Organisation de Coopération et de Développement Economiques
Organisation for Economic Co-operation and Development

30-Aug-2006

English - Or. English

NUCLEAR ENERGY AGENCY
COMMITTEE ON THE SAFETY OF NUCLEAR INSTALLATIONS

NEA/SEN/SIN/SMAP(2006)2
Unclassified

Task Group on the CSNI Safety Margins Action Plan (SMAP)

SMAP TECHNICAL NOTE

Task 2: Assessment Process for Safety Margins

JT03212845

Document complet disponible sur OLIS dans son format d'origine
Complete document available on OLIS in its original format

English - Or. English

FOREWORD

Recent NPPs operating experience shows that in some cases operational and design modifications may lead the plant far away from the original design. Power uprates, life extension or increased fuel burnup as well as cumulative effects of simultaneous or subsequent design changes in a plant, which can be larger than the accumulation of the individual effects of each change, can challenge original safety margins while fulfilling all the regulatory requirements. It has been recognised that currently used methods for safety analysis may not be sufficient to guarantee that enough safety margin exists.

To address this problem the CSNI approved in December 2003 an Action Plan on Safety Margins (SMAP) and established an international Working Group aimed at developing a framework for integrated assessments of the changes to the overall safety of the plant as a result of simultaneous changes in plant operation/conditions. The SMAP plan consists of five tasks:

Task 1 : Definition of Safety Margins and Related Concepts
Task 2 : Assessment Process for Safety Margins
Task 3 : Safety Margin Evaluation Methods
Task 4 : Quantification of Safety Margins
Task 5 : Preparation of a CSNI Guidance Document.

This Technical Note, which represents a SMAP working document, is the result of the SMAP Task 2 - Assessment Process for Safety Margins consisting of

Sub-task 2A: Develop guidance for building a complete set of representative initiators and associated event trees using existing PSA trees and design basis events as a starting point
Sub-task 2B: Determine changes in representative initiators and associated event trees due to changes in the plant (e.g., ageing, uprates, and operations modifications).
Sub-task 2C: Determine what kind of information should be obtained to quantify generalised safety margins.

Although this paper is an outcome of extensive discussions of the whole SMAP Group, the special appreciation belongs to J. Hortal (CSN, Spain) and M. Gavrilas (US NRC) who provided the first outline of this document, as well as to R. Lopez Morones (National Commission of Nuclear Safety & Safeguards, Mexico) who provided valuable help in finalizing it.

TABLE OF CONTENTS

1. INTRODUCTION

The main objective of the Task Group on the CSNI Safety Margins Action Plan (SMAP) is to develop guidance on how to assess safety margins in nuclear power plants. The proposed SMAP relies on the premise that an adequate combination of deterministic and probabilistic methods can provide the best achievable framework for solving the safety margin assessment. Five major tasks were defined to fulfill the SMAP objective. Task 1 is the *Definition of safety Margin and related concepts* which is already finished [1,2] and task 2 is the subject of the present report, related with the use of extended PSA methods to assess safety margins. The third task is devoted to identify the *safety margin evaluation methods* and to provide guidance on how to apply computational analysis capabilities to estimate several aspects of the safety margins evaluations such as plant behavior, determination of suitable criteria for barrier integrity and frequency quantification, taking into account existing uncertainties. The last tasks are related with the *quantification of safety margins* under a deterministic and probabilistic approach and the *preparation of a CSNI guidance document*.

The second task of the SMAP was divided in three subtasks. The first subtask (2A) provides guidance on how to build a complete set of representative initiators and associated event trees, using existing analyses of design basis accidents and PSA models as starting points, to get the capability to address the assessment of safety margins. The second subtask (2B) gives guidance on how to represent proposed plant modifications on this plant safety description in order to compare the "before" and "after" status of the safety margins. The kind of information that will be needed to quantify the effects of the plant modifications on the safety margins is the purpose of sub-task 2C.

2. GENERAL CRITERIA FOR DEVELOPING AN INITIAL SET OF RISK SPACE EVENT TREES

As explained in [1], the assessment of generalized safety margins requires consideration of all possible scenarios having non-negligible frequencies of occurrence; this almost complete set of scenarios was named the *risk space* and is described through a set of PSA-like event trees which provides capability to analyze multiple safety objectives. Sub-task 2A of SMAP can be understood as the development of an initial set of event trees and their corresponding fault trees, able to provide the risk profile of the plant before significant design and/or operation changes are implemented. Thus, this set of event and fault trees for the "before" status of the safety margins can be properly named the ***base case risk space***.

The development of a base case risk space is in some aspects similar to the event tree delineation in classical PSA, but the capability to address different safety objectives and to evaluate generalized safety margins introduce additional requirements that result in important methodological differences. General guidance for the determination of the base case risk space can be found along this section. Methodological details may depend on the specific plant technology and should be developed on a case specific basis.

2.1. The capability of traditional methodologies to characterize the plant safety

As discussed in Reference [2], licensing is usually based on a conservative analysis of plant physical responses to specific challenges. For example, the licensee in the U.S. is required to demonstrate the plant's capability to achieve "success" despite:

- physically conservative assumptions in the analysis,

- concurrent loss of offsite power,
- concurrent limiting single failure.

Showing this capability for a comprehensive set of demanding challenges of different severity and likelihood (i.e., the design basis transients and accidents) is a deterministic demonstration of plant safety. The acceptance criteria which define the "success" are dependent on the type of challenge being analyzed. The complement of equipment needed is then subject to many programmatic requirements, including special treatment requirements. The required capability is maintained operationally through compliance with technical specifications that deal for example with functional availability, surveillance and testing.

A key element of the demonstration is the idea of "margin." The physically conservative assumptions mentioned above include such things as conservative values of decay heat, conservative assumptions regarding the timing of events, generally unfavorable assumptions about actuation set points, and safety limits related to design limits that reflect factors of safety. Partly because of this conservatism, licensing demonstrations of plant capability in responding to design-basis events are considered to be robust.

Unfortunately, while significant plant capability is included within the scope of this demonstration, single-failure-proof response to design-basis events would not by itself guarantee a risk profile that would be considered satisfactory. Experience has shown that out-of-design situations are not as unlikely as expected and the deterministic approach was complemented with risk insights based on the PSA technology. PSA takes credit for success paths that are not part of the design basis capability, including interventions of non-safety equipment, or safety-class equipment in situations not necessarily contemplated in the design basis as well as operator actions beyond those (few) contemplated in accident analysis. At the same time, it allows for the possibility that the equipment relied upon in the safety analysis will not perform its intended function. In some cases, the most conservative assumptions used in the design basis accidents are also relaxed. As a consequence, some success paths in risk analyses are not as robust as the design basis success paths, but they give a valuable contribution to the risk profile description.

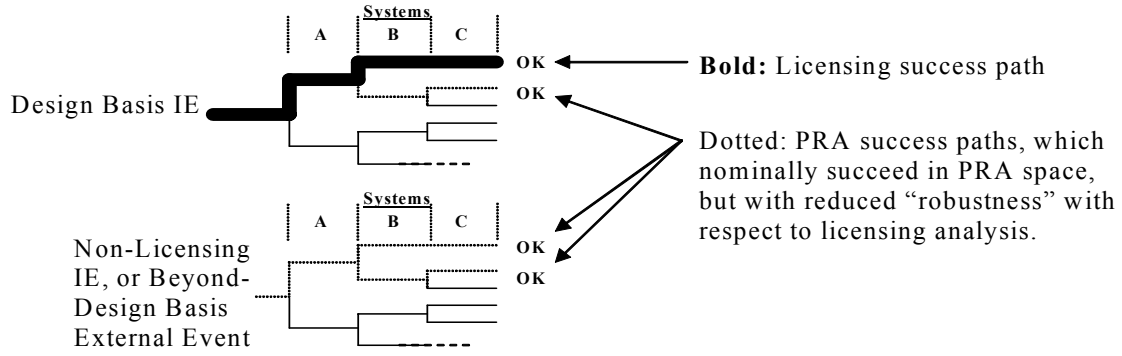

Figure 1. Design Basis (Licensing) Approach vs. PSA Approach.

Figure 1 shows that typical PSA event trees include some sequences whose success directly derives from the licensing analysis, because a more demanding design basis accident has been shown successful. These sequences are normally located at or near the upper part of the event tree, since they imply a low number of safety equipment failures. The same figure shows that the event tree also includes many other success paths farther down.

2.2. The need to expand PSA models

Existing PSA approaches, although providing valuable insights to the plant risk profile, have some limitations. On one hand, the idea of "margin" is not used and, consequently, the PSA results and how they are impacted by plant changes, should be considered in a mostly qualitative way. On the other hand, among the several safety objectives addressed by licensing analyses, only the potential for severe accidents and their consequences is analyzed in PSA. Less severe barrier failures leading to relatively small but potentially frequent radioactive releases are not included under the scope of PSA. It is then necessary, according to the proposals in SMAP task 1 [1], to extend the PSA methods to the risk space in order to get capability to address at least the same safety objectives considered in licensing analyses and to evaluate generalized safety margins.

To determine the safety margins for a set of events, it is necessary to have a quantitative measure of the plant response given by deterministic simulation models.The responses will vary over the spectrum of events. To cover the spectrum of events considered in the risk space, all systems in the event trees must be included in the deterministic model. The model must also compute the safety variables relevant to the safety inquiry (i.e., fuel temperature, clad oxidation, containment pressure, etc.). Once the deterministic model is developed, all relevant sequences can be simulated.

While the focus of this step is on the deterministic analysis, it also includes an iterative process with the risk space event tree delineation. The results from the deterministic analysis may trigger new barrier damage mechanisms or safety system failures, which then require iteration back to the event trees and event re-qualification.

Quite generally, the safety objectives to be analyzed with a risk space analysis model determine the level of detail in the definition of end states. In other words, there is a close link between safety objectives and end states since the end state of a sequence is determined by the safety objectives that have been exceeded. Moreover, safety objectives and end states also condition the level of modeling detail both in dynamic (sequence simulation) and reliability (fault tree) aspects. The selection of initiating events for the event trees is also conditioned by the scope of the analysis. These interactions are discussed in more detail in next section. For example, if one is quantifying core damage frequency, then it may suffice to characterize scenarios simply in terms of whether they lead to core damage.

