INL/EXT-21-64565-Revision-0

Treatment of uncertainties for security-related design aspects of advanced reactors when using a risk-informed licensing approach

September 2021

Curtis L Smith, Robert W Youngblood III, H. Christopher Everett, Michael Mankosa, Cesare Frepoli, Ram Sampath, Tina M Miyake, Jarrett Valeri

Changing the World's Energy Future

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/EXT-21-64565-Revision-0

Treatment of uncertainties for security-related design aspects of advanced reactors when using a risk-informed licensing approach

Curtis L Smith, Robert W Youngblood III, H. Christopher Everett, Michael Mankosa, Cesare Frepoli, Ram Sampath, Tina M Miyake, Jarrett Valeri

September 2021

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

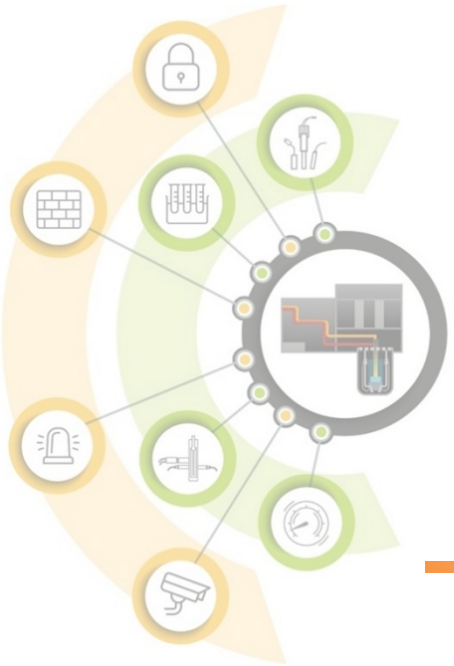

Treatment of uncertainties for security-related design aspects of advanced reactors when using a risk-informed licensing approach

Prepared for
U.S. Department of Energy

Curtis Smith[1], Robert Youngblood[1], Chris Everett[1], Tina Miyake[1], Jarrett Valeri[2], Michael Mankosa[2], Cesare Frepoli[2], Ram Sampath[3]

1. Idaho National Laboratory
2. FPoliSolutions
3. Centroid Lab

September 2021
INL/EXT-21-64565

ABSTRACT

Sabotage of nuclear plants and theft of special nuclear material are different from many other issues potentially affecting public health and safety, and some of those differences drive the content of the present report. A high-level indication of these differences is provided in the U.S. Nuclear Regulatory Commission's Safety Goal Policy. Promulgated in the mid-1980s when it had become reasonably clear that risk analysis had improved to the point where it was possible to understand the risks associated with plant operation, the Safety Goal Policy articulates qualitative safety goals and quantitative health objectives that are meant to guide regulatory and risk management activities, with the following key exceptions noted in the original policy statement:

The possible effects of sabotage or diversion of nuclear material are also not presently included in risk-informed approaches. At present, there is no basis on which to provide a measure of risk on these matters, however theft, diversity, and safeguards are considered under the deterministic requirements such as those found in 10 CFR Parts 73 and 74. It is the Commission's intention that everything that is needed will be done to keep these types of risks at their present very low level; and it is the Commission's expectation that efforts on this point will continue to be successful. With these exceptions, it is the Commission's intent that the risks from all the various initiating mechanisms be considered to the best of the capability of current evaluation techniques.

The present report discusses extensions of classical risk management to address some of the special issues that arise in the context of security. Although the present emphasis is on physical security, some attention will be paid to cybersecurity.

A particular focus of the report is on quantitative framework to manage and address uncertainties. This framework is demonstrated via a couple of hypothetical examples.

Page intentionally left blank

CONTENTS

FIGURES

TABLES

ACRONYMS

AOE	anticipated operational events
AOO	anticipated operational occurrences
BDBE	beyond design basis events
CAD	computer-aided design
CCDF	complementary cumulative distribution function
CDF	Cumulative distribution function
DBA	design-basis accidents
DBE	design-basis events
DET	dynamic event tree
DG	diesel generators
DID	defense-in-depth
EAB	exclusion area boundary
EMDAP	Evaluation Model Development and Assessment Process
EMDP	Evaluation Model Development and Assessment Process
EMRALD	Event Modeling Risk Assessment using Linked Diagrams
EPRI	Electric Power Research Institute
ES	event sequences
ESF	event sequence families
ESS	engineered safety system
F-C	frequency-consequence
FLEX	diverse and flexible coping strategies
F-N	frequency-number
FSF	Fundamental Safety Function
FTA	Fault Tree Analysis
IAEA	International Atomic Energy Agency
IE	initiating event
INL	Idaho National Laboratory
JPL	Jet Propulsion Laboratory
KPP	key performance parameters
LBE	licensing-basis events
LMP	Licensing Modernization Project
LOF	loss of flow
LOOP	Loss of offsite power

LOS	loss of offsite power
LPS	low pressure safety injection
LWR	Light Water Reactors
MOE	measures of effectiveness
MOP	measures of performance
MSF	main safety functions
NEI	Nuclear Energy Institute
NRC	Nuclear Regulatory Commission
NSRST	non-safety related with special treatment
NSSS	Nuclear Steam Supply System
NST	non-safety related with so special treatment
PRA	probabilistic risk assessment
PRISM	power reactor innovative small module
PSF	PRA Safety function
PWR	Pressurized water reactor
QHO	quantitative health objectives
QSG	qualitative safety goals
RAP	reliability assurance program
RCS	reactor coolant system
RFDC	required functional design criteria
RHWG	Reactor Harmonisation Working Group
RIMES	Risk-Informed Management of Enterprise Security
RIPB	risk-informed, and performance-based
RISE	Risk-Informed System Engineering
RSF	required safety functions
RWST	refueling water storage tank
SBO	station black-out
SCSSC	safety classification of structures, systems, and components
SF	safety function
SI	safety injection
SMR	small modular reactor
SNL	Sandia National Laboratory
SNM	special nuclear material
SPH	Smoothed-particle hydrodynamics
SR	safety related

SRDC	safety-related design criteria
SRP	Standard review plan
SSC	structures, systems, and components
STPA	System-Theoretic Process Analysis
TI-RIPB	technology-inclusive, risk-informed, and performance-based
TPM	technical performance measures
TS	Technical Specification
UCA	unsafe control actions
UU	unknown and/or underappreciated

Treatment of Uncertainties for Security-Related Design Aspects of Advanced Reactors When Using a Risk-Informed Licensing Approach

1 INTRODUCTION

1.1 Background

Next generation reactors will be able to use risk to license many aspects of the design, including security, which potentially saves costs. Details of this risk-informed approach, or "next-generation regulation" is found in the Nuclear Regulatory Commission's (NRC) SECY-19-0117[a] where probability is widespread through the guidance via a safety case. Further, probabilistic concepts are built into metrics such as the frequency-consequence curve (see Figure 1-1). However, one of the weaknesses found in risk-informed licensing is that they must consider the uncertainty inherent in the probabilistic concepts. Thus, nuclear designers should be "considering uncertainty" but the approach of how to do this in a real way is not well understood nor are many of the existing tools and methods set up to facilitate a technically-defensible treatment of uncertainties found when evaluating potential scenarios and the determination of associated consequences.

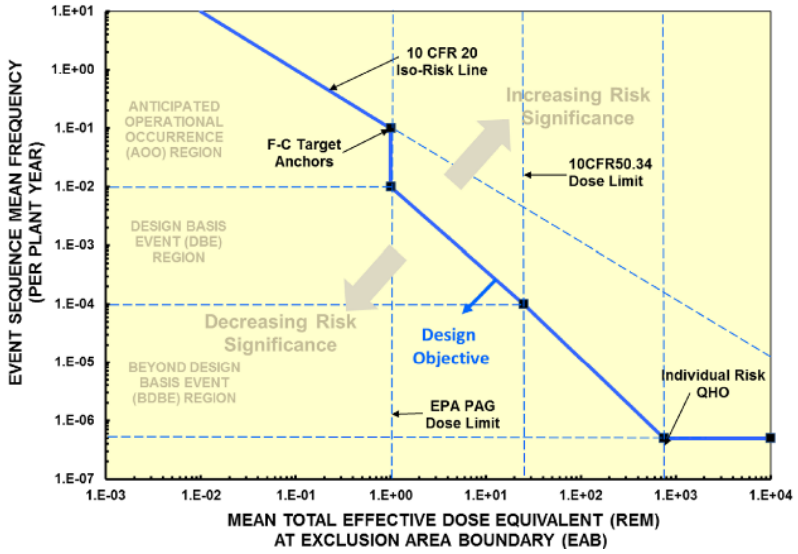

Figure 1-1 Risk-informed guidance for advanced reactor licensing per NEI 18-04.

If the nuclear industry better understand and reduce uncertainties, that knowledge will translate into a direct reduction of unnecessary conservatisms and design margin. This enhancement then translates into better economics. This report describes risk-informed thinking and the associated technical approaches for the safety case for advanced reactors. We describe scenario-based information on potential upset conditions, how the reactor responds to those conditions, and the associated consequences. We further describe methods and tools to address the security-design uncertainty.

[a] Technology-Inclusive, Risk-Informed, and Performance-Based Methodology to Inform the Licensing Basis and Content of Applications for Licenses, Certifications, and Approvals for Non-Light-Water Reactors.

2 RISK-INFORMED APPROACHES

2.1 Overview

The scope of consideration in this report corresponds to public health and safety and protection of the environment, due to threats addressed in 10 Code of Federal Regulations (CFR) 73 (physical attack, cyber attack, hybrid attack), because adequate management of the risks called out in that rule is already a challenge. However, other considerations need to be addressed for plant purposes, such as many kinds of costs, including generation risk, cost to replace equipment damaged by sabotage, etc.

Following are key points discussed in this section:

- In general, understanding the scenario to comprehend in risk management. It is important to understand the scenarios and their consequences, even if one cannot estimate likelihood.
- Classical risk management requires understanding scenario likelihood. Currently, various groups are studying scenario likelihood, but some argue that understanding likelihood is infeasible, and some workers have developed methods for risk management that try to make sensible recommendations even without explicit consideration of likelihood.
- There is an interesting mapping from a traditional frequency/consequences perspective (as reflected in the Farmer Curve and the Licensing Modernization Project) to a comparatively recent Sandia approach to management of risks from attacks on nuclear plants (RIMES). Instead of working with [scenarios, frequencies (reflecting attack likelihood), consequences], RIMES works with [attack scenarios, attack difficulty, consequences of successful attack].
- There are many approaches to risk management under conditions where quantitative information is seriously lacking. (Info-gap, robust decision-making, etc.) The key features of the RIMES approach respond to some of the considerations that drive those approaches.
- When uncertainties are large, and the future is uncertain, it is a good idea to take actions that will work out reasonably well for a broad range of possible futures.

2.2 What Does “Risk-Informed” Mean?

From a speech in 1997, Chairman Shirley Ann Jackson addressed the Plant Life Management and Plant Life Extension International Conference and Exhibition [1]

> “A "risk-informed" approach means that, in the decision-making process, quantitatively derived risk information is considered along with other factors such as the need for defense-in-depth, good engineering practice, and operating experience. Risk information does not become the sole basis for a decision (that is, the decision is **not** "risk-based"), but rather provides a systematic way of identifying and comparing what is important and where uncertainties exist.”

This quote, having originated in 1997, is very close to being the original source for the definition of “risk-informed.” By the time that speech was delivered, the NRC had been making extensive use of risk analysis for many years and had developed a policy statement (1995) [2] according to which “The use of probabilistic risk assessment (PRA) technology should be increased in all regulatory matters to the extent supported by the state of the art in PRA methods and data, and in a manner that complements the NRC's deterministic approach and supports the NRC's traditional defense-in-depth philosophy." But there was still controversy regarding how to use PRA appropriately. At that time, the NRC was still formulating details about how, or even whether, it would make use of risk information in regulations.

In part, Chairman Jackson's statement reiterated the official policy, but beyond that, it named and mandated a conceptual approach, which is now used not only at NRC but also at other federal agencies. Unfortunately, nearly everyone who writes about "risk-informed" redefines the concept. A recent NRC rewrite of it appears in NRC's white paper, "Risk-Informed and Performance-Based Regulation:" [3]

> "Risk-Informed Approach": A "risk-informed" approach to regulatory decision-making represents a philosophy whereby risk insights are considered together with other factors to establish requirements that better focus licensee and regulatory attention on design and operational issues commensurate with their importance to public health and safety. A "risk-informed" approach enhances the deterministic approach by: (a) allowing explicit consideration of a broader set of potential challenges to safety, (b) providing a logical means for prioritizing these challenges based on risk significance, operating experience, and/or engineering judgment, (c) facilitating consideration of a broader set of resources to defend against these challenges, (d) explicitly identifying and quantifying sources of uncertainty in the analysis (although such analyses do not necessarily reflect all important sources of uncertainty), and (e) leading to better decision-making by providing a means to test the sensitivity of the results to key assumptions. Where appropriate, a risk-informed regulatory approach can also be used to reduce unnecessary conservatism in purely deterministic approaches, or can be used to identify areas with insufficient conservatism in deterministic analyses and provide the bases for additional requirements or regulatory actions. "Risk-informed" approaches lie between the "risk-based" and purely deterministic approaches. The details of the regulatory issue under consideration will determine where the risk-informed decision falls within the spectrum.

Risk-informing leads to better risk management outcomes and achieves those outcomes in a more efficient way than can be achieved within "deterministic" approaches. WASH-1400 [4] illustrated the point that if we try to derive priorities through consideration of a set of design-basis events, as was done early in the history of nuclear regulation, we may fail to consider some risk-significant scenarios, while allocating excessive resources to prevention or mitigation of others. Risk-informing safety has yielded significant benefits both in safety and in efficiency. Analogous benefits should exist for security, but the existing approach to security is not as risk-informed as it arguably could be.

There are difficulties associated with risk-informing security, in the classical sense of "risk-inform." One problem has to do with the use of the concept of "importance:" Jackson's original definition invokes the idea of "important," and the NRC's white paper rewrite refers to "importance to public health and safety." Unfortunately, many applications of WASH-1400-style PRA argue that certain structures, systems, and components (SSCs) are "important" in certain ways: either they are involved in classes of scenarios that are major contributors to risk metric values, or those risk metric values are sufficiently sensitive to assumptions about SSC performance that attention to their performance is warranted. Arguably, there are issues associated with the way many applications try to apply the concept of importance, but the issue that concerns us here is that many require some sort of quantification of scenario frequencies, including the likelihood of an attack. However, this is one area where the concept of "relative risk" may play a role in helping decide on importance.

The scope of security events can be considered via the rule language from 10 CFR 73.1. The scope of NRC's concerns includes public health and safety as well as environmental protection, including both plant safety and protection of special nuclear material. These and other consequence categories are shown below on Figure 2-1, which was drafted as a tool to help people involved in risk analysis for security think clearly about the scope of their undertakings.

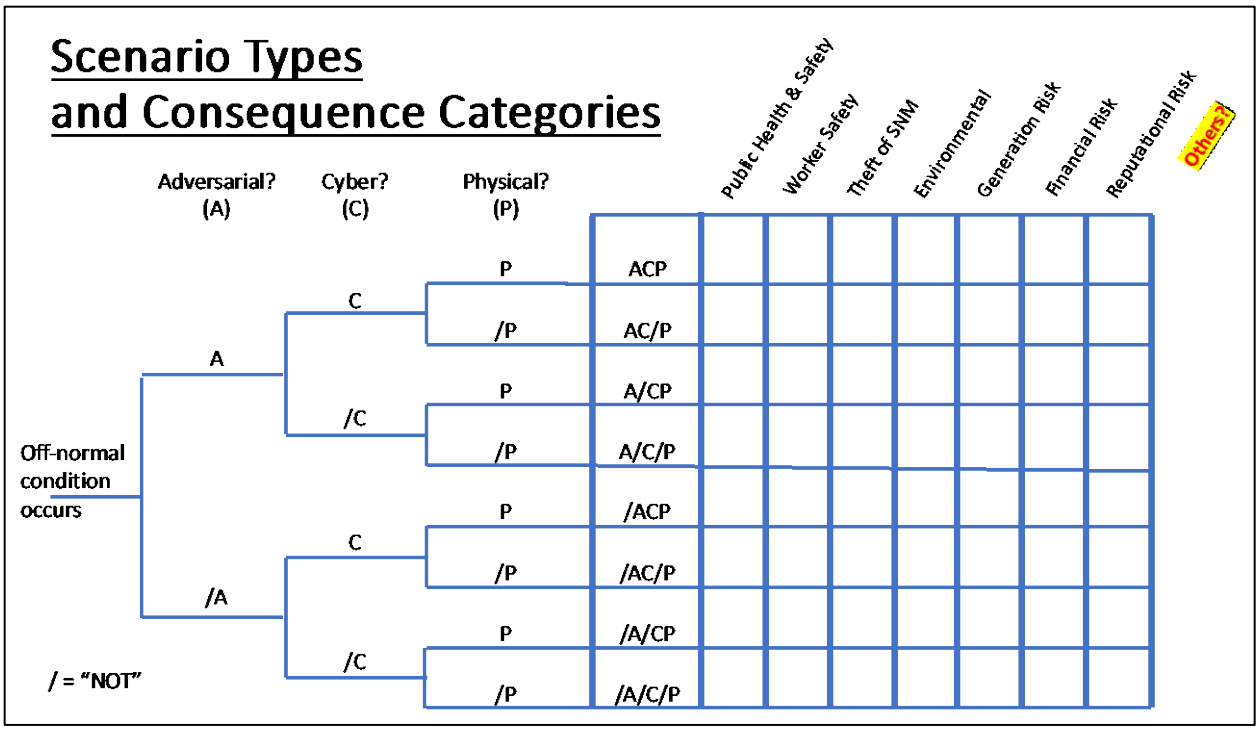

Figure 2-1 Scenario types and consequence categories of potential interest in risk analysis for security.

Table 2-1 below gives examples of the various sequence types called out in Figure 2-1, purely for purposes of illustration. In principle, the present scope of discussion includes radiological consequences, environmental consequences, or theft of special nuclear material (SNM), caused by the following sequence families: ACP, AC/P, A/CP. These three sequence families have in common that they are initiated by adversarial action, and involve cyber aspects, physical aspects, or both.

In preparing this report, we have drawn on publicly available material pertaining both to cybersecurity and physical security. Examples are for hypothetical facilities that do not exist. Scenarios, and the data behind them, are also hypothetical.

