-191. The uncertainty bands of NPSH margin in Figure 11 and Figure 12 are, however, remarkable.

5 Conclusions and Future Work

Integrating risk and safety margins accomplishes the objective of establishing a framework that can be used to evaluate the impact of a broad range of plant modifications. Thus, the framework can be used to quantitatively respond to concerns that recent and proposed plant modifications have eroded plant safety margins. The method augments existing decision-making tools when adequate margin cannot be shown through design-basis analyses. This is the case, for example, when new phenomena surface that bring into question the sufficiency of specific safety margins. In such cases, the issue of margin sufficiency arises as "epistemic uncertainty turns into certainty." [24] Furthermore, for reactor designs that differ radically from the current LWR fleet, there is reason to expect broad uncertainty distributions for both load and strength that will require specialized treatment. [25]

Most importantly, integrating risk and safety margins has the accuracy and precision necessary to evaluate the overall impact of a modification that has simultaneous positive and negative safety consequences. This is different from the current realistic treatment of PRAs and the conservative treatment of design-basis analyses. A realistic calculation without regard to uncertainties can be misleadingly optimistic. However, treating the uncertain parameters conservatively can make the picture overly pessimistic. By integrating risk and safety margins, the uncertainty becomes part of the calculated risk metric such that neither benefits nor detriments are exaggerated.

The integration of risk and safety margins is made possible by firming up the definition of safety margin. The adopted definition acknowledges the fact that the reliability of a system or component is not only dictated by the separation between the means of the strength and load

probability density functions but also by their standard deviations. These are, in turn, determined by subsidiary uncertainties and variabilities. Furthermore, having safety margin means leaving room for "unknown unknowns." This requires imposing a safety limit that is conservative relative to the strength probability density function. As defined, safety margin conforms to its traditional role and its current, established usage.

The adopted definition of safety margin makes it possible to obtain the conditional probability of loss of function in any given event sequence. Specifically, the framework assumes that the probability of exceedance of the safety limit is equivalent to the probability of loss of function. This is potentially the most controversial aspect of the framework for integrating risk and safety margins. However, from the perspective of a regulator, it is both wise and prudent to account for "unknown unknowns" not only in design-basis analyses but also in risk-informed regulatory decision-making. If sufficient knowledge exists about both the load and strength density functions, the probability of loss of function can be computed precisely from their convolution, and the balance of the framework still applies.

For any event sequence, the unconditional frequency of function loss is determined by multiplying the unconditional probability of loss of function and the frequency of occurrence of the event sequence. When one looks at all the event sequences that are impacted by a modification, an aggregate metric is obtained. The sum of unconditional probabilities of loss of function for all affected event sequences can be related to CDF if the loss of function can be reasonably tied to core damage. This makes it possible to focus a safety inquiry to examine only a limited portion of the risk space, yet calculate metrics that are suitable for use against existing risk acceptance criteria. When a modification also changes the radiological consequences of an event sequence, successive barriers are used to determine a higher level risk metric that is comparable to the Commission's safety goals. The framework can thus be as narrow or wide as needed given the scope of a specific plant modification.

The framework merges information from all the disciplines that are important in nuclear regulatory decision-making—deterministic calculations, PRAs, materials science, and engineering. The integration uses existing, tested tools and methods, and the integrated framework has the potential to evolve as constituent parts change. Integrating risk and safety margins provides an opportunity to use state-of-the-art techniques from the various disciplines to generate a single metric for the decision-maker. It extends the depth of design-basis

analyses to risk analyses. Best-estimate techniques plus uncertainties are preferred, but simpler approaches can be used when uncertainty is not a determinant factor and when it can be shown that a conservative approach yields conservative predictions. The framework takes advantage of significant advances in risk assessment techniques. When necessary, it is possible to refine both the success criteria and the end states relative to traditional PRAs. The metrics are devised to use existing risk acceptance criteria.