2.3. Addressing multiple safety objectives

A key point to be discussed is whether or not it is possible to develop a unique base case risk space, i.e., a unique initial set of event trees for a given plant. As above noted, without taking some precautions, a particular choice of initiating events and event tree headers may condition the type of safety objectives that can be analyzed and, therefore, the type of safety inquiries that can be solved. For example, the event trees in a typical level 1 PSA are intended to address the safety objective of maintaining coolable geometry in the reactor core, which is ensured if the sequence success criteria (which are coincident with the LBLOCA design basis acceptance criteria) are not exceeded. However, the same trees cannot be used without extensions, changes or further development to address other safety objectives such as different barrier failures (either failures in other barriers or other types of fuel integrity losses) or their radiological consequences. If one is quantifying the frequency and severity of radiological releases in a severe accident, more detail is necessary even in the core damage model, because the phenomenology of the containment depends on certain characteristics of the scenarios leading to core damage. The application determines the end states, and the end state definitions then determine the success criteria that are the basis for classifying scenarios. As another example, consider the full spectrum of fuel failures from pinhole leaks to catastrophic fuel melt and major core damage. If the application is the assessment of intermediate radionuclide releases, as those allowed for design basis accidents of very low frequency (Condition 3, as

described in [2]), then the end states will be different and possibly more refined compared to those for core damage. An end state could be incipient cladding embrittlement. These newly defined end states then will dictate the appropriate acceptance criteria and the risk space model's success criteria and deterministic attributes. The success criterion could be the decoupling criterion described in Reference [2], namely that "the total number of rods affected by DNB must be less than 10%."

Potential restrictions in the type of safety objectives that can be analyzed with a particular set of event trees come from three main sources:

- Selection of safety functions and associated systems
- Grouping of initiating events and subsequent transient paths
- Dependency of fault tree structure on sequence success criteria

These sources are now discussed in more detail.

2.3.1. Selection of safety functions and associated systems

Event tree headers represent protective features aimed at mitigating the consequences of an accident or at preventing the degradation of the accident evolution. Independently of further refinements due to feedbacks during the process, one of the first stages in the development of event trees is the selection of adequate headers. They should be able to logically characterize different accident paths in order to quantify their frequencies as a function of the frequency of the initiating event and the probabilities of involved headers.

In principle, only those safety functions which are relevant for the safety objective to be assessed are used as event tree headers. For example, in level 1 PSA, containment safeguard headers are not considered (unless they have also some effect on the reactor core response) since the containment failure is not analyzed in level 1. In addition, the selected headers are only detailed to the extent needed for the assessment of the safety objective. For example, the "Reactor Trip" header in a PWR level 1 PSA represents the failure or success of the trip breakers without taking into account the details of the reactor trip signal generation. It is implicitly assumed that, even if the expected primary trip signal fails, other trip signals will be activated and the probability of failure to generate a reactor trip signal is negligible. The consequence of a reactor trip function failure is to delay the reactor trip, probably causing local fuel cladding failures due to critical heat flux (CHF) phenomena or fuel centerline melting in some fuel rods. These failures, however, do not affect the subsequent evolution and the fulfillment of the PCT or oxidation criteria which define the sequence success.

Stated in other way, it is not enough to know whether a given system "succeeded;" it is necessary to know whether more trains than the minimum were successful, and whether they were actuated at the last possible moment or with time to spare. Iteration with the deterministic evaluation is needed to determine the appropriate level of detail in modeling system "failure."

As another example, consider an auxiliary feedwater system with three trains. Only one train might be needed for success (no core damage) in the context of standard PSA. However, when considering safety margins, it is important to consider the probability and implications of two or three trains being successful so that proper weighting of the frequencies can be performed; this is consistent with obtaining true risk metrics. Also, one train may take longer to activate. This could be an important consideration for plant changes that compromise the time to manually actuate a given train.

It is, therefore, essential in the development of the event trees describing the risk space to keep in mind all the potential safety objectives to be covered and to include all the pertinent safety functions and

systems in the event tree headers. To the extent possible, it is recommendable to develop a single set of event trees, with adequate level of detail for any safety objective, even though in a particular analysis some headers could be superfluous. However, if an excessive level of detail leads to the generation of unpractical event trees, the development of different sets of event trees with specific headers for different kinds of safety objectives, can be considered. In this case, the consistency between the different sets should be carefully ensured and the use of contradictory assumptions in different event tree sets must be avoided.

2.3.2. Grouping of initiating events and subsequent transient paths

Initiating events are the root of event trees. Therefore, to speak about number and type of initiating events is equivalent to speak about how many event trees are needed to describe the risk space.

An important step in the development of the risk space event trees (as it is also in the development of PSA event trees) is, therefore, the identification and grouping (or binning) of initiating events. The result of this step is a set of events severe enough to challenge the safety functions, frequent enough to contribute at some significance level to risk metrics, and different enough from each other to affect the modeling of plant response, so that a different event tree model is used for each initiating event. In essence, what is called an initiating event is really a set of events that has been grouped for purposes of that analysis.

The typical PSA grouping of initiating events is not necessarily optimal for margin assessment and it could even vary when assessing different safety objectives.

As stated in NASA's PRA Procedures Guide [3],

> *"When different IEs are combined into a representative group, the frequency or probability of occurrence for the representative group is the sum of the individual frequencies or probabilities of each group member. Since not every member of the group will cause the exact same response from the system, typically the impact of the representative group is modeled as the most severe perturbation caused by individual group members. This technique is conservative."*

The conservatisms that NASA identifies can be reduced by further dividing the "representative group" into smaller groups. However, the conservatisms can not be entirely eliminated. The cost is an increased computational effort.

Therefore, the definitions of certain initiating events (groups) may also need refinement in light of the more stringent requirements placed on the margins analysis. A broadly defined event category may suffice for purposes of evaluating CDF, but there may be more or less margin in selected subclasses of events within this category. Moreover, the need to characterize change in margin may lead to the need to focus on a narrowly defined event category that practically requires introducing a new initiating event.

An example of possible new initiating events comes from the NRC study of alternatives to the current DBA large-break LOCA with its associated margin requirements (e.g., PCT less than 1204ºC). Finer division of initiating events with different pipe sizes, pipe locations, and failure characteristics may be important in risk-informing the restructuring of 10 CFR 50.46.

Whether two initiating events can be considered similar (then grouped) or not, depends on the safety objective being analyzed. A typical initiator in level 1 PSA is the turbine and reactor trip, but it is usually intended to cover a number of generic transients leading to reactor trip which may include very different transients like loss of feedwater heating, loss of primary flow or control rod malfunctions. The plant response to all these transients is adequately covered by the turbine trip when the safety objective is the

assessment of severe core damage. However, if the safety objective is the analysis of DNB as a proxy for local cladding failures, differences in the transient before and immediately after the reactor trip are very important and the sole consideration of the turbine trip leaves essential parts of the analysis out of scope.

As indicated in the above NASA's quotation, the consequences of a group of initiators should be modeled as those of the most severe member of the group. This means that the selection of the representative initiator of a group (i.e., the case to be analyzed) should be based on enveloping criteria, i.e., trying to maximize the consequences with respect to the safety objective. A consequence is that, even if the same grouping is valid for different safety objectives, the analysis case of each group can vary when the safety objective changes. This can make difficult the desirable aim of using the same set of event trees for any safety objective.

In order to cover the whole risk space, the set of initiating events must be always complete in the sense that any possible real initiator should be adequately represented by an analyzed initiator whose frequency is the collective frequency of all the events under its scope. These criteria have been followed in current approaches to nuclear safety. This suggests that the set of initiating events that should be used to develop the event trees describing the risk space should take as main references both the Design Basis Events and the PSA initiators.

Both sets of initiating events try to be complete, except for some low frequency events which have been intentionally left out of the design basis scope as, for example, ATWS. However, the two sets are quite different due to the differences in safety objectives, which lead to differences in the level of detail of event grouping. It could be said that the set of PSA initiators is more complete but, on the other hand, the set of design basis events is more detailed.

In the development of the risk space event trees, the completeness of the PSA initiators should be combined with the level of detail of design basis events, although care should be taken to avoid overlapping or redundancies in the resulting set of initiating events.

Like in traditional PSA, all the operating conditions should be represented in the set of initiators. Every combination of initial conditions and initiating event having a non-negligible probability should be covered. Although steady state initial conditions are much more frequent, the possibility of the initiating event occurring during an operational transient (typically, a load change) should be considered[1], especially if the plant transient makes the initiator more likely or if the non-steady initial state can contribute to significantly worse consequences of the event. Initiating events in shutdown modes will also be considered.

A similar discussion applies to the grouping of individual accident paths into event tree sequences. An event tree sequence does not represent a single transient but a set of transients composed by the same combination of event tree header states, starting from any initiator in the group of the event tree initiator, from initial conditions within a given neighborhood of the analyzed case and possibly with timing differences in safety function interventions. Again, the grouping criteria can be conditioned by the safety objective. For example, a success sequence in traditional PSA contains only success paths with respect to severe core damage criteria but may contain significant amounts of both success and failure sequences with respect to other safety objective (e.g., avoiding excessive fuel centerline temperature). If it is found that the same sequence groups transients exceeding the safety objective together with others remaining away from it, a further division of that sequence by introducing new headers could be considered. This is further discussed in section 2.4. below.

[1]This practice is usual in DBA analyses.

2.3.3. Dependency of fault tree structure on sequence success criteria

In current PSA methodologies an event tree header, representing the intervention of a system or an operator, is considered successful if the safety function is adequately performed, i.e., if well defined safety function success criteria are fulfilled. Safety function success criteria are defined in terms of the outcome obtained from the systems performing that function and the time delay of the safety function onset.

However, the safety function success criteria depend on the sequence success criteria. In other words, whether a safety function is successful or not, is decided according to its capability to prevent the exceedance of the sequence success criteria. This clearly indicates that a change in the sequence success criteria may affect the safety function success criteria.

In some cases, the success or failure of a safety function is equivalent to the logical state of a system or component. For example, if a pressure relief function is performed by a valve, the safety function will be successful whenever the valve opens upon request from an automatic signal or an operator action, at least under a wide range of operating circumstances. In other cases, the state of the system does not provide enough information because the safety function success depends also on the dynamic state of the plant. This could be the case of a water injection system whose success is defined in terms of injection flowrate, which depends not only on the state of pumps and valves in the system but also on the counterpressure of the system receiving the flow.

The potential failure of each safety function is modeled by a fault tree, i.e., a boolean description of the safety function failure in terms of more elementary events. Since the safety function success depends both on the sequence success criteria and on the plant dynamics, the fault trees have, in the general case, the same dependencies. The fault tree structure not only reflects the number and type of working devices, components or structures needed to get success but also the dynamic conditions that can affect the success of the safety function which are usually modeled as *house* events (user-defined sequence-specific boundary conditions).

A set of event trees and fault trees where these dependencies are present are only usable to assess the sequence success criteria for which they were developed. If a different safety objective is being assessed, a careful revision of the fault tree structure is needed. This poses a serious difficulty on the use of existing PSA event trees and fault trees to evaluate generalized safety margins in the risk space.