Later, we will discuss an important recent development, the Licensing Modernization Project (LMP) [5]. The purpose of the LMP is concisely stated in its abstract:

> This guideline presents a modern, technology-inclusive, risk-informed, and performance-based (TI-RIPB) process for selection of Licensing Basis Events (LBEs); safety classification of structures, systems, and components (SSCs) and associated risk-informed special treatments; and determination of defense-in-depth (DID) adequacy for non-light water reactors. This guidance document provides one acceptable means for addressing the aforementioned topics as part of demonstrating a specific design provides reasonable assurance of adequate radiological protection.

That is, the LMP is not about risk analysis per se: it is about formulating the safety case. Nuclear Energy Institute (NEI) 18-04 says almost nothing about security, but what it does say is interesting:

> Task 7e: Risk-Informed, Performance-Based Evaluation of Defense-in-Depth
>
> In this task, the definition and evaluation of LBEs should be used to support a risk-informed, and performance-based (RIPB) evaluation of DID. This task involves the identification of risk-significant sources of uncertainty in both the frequency and consequence estimates, and evaluation against DID criteria. Outcomes of this task include possible changes to the design to enhance the

plant capabilities for DID, formulation of conservative assumptions for the deterministic safety analysis, and input to defining and enhancing programmatic elements of DID.

…

This may be a point for designers to assess plant features for effective satisfaction of regulatory requirements such as 10 CFR 50.155, "Mitigation of Beyond-Design Basis Events," and 10 CFR 73, "Physical Protection of Plants and Materials."

Table 2-1 Examples of instances of sequence types called out in Figure 1.

Sequence Type	Examples
ACP	Cyber-Physical Attack (Hybrid) (could include insiders)
AC/P	Cyber Attack
A/CP	Physical Attack
A/C/P	Adversarial action of some kind, but not cyber, and not physical.
/ACP	Not adversarial, but cyber and physical. Instrumentation and control malfunction causes heavy load drop? Instrumentation and control malfunction concurrent with earthquake.
/AC/P	Non-adversarial instrumentation and control malfunction
/A/CP	Earthquake, Fire, Building Collapse, Truck Hits Transformer By Accident, Accidental Heavy Load Drop, …
/A/C/P	Not adversarial, not cyber, not physical. This category would include normal operation and many internally-initiated accident sequences.

In other words: examine the performance of the LBEs (the success paths that are credited) in security space. Implicit in this is the point that if we know what SSCs we need to make the plant "safe," then we know what SSCs we must defend from attack (cyber or physical) or failure to keep the plant safe.

Theft of SNM is a different matter. Since the success paths involved in NOT suffering theft of SNM are not identical to the success paths relevant to safety, the theft concern calls for a separate hazard evaluation.

2.3 Key Concepts of Risk Management

Risk is a term with a variety of meanings and is subject to many interpretations. Common definitions of risk include the possibility of loss or injury [6]; someone or something that creates a hazard [6]; a situation involving exposure to danger [7]; and an intentional interaction with uncertainty [8]. Consequently, in the absence of an unambiguous shared understanding of the term and the concepts that surround it, there is a possibility that those who are tasked with identifying, understanding, managing, and communicating risk will to some extent talk past each other and/or work at cross purposes, inhibiting the ability of an entity to achieve its goals safely, reliably, and affordably.

The purpose of this section is to present key definitions and concepts relating to risk and risk management that can be used as a basis for establishing, structuring, or refining risk management processes, tools, and activities, and for extending them to the distinct, but related, issue of security. The concepts presented herein reflect much of the thinking that has gone into the development of risk management in other domains, for example that found in the *NASA Risk Management Handbook* [9]. To be precise, a somewhat mathematical approach to exposition has been taken here, involving mathematical functions, coordinate axes, and the like. This is not intended to imply that risk management is necessarily quantitative. Rather, it is intended to communicate the concepts with sufficient clarity that they can be adapted to an appropriate level of rigor of risk management, whether quantitative or qualitative. The goal is to facilitate the establishment of risk management processes that are internally coherent and effectively integrated into the management activities they support.

2.3.1 An Objectives-Focused Definition of Risk

Risk is defined here as "the potential for shortfalls with respect to achieving explicitly established and stated objectives" [10]. This definition focuses attention on the "explicitly established and stated objectives" are being pursued by the entity in question, be it a process, a department, a facility, or an organization. These are the objectives that define success for the entity: If the entity meets its objectives then it is successful; if it fails to meet its objectives then it is not successful, or at least not fully successful. Defining risk in terms of the objectives that define an entity's success ensures that entity's risk management activities, like the entity's activities in general, are focused on success.

> Risk
>
> "Risk is the potential for shortfalls with respect to achieving explicitly established and stated objectives." [10]

Every entity manages its objectives. To do so, the manager of the entity establishes a set of performance measures, defined in [10] as "[the metrics] used to measure the extent to which a system, process, or activity fulfills its intended objectives." In a system engineering context, performance measures they take the form of technical process and product measures such as measures of effectiveness (MOEs), measures of performance (MOPs), key performance parameters (KPPs), and technical performance measures (TPMs), which are assessed or tracked to determine the status of the technical effort [11]. These same performance measures are also used to assess the risks to the accomplishment of its objectives. Roughly speaking, performance assessment is backwards-looking, providing managers with an understanding of how well they have succeeded so far in their efforts to accomplish their objectives. Risk assessment is forward-looking, providing managers with an understanding, however imperfect, of the challenges they are yet facing. Both activities rest on the same set of performance measures, which codify success as derived from the entity's objectives.

2.3.2 Types of Risk

The character of an entity's risks, and of the risk management processes, are a function of the corresponding objectives. Risks to the objectives of an agency or corporation are *enterprise risks*; risks to the objectives of a facility, support organization, or other form of infrastructure are *institutional risks*, and risks to the objectives of an operation or activity are *program risks*.

Within a given entity, the entity's objectives can be categorized by domain. For example, [10] defines the mission execution domains of *safety, technical, cost, and schedule*. Correspondingly, the risk associated with a particular objective can be categorized in terms of the category of the objective. Thus, the potential shortfalls with respect to a safety objective constitute a *safety risk*, the potential shortfalls with respect to a technical objective constitute a *technical risk*, and so on.

Common Types of Risk

Organizational Risk Types

Entities can be classified in terms of organizational type, for example:

- Enterprise (Strategic, Agency) risk
- Institutional risk
- Program (Project, Operation) risk

Domain-Specific Risk Types

Objectives can be classified in terms of the technical or programmatic domain to which they belong. Risk can be correspondingly classified, form example:

Enterprise	Institutional	Program
• Strategic risk	• Staffing risk	• Safety risk
• Operations risk	• Training risk	• Technical risk
• Reporting risk	• Security risk	• Cost risk
• Compliance risk	• Maintenance risk	• Schedule risk

As one would expect, the set of domains into which an entity's objectives partition is a function of the entity. Enterprise risks entail different domains than institutional or program risks, as illustrated in the table above.

2.3.3 Risk Characterization

The "potential for shortfalls" to which a given objective is exposed is typically the result of a variety of things that can go wrong during an entity's efforts to achieve it. Each possible way that things can go wrong has its own specific probability of occurrence and produces its own specific magnitude of shortfall relative to the objective. Similarly, there are also a multitude of ways that things can go right, each of which has its own specific probability of occurrence and its own successful outcome in which the objective has been achieved. Collectively, the totality of ways that things can go wrong or right define a probability density function over the space of possible outcomes, as shown in Figure 2-2. In the figure, the outcomes to the left of Objective A represent shortfalls relative to Objective A, and the area under the probability density function in that region quantifies the probability that a shortfall will occur. In other words, the blue shaded area of the figure represents the risk to Objective A, which is the cumulative result

of all the ways that things can go wrong with respect to its achievement. This type of characterization also captures the uncertainty related to the performance measure of interest.

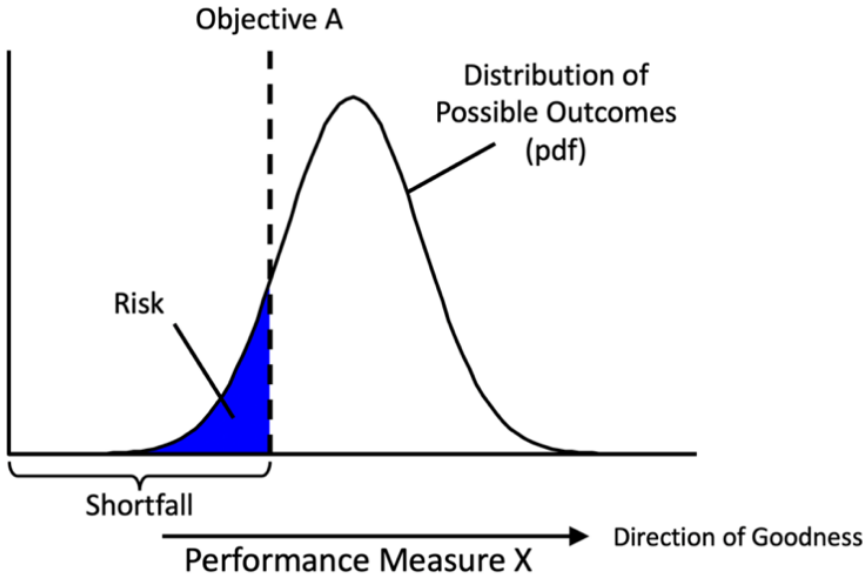

Figure 2-2 Top-level anatomy of risk.

This picture of risk as a probability of failing to meet an objective is consistent with the situation of a "binary" objective that is either met or not met[2]. It represents a model of success that is especially applicable to requirements, which are formally verified as met or not met according to their established verification protocols.

Of course, this binary picture of risk, namely that risk is the probability of falling short of the objective, is not the whole story about what can go wrong in pursuit of an objective. The magnitude of the shortfall can also be relevant, especially when the objective is a so-called "soft" objective where there is some arbitrariness to where the "line in the sand" is drawn. For example, a $100M reactor design project that overruns by $1M is likely to be perceived as more successful than one that overruns by $50M. Similarly, given a reliability objective of 0.999, an achieved reliability of 0.995 is likely to be perceived as more successful than an achieved reliability of 0.5. Similarly, the margin by which an objective is successfully achieved can also be relevant.

[2] This contrasts with an objective such as "maximize revenue," where more is better but there is no defined threshold separating success from failure.

2.3.4 Risk versus Individual Risks

The picture of risk presented in Figure 2-2 is essentially complete as a theoretical expression of risk. However, from the point of view of *managing* risk it is necessary to understand *how* things can go wrong, so that the entity can intervene to the extent practicable, both to prevent things from going wrong and to mitigate the shortfall should they go wrong. To do this, a given risk must be characterized in a way that reveals what its causes are, how likely those causes are to be operative, and how those causes can propagate through the entity and/or its work products, defeat barriers, and compromise the achievement of objectives. To this end, [10] elaborates on its definition of risk by *operationally characterizing* risk as a set of triplets: the scenario(s) leading to the shortfall, their likelihoods, and their consequences (each of which, by definition, counts as a shortfall).[3] Each scenario in this set of triplets is an *individual risk*, and represents a portion of the risk per Figure 2-1 that has been identified and analyzed, rendering it explicitly addressable by the entity's risk management function.

> Operational Characterization of Risk
>
> "Risk is operationally characterized as a set of triplets:
>
> - The scenario(s) leading to degraded performance with respect to one or more performance measures (e.g., scenarios leading to injury, fatality, destruction of key assets; scenarios leading to exceedance of mass limits; scenarios leading to cost overruns; scenarios leading to schedule slippage).
> - The likelihood(s) (qualitative or quantitative) of those scenarios.
> - The consequence(s) (qualitative or quantitative severity of the performance degradation) that would result if those scenarios were to occur.
>
> Uncertainties are included in the evaluation of likelihoods and identification of scenarios." [10]

As mentioned above, typically many things can go wrong during an entity's efforts to achieve its objectives, and in principle, this is infinitely divisible based on finer and finer distinctions between what are otherwise similar scenarios. Ss a practical matter, a given individual risk will generally envelope classes of similar scenarios that share similar causes, propagation pathways, and/or consequences. The likelihood of the individual risk will then be the likelihood that *some* scenario within this envelope will occur. In other words, individual risks are proxies that are constructed to represent a bounded multitude of underlying undesired possibilities. Construction of individual risks is a human activity, and as such is subject to human judgment, so different people may construct different sets of individual risks to characterize the same underlying risk.

Additionally, the "risk triplet" characterization of risk in [10] explicitly recognizes that a scenario leading to a shortfall with respect to one objective may very well also lead to a shortfall with respect to another objective. In practice this is usually the case, due to the high degree of correlation among the corresponding performance measures. Therefore, individual risks can, and often do, entail shortfalls relative to multiple objectives. The triplet characterization of individual risks captures these correlations, illustrating the fact that risk management cannot be effectively stove-piped by objective. It also illustrates that unlike risk proper, individual risks are not domain-specific. For example, it would make little sense to partition an entity's individual risks into safety risks, technical risks, cost risks, and schedule risks. Instead, a single individual risk can potentially threaten objectives in all these domains.

[3] NPR 8000.4B further specifies that uncertainties are included in the evaluation of likelihoods and identification of scenarios.

2.3.5 Incompleteness of Individual Risk Characterization

Many methods can be used to identify individual risks, and *risk identification* is a fundamental risk management activity. However, like most if not all discovery activities, risk identification is vulnerable to *incompleteness*, which means that there is never any guarantee that an entity's individual risks will collectively address the totality of risk to which the entity is exposed. An entity can and should minimize incompleteness by applying formalism, structure, and expertise to the task of risk identification, but it can never eliminate it. Entities such as those involved in the design, construction, and operation of advanced reactors, which operate at the cutting edge of technological innovation and human accomplishment, seldom have extensive experience bases to draw upon when identifying individual risks and are particularly vulnerable to incompleteness of risk identification. Even entities with long operating histories, and which have learned from real-world experience what their dominant individual risks are and how to control them, are still vulnerable to infrequent scenarios having mean times between events on the order of their operating experience, or which reflect unknown or unaccounted-for changes in the entities, their activities, or the environments in which they operate.

Consequently, characterization of risk based solely on identified individual risks is systematically non-conservative, and decisions based on risk characterized solely in this manner tend to unknowingly expose an entity to excessive risk, setting it up for failure. The situation is illustrated in Figure 2-3, which shows the risk to notional Objective A (from Figure 2-2) partitioned into the *known risk,* i.e., the risk that is collectively accounted for by the entity's individual risks, and the *unknown and/or underappreciated (UU) risk* that eludes operational characterization but is nevertheless just as real as the known risk. Thus, it is important, when managing risk, to recognize the potential for UU risk and include it in the overall management of risk.

The assessment and management of UU risk is particularly challenging, since it cannot be characterized and examined in terms of a set of risk triplets. Instead, entities can gain a perspective on its presence by having an appreciation of its historical magnitude in the types of activities in which the entity is engaged, and how factors such as complexity, novelty, and schedule pressure can mitigate or exacerbate it [12]. They can augment purely analytical risk identification methods with testing aimed at exposing undiscovered vulnerabilities. They can attempt to identify emerging UU risks through the liberal implementation of health monitoring mechanisms. They can increase the robustness of systems through the application of generous margins. And they can implement broad coverage risk management controls capable of addressing large classes of adverse scenarios.

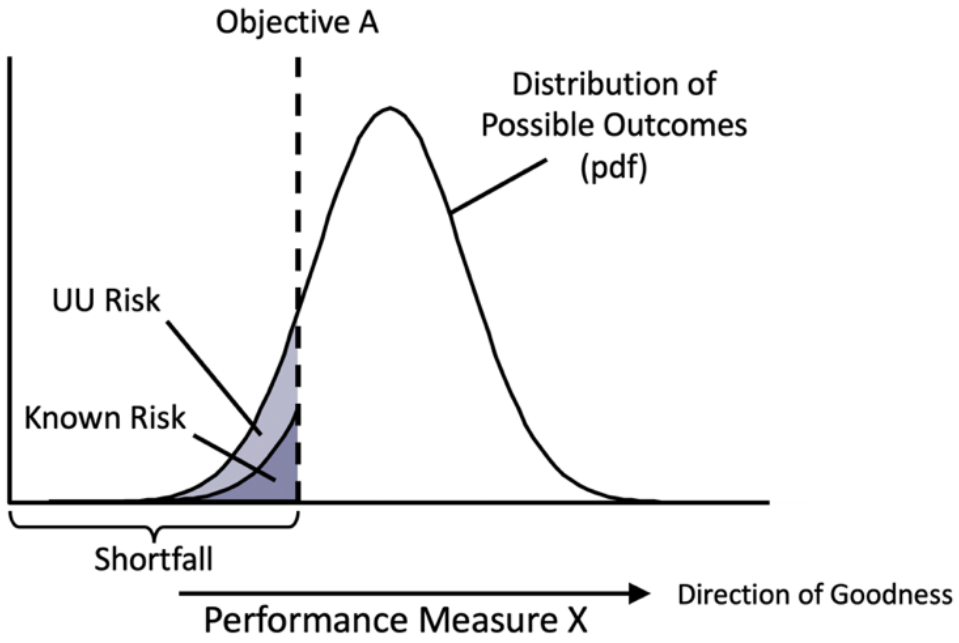

Figure 2-3 Anatomy of risk: known risk and UU risk.

2.3.6 The "Risk Model" – Putting It All Together

A full treatment of risk accounts for all sources of risk, whether explicitly identified and characterized as sets of individual risks or nebulously inferred from experience. It addresses the risk to which each of the entity's objectives are exposed, based not only on its identification and analysis of individual risks that contribute to it (i.e., the known risk) but also on an assessment of the potential for risk due to unidentified and/or underappreciated causes (i.e., the UU risk). This totality of risk information constitutes the entity's *risk model* (i.e., the authoritative representation of the risks faced by the entity), developed to a level of detail that enables it to effectively inform decision-making. Figure 2-4 notionally illustrates Entity *i*'s risk model as it relates to Objective B.