The proof-of-concept demonstration shows the application of the framework to an issue of current regulatory interest. Although the simplified, abstracted model used to determine the effect of increasing sump debris screen size cannot be used to draw any safety conclusions with regard to GSI-191, the example does illustrate some of the advantages of the framework. First, the framework is comprehensive. Applied in its fullest, it integrates every factor with as much detail as available into the decision-making metric. Only incomplete knowledge limits the realism of the framework. Where appropriate, it is possible to report the contributions of aleatory and epistemic uncertainties separately. However, largely because the framework is based on elements that have already been tested within the nuclear industry, the current practice of trading realism for expediency can continue where sufficient margins exist. In lieu of sufficient analytical or experimental data, reasonable conservative assumptions can be made with regard to both probabilities and consequences. Assumptions and simplifications that are routinely used in regulatory analyses remain applicable, and, in its simplest form, the framework reduces to existing regulatory practice.

Further work would be useful in several areas. First, it would be informative to explore if the framework can be advanced to capture changes in operating conditions far from the safety limit, where far is determined relative to the load probability density function. Second, as the report mentioned, integrating risk and safety margins applies in a technology-neutral context, but development of the framework in that direction has exceeded the scope of the current work. It may be fruitful to explore the link between integrating risk and safety margins and the frequency-consequence curve approach proposed for licensing reactors that differ radically from the operating fleet. Because the framework to integrate risk and safety margins is firmly grounded in current regulations, the link may provide a nexus between the present and the novel licensing approaches. Furthermore, because the framework presented in this report quantifies risk for any barrier, it makes it possible to reach regulatory decisions based on the integrity of the first barrier. Where appropriate acceptance criteria exist, such decisions are

desirable because they carry the lowest amount of uncertainty. Also, because risk can be quantified for various barriers and because mitigation and prevention are intrinsically treated within the framework for integrating risk and safety margins, it would be productive to explore the development of relationships that can provide more formal insight into defense in depth.

Acknowledgments

The author gratefully acknowledges feedback received in the course of this work from colleagues in the Office of Regulatory Research and in the Office of Nuclear Reactor Regulation. Discussions with Marty Stutzke, Hossein Nourbakhsh, Charlie Tinkler, Allen Notafrancesco, David Bessette, Mark Kirk, Steve Long, and Gareth Parry have been very valuable in consolidating the concepts addressed in this report. Comments made by Farouk Eltawila and Brian Sheron have impacted the direction and content of this work. The framework has benefited greatly from brainstorming and background material assembled by the SMAP Task Group. Most notably, Javier Hortal (Consejo de Seguridad Nuclear, Spain), Andrej Prosek (Jozef Stefan Institute, Slovenia), Oddbjörn Sandervag (Swedish Nuclear Power Inspectorate), Ramon Lopez Morones (Comisión Nacional de Seguridad Nuclear y Salvaguardias, Mexico), Michel Réocreux (Institut de Radioprotection et de Sureté Nucléaire, France), Bernard Krzykacz-Hausmann (Gesellschaft für Anlagen- und Reaktorsicherheit, Germany) and Martin Zimmerman (Paul Scherrer Institut, Switzerland) have expressed important ideas that have been incorporated into the integrated risk/safety margins framework described in this report.

APPENDIX A: MEASURING LOSS OF SAFETY MARGIN THAT DOES NOT INVOLVE EXCEEDING THE SAFETY LIMIT

The quantification of safety margins discussed throughout this report fails to capture a loss that does not lead to exceeding the safety limit. However, when a plant modification moves operating conditions closer to the safety limit, less distance is available in case of an unexpected event sequence, and thus less margin exists to deal with unknowns. Initial efforts in the development of this framework targeted such modifications. The approach of quantifying changes that occur "under" the safety limit was abandoned because no infrastructure exists that would allow regulators to draw conclusions from such determinations. Some of this early work is discussed in Reference 13.

The definition of safety margin adopted in Chapter 2 remains applicable. For simplicity, let us assume that the barrier (e.g., the fuel) has only one damage mechanism (e.g., embrittlement). Let us also assume that a single safety variable (e.g., PCT) determines the loss of function for this damage mechanism and that it is possible to identify an operating state in which the PCT assumes its lowest value. For currently operating plants, that state would be cold shutdown as defined in the plant's technical specifications. When a plant is in cold shutdown, PCT is at 100 ºC (212 ºF) and has the largest possible safety margin. In cold shutdown the plant has 100 % margin.