There are two possible ways to solve this problem. The first one is to develop different sets of event trees / fault trees for each safety objective to be analyzed. The second one is to define the event tree headers at system level rather than at safety function level. This way, the fault trees would model only the system states and the sequence success will no longer be an exclusive function of the logic combination of header states but the result of the deterministic calculation of a well selected transient used as the sequence representative[2]. Taking into account that the dynamic verification is needed for the evaluation of the generalized margins, the second option seems more recommendable.

2.4. Sequence delineation and classification.

The sequence delineation consists of identifying which event tree headers are hit as a consequence of the plant evolution after the initiator and subsequent events. Under each identified header, the sequence is split to take into account the failure or success of that header. Most PSA event trees contain only binary branching points, i.e., only a success and a failure branch are considered at each branching point. There is

[2]In some cases the final state with respect to a particular safety objective can be determined a priori and the deterministic calculation is not needed.

no fundamental restriction for the use of higher level branching in order to consider, for example, different types of header failures which could give rise to different accident evolution thereafter. However, this alternative makes the automatic processing of the event tree more difficult.

As a general criterion, event tree sequences should be developed along the time until a long-term safe steady state has been reached or until the sequence success criterion has been exceeded. In addition, some operational conditions necessary leading to the exceedance of the success criterion can be used as surrogates for the damage state. For example, if the available water inventory for safety injection is not enough to guarantee long term core cooling, the sequence can be classified as exceeding the severe core damage criteria. If several safety objectives are being considered, each one associated to a success criterion, the sequence development should continue until the steady state is reached or until all the success criteria or their surrogates have been exceeded. All the headers visited along the sequence should result in a branching point. The consequences of failing to open a required branching point are discussed below. It is current practice in PSA to limit the sequence development to a given mission time. If a damage state or its surrogates have not been reached at the mission time, the sequence is considered successful with respect to that damage state, even if the final state is not a steady state.

The main difficulty in sequence delineation is to identify when a header is challenged and, in consequence, when a branching point should be introduced. Without relying on simulation resources, it is, in general, difficult to identify when a setpoint or an operating instruction are reached. This difficulty is even higher when the safety objective being analyzed depends on events and phenomena occurring in the short time scale. Lack of accuracy in the identification of branching points can result in important errors in the analysis results.

As an example, let us consider a mitigating system represented by a particular header. Let us also assume that, in a particular sequence where this system would have enough capability to guarantee the sequence success, it is erroneously determined that the corresponding setpoint is reached. A branching point is then introduced. The success path assumes that the system performs its safety function and the end state will be successful with a conditional probability close to 1. Instead, the failure path leads to an unsuccessful end with a high conditional probability, although its product by the system failure probability results in a small probability of the damage state. However, if the setpoint is not actually reached, the system will not come into play and no branching point will be added. This means that the dynamic evolution is similar to the failure path, but the probability of this evolution is now 1. A corollary is that the introduction of fictitious headers representing safety functions gives unconservative results. On the other hand, if a safety function actually demanded during a sequence is not considered in the event tree, the consequences assigned to that sequence will be overestimated.

As explained in section 2.3.2 above, since a sequence is actually a collection of different accident paths, it could happen that both success and failure paths are grouped under the same sequence. Formally, one ought to reflect in the analysis the probability that a given success sequence might contribute to failure frequency, even if the hardware and operators perform nominally in the scenario. This probability is given by the “fraction” of sequence paths whose end state actually differs from the sequence end state. If this probability is omitted, then it is implicitly argued to be negligible. The following is the criterion for approximating the failure probability of a success path that is being invoked:

The probability of functional failure of a success sequence, conditional on hardware and operator success, can be neglected in the calculation, provided that it is dominated by the probability of operator or hardware failure. Symbolically,

$$P(\text{limits exceeded} \mid \text{hardware \& operator success}) << P(\text{hardware or operator failure}). \quad (1)$$

Few existing PSAs systematically address this point. PSA guidance (including the recent ASME standard [4]) typically recommends avoiding the systematic conservatisms called out above in analyzing success criteria.

As a general criterion, if a success sequence includes a significant number of paths leading to damage state, new headers should be introduced in order to discriminate actual success and damage paths. In addition, it is recommendable to develop quantification methods and algorithms allowing for a better account of the fraction of each sequence frequency actually contributing to the exceedance frequency of the safety objective. On the other hand, a sequence nominally ending in a damage state needs not be further developed even if it contains some success paths with respect to that damage state; nonetheless, it is recommendable to adequately discriminate these success paths since failure to do it may result in an unacceptable lack of optimization of the event tree. In both cases, no action is needed if the collective frequency of the wrongly classified paths is only marginal according to expression (1).

The final result of the delineation process is the set of event trees describing the risk space. As described above, the process starts from the existing PSA event trees and the Design Basis Accidents. However, it is not clear that PSA sequences and DBAs will be subsets of the resulting event tree sequences. Any PSA sequence will be covered by one or more risk space sequences but the possible introduction of new event tree headers in the risk space will make the correspondence difficult to identify. On the other hand, some of the risk space sequences will be covered by Design Basis Accidents but probably none of them can be identified with the analyzed cases in the plant Safety Analysis Report. This reflects only the different approach used, not an inconsistency between the risk space description on one side and the PSA and DBA on the other.

2.5. Treatment of stochastic phenomena

The plant behavior during the development of an accident scenario cannot be deterministically predicted. Apart from the possibility of random safety function failures, which are modeled by fault trees whose top events are used as event tree headers, the accident evolution may result in conditions allowing the occurrence of stochastic phenomena. A typical example is the hydrogen combustion during a severe accident. While this phenomenon cannot occur if the relative concentrations of gases in the containment do not match some conditions, the fulfilment of such conditions does not mean that the phenomenon will necessarily occur. Rather, the hydrogen combustion is characterized by a probability which is a function of the containment conditions.

Typically, stochastic phenomena are not treated in level 1 PSA but they are a common ingredient in level 2. Given the usually high dependency of the probabilities of stochastic phenomena on the dynamic conditions, fault trees are not the best option to deal with them. In level 2 PSA, the influence of safety functions and systems not included in level 1 is modeled as an extension of the level 1 event trees with their corresponding fault trees for the new headers. However, a different tool called Accident Progression Event Tree (APET) is typically used to address stochastic phenomena. They are an alternative format for

event tree representation where the branching point associated to a header (now called "node") is allowed to generate more than two output branches whose probabilities are no longer derived from fault trees but from other types of models better accounting for the dynamic dependencies.

In the description of the risk space, the use of one or more sets of event trees / fault trees, similar to those of level 1 PSA, has been proposed but, at the same time, the need for a detailed dynamic verification of the sequences has been stressed. Also, it has been said that some event tree headers in level 1 PSA contain dynamic dependencies which are usually modeled as *house* events in the header fault tree. Extensive dynamic sequence verification allows for the explicit and detailed accounting of these dependencies. This way, the system states are still modeled through boolean fault trees while the dynamic dependencies are removed from the fault tree structure. An interesting consequence is that stochastic phenomena can now be modeled as event tree headers where the explicit dependency on system behavior is weak or null but the dependency on sequence dynamics is high. This way, in the risk space description, event tree headers are not restricted to safety functions or systems performing safety functions but they may also include stochastic phenomena which can significantly alter the course of the accident.

In summary, the extension of the event tree header concept and the dynamic sequence verification allow to unify the typical methods of level 1 (pure boolean event/fault trees) and level 2 (APET) PSA, giving rise to the concept of *Dynamic Event Tree*.

3. MODELING PLANT CHANGES IN THE RISK SPACE

The second step in the assessment process for safety margins is the development of a methodology for determining changes to the base case risk space in order to get a set of event trees and fault trees that represent proposed plant modifications to be analyzed. This task is addressed in this section and their final outcome is to provide general guidance for the determination of the ***modified case risk space***, i.e., for the "after" status of the safety margins. Plant modifications may impact some aspects of the base case risk space ("before" status of the safety margins) which may include tree catalogue and structure, consequences of tree sequences or frequencies. The assessment of plant changes will not generally require a reanalysis of the whole base case risk space but only of those aspects previously identified as affected by the change, from which the change in safety margins can be evaluated.

There are a variety of changes to the plant that are of interest to safety stakeholders. It is important to keep these in mind when exploring the changes that will be needed in a deterministic-probabilistic approach to properly reflect the plant changes. Changes, individually or as a group, include the following as examples:

- extended power uprates
- aging of plants, especially in the context of regulatory approval for extended operation
- increased fuel burnup
- increased refueling cycles, for example, up to 24 months
- converting to MOX fuel
- external pressures on the nuclear power industry to be cost-competitive
- relaxation of regulatory requirements, including technical specifications by risk-informing the process

The first step in computing the changes to safety margins is to define the safety inquiry. The capability of the risk space model to be used for the inquiry, is assessed and, if necessary, it is modified to

be consistent with the scope of the inquiry. The level of detail needed is also assessed. For example, the inquiry may be limited to a determination of the change in the margin to extensive fuel failure, as represented by peak clad temperature (that would otherwise lead to major core damage), brought about by a design change in an accumulator. Another example inquiry might be to determine the effect of a power uprate specifically in the change to peak clad temperature. Thus, impacts of the change on other barriers or concerns regarding the change in source term would be excluded from the inquiry, greatly simplifying the study. On the other hand, an inquiry may have as its goal the change in the global plant margin, for a combination of changes to the plant, considerably expanding the process. The global plant margin requires an understanding of all barrier failures and the transport of the radionuclide source term to the environment.

Figure 2 shows a flow diagram describing the steps needed to compute changes to plant safety margins for a given safety inquiry. This approach to a safety inquiry is presented as an example. It focuses on addressing barrier states and associated safety metrics. A more generalized approach can be developed which is more global and would address such issues as changes in dose to the public. Each country licensee/regulator can develop a customized approach to best meet its needs.