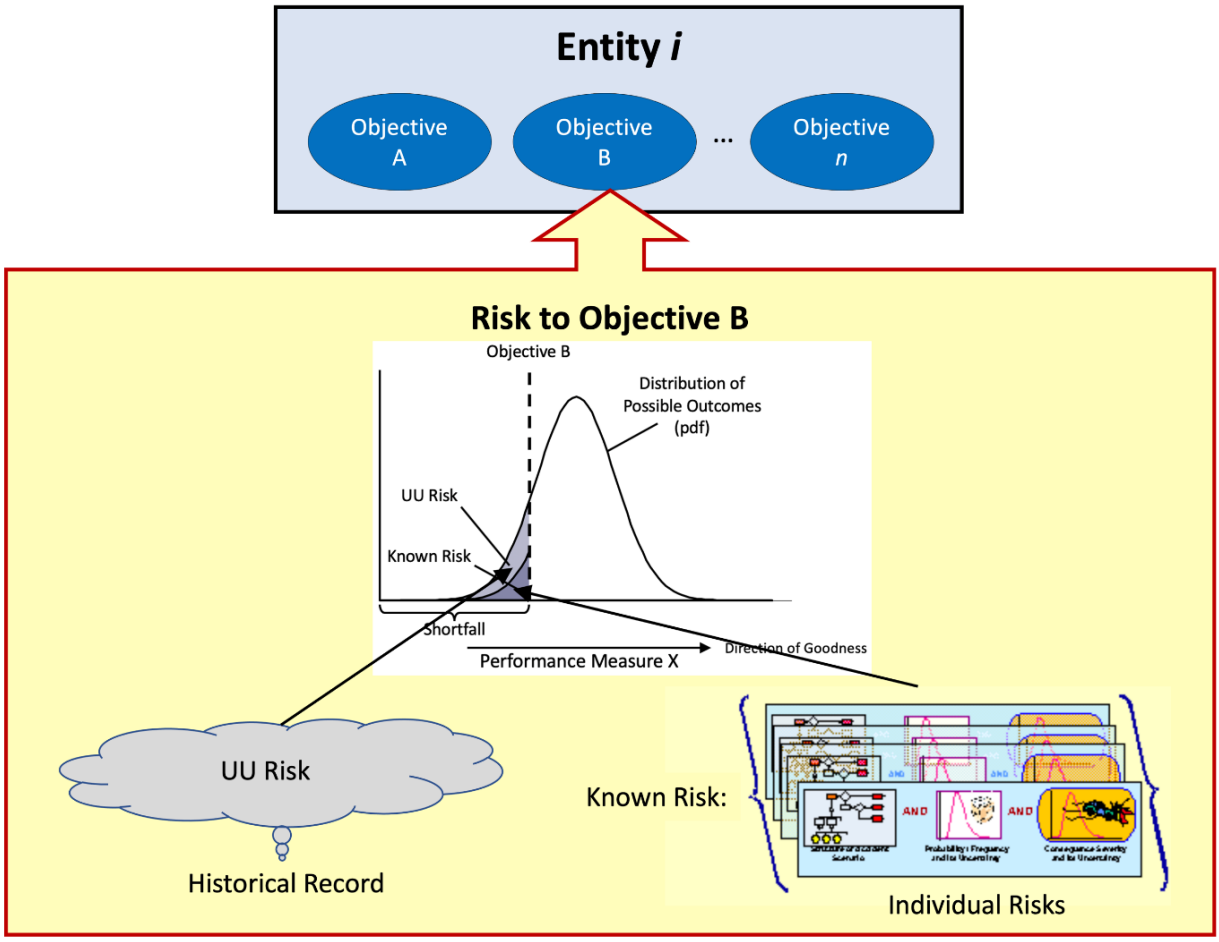

Figure 2-4 The risk model (of Objective B).

2.3.7 Risk Posture and Risk Attitudes

Up to this point we have focused on the objective characteristics of risk as described by an entity's risk model. In this section we address the entity's *attitudes* towards the risks it is exposed to, namely whether they are acceptable (and if so, by whom), or whether risk management actions should be taken to reduce or better understand them.

An entity's attitudes towards risks are an expression of its *risk posture*. NASA's *Risk Leadership* course [13] describes *risk posture* as:

> "Risk posture is our willingness to accept risk of not achieving objectives at high and low levels. It's customized to mission type and should be assessed on a regular basis throughout the lifecycle. Risk posture determines which requirements are used as defined, which are relaxed, and which are tailored – and ultimately affects the performance of designs and our chances for mission success.

Another way to think about risk posture is: 'Given the characterized risk, does the benefit to mitigating the risk outweigh the cost of the mitigation?'"

Risk posture forms the basis for a structured approach to risk management. It is understood in this paper as the framework that shapes, or even determines, an entity's attitudes towards the specific risks it faces. As such, an appropriate risk posture for an entity must be established by an entity based on its objectives, constraints, mission type, etc., *prior* to the management of its actual risk. The term *risk attitude* is used in this paper to describe the level of discomfort that an entity has towards the risks it is facing, given its risk posture.

Risk Posture, Risk Tolerances, and Risk Attitudes

Risk Posture:

An entity's risk posture is an expression of its willingness to accept the risk of not achieving its objectives. It is defined up front and serves as the attitudinal framework for responding to the actual risks that the entity faces. As discussed in [13], risk posture determines which requirements are used as defined, which are relaxed, and which are tailored – and ultimately affects the performance of designs and our chances for mission success.

Risk Tolerance:

An entity's risk tolerances relative to its objectives are the amounts of risk it is willing to expose itself to in pursuit of those objectives. Risk tolerances are fundamental elements of an entity's risk posture.

Risk Attitude:

An entity's risk attitude represents the level of discomfort it has towards an actual risk to which it is exposed. Risk attitudes are expressions of an entity's risk posture as applied to a specific risk, and motivate its response to that risk.

In general, entities must be willing to accept non-zero likelihoods of falling short of the objectives they pursue. These limits of acceptable risk, which are specific to each of an entity's objectives, are a fundamental element of an entity's risk posture, and are referred to as an entity's *risk tolerances*. A risk tolerance is the probabilistic balance point between the net value that would be gained by achieving the objective and the net cost that would be incurred by failing in the attempt. Pursuits that expose an entity to levels of risk that exceed its tolerances are to be avoided based on not being worth the risk of failure. Conversely, pursuits that expose the entity to levels that are within its risk tolerances entail rewards that justify the risk.

It is a key point of risk management that *risk characterization* is value-neutral, and it is only after risk has been characterized that a subjective elicitation of the entity's attitude towards that risk can be performed, given its risk posture. Risk management brings together an entity's understanding of the risk it is exposed to with its risk posture given the objectives its pursuing. This results in attitudes towards risk that have a rational basis in the entity's risk posture.

2.3.8 Risk Posture for Binary Objectives

A binary objective is an objective that is either met or not met, with no degrees of shortfall in the case of not being met. The risk posture for a binary objective addresses the probability of loss, with the risk tolerance expressing the limit of acceptable probability of loss the entity is willing to accept, if it must, in pursuit of the objective. For example, in security context, plant safety is typically treated as a binary objective, in which the plant (or a portion) either survives or is lost, and the risk tolerance for safety is characterized in terms of probability of loss.

We may also consider goals related to a specific objective. For example, NASA has requirements for human rating of space systems that include the establishment of risk tolerances for crew safety in terms of Agency-level safety goals and thresholds that define long-term targeted and maximum tolerable levels of risk to the crew as guidance to developers in evaluating "how safe is safe enough" for a given type of mission [14].

An entity's risk posture serves to anchor its attitude towards the assessed risk it is facing, as illustrated in Figure 2-5. In the figure, amounts of risk that exceed the entity's risk tolerance are intolerable, as shown in the red portion of the bar graph. Risk in this region renders the objective not worth pursuing. Below this level, the entity's attitude towards risk is divided into a green region where it is considered acceptable and a yellow region where it is not intolerable, but high enough to be considered marginal.

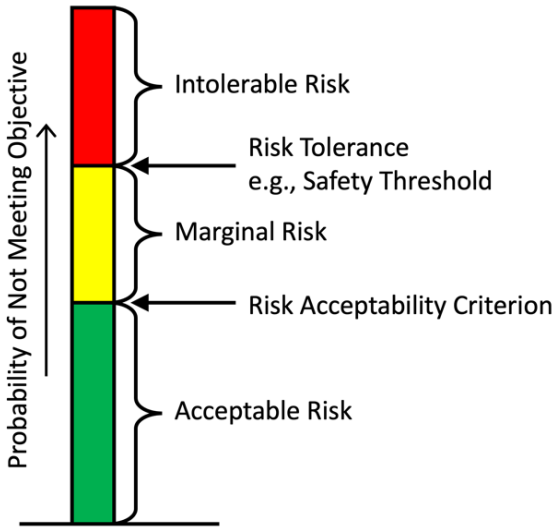

Figure 2-5 An entity's risk attitudes are anchored to its risk tolerances.

The risk posture that an entity has towards various levels of risk inform the management and acceptance of its risk. For example, [10] defines risk acceptability criteria below which an entity has the authority to accept the risk. Similarly, GEIA-STD-0010, *Standard Best Practices for System Safety Program Development and Execution* [15], describes the regions into which risk is partitioned as *risk categories* and suggests using them to determine who has the authority to accept the risk, with yellow risks requiring acceptance at higher organizational levels than green risks.

2.3.9 Risk Posture for Continuous Objectives

When the potential for a shortfall with respect to some objective spans a range of values, the magnitude of the shortfall factors into the entity's risk posture. For example, a high probability of a small cost overrun may represent the same expected loss as a low probability of a large cost overrun, but the consequences of these two scenarios may be qualitatively different; therefore, so may be the entity's attitudes towards them. One way to fully express an entity's risk posture for continuous objectives is using "F-N curves" (frequency versus a number of negative outcomes such as fatalities), as shown in Figure 2-6. In the figure, the entity's risk posture is expressed not just in terms of the (binary) probability of failing to meet the objective, but in terms of the (continuous) probability of exceeding a given shortfall magnitude, over the range of possible magnitudes. The point on the Y-axis where the risk tolerance curve intersects it is the risk tolerance for experiencing a shortfall at all, regardless of magnitude.

The blue line in Figure 2-6 is a notional example of an entity's risk to an objective, in this case Objective A from Figure 2-2. Mathematically, it is the complementary cumulative distribution function

(CCDF) for the shortfall region of the Figure 2-2 probability density function. In other words, it is the running integral of the shaded region, beginning on the left and integrating to a shortfall of zero, but flipped in Figure 2-6 so that the magnitude of the shortfall increases from left to right instead of right to left. The fact that there is a portion of the CCDF in the red region shows that the risk is intolerable, in this case not so much because of the potential for large shortfalls, but because of the excessive potential for small shortfalls.

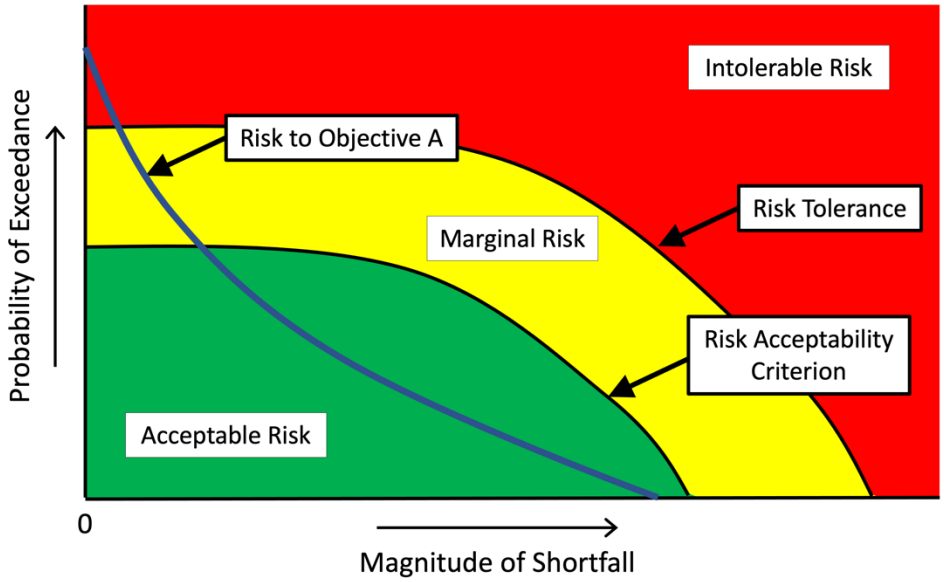

Figure 2-6 Risk posture for a continuous objective.

2.3.10 Accounting for Incompleteness of Individual Risk Characterization

As discussed earlier, identification of individual risks is intrinsically vulnerable to incompleteness, so the known risk will, in general, be only a fraction of the actual risk, the balance being UU risk. This presents a challenge for risk management because structured information of the kind necessary to characterize risk for comparison to the risk postures of Figure 2-5 or Figure 2-6 is typically only available for known risks, and it would be a serious mistake for an entity to manage known risk against its posture towards risk generally, neglecting to account for UU risk. Such practice would likely result in unknowing acceptance of intolerable levels of risk, putting the entity's efforts in jeopardy.

In the case of a binary objective like crew safety, a risk posture that is limited to known risk can be generated from the entity's risk tolerance proper by reserving a risk margin (or risk reserve) to accommodate the potential UU risk. This smaller risk tolerance, now applicable to known risk only, is the new anchor for defining risk attitudes such as *marginal* and *acceptable*, now also applicable to known risk only. The situation is illustrated in Figure 2-7, which would be used to assess the entity's attitude towards the *known* risk to the objective, calculated as the aggregate risk from all the identified *individual risks*. Risk margins/reserve are addressed in detail for safety risks in [12, 16].

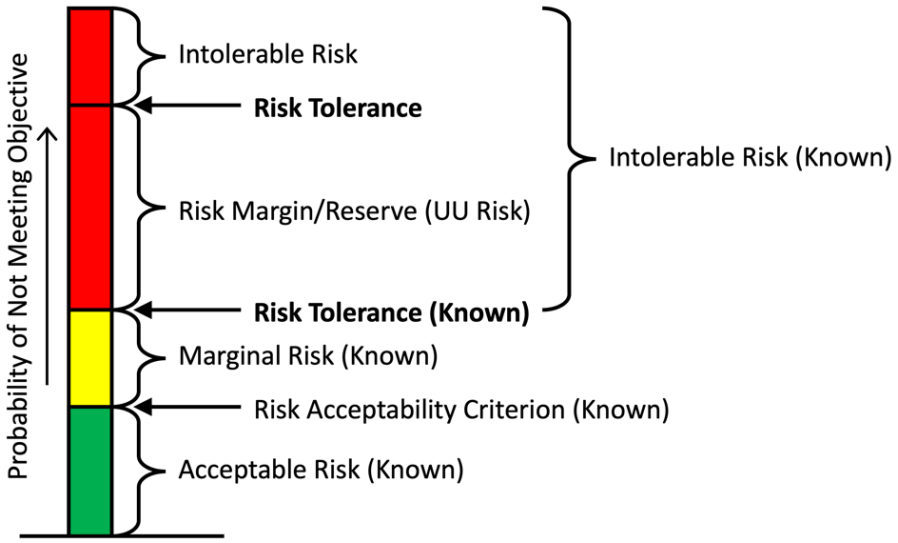

Figure 2-7 Risk posture for a binary objective, accounting for UU risk.

The situation is much the same in the case of continuous objectives, in which the entity begins with the risk tolerance curve of Figure 2-6, but then generates a risk tolerance curve specific to *known risks* by reserving a margin to account for UU risk. The continuous case is slightly more complicated theoretically, in part because the potential for unknown and/or underappreciated sources of risk is a function of likelihood since, in general, less is known about risks with lower likelihoods than risks with higher likelihoods. One property of F-N curves that can be of assistance is the fact that the expected value of the shortfall is obtainable as the area under the curve. So, for example, if an entity assessed that UU risk was likely to be on the order of half the total risk, then the F-N curve for known risk could be constructed such that the area between the two risk tolerance curves (risk proper and known risk only) is the same as the area under the known risk curve. The situation is illustrated in Figure 2-8. In the figure, the blue line now represents the *known risk* to Objective A rather than the total risk and is constructed as a CCDF from the *known risk probability density function* in Figure 2-3. Like the notional total risk CCDF of Figure 2-6, the fact that there is a portion of the CCDF in the red region shows that the known risk is intolerable.

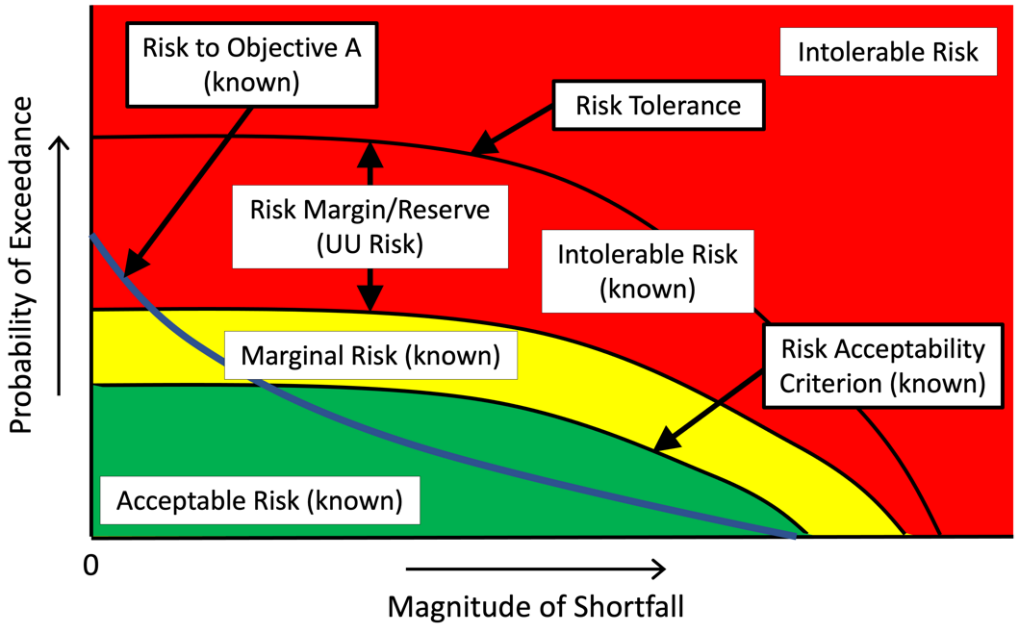

Figure 2-8 Risk posture for a continuous objective, accounting for UU risk.

2.3.11 Risk Posture Towards Individual Risks

Historically, a great deal of risk management activity has centered around entities' risk postures towards individual risks, rather than risk proper (i.e., total risk). Typically, an entity's attitudes towards the individual risks they are exposed to are expressed by plotting the risks on a *risk matrix* based on the potential consequences of the risk and the likelihood that these consequences will be realized. The risk matrix itself is divided into several regions (e.g., *red, yellow, and green*), each of which corresponds to a different risk attitude. The entity then responds to each risk in a manner that is consistent with the region of the risk matrix into which the risk falls. Figure 2-9 presents a typical 5×5 risk matrix in which each individual risk is binned in a matrix element based on its assessed likelihood and the magnitude of its assessed shortfall. Other designs and configurations also are used, based on project goals and risk posture. Augmentations of the risk matrix are also sometimes used, such as at the Jet Propulsion Laboratory (JPL), which has implemented an administrative category of *epsilon (ε)*, representing s likelihood for which there was no reportable residual risk [17].