In an accident scenario in which PCT increases to 1200 ºC (2200 ºF), the plant has lost safety margin. In a manner consistent with the approach adopted by the framework, the PCT value for an event sequence is the bounding value at the requisite confidence as calculated from the appropriate best estimate plus uncertainty methodology (e.g., the 95-95 value of SUSA). Therefore, for any accident sequences in which the value reaches or exceeds the safety limit, there is 0 % PCT margin.

At any temperature between cold shutdown and the safety limit, the amount of margin is calculated by assuming linear loss of margin as PCT increases. For example, assume that during full power, normal operation the plant PCT goes to 816 ºC (1500 ºF). At this PCT, the plant has lost 65 % of its safety margin. The amount of safety margin available in every event sequence is then combined into an aggregate safety margin.

Assume that a plant has the risk space depicted in Figure 13. For each event sequence, the bounding (e.g., 95-95) value of PCT is calculated from a best estimate plus uncertainty analysis. The safety margin is then computed assuming that full margin exists at cold shutdown conditions, and that no margin exists when the safety limit is reached. The percent margin available in each event sequence is listed together with the probability of occurrence and PCT in Figure 13.

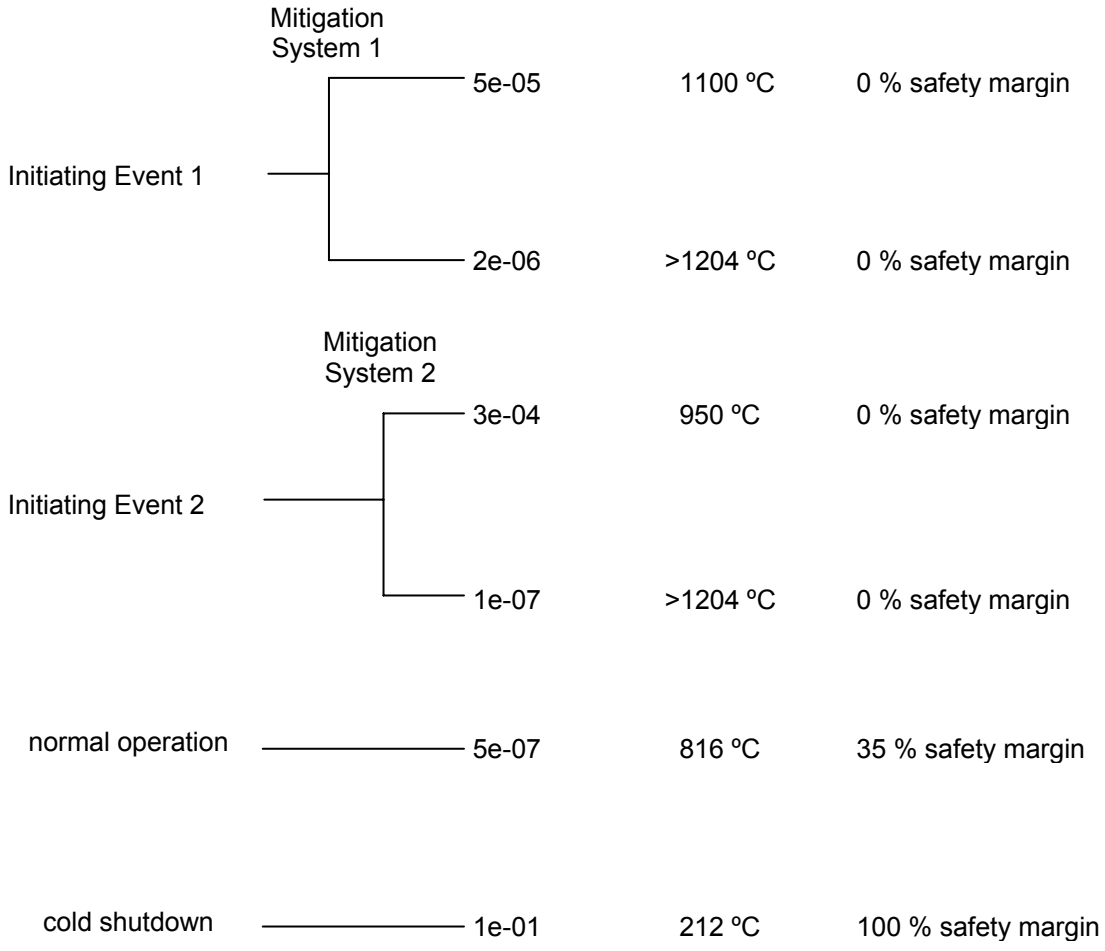

Figure 13 Risk space of a representative reactor

The data is also listed in Table 6. Cold shutdown and normal operation are included to account for all the possible states and obtain a true expectation value for safety margin. The probability of occurrence of each event sequence (including the probability of being in normal cold shutdown or at normal power operation) is used to weigh the safety margin available into an aggregate metric. The last entry in the table is the product between the available margin and

the probability of occurrence of each event sequence. Because the probabilities of occurrence add up to one, the sum of the products over all event sequences is the expected safety margin. The expected margin for the plant of Figure 13 is 42 %.

Table 6 Calculation of Average Safety Margin for the Risk Space of Figure 13

Probability of Occurrence of the Event Sequence	PCT (ºC)	Percent Margin in the Event Sequence	Product of Probability of Occurrence and Safety Margin in the Event Sequence
5.00E-05	1100	9%	4.71E-06
2.00E-06	1204	0%	0.00E+00
3.00E-04	950	23%	6.90E-05
1.00E-07	1204	0%	0.00E+00
9.00E-01	816	35%	3.16E-01
1.00E-01	100	100%	1.00E-01
Expectation Value For Safety Margin			42%