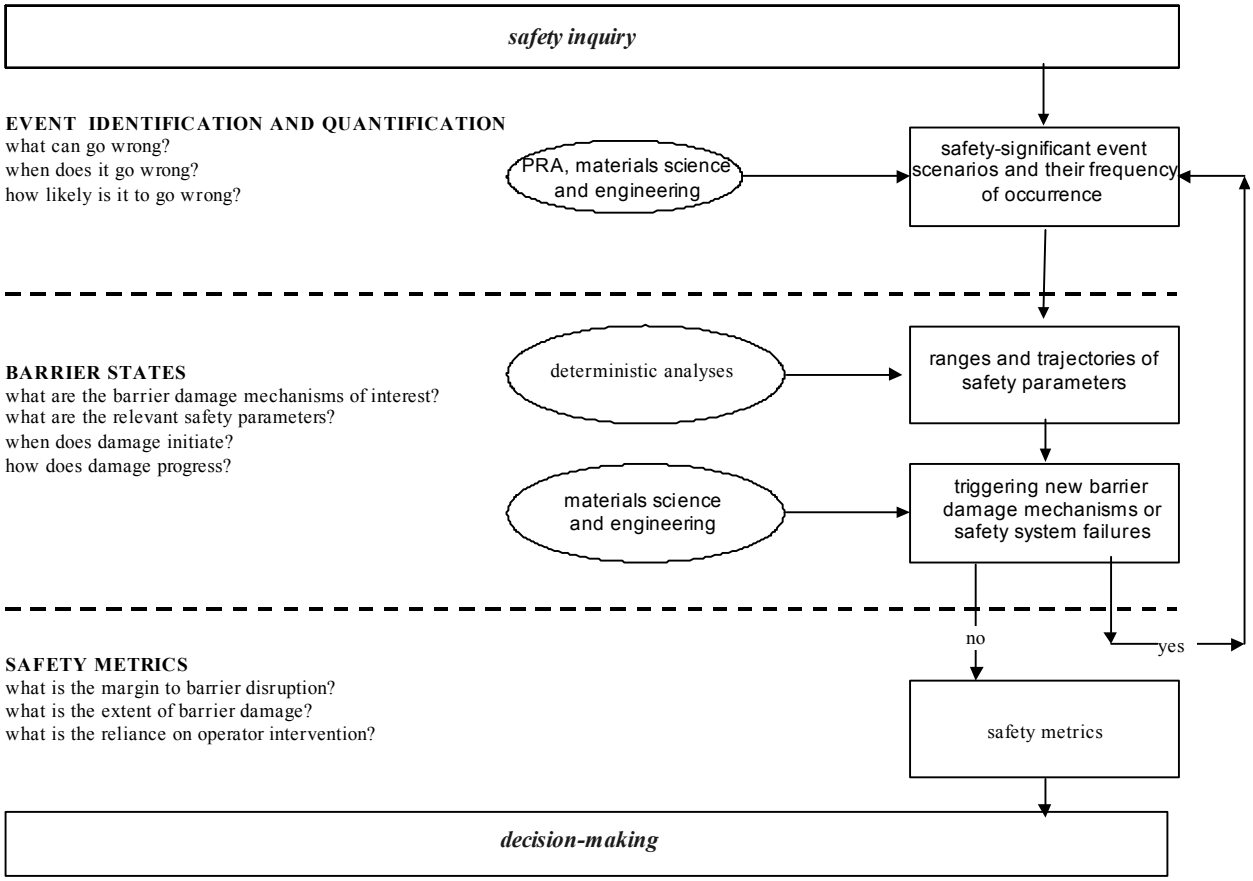

Figure 2. Implementation of the Proposed Safety Framework

One could say that changes in the licensing basis could be addressed by performing a complete safety analysis before and after the changes of interest are made. While formally correct, this way of thinking about changes in margin is not practical; it is necessary to focus effort on changes in certain areas. This approach also has the disadvantage of large uncertainty in the high-level figures of merit that are the basis for the decision.

Taking power uprates as typical examples of candidate plant changes for safety margin assessment, licensing analysis is aimed at confirming that the before-uprate licensing success paths are still success paths after the uprate and that no sequence matching the design basis assumptions escapes the envelope provided by the licensing success paths.

In the risk space, since there could be a large number of success sequences potentially affected by the change, examination of all of them could prove burdensome. It is possible to focus the examination very significantly if sequences are ranked by marginality, or perhaps just by grouping them by similar characteristics. Attention can then focus on one or on a relatively low number of sequences. The idea of sorting success sequences according to their margin, and focusing on the most marginal, is shown in Figure 3. Following are examples of ways in which a success sequence could be "marginal":

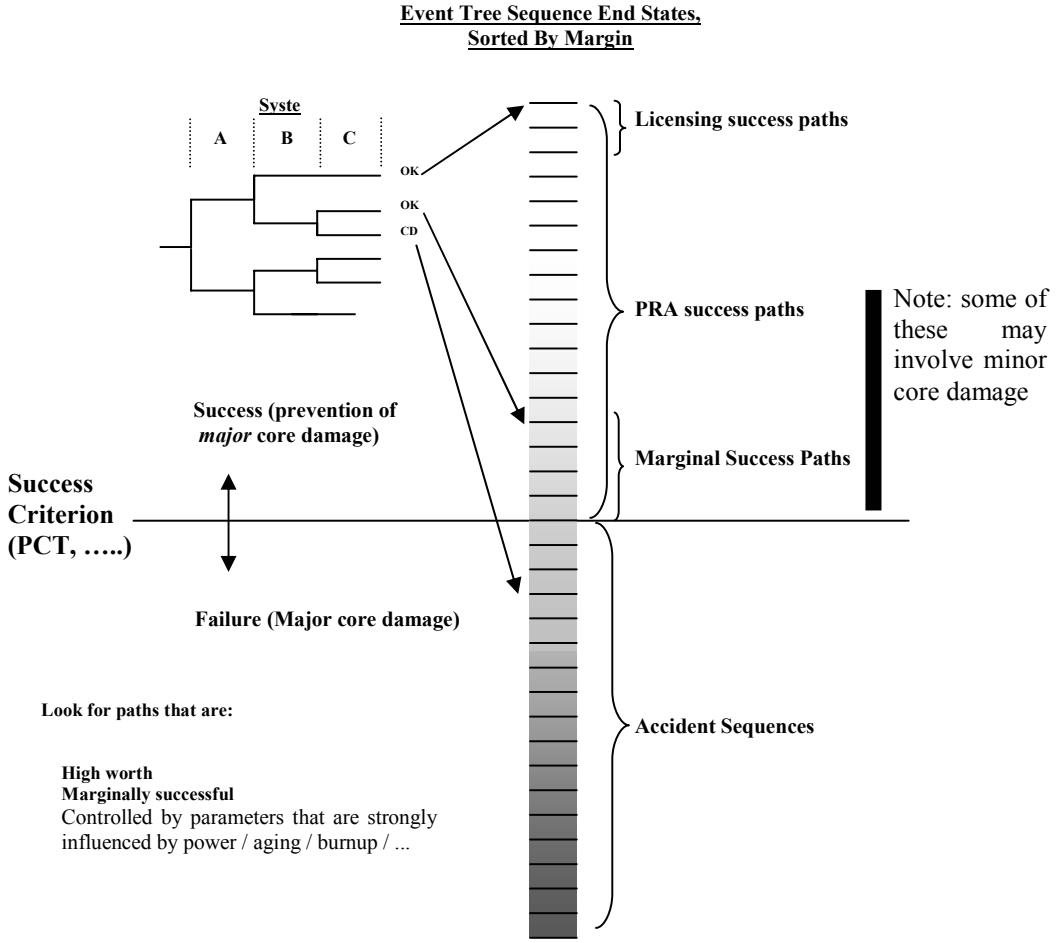

Figure 3. Variable Margin in Risk Space Success Sequences

- Physical limit criteria barely met.
- Barely enough flow (poison, ...).
- Time allowance on operator action allows action nearly too late to be "success."
- Barely enough relief capacity.
- Non-conservative assumption regarding leak rate.

With a focus on events in this range of interest, the next step is an examination of the key elements of the risk space description which can be affected by the change.

In Section 2.1 of this report, the correspondence between the risk perspective and the licensing perspective was established in terms of the success paths credited by each in their respective evaluations. Because of this correspondence, it is possible to make some use in risk space of one framework used in deterministic licensing to address plant changes. In the U.S., to determine if regulatory review of a proposed plant change is required, licensing success paths are evaluated in terms of a set of conditions spelled out in 10 CFR 50.59. As a starting point, it is useful to apply these conditions to the success paths credited in the risk space model.

3.1. 10 CFR 50.59 Basis

In the U.S. Code of Federal Regulations, 10 CFR 50.59 addresses licensee's responsibilities in assessing the safety implications of changes the licensee wants to make to the plant [5]. It is helpful to use the approach in 10 CFR 50.59 as a starting point for addressing changes in safety margins.

The following excerpt from 10 CFR 50.59 contains a useful conceptual checklist of ideas for the present purpose:

> *(2) A licensee shall obtain a license amendment pursuant to § 50.90 prior to implementing a proposed change, test, or experiment if the change, test, or experiment would:*
>
> > *(i) Result in more than a minimal increase in the frequency of occurrence of an accident previously evaluated in the final safety analysis report (as updated);*
> >
> > *(ii) Result in more than a minimal increase in the likelihood of occurrence of a malfunction of a structure, system, or component (SSC) important to safety previously evaluated in the final safety analysis report (as updated);*
> >
> > *(iii) Result in more than a minimal increase in the consequences of an accident previously evaluated in the final safety analysis report (as updated);*
> >
> > *(iv) Result in more than a minimal increase in the consequences of a malfunction of an SSC important to safety previously evaluated in the final safety analysis report (as updated);*
> >
> > *(v) Create a possibility for an accident of a different type than any previously evaluated in the final safety analysis report (as updated);*
> >
> > *(vi) Create a possibility for a malfunction of an SSC important to safety with a different result than any previously evaluated in the final safety analysis report (as updated);*
> >
> > *(vii) Result in a design basis limit for a fission product barrier as described in the FSAR (as updated) being exceeded or altered; or*
> >
> > *(viii) Result in a departure from a method of evaluation described in the FSAR (as updated) used in establishing the design bases or in the safety analyses.*

This excerpt is aimed at deciding whether a proposed "change" alters the situation in some way that is not adequately addressed in the current licensing basis. This thought process is generic and comprehensive, and can be applied to uprates, life extensions, etc.

Although one might consider the above to be an essentially deterministic thought process, the essential thoughts are more broadly applicable. It is useful to apply these essential thoughts to a broader class of events than that contemplated in the licensing basis. In the usual application of the above excerpt, the domain of the thought process is the accident analysis in the FSAR. In the present review, which is concerned with the risk impact of changes, the domain of the thought process should be the success paths credited in the risk model. Table 1 below compares the elements of the deterministic licensing approach with the comparable elements of the risk model. The 10 CFR 50.59 questions raised about each element in the middle column need to be extended to the elements in the right-hand column, and slightly adjusted for context.