An entity's attitude towards an individual risk is a function of the objectives that are threatened by that risk, so in principle, each objective should have its own, unique risk matrix. Decisions about what magnitudes of shortfall fall into which consequence severity bins, what ranges of likelihood map into which likelihood bins, and what risk attitudes correspond to the various intersections of likelihood and consequence severity, should be made on an objective-specific basis. However, as a practical matter, it is customary to use a single risk matrix for all an entity's individual risks, but with objective-specific mappings of likelihood and consequence severity into the matrix (i.e., a multi-axis matrix). For example, a consequence severity of "5" could apply to loss of life or crippling injury for a safety objective, a monetary loss greater than some value, or the loss of some core institutional capability. However, whatever the mapping of consequences and likelihoods into the matrix, care must be taken to ensure consistency among the different risks in terms of the entity's risk attitudes towards them. For example, a consequence severity for a 20% cost overrun that is higher than the severity of failing to achieve the entity's primary objective would be an indication that there may be an inconsistency in how the severities were decided.

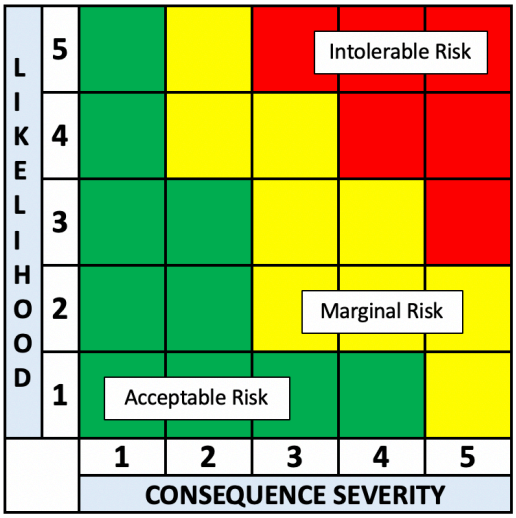

Figure 2-9 A risk matrix.

An individual risk characterizes only a piece of the risk to an objective, specifically a piece of the known risk. As such, an entity's attitudes towards individual risks cannot be anchored to its risk tolerances in the same way that total risk (or even total known risk) can. Instead, risk matrices are typically anchored not to intolerable levels of risk, but to levels of risk below which individual risks are considered insignificant. This is the green region of Figure 2-9. Bordering the green region is a yellow region of marginal risk where the entity's attitude might be to actively manage individual risks, but with the authority to accept them if the entity can make the case that they have been reduced as much as practicable. Beyond this is a red region of intolerable risk where the risk is seen as a significant vulnerability. The entity may be required to elevate the management of risks that it cannot reduce to at least the yellow region.

2.3.12 The Complementary Nature of F-N Curves and Risk Matrices

As discussed above, F-N curves and risk matrices have different, complementary anchorings. F-N curves are anchored to the tolerability of the total risk, whereas risk matrices are anchored to the adequacy of the risk management effort to address discrete vulnerabilities. An entity whose F-N curves show its risks to be acceptable can validly claim that its objectives are worth pursuing despite the risks. An entity whose risk matrices show its individual risks to be acceptable can validly claim to have adequately controlled the known vulnerabilities.

The complementary nature of F-N curves and risk matrices is closely related to the notion of "adequate safety" in [16]. In [16], adequate safety is defined as (1) achieving a system that meets or exceeds the minimum tolerable level of safety, and (2) achieving a system that is as safe as reasonably practicable. The status of these two components of adequate safety is communicable via F-N curves and risk matrices, respectively. Figure 2-10 illustrates the concept, expanded from safety to risk generally.

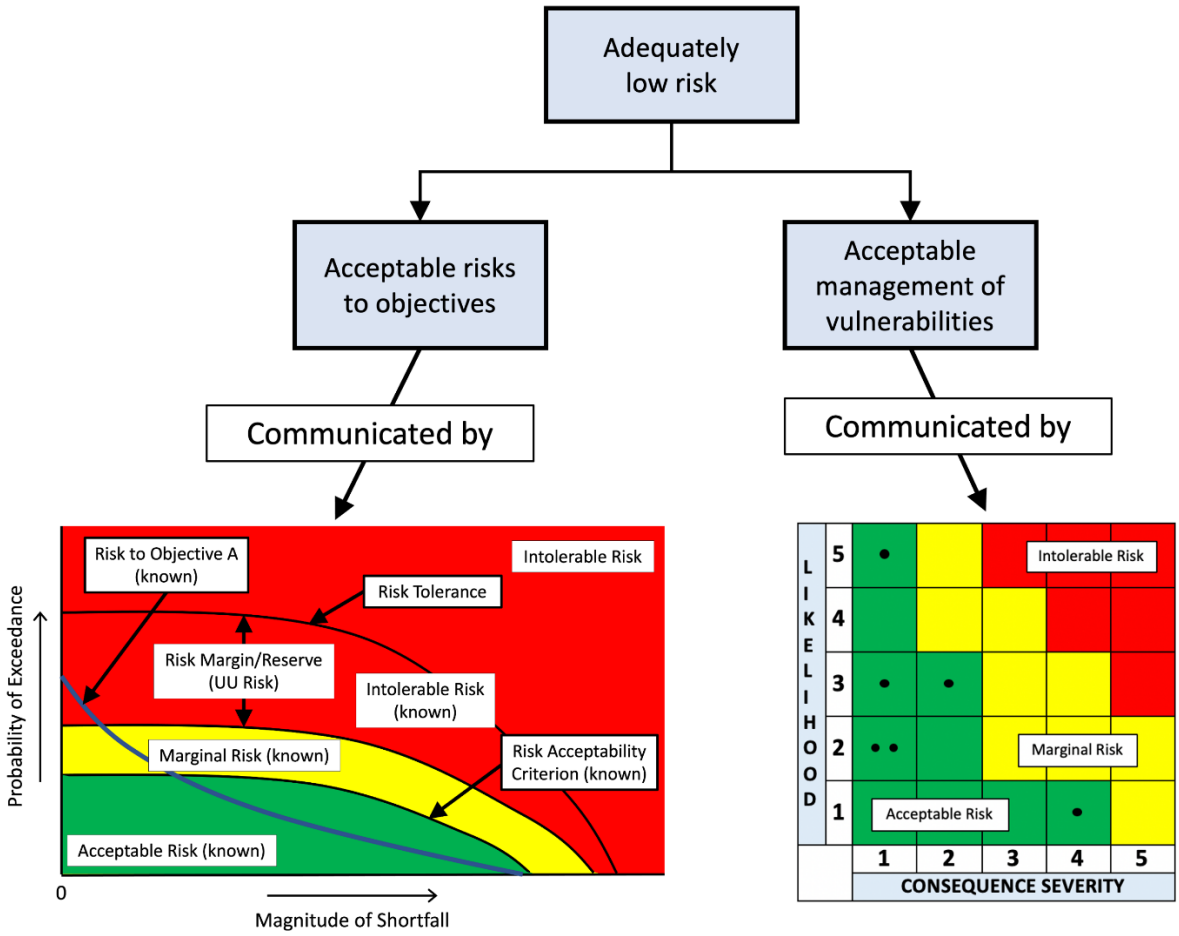

Figure 2-10 Complementary functions of F-N curves and risk matrices.

Figure 2-10 reinforces the point that F-N curves and risk matrices are planning and communication tools, *not* analysis tools. They map the entity's already-assessed risks against its risk attitudes, graphically making the case for their acceptability. When risk is in the red, a risk management response is needed (potentially including elevation) to rectify its intolerability. When risk is in the green, stakeholders can be assured of adequately low risk. All this is contingent, of course, on proper *prior* analysis of risk (including risk margins to account for UU risk) and *prior* establishment the entity's risk posture.

2.3.13 Application of Risk Management Concepts to Advanced Reactor Risk Security

The NEI Technical Report 18-04 [5] presents a modern, TI-RIPB process for selection of Licensing Basis Events (LBEs); safety classification of SSCs and associated risk-informed special treatments; and determination of defense-in- depth (DID) adequacy for non-LWRs. Part of this process involves the evaluation of radiological dose risk, both in terms of individual risks and cumulative risk.

2.3.14 Treatment of Individual Risks in NEI 18-04

Figure 2-11 shows the frequency-consequence evaluation correlation (i.e., the "F-C Target") used in the NEI 18-04 process to evaluate individual risks. It is NEI's implementation of the risk matrix concept

of Figure 2-9. Although not strictly a matrix, its use is equivalent, its purpose being "to evaluate the risk significance of individual [anticipated operational events (AOOs), design basis events (DBEs), and beyond design basis events (BDBEs)," where AOOs, DBEs, and BDBEs represent a partitioning of individual risks, i.e., licensing basis events (LBEs), based on event sequence frequency. The F-C Target values themselves are analogous to the boundary between the red and yellow regions of Figure 2-9, with the stated caveat that it should not be considered as a demarcation of acceptable and unacceptable results. Rather, the F-C Target "provides a general reference to assess events, [structures, systems, and components (SSCs)], and programmatic controls in terms of sensitivities and available margins."

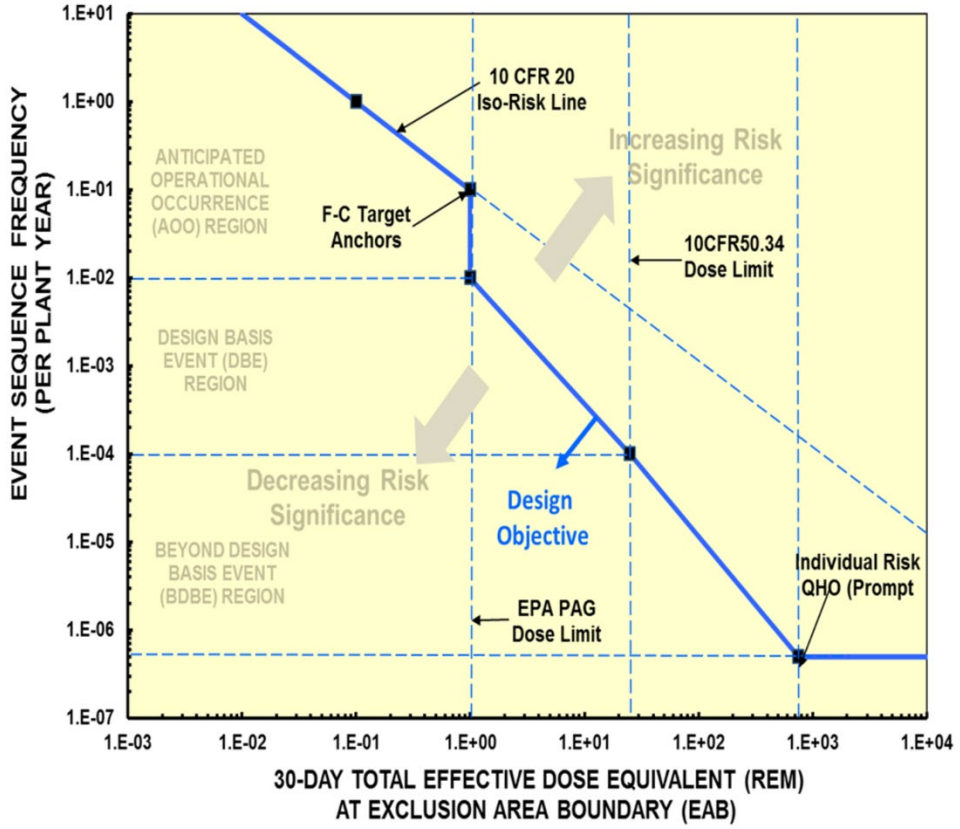

Figure 2-11 NEI 18-04 frequency-consequence target.

2.3.15 Treatment of Aggregate Dose Risk in NEI 18-04

NEI 18-04 specifies the evaluation of aggregate risk, accounting for all LBEs, against three cumulative risk targets:

1) The total mean frequency of exceeding a site boundary dose of 100 mrem from all LBEs should not exceed 1/plant-year. This metric is introduced to ensure that the consequences from the entire range of LBEs from higher frequency, lower consequences to lower frequency, higher consequences are considered. The value of 100 mrem is selected from the annual cumulative exposure limits in 10 CFR 20.
2) The average individual risk of early fatality within 1 mile of the Exclusion Area Boundary (EAB) from all LBEs based on mean estimates of frequencies and consequences shall not exceed 5×10-

7/plant-year to ensure that the NRC safety goal quantitative health objective (QHO) for early fatality risk is met.

3) The average individual risk of latent cancer fatalities within 10 miles of the EAB from all LBEs based on mean estimates of frequencies and consequences shall not exceed 2×10-6/plant-year to ensure that the NRC safety goal QHO for latent cancer fatality risk is met.

The first target implies a binary objective of not exceeding 100 mrem at the site boundary, and is an instance of the situation illustrated in Figure 2-5. The second and third targets address the objectives of protecting the public against radiologically caused early fatalities and latent cancer fatalities, respectively, both of which are functions of dose magnitude. In these cases, Figure 2-6 applies, where the aggregate risks can be calculated as the areas under their respective F-N curves. Like the F-C Target for individual risks, the cumulative risk targets are not intended to define acceptable vs. unacceptable, but rather to "[focus] attention on matters important to managing the risks."

NEI 18-04 does not explicitly apply any risk margins per Figure 2-7 and Figure 2-8. Instead, in the context of DID adequacy, its process includes an evaluation of the safety analysis in terms of its completeness, as discussed in detail in the report.

2.4 Setting Priorities

2.4.1 Delineation of Goals and Thresholds

In some domains (process safety, space flight, etc.) and in some countries, attempts are made to articulate "threshold" values of risk. The response to exceeding such a value may vary across applications, but in at least some of those domains, the operational significance of a "threshold" value would be that an activity posing risk exceeding that threshold is simply not permitted, on such grounds as that the risk posed by that activity is not justified by that activity's benefit. However, there are very real issues associated with applying uncertain risk numbers in such a simplified, all-or-nothing way, even for topic areas where uncertainties are only moderate. In security, certain uncertainties are especially daunting, applying threshold values of risk metrics as a sole decision approach is arguably a questionable idea in the security arena.

However, the use of safety *goals* by the NRC has been very significant, and the idea of "goals" has also found its way into other domains. The following discussion is taken from Mubayi and Youngblood [18]:

In 1986, the NRC issued a Safety Goal Policy Statement that established two qualitative safety goals (QSGs) and two QHOs. The adoption of these goals was meant to indicate to the public the NRC's expectation that nuclear power plants should be designed and operated in a safe manner so members of the public should bear no significant risk of adverse health consequences resulting from living near a nuclear plant. The QSGs are as follows:

> 1. Individual members of the public should be provided a level of protection from the consequences of nuclear power plant operation such that individuals bear no significant risk to life and health.
>
> 2. Societal risks to life and health from nuclear power plant operation should be comparable to or less than the risks of generating electricity by viable competing technologies and should not be a significant addition to other societal risks.

Not only do the QSGs point to a general goal, they also tie that goal to the societal benefits of the activity (nuclear power production): risks from nuclear plants should be less than those associated with viable competing technologies. An implicit corollary of this point is that the QSGs are meant to include

radiological risks to health and safety from all accident sequence types, arguably including accidents initiated by adversarial action.

The QHOs provided a quantitative guideline for what the NRC considered as no significant risk to individual members of the public from plant operation. They were developed in terms of the risk of radiation exposure to members of the public from accidental releases of radioactivity that could lead to either a prompt fatality or an induced cancer fatality. The QHOs are as follows:

> 1. The risk to an average individual in the vicinity of a nuclear power plant of prompt fatalities that might result from reactor accidents should not exceed one-tenth of one percent of the sum of prompt fatality risks resulting from other accidents to which members of the U.S. population are generally exposed.
>
> 2. The risk to the population in the area near a nuclear power plant of cancer fatalities that might result from nuclear power plant operation should not exceed one-tenth of one percent of the sum of cancer fatality risks resulting from all other causes.

The first QHO is calculated as an average individual risk over the population residing within 1 mile of the plant. The second QHO in practice is calculated as an average individual risk over people living within 10 miles of the plant. Although the second QHO is referred to as a societal health objective, it is estimated in Level 3 PRAs as an individual latent cancer fatality risk. Both QHOs are focused on individual risk alone.

Explicit comparison of the risk posed by a given plant with the QHOs obviously needs to be based on quantitative risk analysis. This has been done many times for internal events risk and for external events risk. It must be reiterated that the goals are not regulatory requirements. In the context of PRAs performed after the Three Mile Island Unit 2 accident, the QHOs performed a valuable function of indicating the level of risk considered significant by the Commission. The QHOs were not requirements, but they were meant to indicate to reactor designers and operators as well as the public what level of safety the regulations were trying to achieve.

In practice, the QHOs play a role in reasoning about the significance of contributors to risk (see, Regulatory Guide 1.174 [19]). Since the QHOs are meant to be compared with total plant risk, any single issue that contributes a significant fraction of QHO risk must arguably be considered significant and potentially worth reducing, and that idea is built into the safety-goal-based screening criteria used in the U.S. to decide whether to pursue a candidate generic safety issue. So, for example, one could in principle consider a goal for adversarially initiated radiological releases corresponding to some small percentage of the QHO value.

Another important feature of that screening process is that comparisons with the QHOs tend to be predicated on credit for evacuation of the population that would otherwise be exposed to the radioactive material released in a severe accident. To consider a hypothetical case for purposes of illustration, if all members of the public near the plant are evacuated without incurring exposure, then the radiological consequences to the public are zero. Crediting evacuation in reduction of estimated radiological consequences has an enormous effect on the perception of accident risk. However, the non-radiological consequences of long-term evacuation are very significant, as illustrated in the case of the large number of people evacuated from the vicinity of the Fukushima plants because of plant damage caused by the Great Tohoku Earthquake and Tsunami. This point receives a great deal of discussion in Ref. [18], and there are multiple reasons to consider it for advanced reactors. One reason is that although the U.S. goal structure does not currently reflect consideration of long-term evacuation, the U.S. might someday decide to consider it. Another reason is that vendors wishing to sell into foreign markets might be expected by their prospective customers to address this point, as discussed below.

A European idea of interest in this connection is the idea of “practical elimination” [20] of scenarios associated with releases having significant societal implications. The idea of “practical elimination” is advocacy of modification of plant designs to essentially “eliminate” failure modes capable of leading to a

significant radiological release. Apart from the arguable safety advantages of such a design approach, it avoids the necessity of arguing extremely small probabilities in domains where there is significant phenomenological uncertainty. Some categories of advanced designs have advantages in this area [21]. Actually, there is sufficient diversity of views and situations in Europe that while "physical impossibility" is the preferred approach, "[d]emonstrating practical elimination via 'extreme unlikeliness with a high degree of confidence'" is also allowed for. [20]

This topic arises here partly because of the essential difficulty of quantifying attack frequency. Some would say that all risk metrics are difficult to quantify, and the difference between attack frequency and (say) earthquake frequency is just a difference of degree. It is true that earthquake hazard is uncertain, but attack likelihood is not only uncertain, it may change as a function of current events*;* adversaries' motivations can change, and adversaries are moreover adaptive to plant changes*.*

In practice, in U.S. regulatory application, it is recognized that uncertainties in modeled values of risk metrics are significant. Satisfaction of the goals cannot be proven in any case, and may be far from assured. Instead, the goals are used to set priorities, and sometimes as a stopping rule (as part of an argument that a given situation is safe enough).