Because the metric includes margin available during the most likely plant state, normal operation, the expected safety margin is insensitive to very rare events. Conversely, increasing the amount of time the plant spends in cold shutdown has a positive impact on this metric, which is consistent with physical expectation. Thus this approach is most useful if the plant normally operates close to a particular safety limit. This is usually the case only for damage mechanisms that are less consequential but may occur more frequently.

The approach can be generalized to multiple damage mechanisms. For each damage mechanism, the analyst must identify the plant state that offers the largest amount of margin, and the safety limit. Because this is a “point-of-information” type metric, or one that the designer can use in optimizing his selections, the assumption of linear erosion of margin as the safety variable approaches the safety limit appears to be acceptable. This assumption can be justified wherever there is no physical change that occurs as the safety variable ranges from its safety value to the safety limit. For a plant that has multiple barriers, the metric can be devised for the ultimate barrier, in a manner that complements the frequency-consequence curve approach. Alternatively, the designer can prioritize design objectives and deal with several margins simultaneously through an importance-weighing process such as that described in Reference 13.

To get a relatively precise estimate of the safety margin expectation value, one will have to develop a PRA and carry out best estimate plus uncertainty calculations for a subset of event

sequences. This process can easily become expensive. Therefore, it is hard to justify such an analysis unless it becomes the basis for a regulatory decision. The absence of acceptance criteria, e.g., what erosion of safety margin is acceptable, precludes using such an approach in regulatory decision making. It is for this reason that the work on margin erosion was suspended and the framework focused instead on quantifying the probability of loosing function caused by the loss of safety margin.

APPENDIX B: THE EQUIVALENCE OF FAILURE RATE AND FAILURE FREQUENCY FOR VERY RARE EVENTS

The distinction between failure rate and failure frequency is the condition of finding the number of failures per unit time. More precisely, the failure rate at time t, $\lambda(t)$, is the number of failures expected between t and t+1 given that no failures had occurred before time t. The failure frequency at time t, f(t), is the expected number of failures between t and t+1 without any presumption on the state at any time. Thus, the requirement of having no failure before time t makes the difference between failure frequency and number of failures per unit time. Therefore, if a system or component can have multiple failures and is reparable, failure rate only applies to the first one.