TABLE 1

Areas to Be Evaluated with Respect to Proposed Changes

Elements	FSAR Elements Addressed in 10 CFR 50.59	Risk Model Elements That Need To Be Questioned Regarding Proposed Changes
Initiating Events	Selected IEs including DBAs, AOOs	Comprehensive set of IEs; comprehensiveness determined implicitly by intent to capture "risk-significant" contributors
Success Paths	Complement of success paths, that is, single-failure-proof, assuming a concurrent loss of offsite power; most credited equipment ends up with special treatment requirements. This determines scope of SSC failure mode questions.	All success paths on each event tree. This should determine scope of SSC and operator action failure mode questions.
Evaluation Basis (how "success paths" are shown to be successful)	Conservative phenomenological evaluation compared to the conservative acceptance criteria.	"Success" may be defined with respect to ultimate failure values of physical variables, instead of design limits. Robustness of success paths is not always carefully considered. Need to ask whether all of these paths' failure probabilities are still dominated by hardware or operator failure probability.
Consequence Evaluation	Radiological consequences of "success paths" and of non-mechanistic source terms are analyzed using prescribed methods and compared with regulatory limits.	Extending from the consequences under the FSAR elements, radiological consequences of all accidents including severe accidents

3.2. Adjustments to 10 CFR 50.59 Concepts

Accident

In PSA, the phrase "accident sequence" refers to a sequence of events leading to an end state in which substantial core damage has occurred. As used in Part 50, Appendix A, the term "accident" (as in "loss of coolant *accident*") refers not to an "accident sequence," but rather to an initiating event that is more severe, and much less frequent, than an "anticipated operational occurrence." Wherever 10 CFR 50.59 refers to "accidents," one needs to apply the thought to "initiating events," or more generally, to any challenge to a plant safety function. Thus, based on 10 CFR 50.59, one needs to know whether there are "new" initiating events to be considered (or whether previously-screened-out initiating events need to be considered), or whether the frequencies of currently analyzed initiating events need to be requantified because the initiating events are more frequent as a result of the change.

SSCs Important to Safety

References to "SSCs evaluated in the FSAR" (essentially, those credited in the accident analysis) need to be broadened to include SSCs (and indeed operator actions) credited in the risk model.

Increase in the Consequences

The magnitude of the radiological consequences of *accident sequences* is of interest. An increase in the magnitude of the consequences could result from an increase in the inventory, release fractions, or physical characteristics (thermal energy) of the release.

Design Basis Limit for a Fission Product Barrier

Here, the interest is not in design basis limits, but rather in the exceedance or alteration of the ultimate failure points of barriers (e.g., over temperature, and overpressure). All the failure modes of each barrier are candidates, depending on the nature of the inquiry.

Based on the above ideas, and rearranging to bring initiating events issues together, mitigating systems failures together, etc., Table 2 suggests a rearrangement of the 10 CFR 50.59 questions that would be useful in addressing changes.

3.3. Discussion of the Issues and Questions Listed in Table 2

Table 2 provides a high-level classification of changes to a plant that could change its risk metrics and/or its margins indices. As discussed above, the table entries are motivated in part by U.S. licensing practice, within which proposed changes to a plant are evaluated in terms of issues such as these to establish whether and how the proposed change impinges on the current licensing basis (the initiating events, the mitigating system success paths, the radiological and barrier consequences of the credited success paths). Evaluating a change in margin is analogous to this at a high level, but involves more technical detail.

It may be that a "success path" comes very close to failure as a result of a change, such that the path may, in fact, fail as a result of variability or uncertainty that would have been considered negligible before the change. Correspondingly, the row in Table 2, "Change consequences of existing event sequences," includes, for example, changes in the thermal-hydraulic trajectory of a success path.

Table 2. Plant Changes That Could Induce Changes in Margins Indices or in (Other) Risk Metrics

Changes in Risk Metrics	Instances
Change consequences of existing event sequences	Cause a physical limit to be exceeded (temperature, pressure, ...) Cause a change in the time evolution of key plant variables in one or more credited success paths, such that the failure potential of those event sequences is substantively affected by variability or uncertainty in the trajectories of those variables Change in plant damage states: creation of new plant damage state (PDS), or parametric change to existing PDS, e.g., a higher source term, or an earlier or more energetic release
Change initiating events: new events or frequency increase in existing events	New IE (class of accidents) (Note: if so, then new event trees would need to be developed, etc.) Example: Previously dismissed (incredible) event becomes credible, e.g., core flow blockage .as a result of passive element failure caused by increased cumulative damage to the element (more neutron fluence more flow)
	Increase the frequency of already-existing IEs. Examples: Higher frequency of small pipe breaks due to more vibration Higher frequency of transients due to less margin in BOP systems
Change Mitigating System Failure Probabilities	Change event tree or fault tree logic Mission success criteria Example: Configuration previously classified as "success" no longer "succeeds" because decay heat is higher Add new failure modes (new basic events) Add new success paths (hardware backfit) Change failure probabilities of basic events or enable CCF Operator actions, e.g., change in failure probability because less time is available Passive components, e.g., increased loading of some kind: vibration, stresses, thermal cycles, ... Active components Increased mechanical or electrical loading Change in operating environment room temperature, or suction source temperature

To the extent possible, the risk space model should be developed in such a way that its structure can be meaningfully quantified for both the "before" and "after" conditions. If a plant change introduces a new initiating event, the new event and its associated event tree will be quantified in the "after" state, since the "before" state will be null. This way, there is a consistent basis for quantifying the margin and the change in margin will be traceable to changes in specific paths.

Similarly, if a header success criterion changes, it will be useful to expand the top event headings such that all paths through the event tree are specified in sufficient detail to support an unambiguous assignment of margin to each path. For example, consider a three-pump system that succeeds with one or more pumps "before" a proposed change, but needs two or more pumps "after." If "success" is defined as in a simple core damage model, one pump's worth of margin is assigned to the frequency of success for the "before" case, even though two or three pumps are probably going to be running, and two pumps' worth is assigned to the frequency of "success" for the "after" case, even though (again) two or three pumps are likely to be running. Formally, then, margin is evaluated for 0, 1, 2, and 3 pumps running, and the appropriate frequency is assigned to each case, before and after. Even if the frequencies of these four configurations change as well as the margin for each configuration, there is a consistent basis for computing "before" and "after" margin. This objective can be better achieved if fault trees are developed at the system level rather than the function level (as they are now developed). The classification of a sequence would be done as a function of its calculated (or estimated) consequences, not as a function of the success headers involved in the sequence. Similar considerations apply to possible variations in event timing "before" and "after."

In the following sections, additional discussion on the items in Table 2 is included, along with examples taken from the possible plant changes listed at the beginning of this section 3.

3.3.1. Change in the Consequences of Existing Event Sequences

In licensing, this row refers to the consequences (e.g., barrier challenge, dose) evaluated for success paths in that context. More generally, the term "consequences" may refer to other risk metrics, or to new metrics defined to address issues of margin.

For purposes of the present report, this row refers to any safety variable being calculated by the dynamic simulation models used in the analysis, either inside or outside the plant. This includes, for instance, an increase in the probability that a previously successful sequence ends in a damage state, even if all the mitigating systems perform as expected. This effect could likely result from power uprates if the capability of some support systems such as cooling or electrical systems is not modified in a consistent way. Changes in radiological consequences due to higher inventories or release fractions are also included in this item.

3.3.2. Change in Initiating Events: New Events or Frequency Increase in Existing Events

The following examples illustrate the potential for changes to alter the structure of the initiating event set, or change the frequencies of initiating events.

1. An impact of aging and associated corrosion of SSCs is that initiators that were previously considered incredible may now become credible. An example might be reactor vessel failure, that is, a new LOCA, due to corrosion of the vessel upper head. A recent Accident Sequence Precursor (ASP) study at NRC [6] indicated that, given the condition of the vessel head at a PWR, there was a probability of core damage of 8E-3.
2. Some plant changes alter the likelihood of an initiating event by causing the plant to operate closer to limiting parameter values, such that an excursion of a given magnitude now causes a plant trip where it did not before.

3. A power uprate could cause main feedwater to require three pumps of a given capacity at full power, rather than only two. This would increase the frequency of the event "loss of a single feedwater pump," and could cause alterations in the control system that also affects initiating event frequency or consequences.
4. The deregulation of the electrical industry in the United States has put pressures on electric grid reliability. These external changes may have an effect on the frequency of the loss of offsite power initiating event.
5. Steam dryers in some U.S. BWRs have been deteriorating, releasing parts into the vessel and downstream from the vessel. These loose parts could interfere with cooling or damage pumps, thereby creating possible new initiating events.

3.3.3. Change in Mitigating System Failure Probabilities

Possibly the most complex part of the changes to a PSA resulting from plant changes resides in the mitigating systems. For a change to a plant or for combinations of those changes, the impact on equipment reliabilities, equipment availabilities, component dependencies, success criteria, modeling completeness, operator response times and common-mode failure logic all need to be considered. For example for power uprate licensing in the United States, the NRC addresses the potential for increased risk and evaluates the impact of uprates on component reliabilities, system success criteria and/or operator response times [7].

At a US PWR plant, deteriorating heat exchanger surfaces on a risk-important system required that the system be taken out of service for maintenance, thus decreasing the system availability. This is another example of how aging can change basic event probabilities in the risk space model.

4. INFORMATION NEEDS TO QUANTIFY SAFETY MARGINS

When decisions are addressed regarding the safety implications of changes to a nuclear power plant, decisions are ideally characterized in terms of their possible outcomes: the scenarios (or events) or changes to scenarios that might ensue as the result of the given change, and the frequencies and the consequences of those scenarios.

Ideally, given perfect knowledge, one could make licensing decisions based on accurately projected outcomes of those licensing decisions. However, in traditional licensing practice, because of uncertainties of many kinds and in light of potentially serious adverse consequences of severe accidents, a different approach to making decisions has been applied. In licensing of commercial reactors, it has been traditional to characterize satisfaction of safety objectives in terms of the consequences of a specific set of postulated events comprised in a spectrum ranging from normal operation to design basis events that pose challenges deemed to be "enveloping." The acceptance criteria on the analyzed consequences of these events (such as Peak Clad Temperature (PCT) < 1,204°C (2,200°F)) are conservative substitutes for objectives relating to protection of public, workers, and environment; in addition, the evaluation methodologies may embed systematic conservatism. Elements of this approach are illustrated in Figure 4, taken from Reference 1. Note that while Figure 4 is about margin and consequence metrics that need to be evaluated, the logic of the overall decision process requires that these evaluations be predicated on scenarios having particular characteristics, including the severity of the challenges and limits on their frequencies.