Common practice in NRC applications is to compare mean values of risk metrics with goals. To calculate these means, it is necessary to have meaningful distributions of the metrics. In the security arena, special difficulties afflict attempts to apply safety-goal-type reasoning. Risk metrics are uncertain, but in security, the situation is worse, because attack likelihood is not only uncertain but also variable.

Although the costs of long-term evacuation are formally within the scope of full regulatory analysis of decision alternatives, they are not factored into the safety-goal-based screening criteria (the criteria applied to decide whether to perform detailed analysis on an emerging generic safety issue). In practice, credit for evacuation in reduction of estimates of accident consequences seems to be wired into the current U.S. safety goal structure.

The safety goals do not address theft of special nuclear material. Trying to formulate a probabilistic goal aimed at promoting a low frequency of successful theft would be affected by the same issues that affect attack likelihood.

It may eventually be possible to articulate a useful qualitative goal on theft of SNM that resembles the European goal of "practical elimination" of severe-accident radiological releases. Such a goal might say something like "the potential for theft of SNM should be practically eliminated as an option for non-state actors."

2.4.2 Reginald Farmer and the Licensing Modernization Project

Within the traditional safety domain, for a given facility, one can identify potential release scenarios and characterize the frequencies and consequences of those scenarios: not to high precision, necessarily, but arguably to a level at which a licensing / siting decision can be usefully informed.

Figure 2-12 is a markup of a figure originally proposed in the 1960s by Farmer [22] as an illustration of an approach proposed to be used in decision-making about siting (well before the NRC's safety goals were promulgated). The horizontal axis (labeled "Curies I-131") refers to the magnitude of an iodine release from a particular envisioned scenario; the vertical axis ("Reactor Years") refers to the average time interval between repetitions of that scenario (in essence, the reciprocal of the average scenario frequency). Characterized by its frequency and the magnitude of its I-131 release, any given scenario maps to a location on this plot. The curve drawn in black represents a notional boundary in this space between scenarios deemed potentially acceptable (relatively low frequency / consequences) and scenarios deemed potentially unacceptable (relatively high frequency / consequences).

The use of I-131 as a proxy for the whole release reflects the presumption of that era that immediate dose from the plume was the dominant consideration. Within this picture, given a facility design, a

candidate site, and a trustworthy model of scenarios, frequencies, and consequences, one could evaluate the suitability of that site for that design based on a figure like the one below, which is a markup (in red) of a figure originally drawn by Farmer.

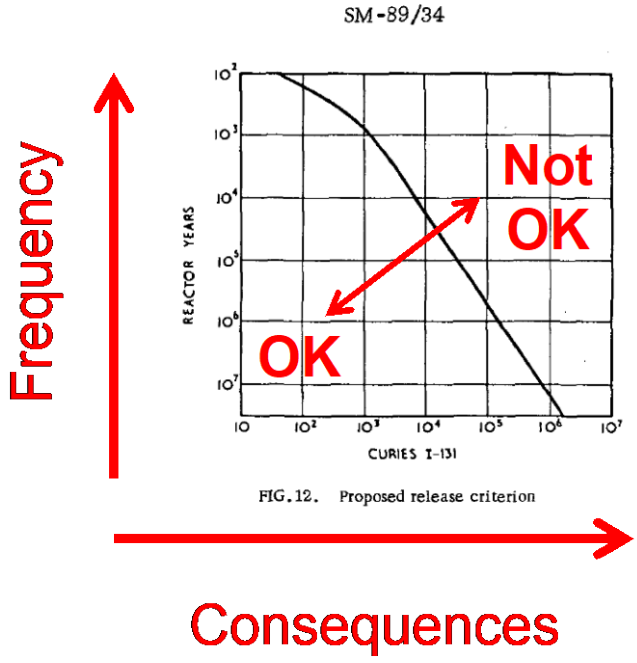

Figure 2-12 Markup of Farmer's original curve.

The curve is, in principle, site-specific, because the potentially affected population is site-specific, and from a societal-risk point of view, the population size is a factor in the relative undesirability of releases of a particular magnitude. Also noteworthy is the increase in the slope of the curve as the consequences become higher; this reflects the attitude that the decision-maker is more averse to high consequences from a single event than might have been expected based on a simple-minded idea of risk as "expected consequences."

It is often remarked that although comparing a set of individual scenarios to such a line is interesting, it does not address the question of total risk (or uncertainty). Farmer was aware of this point, but considered that risk would be dominated by a few scenarios at most, so that viewing suitability based on those dominant scenarios would not be inappropriate.

The idea of plotting frequencies and consequences as in the "Farmer curve" has been widely influential, especially in non-nuclear industries, and something like it has recently been used in the LMP [5]. The LMP figure, shown below (Figure 2-13), is meant to apply a similar (*not* identical) idea in a modern setting.

The thick blue line is drawn by the LMP authors based on the historical NRC views regarding several different categories of events. Again, we see a slope change, reflecting greater aversion to higher consequences.

Shown on the figure is a point corresponding to the "Individual Risk QHO," which is related to the expectation articulated in the Commission's Safety Goal Policy Statement. According to NEI 18-04,

> The F-C Target for the BDBEs [Beyond-Design-Basis Events] range from 25 rem at 10-4/ plant-year to 750 rems at 5×10-7/plant-year to ensure that the Quantitative Health Objective (QHO) for

early health effects is not exceeded for individual BDBEs. The question of meeting the QHOs for the integrated risks over all the LBEs is addressed using separate cumulative risk targets described later in this guidance document.

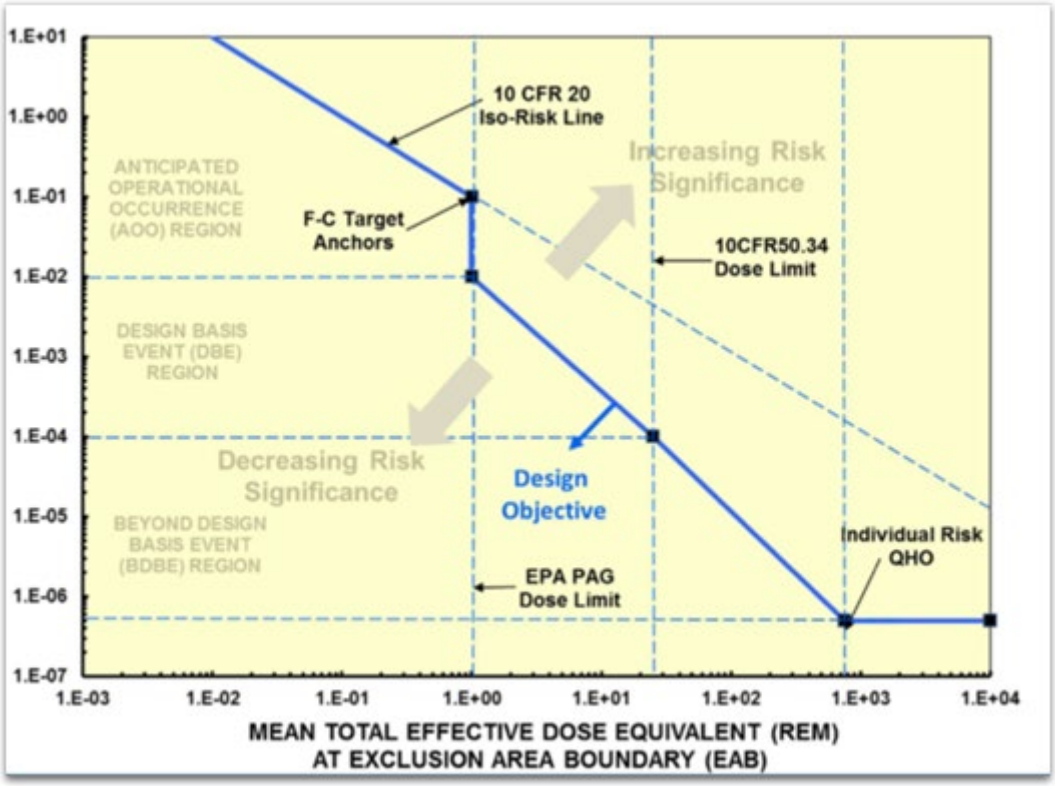

Figure 2-13 From the Licensing Modernization Project [5].

Another NEI comment goes to the point of whether the boundary is pass/fail:

> The F-C Target values shown in the figure should not be considered as a demarcation of acceptable and unacceptable results. The F-C Target provides a general reference to assess events, SSCs, and programmatic controls in terms of sensitivities and available margins.

Different nations apply values like these targets in different ways. In the US nuclear arena, the goals really are goals (objectives that licensees should try to accomplish), not formal pass/fail criteria. Some other nations appear to apply them as criteria in some domains. Also, going beyond health effects of radiation, some other nations recognize the undesirability of large-scale evacuation, which avoids radiological health effects but does so at a potentially tremendous cost. For example, Finland goes so far as to articulate an explicit frequency criterion for large release of cesium. Dose rates from cesium deposited in the environment remain significant for a long time: "significant" not necessarily in the sense of causing health effects in a short time, but in the sense of surpassing typical thresholds on acceptability of public dose rates, and thereby warranting evacuation. In the US, the cost of long-term evacuation is considered in some applications of regulatory analysis, but the NRC's process for screening candidate generic issues is based on the NRC's safety goals, which do not penalize evacuation.

Shown on Figure 2-13 are two gray arrows indicating the directions of increasing risk significance (up and to the right) and decreasing risk significance (down and to the left). On this figure, risk significance is related to the frequency and the magnitude of the dose at the exclusion area boundary. If we articulate an

expectation that a given plant's scenario set should preferably fall below and to the left of a given boundary, then in addition to each scenario having absolute significance, the distance of a given scenario's frequency/consequence point from that line is a measure of the plant's safety performance in that scenario with respect to our expectation.

The LMP process goes on to use this figure as a framework for discussing treatment of SSCs involved in various event categories. For a security analog to follow closely in the LMP's footsteps, we would need to be able to assign frequencies to attack scenarios. Whether it is technically practical to assign frequencies to attacks has been (and remains) a subject of research and debate. The following subsection reviews some related work, and suggests a correspondence between one particular approach (RIMES) and the LMP process.

2.5 RIMES (Risk-Informed Management of Enterprise Security)

Multiple papers [23, 24] have been written about RIMES. The issue addressed by these papers is this. Generic risk management calls for prioritizing security investments based on the risks that they reduce; but if we try to apply this to security, we need to be able to specify likelihoods of attacks.

Risks due to upsets NOT caused by adversaries can be analyzed using engineering models informed by operating history. Some investigators argue that risks due to adversarial attacks can be analyzed in the same way, and perhaps some can. However, the attack likelihood is influenced by diverse considerations, including the difficulties of succeeding with an attack of a particular sort, resources available to the attackers, the attackers' state of knowledge, the magnitudes of the consequences that the attackers are likely to be able to bring about, and numerous other factors relating to the attackers' motives and other influences on their decisions. Even if we could somehow quantify attack likelihood as a function of all these things, most of these considerations fluctuate in time in ways that cannot be anticipated. Alternatively, one might consider quantify the likelihood at a higher level, simply counting different types of "security events" over time like what is done for events that have not been seen at nuclear power plants (e.g., large break loss-of-coolant accidents).

The RIMES team thinks about attacks in the context of two relevant time scales: tactical and strategic. Sometimes, intelligence may indicate that an attack of a particular sort on a particular facility is imminent; this "tactical" situation drives the defenders' response in a particular way (but obviously not during the design of the facility). On the other hand, a much longer "strategic" time scale, for example one appropriate for application to a facility yet to be built, gives rise to completely different possibilities.

Instead of trying to base risk management on consideration of attack likelihood and attack consequences, the response of the RIMES team to the above considerations is to think in terms of attack *difficulty* and attack consequences. These concepts are illustrated in Figures 2-14 and 2-15. [24]

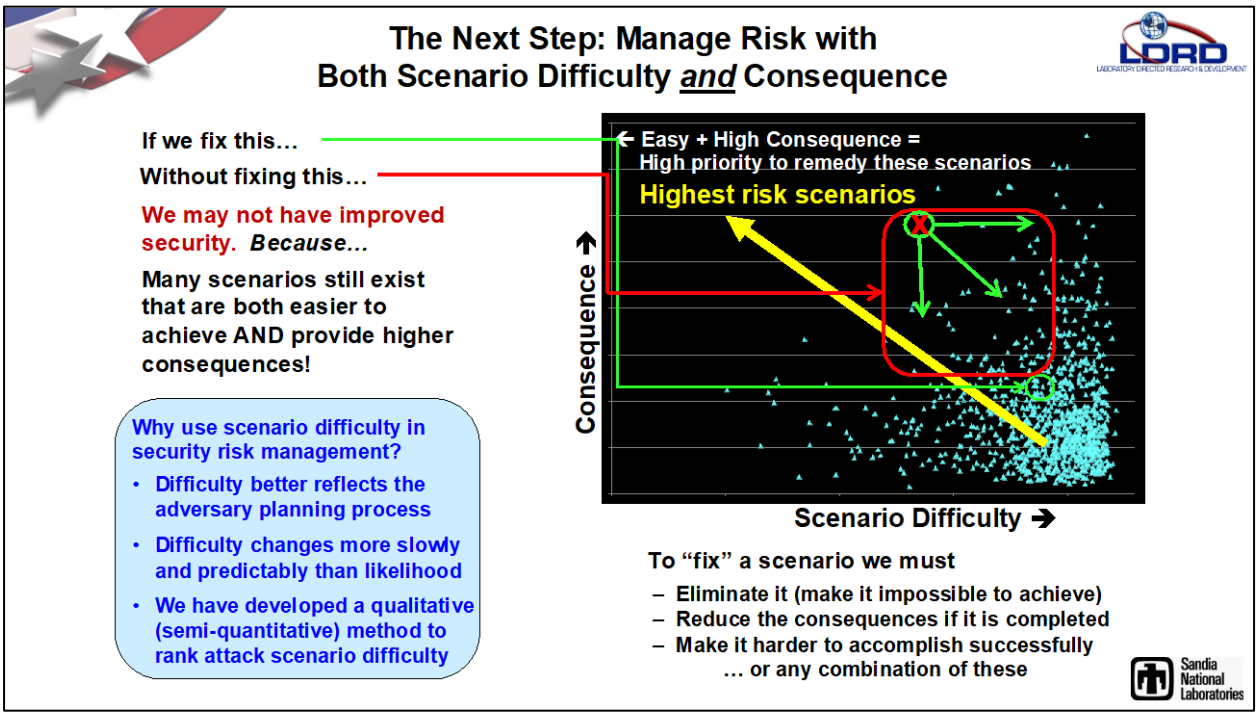

Figure 2-14 From Wyss [24], managing risk with both scenario difficulty and consequence.

It can be argued that from an attacker's point of view, other things being equal, the "best" attacks are ones with the proper balance between difficulty and payoff to the attackers (induced consequences). As noted on Wyss' slides above, addressing scenarios that are dominated by other scenarios from the attacker's point of view (that is: the dominating scenarios are easier for the attacker to accomplish, and/or more impactful) does not necessarily reduce real "risk" in a very significant way: the attacker was probably not going to choose those attacks anyhow (unless the attacker does not know of the better choices). Focusing on dominated attacks would be analogous to reducing the likelihood of minor cut sets in a WASH-1400-style PRA, while leaving the dominant cut sets unaffected.

As suggested in Figure 2-16, we can place the RIMES slide into correspondence with the LMP/Farmer slide by (1) rotating it 90 degrees clockwise, and (2) thinking about attack difficulty instead of attack probability.

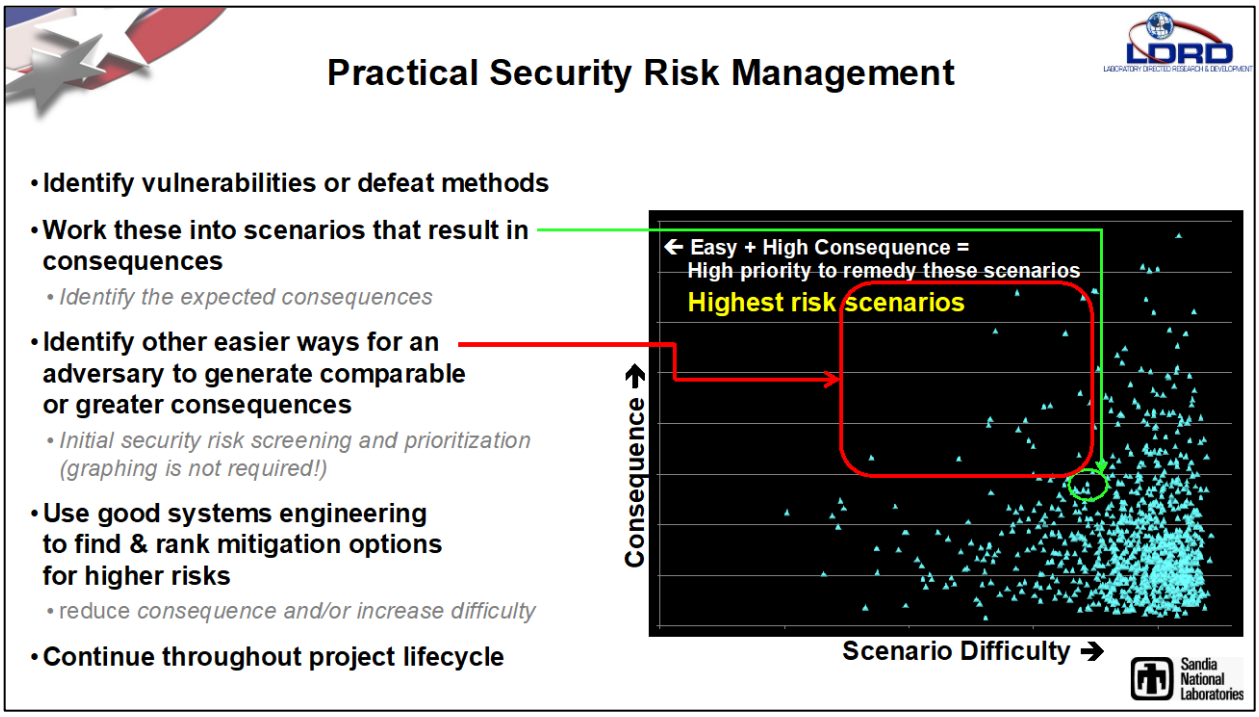

Figure 2-15 From Wyss [24], practical security risk management.