Because failure rate is conditioned on having no failures occur prior to time t, it is often called the conditional failure rate. If the time interval (i.e., t to t+1) for which the failure rate is determined is small relative to the total time of service (e.g., days vs. decades), the failure rate can be thought of as the conditional probability of failure given that no failure has occurred prior to time t.

It can be shown that, when one is interested in rare events, the failure frequency and conditional probability of failure have the same value. If the failure rate is described by a probability density function with f(t) being the instantaneous probability of failure, the failure rate becomes:

$$\lambda(t) = \frac{dF(t)/dt}{1 - F(t)}, \qquad \textbf{Eq. 1}$$

where F(t) is the cumulative probability of failure which relates to the instantaneous probability of failure via:

$$f(t) = \frac{dF}{dt}. \qquad \textbf{Eq. 2}$$

No assumption is made about the form of the failure probability density function.

With reliability, R(t), is defined as:

$$R(t) = 1 - F(t) \qquad \textbf{Eq. 3}$$

with the derivative:

$$\frac{dR}{dt} = \frac{-dF}{dt}, \qquad \textbf{Eq. 4}$$

one can express the failure rate in terms of reliability as:

$$\lambda(t) = \frac{-dR/dt}{R}. \qquad \textbf{Eq. 5}$$

By rearranging into:

$$\frac{dR}{R} = -\lambda(t)dt \qquad \textbf{Eq. 6}$$

one can express the solution

$$\ln R = -\int_0^t \lambda(t')dt' + R_0 \qquad \textbf{Eq. 7}$$

from which:

$$R(t) = R_0 \exp\left[-\int_0^t \lambda(t')dt'\right]. \qquad \textbf{Eq. 8}$$

From the condition of 100-percent reliability at t=0, the initial condition becomes:

$$R(0) = R_0 = 1 \qquad \textbf{Eq. 9}$$

and thus the cumulative probability of failure at time t is:

$$F(t) = 1 - \exp\left[-\int_0^t \lambda(t')dt'\right] \qquad \textbf{Eq. 10}$$

Considering times, t, that are much shorter than the time between failures:

$$\int_0^t \lambda(t')dt' << 1 \qquad \textbf{Eq. 11}$$

one uses the series expansion:

$$e^{-x} \approx 1 - x + \frac{x^2}{2!} - \cdots. \qquad \textbf{Eq. 12}$$

Then, the cumulative probability of failure is:

$$F(t) = 1 - \left[1 - \int_0^t \lambda(t')dt'\right] = \int_0^t \lambda(t')dt'. \qquad \textbf{Eq. 13}$$

which means that:

$$f(t) = \frac{dF}{dt} = \lambda(t) \qquad \textbf{Eq. 14}$$

The above result is general and only requires that the time t at which the failure probability is determined has to be much shorter that the time between successive failures. This is a

reasonable approximation for very rare events (e.g., the severe accident event sequences that are considered in integrating risk and safety margins).

REFERENCES

1 Appendix K, "ECCS Evaluation Models," to Title 10, Part 50, "Domestic Licensing of Production and Utilization Facilities," of the *Code of Federal Regulations*, U.S. Nuclear Regulatory Commission, Washington, DC.

2 "Quantifying Safety Margins: Application of Code Scaling, Applicability, and Uncertainty Evaluation Methodology to a Large -Break Loss-of-Coolant Accident," NUREG/CR-5249, EGG-2659, 1989 (also Nuclear Engineering and Design, 119, 1990).

3 Wall, I. B. , "Realism in Evaluating Nuclear Hazards," American Nuclear Society President's Special Session, June 15, 2004.

4 "Reactor Safety Study, An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants," WASH-1400 (NUREG-75/014), U.S. Nuclear Regulatory Commission, Washington, DC, October 1975.

5 "Safety Goals for the Operations of Nuclear Power Plants: Policy Statement," *Federal Register*, Vol. 51, p. 30028 (51 FR 30028), U.S. Nuclear Regulatory Commission, Washington, DC, August 4, 1986.

6 Bulletin 2003-01, "Potential Impact of Debris Blockage on Emergency Sump Recirculation at Pressurized-Water Reactors," U.S Nuclear Regulatory Commission, OMB Control No. 3150-0012, June 9, 2003.