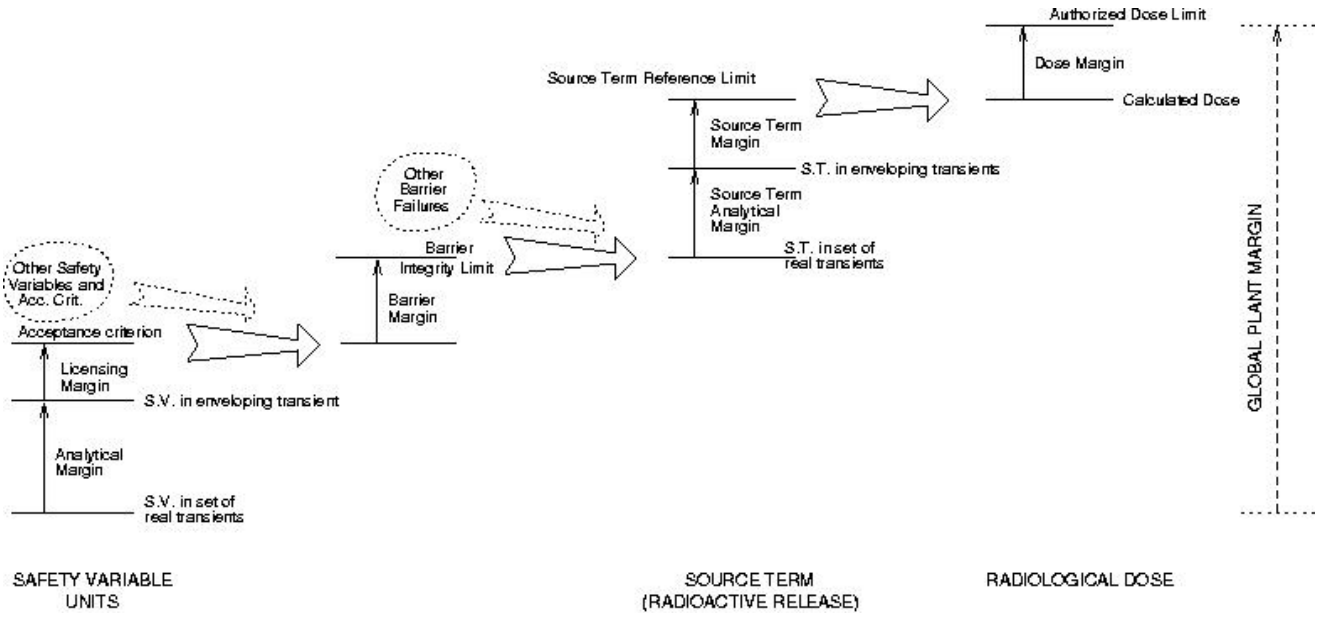

Figure 4. Safety Margins in Typical Design Basis Safety Analysis

This traditional approach can support robust findings regarding adequate protection. However, it is not sufficient to use for assessing the effects of changes in a given plant. It supports "yes" or "no" findings of adequate protection, but it does not support comparing two different plants' level of safety, or assessment of the safety state "before" and "after" certain kinds of changes at a given plant. Correspondingly, in order to go beyond answering yes / no questions of adequate protection, that is, going to the assessment of safety margin and change to that margin, one needs information beyond that needed for the adequate protection question. This is the subject of the section and is directed to the type of information that will be needed to quantify safety margins.

Quite generally, decision-making regarding a particular safety issue, for example a change in safety margin, can be approached by characterizing the scenarios associated with that issue, together with their frequencies and their associated consequences, as discussed in section 3 of this report. Many people associate this "set of triplets" idea –a scenario set with its associated consequences and frequencies– with explicitly risk-oriented decision-making [8], but it also applies at a high level to the traditional decision-making approach as well, and formally applies to essentially any framework likely to be used in reactor safety decision-making. For example, in today's licensing practice, postulated scenarios are analyzed to show that their consequences are acceptable, including margin; the decision rule implicitly reflects the relative likelihoods (frequencies) of these postulated events. (Refer to [2] for examples of postulated scenarios.] Within the more generalized framework addressed here, it is important to analyze a more complete set of scenarios, to quantify their frequencies explicitly, and to understand a broader range of margin-related and consequence-related metrics.

Because the "triplets" idea bridges all foreseeable frameworks, including current practice, the discussion of information needs is organized below in terms of this idea. Information needs are discussed for developing the scenarios and quantifying their frequencies and consequences. Structuring the discussion in this way is not meant to imply that these aspects can be discussed independently; they are highly inter-related. For example, both the frequencies and the consequence metrics influence the structure

of the scenario set. However, this organizing principle fosters understanding of how the existing approach is generalized for purposes of the integrated margins assessment.

The discussion does not presume a specific new framework, but does assume that more detailed and realistic, best-estimate assessments will be needed in order to support improved assessment of margin, consistent with the approach addressed in [1]. The uncertainties associated with both safety margins and event scenario frequencies have to be considered in the decision process.

4.1. Information Needs for Structuring the Scenario Set

4.1.1. Information Needs in Existing Licensing Practice

In today's licensing practice, careful consideration has been given to formulation of categories of events that must be analyzed and for which it must be shown that a plant's design response is adequate [2]. The current process does not require a formally complete set of scenarios: rather, one develops a scenario set intended to envelope the physical challenges associated with probabilistically significant scenarios.

Transients and design basis accidents and their associated acceptance criteria are chosen in the belief that designing to mitigate these scenarios to the stated acceptance criteria will provide more than enough capability to protect the public for any event likely to occur in reality. Many stringent requirements are also placed on plant design, construction, and operation, intended to assure that challenges of that severity are extremely unlikely.

Figure 4 shows the assessment of safety margins in typical design basis safety analysis. The figure shows margins assessed at different points in the scenario. As described above, the basis for a favorable decision is that the margins and consequences are appropriate *for limiting scenarios.* In structuring the scenario set for this purpose, one needs first to establish which scenarios are in fact limiting within the requirements of the evaluation methodology (what needs to be postulated), and then analyze them carefully.

Applying the single-failure criterion in design-basis-accident analysis, one needs to perform failure modes and effects analysis iteratively with thermal-hydraulic analysis, in sufficient detail to identify the limiting single failure to be postulated to occur concurrently with the initiating event. This requires detailed system design information, but only for those systems credited in the licensing basis accident analysis.

Generally analogous considerations apply to what are referred to in the U.S. as anticipated operational occurrences (AOOs). Corresponding to the relative frequencies of such events, the acceptance criteria are correspondingly stringent. The formulation of these scenarios is predicated on their roles in the decision process.

4.1.2. Information Needs in Integrated Margins Analysis

In order to look more quantitatively at margin, it is necessary to establish more than whether a design basis event meets the acceptance criteria. It is necessary to define categories of scenarios in terms of the distinctions needed to support quantification of margins metrics. The categories of scenarios can be determined, for example, from the inquiry process depicted in Figure 2. This depiction of an inquiry process focuses on the safety margins associated with the three barriers to the release of radio nuclides.

If the margins metrics address quantitative distinctions (as opposed to "yes" or "no" findings of adequate protection), then a more complete and detailed scenario set needs to be specified than merely "design basis accidents (DBAs)" and "AOOs." For the sake of realism, the scenario set needs to address systems not credited in the design basis, and consequence categories more detailed than "no damage" and

"complete failure." The scope of information collection is therefore potentially much broader, encompassing not only safety systems but also other systems; and probabilistically significant barrier failures need to be specified in the formulation of the scenario set. Instead of knowing that a barrier withstands a design basis challenge with margin, one now needs to know where it will fail. For example, it is necessary to know not only the design pressure, but also the ultimate pressure capacity of the containment barrier.

In summary, the main questions to be answered for an adequate structuring of the scenario set are:

- What barrier failures can be affected by the change?
- What mechanisms of radiological release are modified by the change?
- What safety variables are indicative of the challenges to identified barriers and release mechanisms?

By answering these questions, a subset of the scenarios that describe the risk space (as defined in [1]) can be identified as impacted by the change. Then, the potential effects of the change on the event tree structures should be identified as a first step in the resolution of the safety inquiry.

4.2. Information Needs for Assessing the Consequences

Formally, the term "consequences" in this context refers to all of the outcome-related metrics in the problem: what happens to the public, what happens to the workers, what happens to the environment, and what happens to the barriers. It also includes quantification of the margin-related metrics, for example, quantification of the margin to failure of the fuel-clad barrier.

4.2.1. Information Needs in Existing Licensing Practice

In existing licensing practice, it is necessary to show that the consequences meet the acceptance criteria. Formally, this means showing that the components perform physically according to the assumptions in the T/H analyses, and that the dose consequences resulting from the postulated events meet acceptance criteria. Depending on the specific events, some barrier damage may occur, but in LWR licensing, this has already been taken into account in the formulation of the acceptance criteria.

Acceptance criteria currently used for normal operation, Class 1 and Class 2 postulated accidents (DBAs), are reviewed in the SMAP Technical Note entitled, "Acceptance Criteria and Related Safety Margin" [2]. These acceptance criteria are conservative; their satisfaction provides defense in depth and a decoupling from one barrier to the next. They cover the frequency spectrum from normal operation to very low-frequency and hypothetical accidents.

In considering the fuel-clad barrier, the consequence model should be able to calculate such things as peak clad temperature, clad oxidation, departure from nucleate boiling (DNB), and fuel enthalpy so that all types of fuel clad failure and the extent of fuel failure can be assessed. Core-wide, as well as hot channel loadings on the fuel clad will need to be addressed. Computer codes used for these purposes include RELAP5 [9] and TRACE [10].

In the US, guidance on treatment of uncertainty and acceptance criteria for accident analysis is provided in Regulatory Guide 1.70 [11] and in the Standard Review Plan (SRP), NUREG-0800 [12]. Chapter 15 of the SRP includes a list of events to be analyzed and acceptance criteria for each type of event. The acceptance criteria are based on applicable General Design Criteria (GDC) and possibly also TMI Action Plan items. Specific criteria necessary to meet the GDC are specified, for example, pressure in the reactor coolant and main steam systems maintained below 110% of design values, minimum DNBR

above the 95/95 limit for PWRs, CPR above the MCPR safety limit for BWRs and assumption of most limiting single active component failure. It is also required that the parameters used in analytical models be suitably conservative. This generally requires biasing the input to allow for uncertainties in power level, safety system actuation delay times, trip setpoints, plant initial conditions, scram reactivity, scram speed, etc. Core burnup is selected to yield the most limiting combination of moderator temperature coefficient, Doppler coefficient, void coefficient, axial power profile and radial power distribution. Mitigating system performance is also biased in a conservative direction to allow for uncertainties.