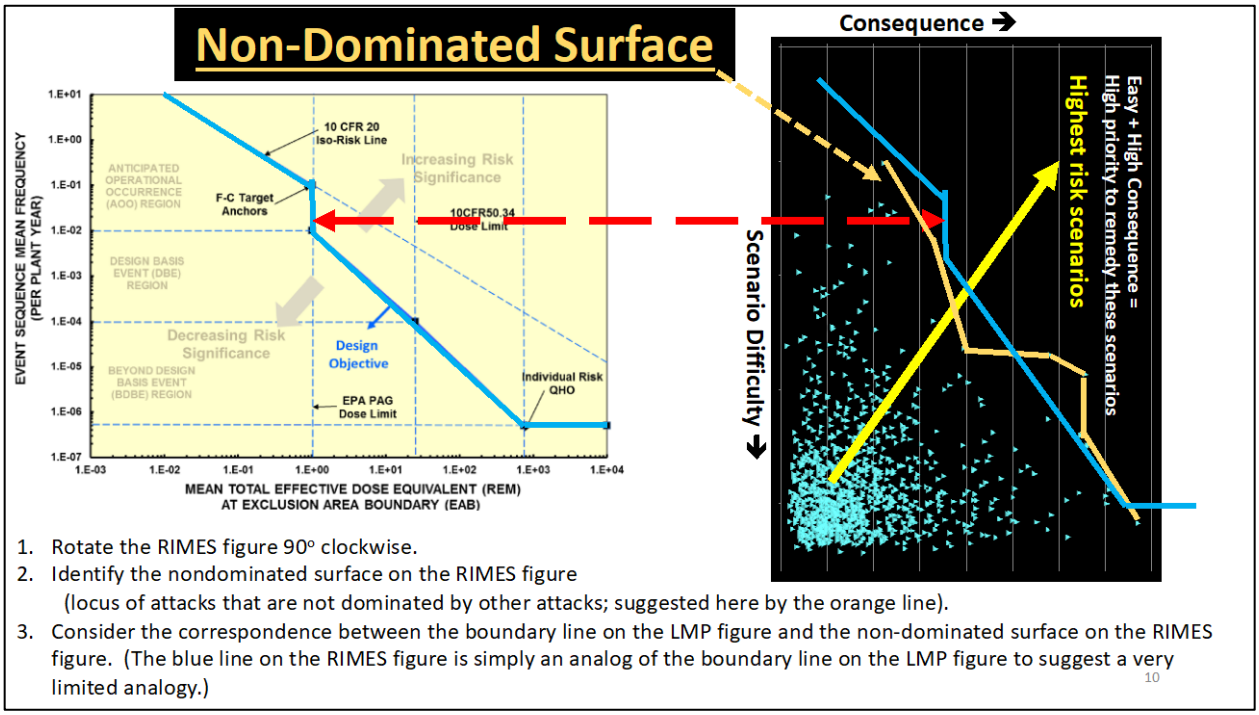

Figure 2-16 Correspondence between RIMES and LMP/Farmer.

It should be visually obvious that there is a meaningful analogy between the difficulty / consequence plot and the frequency / consequence plot. But there are key differences as well, beginning with the point that the LMP figure makes explicit use of probability, while the RIMES plot does not (arguably *cannot*). The boundary line on the LMP figure is meant to have meaning in an absolute sense: scenarios are potentially acceptable, or not, depending on their frequencies and consequences. On the RIMES plot, it is straightforward to compare scenarios with each other, but not with an absolute standard involving frequency / probability.

The WASH-1400 study was published in 1975, and although widely criticized, it nonetheless established the foundational principles of PRA still widely used today. Shortly thereafter, a modified version of the societal risk model was first proposed for nuclear safeguards (security) [18]. Known as the ERDA-7 proposal, this approach was evaluated by Rasmussen, who concluded that safeguards (security) risk could not be quantified using the WASH-1400 developed societal risk approach [19]. Rasmussen said that he did not believe that risks involving malevolent human action could be quantified by traditional risk assessment methods like fault tree and event tree analysis because attack probability estimates could not meet important statistical requirements [19]. Over the years, the ERDA-7 proposal has been subject to reintroduction and modification [1, 20, 21, 22]. Similar to Rasmussen's conclusions, subsequent critical reviews stated an approach like ERDA-7 proposal based on traditional risk assessment not be used for security risk [23, 25]. The ERDA-7 approach is problematic for intentional malevolent acts, the terms in the equation are interdependent, data is, [*sic*] lacking which results in large uncertainties. Instead of using of the ERDA-7 approach, performance-based standards for the effectiveness of security systems as well as addressing consequences were recommended as useful tools [19, 24].[4]

A large body of work over the years has been done in response to considerations analogous to the above. Suppose that despite the above discussion, we have constructed a model of risks from attacks on a particular facility: we have an assessment of scenarios, frequencies (including attack likelihoods), and consequences, and wish to modify the facility to reduce risk. For present purposes, we oversimplify the classical decision rule: the idea is to choose the modification that maximizes expected utility, where the expectation is calculated conditional on all the uncertainties, emphatically including the ones that affect attack likelihood. If the analyst specifies all these inputs, the process will operate, and will recommend an alternative. However, if the uncertainties in attack likelihoods are very, very broad, it may well turn out that the recommendation is sensitive to small changes in those presumptions: if we revise the assumptions about likelihood of particular attacks only slightly, this may be enough to completely upset the original recommendations.

This is one of the class of reasons that drives some workers to argue that when uncertainties are sufficiently profound, the best approach overall is to choose a decision alternative that addresses a significant *range* of possible futures. One way to think about this is the "minimax" idea; one version of minimax is "choosing the option that ***mini***mizes the badness of the ***max***imally bad case."

This idea is similar in spirit to the RIMES idea of addressing attack scenarios that lie on the non-dominated frontier. Reducing the consequences of the non-dominated attacks, and doing this iteratively, precisely minimizes the consequences of the maximally bad case. Interestingly, RIMES' use of the difficulty axis helps limit the "maximally bad" case. In a sense, the worst possible release is "everything there is at the facility;" but if there is a level of difficulty beyond which an attack is effectively impossible, we don't need to analyze the case where "everything" is released.

[4] References in this quote refer to the bibliography in the RIMES paper [23].

2.6 Scenario Identification Methods

2.6.1 Challenges in Scenario Identification for Security

It was noted earlier that the triplet concept of risk (scenarios, frequencies / likelihoods, consequences) is fundamental to management of risks to safety. It is important to understand what can go wrong, how likely it is for that to happen, and how bad the consequences are if it does happen. For applications involving risks to public health and safety, risk analysis needs to furnish risk managers with as comprehensive an understanding of the risk triplets as possible.

Over the last half-century or more, enormous effort has been devoted to modeling (simulating) the evolution of hazardous scenarios and the consequences of those scenarios. In the early days of nuclear power, some work was done to support demonstration of design adequacy within a design-basis framework, but since the Reactor Safety Study (widely considered to be the first plant-scale probabilistic risk analysis), examination of a much broader spectrum of scenarios has become customary. As computers, software, and phenomenological models became better, good modeling became more practical; and as more such work was done, the benefits of doing it became clearer to sponsors, to regulators, to facility operators, and to the community of practice, promoting further improvement, etc. This evolution in modeling quality and modeling effectiveness arguably still has some distance to go, but there is every reason to expect progress to continue.

All that said, one of the key uncertainties that remains in comprehensive risk analysis is the type of uncertainty sometimes called "completeness" uncertainty: we are not sure whether we have inadvertently omitted scenarios that we would have modeled if we were aware of them. To reason appropriately about risk management measures, it is important to have modeled a reasonably complete scenario set. Otherwise, measures of aggregate risk will be understated, and the relative risk significance of various facility elements will be distorted, limiting the quality of risk management decisions that are based on the results.

Unfortunately, there is no tool for proving that a risk model is complete in detail. (At a sufficiently high level, risk models are as complete as the designer's understanding of her own design, but when one gets down to details of system failures, the modeling effort is an inquiry, and it is difficult to know that one has asked all the relevant questions or has answered them correctly.) Humans are involved in determining what needs to be modeled, and the rigor with which it needs to be modeled. For example, the Reactor Safety Study (WASH-1400), pioneering as it was, left out some fairly important things. Its goal was to characterize the risks of the operating fleet, and up to a point, it succeeded: today's estimates of plant risk due to internally initiated accidents are not profoundly different from WASH-1400's (partly because we are nowadays less conservative, so even as new risk contributors have been recognized and modeled, numbers have gone down). But WASH-1400 essentially omitted serious evaluation of risks due to external events, risks associated with shutdown operation, and scenarios involving main coolant pump seal failures. It also underrated the potential for very long station blackout scenarios (like Fukushima). In other areas, it was conservative by today's standards: for example, nowadays, we take more credit for human recovery than they took in those days. But for present purposes, it is the completeness uncertainty that concerns us.

Many methods exist for cuing human identification of scenarios. One of them, STPA [26], is discussed below as part of HAZCADS [27].

2.6.2 Developing Scenarios Within the HAZCADS Approach

A partial overview of HAZCADS [27, 28] is given in Figure 2-17. The process shown in the figure culminates in a scenario set including security concerns that can be used for risk management, but that is only the front end of HAZCADS; the rest of HAZCADS uses that scenario set within a framework called

Top Event Prevention Analysis, which helps the user decide how to manage risk as efficiently as possible. The present section discusses the front end; later we discuss Top Event Prevention Analysis [29-33], which is a general method for using logic model results of many different kinds.

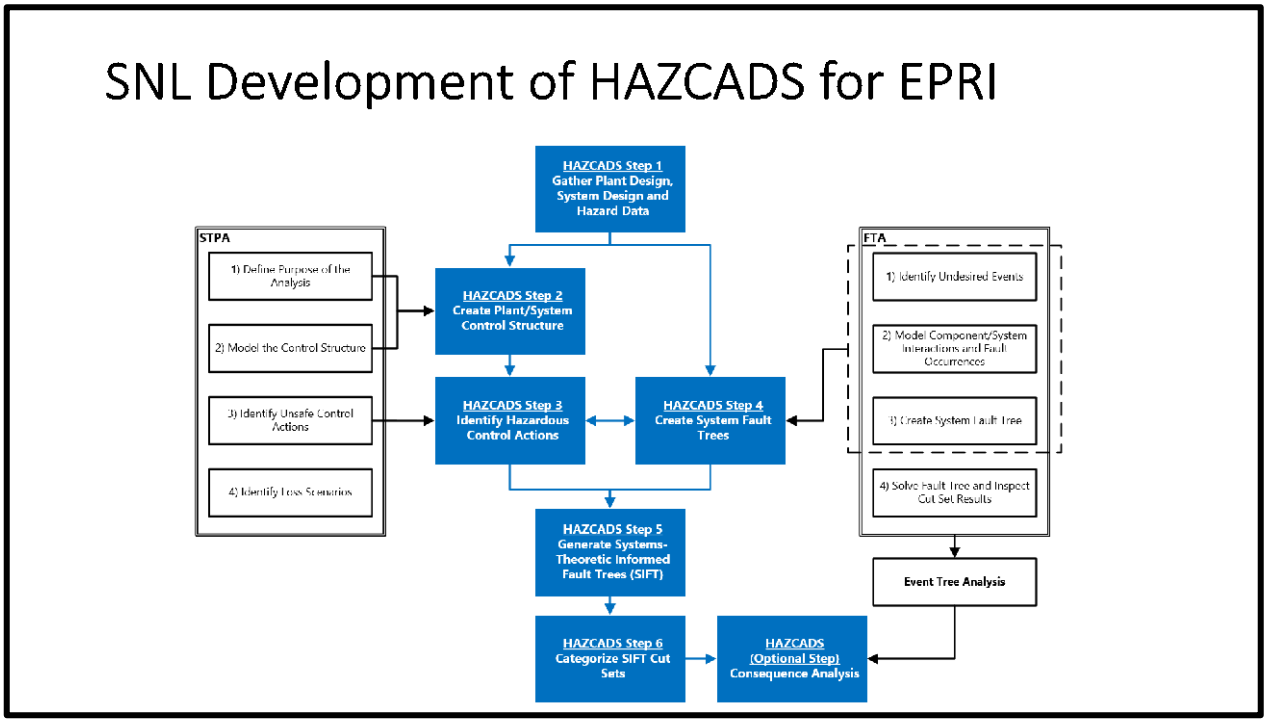

Figure 2-17 Overview of the Sandia National Laboratory (SNL) portion of HAZCADS.

In the figure, the portion on the right, Fault Tree Analysis (FTA) corresponds to part of what is done in ordinary risk analysis. It is a model of the conditions under which a selected "top event" occurs (typically functional failure at the facility level or system level), described in terms of "basic events" corresponding to sets of conditions specified at the component level. Once a particular top event (e.g., "system failure") is specified for analysis, a hierarchical logic model for that top event is developed, expressing the sets of conditions under which the top event occurs, first at a rather general level (e.g., Division 1 fails AND Division 2 fails), and then at increasing levels of detail (Valve 42 in Division 1 fails AND …). It is reasonable to expect a competent fault tree development to result in a generally complete specification of the families of sets of conditions under which the top event occurs, but a typical fault tree may not completely reflect all the possible *causes* of each basic event. In safety applications, little may be lost because of neglecting causes that are extremely unlikely to occur in ordinary circumstances, but the present interest is in understanding what scenarios an adversary might *choose* to cause, and the sorts of component behaviors that are "unlikely" if we consider only stochastic causes may now be exactly what the adversary wants, so their likelihood is limited only by the adversary's difficulty of achieving them. Identifying them before the fact is the point of STPA, the process shown on the left-hand side of the figure.

"STPA," or System-Theoretic Process Analysis [26], is a structured way of thinking about scenarios that culminates in a list of "unsafe control actions" (UCAs). From a certain point of view, it is precisely the objective of adversaries to cause UCAs that combine to produce the consequences that the adversaries are trying to cause. STPA can do a good job of identifying UCAs at a local level; the fault tree reflects the

global structure of the problem. The marriage of the two (local and global) in Step 5 of the figure culminates in the System-Theoretic Informed Fault Tree (SIFT) [*sic*].

This method recommends itself for the present application because its formulation stresses the idea that many (most?) accidents can usefully be viewed as losses of control, with adversarial action as a possible cause. STPA identifies "unsafe control actions," which could be due to component failures, but could also be due to design errors, adversarial actions, cases where design assumptions were violated, etc.

A key premise of the HAZCADS exercise presented in the Electric Power Research Institute (EPRI) report was that although the original fault tree model did not contain cyber causes specifically, its picture of the possible functional deviations was rich enough that it was easy to add basic events corresponding to adversarial action to the model. For example, if one unsafe control action was spurious opening of a particular valve at the wrong time, then ideally, a corresponding basic event corresponding to this happening stochastically would already be in the logic model, and it would be trivial to add another basic event corresponding to an adversarially-caused opening of the valve. If that deviation was not already reflected in the model, then it would need to be added, both as a "random" failure (modeled as occurring stochastically), and secondly as a "systematic event." In the HAZCADS work summarized here, the term "systematic" was applied to events that are not to be thought of as stochastic; this could include conditions such as software errors or adversarial action.

In the HAZCADS report, the Top Event Prevention Analysis generated numerous allocations ("Prevention Sets"), and tested one of them. In that exercise, in each allocation, all the non-systematic events were protected, along with a specific subset of the systematic events. Top Event Prevention Analysis only generates allocations that satisfy prevention criteria, so unprotected systematic events were those that appeared in cut sets only with protected events. Protection of systematic events was assumed to reduce their failure probabilities to a level comparable to (or less than) their random failure probabilities. The Prevention Analysis does not tell the user *how* to protect those events: it just gives the user options for which subset of those events to protect, based on how all the basic events are combined in the minimal cut sets.

2.7 The Concept of Allocation

Earlier, we had occasion to mention the LMP, an important industry document concerned with "a modern, TI-RIPB process for selection of LBEs; safety classification of SSCs and associated risk-informed special treatments; and determination of DID adequacy for non-LWRs." In general, at the design stage, there is considerable flexibility in how best to go about achieving safety objectives. Selecting LBEs is an important investment decision with far-reaching safety and cost implications. In making that selection, one is deciding what one will rely on for safety. SSC performance comes at a price, so the selection needs to be based in part on how efficiently (how cost-effectively) the SSCs that need to perform in those LBEs will get the job done. That is to say: selection of LBEs is (among other things) an *allocation*.

Most risk-oriented discussions of scenario modeling default to failure space: scenarios are expressed in terms of adverse conditions or events. A "minimal cut set" describes a set of basic events whose conjunction (joint occurrence) causes the top failure event in the model (the system / facility *fails*, [i.e., it has an accident]). But the complementary success-space perspective is also important: a minimal "path set" (a minimal "success path") describes a set of basic events whose conjunction corresponds to system / facility *success.* A system cannot succeed and fail at the same time, so the set of events / conditions contained in a minimal cut set must be comprehensive enough to defeat all the success paths, and the set of events / conditions in a minimal path set must be comprehensive enough to defeat all the minimal cut sets.

When we talk about increasing the difficulty of attack scenarios (e.g., in RIMES), we are, in effect, talking about increasing the prospects (the likelihood of success) of one or more success paths.

Preventing any one of a cut set's basic events prevents that particular cut set, and while it is difficult to achieve absolute prevention of an event short of designing it out, preventing several events in a cut set is likely to do a good job of preventing it. Here, by "preventing" a basic event, we mean "reduce event probability to a very low value." Preventing several events in *every* cut set is likely to do a very good job of preventing the top event. Since some cut sets may be purely adversarial, others purely random, and yet others a mixture of random and adversarial, we know that to prevent both ordinary system failures and system failures due to adversarial action, we will have to prevent at least some random events and at least some adversarial events. But, as discussed, we may find that we can identify a ***subset*** of the adversarial events whose prevention, in conjunction with prevention of the random events, suffices to drive top event probability down to a usefully low level. Identifying such a subset is not trivial, but this is a sample of what Top Event Prevention Analysis does, and this was illustrated in the HAZCADS report.

A point not stressed in the HAZCADS report is that Top Event Prevention Analysis was originally formulated to structure a comprehensive safety case [32, 34, 35]. Its application in the HAZCADS report was tailored to focus on prevention of systematic events with an existing plant as an example, but at the design stage of an advanced reactor, the comprehensive focus is exactly what is needed.

Finally, we note that Top Event Prevention Analysis has also been applied to set priorities in protection of vital areas from physical attack [35], in a way that would allow it to be used to revise protection priorities *dynamically* [during an attack], based on the current status of vital areas.