7 Regulatory Guide 1.174, "An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis," Rev. 1, U.S. Nuclear Regulatory Commission, Washington, DC, November 2002.

8 Office Instruction LIC-504, "Integrated Risk-Informed Decision-Making Process for Emergent Issues," Rev. 1, U.S. Nuclear Regulatory Commission, Office of Nuclear Reactor Regulation, Washington, DC, December 20, 2005.

9 Parry, G. W., "The Characterization of Uncertainty in Probabilistic Risk Assessments of Complex Systems," Reliability Engineering and Systems Safety, Vol. 54, pp. 119–126, 1996.

10 Krzykacz-Hausmann, B., Hofer, E., and Kloos, M., "A Software System for Uncertainty and Sensitivity Analysis of Results from Computer Models", Proc. Int. Conf. PSAM-II, Vol. 2, Session 063, pp. 20–25, San Diego, CA, 1994 .

11 O'Connor, P .D.T., "Practical Reliability Engineering," Third Edition, John Wiley & Sons, Indianapolis, IN, ISBN 0-471-92902, p. 95, 1991.

12 "Safety Margins Action Plan Technical Note for Task 3: Safety Margin Evaluation Methods," OECD/CSNI/SMAP, draft for NEA/SEN/SINSMAP, 2006.

13 Gavrilas, M., et al., “A Generalized Framework for Assessment of Safety Margins in Nuclear Power Plants,” Proceedings of the International Meeting on Updates in Best Estimate Methods in Nuclear Installation Safety Analysis (BE-2004), Washington DC, November 14–18, 2004.

14 “Contributed Material for Safety Margins Action Plan Task 5: Preparation of the Guidance Document,” OECD/CSNI/SMAP, draft for NEA/SEN/SINSMAP, April 2006.

15 “Review of Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts,” Panel on Seismic Hazard Evaluation, Committee on Seismology, Commission on Geosciences, Environment, and Resources, National Research Council, ISBN 0-309-56207-4, 1997.

16 Parry, G .W., “The Characterization of Uncertainty in Probabilistic Risk Assessments of Complex Systems,” *Reliability Engineering and System Safety*, Vol. 54, pp. 119–126, 1996.

17 EricksonKirk, M., et al., “Technical Basis for Revision of Pressurized Thermal Shock (PTS) Screening Limit in the PTS Rule (10 CFR 50.61): Summary Report,” NUREG-1806, Draft for Peer Review Panel and ACRS Review, November 2, 2004.

18 “Safety Margins Action Plan Technical Note for Task 2: Assessment Process for Safety Margins,” OECD/CSNI/SMAP, draft for NEA/SEN/SINSMAP, 2006.

19 Title 10, Section 50.59, “ Changes, tests and experiments,” of the *Code of Federal Regulations*, U.S. Nuclear Regulatory Commission, Washington, DC.

20 Izquierdo, J.M., Hortal, J., Meléndez, E., and Sánchez M., “An Integrated PSA Approach to Independent Regulatory Evaluations of Nuclear Safety Assessment of Spanish Nuclear Power Stations,” presented at the IBC’s 7th conference on PSA in the Nuclear Industry, Café Royal, London, November 26–27, 2001.

21 Oberkampf, W., Helton, J., Sentz, K., “Mathematical Representation of Uncertainty,” AIAA Proceedings of Non-Deterministic Approaches Forum, Reston, VA, 2001.

22 Rao, D.V., et al., “Knowledge Base for the Effect of Debris on Pressurized- Water Reactor Emergency Core Cooling Sump Performance,” NUREG/CR-6808, LA-UR-03-880, Los Alamos National Laboratory, February 2003.

23 “Parametric Study of the Potential for BWR ECCS Strainer Blockage Due to LOCA Generated Debris,” NUREG/CR-6224, Los Alamos National Laboratory, October 1995.

24 Parry, G. W., personal communication, November 23, 2005.

25 Pagani, L. P., Apostolakis, G.E., and Hejzlar, P., “The Impact of Uncertainties on the Performance of Passive Systems,” *Nuclear Technology*, Vol. 149, pp.129–140, February 2005.