4.2.2. Information Needs in Integrated Margins Analysis

There are two shortcomings in the current approach to margins assessment, which an integrated framework is intended to overcome. First, although there is confidence in adequate margin because of the conservative nature of the bases, the actual amount of margin is unknown. Further, since for the most part the margin is measured against hypothetical accidents (e.g., large-break LOCAs with conservative assumptions), the benefit of knowing the real risk profile is lost. The framework and methodology presented in [2] address both those issues, in that one is now dealing with best-estimate assessments of realistic events coupled with the event importance through knowing the event's frequency of occurrence.

The extension of the analysis space to cover also beyond-design-basis (e.g., severe accident) scenarios makes it impossible to impose acceptance criteria on the consequences of analyzed scenarios. Rather, some safety objectives are defined and the analysis goal is to quantify on one hand the exceedance frequency of each safety objective and, on the other, the margin with respect to each objective of scenarios that do not exceed it. Acceptance criteria of traditional licensing analyses can be directly used as safety objectives in the integrated assessment but they could also be replaced by more realistic ones, if the latter are found better suited for the higher realism of the integrated approach. Consistency based on physical processes and computational limitations is desirable where existing DBA acceptance criteria permit it.

Consequences need to be addressed not only for the design basis events, but for all categories of scenarios in the risk space model. Potentially, this means more T/H analysis than was needed for design basis analysis. The results can be cast in the form of more realistic measures of consequences, rather than yes/no findings of adequate protection based on plant response in postulated events. In addition, beyond-design-basis phenomena now need to be analyzed.

4.2.3. Barrier Loadings

The assessment of barrier loadings starts with an understanding of fuel behavior. If the inquiry interest is directed to mild transients, codes such as FRAPCON-3 [13] can be used to understand the behavior of a single fuel rod. If the inquiry is focused on more challenging transients and design basis accident behavior, codes such as FRAPTRAN [14] can be used. For accident analysis of the primary system, it is assumed that fuel and clad loadings remain acceptable if selected safety variables remain within specified limits. These assumptions are supported by fuel-type-specific analyses and experiments. More detailed correlations between safety variable values and fuel/clad loadings allow for a better characterization of the barrier damage.

The thermal hydraulics model used should be a complete representation of the plant with the modeling of all of the mitigating systems and operator actions necessary to track the set of scenarios describing the risk space. The model should be able to accommodate the full spectrum of initiating events, including ATWS. It should be able to assess the safety variables relevant to the barrier being addressed. The model (or combinations of models) should be able to track all relevant loadings and phenomena associated with the full spectrum of accidents, including severe accidents. Further, it should be capable of running through to the end of a given scenario until a stable, coolable state is reached or until all the safety

objectives being analyzed are found exceeded. The model should be realistic (best estimate) with a capability of addressing a measure of uncertainty, consistent with the approaches discussed in SMAP task 3. Since an important focus of the safety margins program is the impact of plant changes on the changes to safety margins, the models need to reflect the consequences of plant changes on the safety variables.

Typically, models already exist for addressing at least some of the safety variables of interest for the fuel-clad barrier. Different computer codes are used, primarily depending on the expected consequences of the event. For anticipated operational occurrences (AOOs), system codes like RETRAN [15] and RELAP5 [9] are suitable for the application. A number of vendor proprietary codes are also widely used. For design basis accidents, codes such as RELAP5 and TRAC-P [16] can perform the necessary simulations. Severe accidents are analyzed using codes such as MELCOR [17] and SCDAP/RELAP5 [18].

These models will probably need to be modified to accommodate the full set of mitigation structures, systems, and components (SSCs) and operator actions called on for the sequences under consideration. Further, model inputs may need adjustment to make them "best estimate." The model should include coupling to the containment and should be consistent with the PSA model. Efficiency will be enhanced if a single code can be used for all of the safety margin applications. An alternative solution is to build a chain of codes of incresing capabilities (and complexity) in such a way that, as the accident scenario progresses and the simpler models reach their applicability limits, they automatically launch and initialize a more adequate code to continue with the simulation.

As with the traditional licensing model, the model should be able to calculate peak clad temperature, clad oxidation, departure from nucleate boiling (DNB), and fuel enthalpy so that all types of fuel clad failure and the extent of fuel failure can be assessed. Core-wide, as well as hot channel loadings on the fuel clad will need to be addressed.

4.2.4. Barrier Damage Characterization

It is important to understand the failure profile for a given barrier. Ideally, the failure probability of a barrier as a function of the loading of that barrier is needed to determine the actual margin to failure for that barrier. For example, it is important to know the actual probability of fuel-clad failure when the temperature of the clad is at 1,204°C (2,200°F). This is discussed in more detail in [1]. For a given sequence, the barrier might be breached and the consequences can be determined from knowledge of how much fission-product material is released from the barrier and how much of that material is transferred to the next barrier or how much of that material is transferred to the environment.

The accurate determination of the release of fission products from the various barriers and the ultimate transport to the environment with associated dose to workers and the public can be a complex, resource-intensive, and time-consuming activity, especially when many accident sequences are being considered. There are simplifications, as described in Sub-task 1C [2], which then permit conservative approximations.

It should be noted that an understanding of the consequences of barrier breach is only important if a Global Plant Margin (and the associated Source Term and Dose Margin), as depicted in Figure 1, is desired or if the adequacy of a given safety objective to be used as a decoupling criterion for the analysis should be assessed. To compare plant margins before and after a modification, an evaluation of global margins is not necessarily needed.

The fission-product inventory in the fuel can be determined by computer codes such as ORIGEN [19]. For the normal operation and for the high-probability and low-consequence events, it may also be

important to include activation products that reside in the primary system, in addition to the fission products.

The release from the first barrier, the fuel clad, is a function of the damage to individual fuel rods, the extent of the rod damage across the whole core, and the timing of the failures. Releases can range from minor leaks of noble gases to releases of volatiles and finally non-volatiles. When estimating the amount and nature of the release not only the state of the cladding is important. The retention capability of the ceramic fuel matrix can be substantially reduced if partial melting has occurred during the scenario, in particular in the case of transient overpower before or at the time of the reactor trip.

The release from the second barrier, the primary system, is a function of the pressure and temperature history of the primary system, the location of the failure, and the nature of the initiating event. As for the fuel-clad barrier, the primary system failures can be anywhere from a small leak (even during normal operation) to a catastrophic failure of the vessel. A major failure of the primary system can occur as an initiating event (e.g., a large-break LOCA) or as a consequence of the progression of an accident sequence (e.g., a hot-leg failure after core damage in a PWR, resulting from a station blackout). In addition, releases can also result from valve openings automatically generated or manually executed upon request of emergency procedures. This is the case, for example, of pressure relief through pressurizer PORVs. The radionuclide releases from the primary system will be specific to the sequences being studied. In addition to releases from the primary system into the containment, there may be releases that bypass the containment barrier, for example containment bypass through a steam generator in PWRs.

The release from the third barrier, the containment system, is a function of the pressure and temperature history of the containment, the location of the failure, and the nature of the initiating event. The release from the containment can take the form of a design-basis leak from an intact containment all the way to a major failure that is caused by a severe accident or one that initiates a severe accident. Also, some emergency procedures and severe accident guides include provisions for containment venting under specified conditions, which must be taken into account. As with the primary system barrier, the radionuclide releases from the containment system will be specific to the sequences being studied. The releases could be scrubbed, retained in structures downstream from the containment, or could be released directly into the outside environment.

It is particularly important to understand the radionuclide releases into the environment (whether from the containment or directly from the primary system) as this information defines the initial and boundary conditions for the assessment of the dose that the workers and the public may receive. The information necessary to understand the transport of the radionuclides from one barrier to the next and then finally to the environment, as depicted in Figure 5, is important for an understanding of how much is released. It is of particular importance, if the scope of the safety margins inquiry extends to a determination of the source term and global plant safety margins depicted in Figure 4. This transport is discussed below.

There are various ways to determine the transport of radionuclides from one barrier to the next and from the final barrier to the environment. These methods range from simple approximations to determinations using sophisticated computer codes. The process of release and transport for design basis accidents is discussed in Reference 2.

If an inquiry needs to address occupational or public radiation exposures that result from an event, for example in the control room or at the site boundary, computer codes such as RADTRAD [20] can be used.

The transport can be unique to the events under study, the specific design of the plant, and the operating history of the plant. Approximations are sometimes used that focus on the radionuclide retention, fallout, re-suspension, scrubbing, phenomena interaction, and chemical reactions, especially in the case of

severe accidents. For example, the potential for scrubbing of fission products through the suppression pool is important in determining the final release from the containment for BWRs. Similar mechanisms are assumed in the case of design basis steam generator tube rupture in PWR when the break is covered by liquid water in the secondary side.

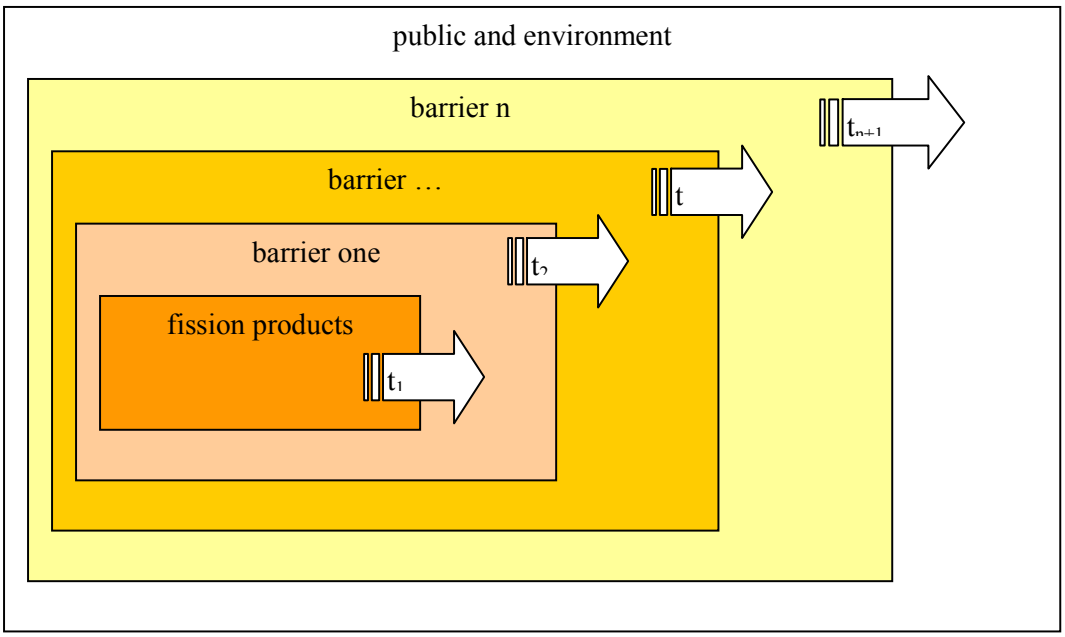

Figure 5. Generalized Portrayal of Fission Product Transport from Barrier to Barrier

The final transport is from the final barrier to the public. For severe accidents, a code such as MACCS [21] is used to calculate dose to the public as well as property damage. Input information includes the groups of fission products released, the timing and duration of the release, the energy of the release and the chemical composition of the fission products. Uncertainty in the environmental conditions such as wind speed and direction, surface and underground water flows, etc., should be taken into account. Bounding values or best estimate high percentiles may be used for these conditions.