2.8 Uncertainty in Risk Management

When analysis is done to support important decisions (including decisions affecting public health and safety), it is imperative to understand the limitations of that analysis, to avoid making questionable decisions. This entails understanding "uncertainty." Two categories of uncertainty seem especially significant in the security arena, and are discussed in the following subsections: completeness uncertainty, which people may try to cope with by making ostensibly conservative assumptions regarding what can or may happen, and uncertainties in certain event probabilities.

2.8.1 Uncertainty in the Adequacy of Coverage of the Event Spectrum: Completeness

To quantify the risk of core damage resulting from outside attack, one could imagine proceeding analogously to modeling core damage risk from so-called "internal" events. One could develop a set of scenarios leading to core damage, based on careful identification of ways in which attackers could cause sufficient failures to cause core damage, and then one could quantify the frequencies of those attacks and the conditional probabilities of their successes, which depend on the capabilities of the attacking force and on the capabilities of the defending force. To proceed by analogy to modeling internal event scenarios, one would first think up the scenarios, and then postulate the characteristics and model the actions of both the attackers and the defenders, and one would need to quantify the frequency at which that attack occurs (more realistically, the probability that the attack will occur within a particular time frame).

Over the years, it has become clear that analysts are willing to model (simulate) pretty much any scenario, once it is specified. Given such a capability, and a facility design, and probabilities for the constituent events that play a role in each attack, one can in simulate many scenario realizations, and thereby learn something about the conditional probability of that attack being successful, what considerations (e.g., what capabilities in either the attacking force or the defending force) had the greatest influence on the outcome, and how to improve the outlook for the defenders.

Many organizations have invested substantially in developing capabilities to do this. There are many, many permutations to consider, but computers are getting better and better. However, at the end of the

day, one still needs to postulate a spectrum of attacks, and to quantify "risk" in the traditional sense, one needs the attack probability or a surrogate for this parameter.

For purposes of this report, certain uncertainties seem particularly significant: uncertainties associated with attack probability, and uncertainties associated with completeness of the scenario set. They are difficult to discuss in any detail with examples because much of what is known about them by the U.S. community of practice is not openly available. But there are fundamental reasons to believe that even if that information were publicly available, we would be devoting significant attention to these uncertainties: attack likelihood is a special kind of issue, and scenario completeness has historically been important in all applications of risk analysis.

Despite generations of work, significant controversy remains regarding assessments of probabilities of severe events. An interesting early work was WASH-740 [36], "Theoretical Possibilities and Consequences of Major Accidents in Large Nuclear Power Plants: A Study of Possible Consequences if Certain Assumed Accidents, Theoretically Possible but Highly Improbable, Were to Occur in Large Nuclear Power Plants." This report, which came out in 1957, paraphrased extensive discussions held by the report's authors with selected colleagues who felt that the probabilities of severe accidents were essentially unknowable. This was considered unsatisfactory by many people for many reasons, and fed into the eventual decision to undertake WASH-1400.

Since many analysts are now accustomed to discussing severe accident likelihoods, the supposed "unknowability" seems surprising, so it is worth recalling the broader framework within which that attitude was voiced. At that time, people were learning to think in terms of the "maximum credible accident." There was broad agreement within the reactor safety community that while a severe accident was not physically impossible, it was extremely unlikely; the people who thought the probability was "unknowable" arguably meant that one could not assign a meaningful number to the probability. But one still needed to be able to argue that plants were safe. Given all that, it was hoped that the needed safety argument would flow from the following:

- Concoct an accident that posed a severe challenge and was seemingly barely credible.
    - The challenge selected was "double-ended guillotine break of the largest pipe in the reactor coolant system, with or without a concurrent loss of offsite power, and assuming the limiting single active failure." System actuation setpoints were to be presumed to have drifted to their most nonconservative allowable values, outside temperatures were to be at a specific extreme (limiting the effectiveness of the ultimate heat sink), system actuation delays were supposed to be at their worst allowable, and so on.
- Show that a proposed design could deal with that challenge, severe as it was.
    - Simulation of plant behavior conditional on that challenge was to be done, and if the core could be shown to suffer only minor damage, the system would be deemed "good enough."
- Declare victory.

An implicit premise was that challenges can be rank-ordered in such a way that if a design could cope adequately with a particular challenge, then it could cope adequately with all less-severe challenges. More colloquially: If the design can cope with the maximum credible accident, then it can cope with anything that will ever really happen, so it's safe. This thought process used probability implicitly, but did not require a high-precision estimate; it required only a judgment-based concept of "credibility," and a way to rank scenarios according to severity.

Three Mile Island was initially much, much less severe than the "maximum credible accident," so it was a real shock to many in the community of practice when core melt occurred. As expressed by

someone who was on NRC staff at the time, "we didn't think core melt could occur in an intact RCS [Reactor Coolant System]." They were anchored to the idea that, provided that a reactor shutdown occurred when it was needed, a really large loss of coolant was necessary to cause core damage; so they focused on the reliability of reactor shutdown systems and coolant injection systems for mitigation of the large double-ended break, for purposes of design, operator training, and technical specifications. Unfortunately, it turns out that mission success criteria over the spectrum of plant initiating events are significantly more complicated than suggested by the maximum credible accident picture. Over time, more and more events were added to the list of analyses that needed to be done for licensing purposes, but it is doubtful that a list prescribed without the benefit of scenario perspective will reliably cover the relevant issue space. Part of the significance of the LMP is precisely that it calls for development of this scenario perspective before determination of design-basis and LBEs.

While not described extensively in this report, for a survey of cyber risk analysis techniques, see [37].

2.8.2 Quantification

It is nowadays customary to discuss two kinds of uncertainty related to event probabilities. One is "aleatory:" even if an event probability is fixed, we do not know when (or even whether) it will occur, so in different time histories, it will occur at different times. The other is "epistemic," related to our state of knowledge: we do not really know the probability.

Most PRA work relies on some knowledge of operating history to inform estimates of event probabilities. But even when knowledge of history exists, someone must interpret that history and argue its applicability to the situation being modeled. This is more difficult than it may sound; there is a strong temptation to resort to formulations such as "we have had zero attacks in N years, so the attack frequency we should use in our analysis has the following distribution: … ." Such a formulation assumes that whatever determinants of attack frequency existed in the past (or not) will also exist in the future. This category of error was discussed in [38] under the rubric of "exchangeability."

In fact, many event frequencies change over time, and this complicates the assessment of them. A huge influence on event frequencies is institutional learning, as (arguably) illustrated in Figure 2-18, based on [39, 40]. The plots on the left show the evolution of the frequencies of pressurized water reactor (PWR) general transients and losses of offsite power. The plots on the right show expanded views of the same quantities. Over the years, the frequencies of certain types of events go down, as plant staffs learn how to avoid those occurrences. Loss of offsite power is a risk-significant initiating event, so it is very significant that its frequency appears to have gone down appreciably. (When WASH-1400 was done, a typical LOOP frequency was 0.2/year.) But even here, the apparent frequency can fluctuate. The 2003 result showed a peak caused by an external event. Even if event arrival rates look stochastic for some purposes, they can be influenced by other events (like "widespread grid disturbances"). Analogously, history-based assessments of flood probabilities have seemed especially fraught in recent years.

In short, there are countervailing influences operating on many event frequencies. Learning operates to reduce some adverse event frequencies [41], while external influences act in ways that can either increase them or decrease them. If event frequencies are sufficiently low, the learning process cannot operate efficiently, and if the events are infrequent but conditionally risk-significant, we need some idea of their frequencies, typically a better idea than we can get from history alone. Regarding events in that category, the good news is that the frequencies are low; the bad news is that because the frequencies are low, our only source of information about them is assessments done by humans.

Although the scientific people interviewed for WASH-740 would probably object to application of expert elicitation to establish probabilities for use in risk analysis, expert elicitation has by now a long history of application in nuclear risk management and elsewhere. In important decision situations, someone must formulate good options [42] and act on at least one of them; "no decision is still a decision." So, sometimes, expert elicitation is undertaken. Details of the merits of different approaches

are beyond the scope of the present treatment; for a snapshot of the picture 40 years ago, see [43], and for two arguably leading current approaches, see [44, 45].

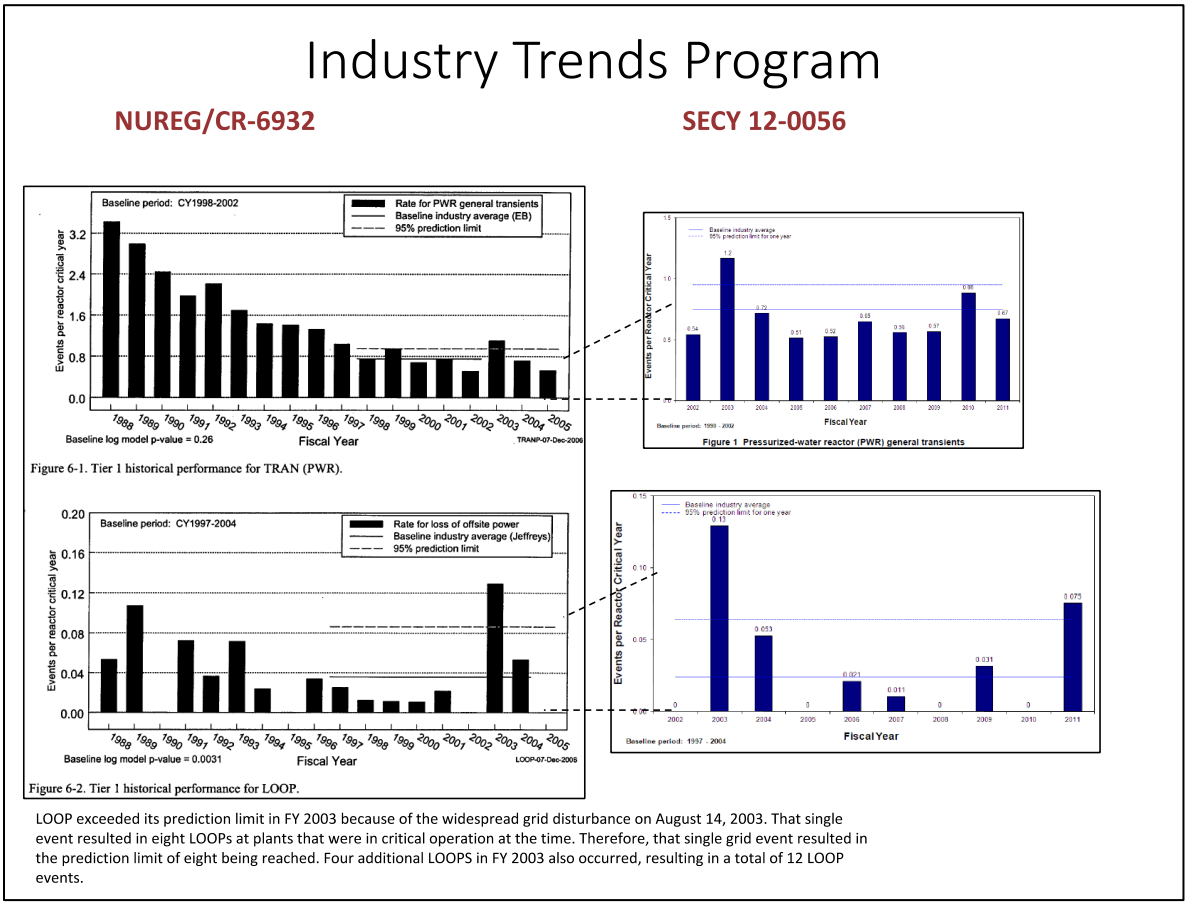

Figure 2-18 Changing event frequencies.

Suppose we have done everything we can to get a reasonably converged estimate of an important quantity, but our state of knowledge is still lacking: the distribution is still too broad to support a decision that is robust (insensitive to changes in model inputs). One response to that situation is to write down very broad probability distributions for the highly uncertain variables, and average over them, and do what the analysis tells us to do; decision analysis will tell us what's best, based on the inputs that we give it. But we need to beware of how robust that answer is: how likely we are to regret the decision we make. A large and active research community has been working for some time to determine demonstrably better ways to approach such decision situations, but so far, consensus has not been reached. For a snapshot of the diversity of opinions in this area, see [46].

For some years now, workers at various institutions have stressed the importance of reflecting safeguards considerations as early in the facility life cycle as possible. This is a familiar idea to anyone in safety: it is much easier to design certain attributes into a paper system than it is to retrofit those attributes into a system that has been built and deployed. Appropriately, [47] cites a Systems Engineering Handbook. There are many Systems Engineering paradigms out there, but they generally stress early focus on stakeholder requirements, and from that point of view, we can say that the idea of Safeguards by

Design is essentially the idea that safeguards priorities should be injected into the Systems Engineering process as early as possible (along with, for example, safety priorities).

2.9 Uncertainty and Biases

Walker and colleagues [48] broadly defined uncertainty as departures from the unattainable goal of a full deterministic understanding of a system. Such a goal is unattainable because uncertainty cannot be eliminated. In the scientific and engineering context, models incorporate both aleatory and epistemic probabilities [49]. If the researcher wants to know the probability of some outcome (e.g., number of defective products out of the total number produced, success versus failure, etc.), aleatory models represent the randomness in the outcome of the process. For example, if the researcher wants to know the probability of observing two heads out of three flips of a coin, the researcher will need to estimate the probability by using a binomial model. Aleatory uncertainty refers to the outcome of a stochastic event being described as a probability. In contrast, epistemic uncertainty reflects the accuracy of our knowledge about the model. The accuracy can be influenced by the uncertainty within the model parameters, complications concerning the model such as scope and completeness, and measurement errors. Even in a simple example of flipping a coin, uncertainty cannot be eliminated without fundamentally changing the nature of the act (e.g., no longer flipping the coin but simply choosing a side to display). Despite the ubiquitous presence of uncertainty, stakeholders must reach conclusions. Models and simulations help reach these conclusions.

Often, people work in teams or groups toward a common goal, which requires effective communication. These teams can be organized into a larger network of teams with people from diverse backgrounds and experience (e.g., operators and human factors experts from a utility, psychologists and engineers at a national lab, and regulators from the NRC working together on a project). This type of heterogenous network must still operate optimally especially when high risk consequences are potential outcomes. We will review a few common biases before discussing how to mitigate certain biases. Our review of biases and how to mitigate them is not exhaustive. Rather, it is an open discussion on how these biases affect our perception of the situation, our decisions based on those perceptions, including what and how we choose to communicate.

2.9.1 Bias and Noise

Kahneman and colleagues [50] argued that understanding error in judgment requires understanding both bias and noise because both contribute to error. To illustrate the difference between bias and noise, they used the example of four teams shooting at a target. Each member of the team takes one shot at the bullseye. All of Team A's shots are clustered tightly around the bullseye. There is little deviation from the bullseye. This pattern approximates the perfect situation in which everyone hits the bullseye. Team B's shots are tightly clustered but are away from the bullseye. This pattern suggests some bias is present. The consistency of this deviation from the bullseye would permit us to make a prediction. If another shot were to be taken, we would predict it would be clustered with the rest of the shots from Team B. Team C's shots are widely scattered. This pattern is evidence of noise. There is a great deal of variability in this team's shooting ability, and we could not make a prediction for the next shot. Team D's performance is not as unpredictable as Team C; however, they are not as reliably predictable as Team A or Team B. Some of the shots are clustered but quite a few are scattered. Team D's performance is both noisy and biased. Consequently, if bias is shared amongst individuals, then it is a systematic deviation from a true value or standard. In contrast, noise represents individual differences which can have different sources including bias.

Regarding decision-making, we would hope that our decisions would be the correct decisions (i.e., Team A performance, unbiased and reliable). However, our decisions can involve both bias and noise. Kahneman and colleagues [50] argue that if we want to improve decision-making than we must address

both contributors to error (i.e., bias and noise). Quite often, we can rely on intuitive, fast thinking or System 1 thinking to reach quick conclusions. This type of thinking can lead to biases. For example, if insurance companies collectively charge a higher premium for cars that will be parked on the street, one could say insurance companies are biased against people who park on the street. In this sense, the systematic bias against people who park on the street represent a pattern of thinking shared by the insurance company and not individual differences between companies. If this bias supports the insurance companies' bottom line, then, from the perspective of the insurance companies, they have reached their goal. However, their thinking is still biased and can be problematic. Biased thinking can lead to noise. It is important to keep in mind that bias and noise are separate but related concepts.

How can bias lead to noise? Bias can produce noise when decision-makers differ in the degree of how biased they are or the nature of their biases [50]. For example, two companies are building new nuclear power plants. Company A states that they will be able to complete the plant in 6 years, and company B states that they will be able to complete constructing their plant in 5 years. If both companies take 8 years to complete their plants, the companies have fallen victim to the planning fallacy. Individuals provide lower estimates than the actual time to complete the project. Thus, the planning fallacy is an example of a judgment bias because the estimate (e.g., 5 or 6 years) is compared to a true value (e.g., 8 years). However, there also is noise because Company B is more biased than Company A. Company B's estimate deviates from the true value by 3 years while Company A's estimate deviates by 2 years. Both projects will be over budget, but Company B will be over by a lot more. Thus, there is noise in the planning estimate because the companies differ in the degree to which they fall victim to this fallacy.

More often, people are faced with the situation in which they cannot be certain how much bias is present. However, we can compare situations, which should lead to identical conclusions, to see if bias and noise are present even without knowing what the true value is. For example, if two people (i.e., a poor and rich defendant) are convicted of the same crime and circumstances surrounding the crime are identical, then we would expect both defendants to be sentenced equally. Often, we do not expect this result because we are aware of bias and noise. If a bias against poor people is present, the criminal justice system will systematically give tougher sentences to poor defendants. If noise is present either random noise or noise produced by other factors, then we would expect different sentences based on who the judge is or the makeup of the jury.

Kahneman and colleagues caution that in diagnosing bias people do not simply decide that a negative conclusion is due to bias. Vagueness will only lead to vague plans to improve the situation, which are generally not helpful. They recommend that the word bias be used when individuals understand what contributed to their faulty thinking. In the substitution bias, the person is making a judgment based on the answer to an easier question. Asking oneself if someone is like a description is easier than asking whether someone is more likely to fit a description. The judgment of similarity is made in place of the judgment of probability.

People frequently disregard base-rate information [50]. For example, if you ask someone to make a judgment about the safety of nuclear power plants, salient examples such as Chernobyl will probably influence their judgments. They would probably not consider what is the base rate of accidents for nuclear. They would probably not consider the actual mortality statistics for nuclear. If you asked the same person to make a similar safety judgment about driving cars, then they would most likely consider their own driving experience rather than base-rate information for car accidents. Both situations involve relying on the most accessible information and what the person considers the most representative information. Based on this intuitive thinking style, the person would most likely believe that nuclear power is dangerous while driving is not dangerous at all. Both conclusions would not be accurate.