An integrated margins quantification approach would use the same barrier failure mechanisms that are in [2]. Consistency with the approach described in this reference is a requisite for regulatory decision-making based on an integrated safety margins methodology.

4.3. Information Needs for Analyzing the Frequencies

4.3.1. Information Needs in Existing Licensing Practice

In traditional licensing practice, scenario frequencies are assigned to qualitatively defined categories based on the kinds of events that cause them. For example, design basis events generally require gross failures of passive components that are not expected to fail, and this is all that the decision-maker needs to know about that category of events. In fact, some DB events border on being non-mechanistic. In contrast, AOOs can be caused by failures of active components. The acceptance criteria are defined accordingly. For traditional licensing purposes, this level of frequency information is considered sufficient, partly because numerous requirements on licensees are imposed in order to assure that frequencies assumed to be low are, in fact, low. Examples of frequency ranges are provided in Reference [2].

4.3.2. Information Needs in Integrated Margins Analysis

It is necessary to explicitly quantify the frequencies of the categories of scenarios defined when the scenario set is structured. For rare events, it will be necessary to synthesize the frequencies using state-of-knowledge logic models. The role of simulators could turn out to be important. Depending on the formulation of the margin-related metrics, and on the design of the plant, margin may depend on the timing of operator actions in some scenarios.

For the determination of safety margins it is assumed that a rigorous plant-specific description of the risk space through a complete set of event trees has been developed following the criteria of section 2 in this report. Initiating event frequencies and header probabilities are the essential ingredients for frequency calculations. For those plants where a PSA is already available for the assessment of severe core damage (level 1) and large early releases (level 2), many frequency and probability data are already available. Initiating events in the risk space will be most likely similar to PSA initiators or would result from a finer subdivision of them. System related probabilities used in PSA fault trees would be mostly applicable in the risk space and probabilities of some stochastic phenomena may be taken from level 2 PSA.

Until now, the focus has been on a single scenario and the deterministic evaluation of the safety margin for that scenario. However, in the risk space, all the relevant scenarios for a particular inquiry and the frequency of those scenarios must be known for the risk-informed safety margin to be determined. By using methods to some extent similar to those of PSA, the set of scenarios and their frequencies can be determined. While this can be a very large number of scenarios, simplifications can be made.

From the risk space event trees and for each safety objective being analyzed, events can be binned into three groups. The first group is made up of events that are clearly successes and would have a very large safety margin. These events do not need to be analyzed further. The second group is made up of events that are clearly failures and would have no safety margin. These events likewise do not need to be analyzed further. The events in the third group need to be analyzed more carefully: they have the potential to challenge the barrier (or the safety objective) under study. The particulars of this process have been already described in [1] and in this report. Particular information needs for the PSA contribution to the process are discussed below.

Event trees display the top events, some of which are directly related to the structures, systems, and components (SSCs) that should mitigate the consequences of the event. Fault trees, made up of these SSCs, link to the top events and determine the top-event split fractions. For the mitigative SSCs that are impacted by plant changes, data are needed on any changes to the performance of these SSCs and/or changes to the requirements for these SSCs. For example, deterioration of the surfaces of a heat exchanger could degrade the heat exchanger performance while imposing a power uprate could challenge the heat exchanger's specifications. These changes may impact the system failure probabilities. Care should be taken in making sure that the system failure probability changes are consistent with any changes in the deterministic analysis. Also, it is possible that plant changes might require additions to, or a rearrangement of, the event trees.

An important part of event-scenario frequency determination is the proper modeling of operator intervention. It is particularly important to capture this impact if a plant modification has the simultaneous consequence of eroding margin to failure of a barrier, while at the same time shortening the time available for successful operator intervention.

Finally, plant changes may impact also the dynamic conditions in sequences where stochastic phenomena may appear. The probabilities of these phenomena must be recalculated accordingly.

4.4. Further Points on Acceptance Criteria and Safety Objectives

In previous sections, acceptance criteria for licensing accident analysis and safety objectives for Integrated Margin Assessment have been discussed in the context of the consequences and frequencies of the scenarios that make up a particular safety margin inquiry. It was also stated that for an integrated safety margins inquiry, acceptance criteria of the accident analysis become safety objectives, eventually improved to make them more realistic, consistent with the best-estimate approach of the integrated inquiry.

In [0], high-level acceptance criteria are discussed for ultimate measure of safety relative to dose to the public. Low-level acceptance criteria are also discussed for the individual barriers, which permit a decoupling from the high-level acceptance criteria because of the conservative nature of the criteria and the deterministic analysis. That is, if the fuel-clad criteria are met (together with the primary system and containment criteria) then the radiological releases are limited and acceptable. There is no need to directly address the high-level criteria. For the realistic methodology for assessing safety margin, discussed here, the same approach for decoupling can be used. Properly formulated, realistic safety objectives can be developed at the barrier level that allows for decoupling and assurance that the high-level safety objectives can be met.

If a requirement of the inquiry is that the plant global plant margin be measured, then the safety objectives for the plant global margin also need further attention. In the U.S., the "subsidiary objectives" in the Safety Goal Policy Statement [22] can be seen as examples of acceptance criteria, defined in terms of exceedance frequency, when considering global plant margin with respect to the safety objectives of avoiding severe core damage or large early releases. These subsidiary objectives have been incorporated into the risk-informed approach for considering plant changes as described in Regulatory Guide 1.174 [23]. Changes in safety margin relative to the subsidiary objectives could be an important metric, if one were considering global plant margin. It should be kept in mind, however, that these subsidiary objectives refer only to very specific safety objectives. Similar criteria should be defined with respect to other safety objectives like keeping annual doses below specified limits or complying with the site criteria defined by 10 CFR 100.

However, it should be taken into account that safety objectives related with the global plant margin do not stand on their own. In order to preserve defense-in-depth, the lower-level figures of merit need to be considered and the associated decoupling that they would preserve. In other words, deriving low-level acceptance criteria from high-level acceptance criteria is only possible in individual safety inquiries. This is because there is no direct top-down correspondence in regressing either frequencies or low-level criteria for individual safety margins form high-level criteria, e.g., 10 CFR 100 limits or the safety goals.

REFERENCES

[1] SMAP Sub-task 1B Technical Note entitled "Definition of Generalized Concepts of Safety Margins and Characterization of Safety Margin Sources," NEA/SEN/SIN/SMAP(2005)4, September 2005.

[2] SMAP Sub-task 1C Technical Note, entitled, "Acceptance Criteria and Related Safety Margin," NEA/SEN/SIN/SMAP(2005)3, August 2005.

[3] "Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners," NASA HQ, February 5, 2002.

[4] ASME Standard for Probabilistic Risk Assessment for Nuclear Power Plant Applications, ASME-RA-S-2002, April 5, 2002.

[5] Code of Federal Regulations, 10 CFR 50.59, "Changes, Tests and Experiments," July 1996.

[6] Davis-Besse - Final Accident Sequence Precursor Analysis of February 2002 Operational Condition, Office of Nuclear Reactor Research, March 14, 2005.

[7] Attachment 1 to Matrix 13, Review Standard for Extended Power Uprates, U.S. NRC, Office of Nuclear Reactor Regulation, December 2003.

[8] Kaplan, S., and B.J. Garrick. 1981. On the Quantitative Definition of Risk. Risk Analysis 1(1): 11–27 (1981).

[9] RELAP5/MOD3.3 Code Manual, Volume 1: Code Structure, System Models, and Solution Methods, NUREG/CR-5535, Vol. 1, Rev. 1, July 2003.

[10] TRAC-M/Fortran 90 (Version 3.0), Theory Manual, LA-UR-00-910, July 2000.

[11] US Nuclear Regulatory Commission, Standard Format and Content of Safety Analysis Reports for Nuclear Power Plants, Regulatory Guide 1.70.

[12] Standard Review Plan for the Review of Safety Analysis Reports for Nuclear Power Plants, LWR Edition, NUREG-0800, June 1987.

[13] NUREG/CR-6534, Vol. 4, "FRAPCON-3 Updates, Including Mixed-Oxide Fuel Properties," May 2005.

[14] NUREG/CR-6738, Vol. 1, "FRAPTRAN: A Computer Code for the Transient Analysis of Oxide Fuel Rods," August 2001.

[15] RETRAN-02 – A Program for Transient Thermal-Hydraulic Analysis of Complex Fluid Systems, Volume 1: Theory and Numerics (revision 4), Electric Power Research Institute Report NP-1850-CCM-A, 1988.

[16] TRAC-PF1/MOD2 Code Manual User's Guide, NUREG/CR-5673, LA-12031-M, July 1992.

[17] MELCOR Computer Code Manuals, Volume 1: Primer and Users' Guide, Version 1.8.5 May 2000, NUREG/CR-6119, Vol. 1, Rev. 2, SAND2000-2417/1, October 2000.

[18] SCDAP/RELAP5/MOD3.3 Code Manual, NUREG/CR-6150, Vol. 1, Rev. 2, INEL-96/0422, January 2001.

[19] Croff, A.G., "ORIGEN2—A Revised and Updated Version of the Oak Ridge Isotope Generation and Depletion Code," ORNL-5621, July 1980.

[20] Humphries, S.L., et. al., RADTRAD: A Simplified Model for RADionuclide Transport and Removal And Dose Estimation, NUREG/CR-6604, December 1997.

[21] Chanin, D.I., J. L. Sprung, L. T. Ritchie and H. N. Jow, "MELCOR Accident Consequence Code System (MACCS): User's Guide," NUREG/CR-4691, Vol. 1, Sandia National Laboratories, February 1990.

[22] "Safety Goals for the Operations of Nuclear Power Plants; Policy Statement," U.S. Nuclear Regulatory Commission, August 4, 1986.

[23] Regulatory Guide 1.174, "An Approach for Using Probabilistic Risk Assessment In Risk-Informed Decisions On Plant-Specific Changes to the Licensing Basis," U.S. Nuclear Regulatory Commission, July 1998.