The nuclear example can be used to illustrate conclusion and excessive coherence biases. Conclusion bias or prejudgment involves starting with a desired conclusion. The judgment process then works toward that desired conclusion. Here, System 1 thinking suggests what is the desired conclusion and the person

skips the information collecting phase altogether or uses the slower System 2 thinking to justify the desired conclusion. For instance, if our gut reaction to nuclear power is that it is dangerous, then we can simply refuse to consider evidence and stick with our prejudgment. Or we can marshal our deliberative thought processes to justify this conclusion (i.e., only consider evidence that supports the prejudgment), and, thus, fall victim to the confirmation bias. When we are vulnerable to confirmation bias, we demonstrate how we can form quick impressions and refuse to alter our impressions even when presented with disconfirming evidence.

2.9.2 Confirmation Bias and Uncertainty Paradox

As noted above, confirmation bias can serve to filter out disconfirming evidence in favor of a desirable conclusion or outcome [50, 51]. This bias can present even when the decision-maker is approaching the endeavor honestly and constraints are in place to help avoid this bias (e.g., peer review). Digdon [51] notes that confirmation bias has the potential to be problematic to many disciplines. However, the appropriate safeguards against confirmation bias depends on the field of study. In psychology, the defense against confirmation bias is the rigorous peer review process.

As Digdon [51] noted that the level of uncertainty in the argument was glossed over unintentionally. In general, accurate reporting of the uncertainty in research is critical because uncertainty is inherent in research. Digdon does not specifically name the uncertainty paradox. However, Digdon's analysis is consistent with the perils of this pattern of thinking.

Patterns of problematic thinking (e.g., confirmation bias, substitution bias, uncertainty paradox) can lead to bad decisions. Thus, how do stakeholders avoid making bad decisions when many biases can be shared among many people? Reducing bias will take effort but also knowledge and practice [52].

2.9.3 Reducing Bias and Noise

Lee and colleagues [52] demonstrated that the anchoring bias and the representativeness bias can be reduced if participants are educated about the biases and allowed to practice mitigating the biases in a digital game. The anchoring bias is when people use the first piece of information as an anchor to make their estimates (e.g., How much would you pay for this bottle of wine?). If a number—even if it is unrelated—is presented at the time of judgment, that number becomes an anchor that influences the estimate. The higher the number, the higher the estimate. The representativeness bias was mentioned earlier. If someone asks how safe is driving a car, most people do not consider base-rate information. They might focus on what they consider is the most representative information to make their probability judgments (i.e., "I've never had an accident…").

The implicitness of the cognitive bias makes them particularly insidious in eradicating them from our thinking . The combined condition (i.e., lecture plus digital game) was the best at mitigating the cognitive biases. The short lecture provided information on the cognitive biases. Lee and colleagues argued the digital game helped learners transfer the knowledge by practicing what they learned.

Kahneman and colleagues [50] proposed that groups use a decision observer to help reduce bias in a team setting. The job of the decision observer is to identify biases based on the following checklist.

- Approach to judgment
    - At any point did the discussion of the evidence involve using the answer to an easier question for the harder question (substitution bias)?
    - Was any piece of relevant information ignored?
    - Was any piece of information inappropriately weighted?
    - Did the group attempt to think of the issue from an outside perspective?
    - Were comparisons made? How were they made?

    - Could group members share biases?
    - Was there an expertise or viewpoint missing from the group?
- Prejudgments and premature closure
    - Do any of the decision-makers benefit from a particular conclusion?
    - Did any of the group members come in with a preconceived notion about what the right decision is?
    - Were dissenters able to express their views?
    - Is there a possibility that group members were unwilling to change course on a decision because resources were already allocated in that direction?
    - Is it possible that the chosen course of action was due to it being discussed earlier in the discussion than later?
    - Were all alternatives discussed? Was all relevant information discussed?
    - Were group members able to discuss anything uncomfortable or unpopular opinions?
- Information Processing
    - Are group members being influenced by the recency, dramatic nature, or personal history of a piece of evidence regardless of its diagnostic value?
    - Is the evidence anecdotal or empirical?
    - Is there an anchoring bias?
    - Did group members ignore that regression to the mean could be taking place?
- Decision
    - Did anyone fall victim to the planning fallacy?
        - Were sources questioned about their validity?
        - Were alternative views considered?
    - How aligned are the stakeholders in how much risk they are willing to take on?
    - Do the decisions reflect the organizations short-term and long-term priorities?

Kahneman and colleagues recommended that decision hygiene was the best way to reduce noise. Six principles that describe decision hygiene are the following:

- Judgment should be focused on accuracy and not expressing one's individuality.
- The judge should view the case as another member of a reference class. The judge should not view the case as a unique case within a narrative. This recommendation is to avoid introducing noise. Also, this statistical perspective helps judges moderate their predictions.
- Judgments should be individual tasks to avoid excessive coherence. For example, when nurses assign an Apgar score to a newborn, the nurses judge the infant on several criteria, which add up to the overall Apgar score. Nurses do NOT just assign an overall score without making the individual judgments.
- Avoid making prejudgments by restricting biased information from decision-makers.

- Use multiple judges and aggregate their judgments.
- Use relative judgments and relative scales rather than absolute judgments.

Decision hygiene should begin with a noise audit as the initial step [50]. The noise audit team can consist of only internal, only external, or a combination of internal and external team members. The consultants will collect and analyze data and provide the final report. The project team also will need subject matter experts who will compile relevant cases for the judges. A high-level manager should be involved for practical reasons (e.g., the client is assigning this project importance, etc.). The client must be open-minded in discussing the results and what changes can be implemented. Judges are professionals similar to the employees who work for the company in the unit being audited. Each unit being audited will have several judges. Prior to administration of the study, the project team should meet with the client to reach a standard for unacceptable noise and provide an opportunity for the client to raise objections about the study plan. Again, the consultants will collect and analyze the data. And, if the materials support identification of biases, consultants should also focus on that.

2.9.4 Communicating Uncertainty

Researchers and decision-makers should make every effort to avoid bias by identifying if any biases are present and mitigating the bias through effort and practice [52]. Groups should follow the Kahneman and colleagues' recommendations. However, bias and noise cannot be eliminated. Thus, researchers must accurately convey this information to decision-makers. Thus, the first step in communicating uncertainty is not too gloss over the uncertainty present in a situation.

An obvious but helpful perspective is to remember that communication is interactive. Consequently, any recommendations should take that dynamic quality into account. Gesser-Edelsburg and Shir-Raz [53] noted that a key issue is when multi-disciplinary stakeholders are discussing the topic and possess differing risk perceptions. If the committee has not reached consensus on risk severity, how can the risk be communicated and discussed accurately with others? Following Fischhoff's [54] recommendations in easing the fears and reluctance of both researchers and decision-makers and acknowledging the dynamic quality of communication can address this issue.

Fischhoff [54] argued that expert knowledge will only confer practical significance if decision-makers are able to understand the quality of that information. Part of assessing the quality of information is understanding the amount of uncertainty present. Thus, in the interaction between researchers and decision-makers, researchers have a duty to provide candid disclosures of how certain they are about a particular piece of evidence. And decision-makers have a duty to provide candid disclosures about what information they need, whether they understand the information, and what they intend to do with the information.

Fischhoff [54] proposes that identifying sources of expert reluctance in conveying uncertainty will help promote better communication of uncertainty because we can address root issues that are impeding communication. Fischhoff argues that sometimes experts will withhold uncertainty information because they believe the uncertainty will suggest a certain level of precision that might not be warranted. Fischhoff urges that experts still share the uncertainty information to prevent decision-makers from guessing. If the decision-maker's estimate is incorrect or they do not understand what is causing the uncertainty, their decision or plan of action could be incorrect. Other times, the expert is not reluctant to share uncertainty information. They assume the decision-makers already know this information. Fischhoff recommends that experts convey uncertainty information without assuming the audience knows what the expert knows. Thus, a trial run with a practice target audience can help identify unclear parts of the message or what the target audience might not know. The practice audience should be similar in makeup to the target audience. The communication should be tailored to the audience and not a general audience. A personal briefing will support the opportunity to ask clarification questions.

Another challenge to communicating uncertainty is if the expert believes that the audience will not understand the information [54]. Experts can overestimate the amount of information needed to for the certainty information to be understood by the audience. Depending on the level of knowledge in the audience, experts can provide enough information that the gist or key points of the message are understood. This type of tailoring the message does not mean the expert should be vague. The message should still be clear. If using a few concrete examples help, then use examples.

Another source of reluctance in communicating uncertainty is the fear of reprisals for conveying uncertainty information [534]. Experts should be supported in reporting uncertainty. Even if the information is not what decision-makers want to hear, accurate reporting supports understanding the risks that are being taken.

The duty of the decision-maker is to help the expert understand what they need from them and that uncertainty information is valued [54]. Decision-makers should not assume that experts understand the reasoning behind the decision-maker's actions or decisions. They should provide direction to the experts in terms of the decision-maker's goals, options, and beliefs. The decision-maker could have unrealistic expectations or goals. However, the expert cannot help the decision-maker if they do not discuss their reasoning. Experts should have decision-makers frame their decisions in their own words. Then, experts should ask them to elaborate or listen. If an issue was raised or an issue that was expected to be raised and was not, experts should probe decision-makers on these instances. Fischhoff argues that making the decision explicit will help clarify how uncertainty might be affecting the decision. The decision-makers can assure the experts that they value reporting of uncertainty and that they will support the experts when they do.

After the Columbia disaster, NASA [55] produced a set of best practice recommendations for modeling and simulations. They wished to reduce the level of risk in decisions based on models and simulations by increasing the transparency in the decision-making process. In the report, NASA describes the information that should be conveyed but not how the message should be communicated. A general recommendation is stated that stakeholders seek out discipline-specific guidelines in communicating uncertainty. Thus, we will first describe some recommendations from the NRC followed by a nuclear-specific example.

2.9.5 Risk Communication within the NRC

The NRC [56, 57] notes that engineering risk models involve expertise in multiple systems, different operating environments, and scientific phenomena. NRC risk analysts must be able to interact with a range of individuals to acquire the needed information. The following are the recommendations for working with technical staff of different backgrounds:

- The risk analyst should be able to explain why they need the audience's expertise and how it will influence the decision.
- The risk analyst should use plain language (i.e., avoid jargon).
- The risk analyst should understand how different expertise might influence the communication.
- The risk analyst should include stakeholders early in the process.
- The risk analyst should be aware that there will be a differing level of expertise when it comes to statistics.
- The risk analyst should be prepared to answer any questions on the NRC's risk-informed, performance-based regulatory approach. In addition, they should be ready to answer any questions about specific risk numbers or assumptions.
- Diversity in views should be encouraged.

In terms of communicating with nontechnical staff, the NRC recommends a more general approach. They recommend that the risk analyst provide an overview of the project. And the risk analyst should use qualitative terms and explain how risk is used in the decision-making process. Again, plain language should be used rather than jargon. Specific examples should be used to explain parameters. Pay attention to nonverbal cues to assess if your audience is understanding the message. The risk analyst should provide time for the nontechnical audience to ask questions.

2.9.6 Nuclear Industry-Specific Example

In recommending best practices in communicating uncertainty in modeling, simulations, and research in the nuclear industry, we need to know what message we are trying to communicate, why we are trying to communicate the message, and to whom or between which parties. For this report, the parties concerned are the utilities, researchers, and the regulatory bodies. Fischhoff and Davis [58] argue that conveying scientific uncertainty is a challenging balancing act of finding the appropriate level of detail and identification of relevant uncertainties. The message must integrate the decision-relevant elements. To this end, researchers must seek the relevant uncertainties to be included in the model or simulation.

Fischhoff and Davis propose that relevance is determined by the decision that needs to be reached. They argue that a decision might fall into one of three categories:

- Deciding if it is time to take action.
- Deciding which option is best.
- Deciding what is possible.

Fischhoff and Davis argue that communicating the amount of uncertainty in these decisions require specific actions. First, the uncertainty must be characterized by researchers identifying the relevant issues affecting the decision. Second, evaluating the uncertainty by distilling it into useful terms. And conveying uncertainty by providing an appropriate level of detail for the decisions being made.

Threshold decisions—decisions of whether to act or not—involve settling on a value to set the threshold [58]. For example, at which thresholds should we set fire warnings? Or, in the nuclear industry, a challenging and important problem is identifying which sections and components of the piping system to inspect during scheduled maintenance. The system is extensive with many bends and turns, which is part of what makes this so challenging. Potential components could have degraded in the harsh environment of the nuclear power plant. Safety is paramount in the nuclear industry. However, any downtime—scheduled or unscheduled—is costly for the utility. Thus, a maintenance plan that minimizes downtime is economically beneficial to the company. The question of whether a company should schedule a maintenance downtime to inspect certain sections of the piping system is an example of a threshold decision. Should the action be taken? But, to answer the question, the company must have some value or threshold that is met that would instigate the action.

Gribok and colleagues [59] have been engaged in an effort to develop fiber sensor technology that will increase the efficiency of detecting degraded pipes. Their prototype has the potential to reduce the cost of implementing their fiber sensor technology, which, itself has the potential to reduce costs for utilities. Their work partially relies on modeling and simulation and, like all research, uncertainty is inherent in their work. The conversation between the regulatory bodies and the utility for the deployment of new technology such as fiber sensor technology will inevitably involve conversations on uncertainty, especially, in a risk-adverse culture as nuclear. How can researchers best facilitate that conversation?

We can imagine the above scenario in which a utility wants to use the new fiber sensor technology. However, they may need the approval of the NRC. In the discussion over the new technology, the experts who developed the technology (i.e., Gribok and team) should be assured that the utility and the NRC values full disclosure of any uncertainty within the model and simulations. Moreover, the utility should make clear what their goals and expectations are. According to the NRC [56, 57], in discussing their

technology, Gribok and colleagues should address their audience using plain language (i.e., explain any technical jargon). They should start with the overall picture (e.g., What problem is the fiber sensor technology addressing?, etc.). They can also direct people to reports that are available. Keep presentations brief and use succinct phrasing. Clean graphs that are not filled with irrelevant information will help people understand statistical information. Use concrete examples to make the numbers more understandable. The NRC recommends addressing the following questions concerning uncertainty:

- What are the weaknesses in your data set?
- What assumptions did the researchers make during the study?
- How sensitive are the estimates to assumption violations or changes?
- How will the decision be affected if the estimates change?
- What did the research team do to reduce uncertainty?

Bias and noise are important factors to consider when stakeholders want to reach effective decisions. Kahneman and colleagues' framework for reducing bias and noise can be very helpful for groups in navigating this complex roadblock. [50] As mentioned earlier all bias and noise cannot be eliminated. Experts must convey this information to the rest of the team, and, in turn, the rest of the team must value the candor with which the uncertainty information was shared.

3 QUANTITATIVE MODELING APPROACHES

3.1 Background

The Idaho National Laboratory (INL) is beginning work on a suggested method for the treatment of uncertainties within the risk-informed licensing approach for advanced reactors [62]. The problem statement laid out states that advanced reactor designers are required to consider uncertainty, but "the approach of how to do this in a real way is not well understood nor are many of the existing tools and methods set up to facilitate a technically-defensible treatment of uncertainties found when evaluating potential scenarios and the determination of associated consequences."

The outcome of this research is a generic methodology that could be adopted by interested advanced reactor vendors to guide their efforts in design and maintenance of the power plant licensing basis.

First the meaning and scope of defining the 'safety case' is discussed in Section 3.2. The solution approach for this project is introduced in Section 3.3. The NRC position in [60] endorses the position in draft regulatory guide DG-1353 [61] that the licensing basis development of advanced reactor designs should follow the methodology in NEI-18-04 [62]. To achieve the desired goals of this project, we will construct the methodology around the concept of a probabilistic digital twin. Section 3.3 of this report introduces the concept of the probabilistic digital twin, and Section 3.4 describes FPoli RISE technology developed for handling risk-informed solutions. Section 3.5 provides a review of available guidance material to support the development of uncertainty methodology for PRA and DBA analyses.

Section 3.6 presents an example application of the methods discussed in the earlier sections using a simplified model of a PWR in a station blackout (SBO) scenario. Section 3.7 describes an example application of dynamic PRA to situations which are less amenable to a traditional PRA approach. Finally, Section 3.8 describes an example application of the graded approach for a hypothetical advanced reactor.

3.2 The Safety Case

The safety case for nuclear systems is a documented expression of safety, demonstrating the protective measures against uncontrolled radiological releases [68]. Safety equates to design resilience which must be coordinated with function and performance to form the foundation of a viable product. Given the real and perceived consequences of failure relating to providing power from nuclear energy, design resilience or safety takes precedence over function and performance. As such, the introduction of safety characteristics early in the design process, especially for new designs, is essential to reduce uncertainties associated with safety throughout the product cycle.

The safety argument or claims can be qualitative or quantitative. Quantitative arguments can be deterministic (conservative) or probabilistic. The designer determines the most appropriate 'packaging' of such evidence with the goal of providing a logical, traceable, and scrutable formal construct to regulators and ultimately to the public. Since the dawn of the nuclear industry, regulatory frameworks have been constructed to facilitate this complex process. Among these, the standard review plan (SRP), NUREG-0800 [69], was built on the experience of operating Light Water Reactors (LWRs), the predominant technology of the operating fleet in the U.S. The SRP is not a regulation by itself but it outlines acceptable elements of a review and compliance is generally expected. The SRP suggests a specific list of initiating events to consider and classifies those in two broad categories based on anticipated frequency of occurrence, as either anticipated operational occurrence (AOO) and postulated accidents, or Design Basis Accidents (DBA). Historically, the SRP and most NRC regulations and guidance are largely deterministic.

In general, the regulatory requirements and acceptance criteria were designed to ensure there are no undue risks to public safety due to the operation of nuclear facilities, The concept of "risk" to the public,

as described by the NRC, can be characterized by asking, “What can go wrong?”, “How likely is it?” and “What are the consequences?”. These questions can be answered for each hypothetical scenario. The deterministic approach is well suited to respond to most of these questions but does not address the “How likely is it?” question. Therefore, traditionally, PRA analyses are conducted on the backend of the process to ‘verify’ the safety case of a design.

In recent years, rulemaking and regulatory activities have become more “risk-informed” and “performance-based”. This trend is justified as a mean to strengthen the regulations, ensure that resources are properly allocated, and create a more technology-neutral regulatory environment. The LMP, the roadmap formalized in NEI 18-04, and ultimately the proposed regulatory framework in the 10 CFR Part 53, is consistent with these trends. This approach asks for a maximal role of PRA early in the design process.

A risk-informed approach is currently considered by the developers of advanced reactors. However, the industry recognizes the challenges in reliance on a maximal PRA role in their design and licensing efforts. Rather a degree of flexibility is deemed necessary, and the conversation is around a more pragmatic graded approach.

3.3 Uncertainty Management in Risk-Informed Applications

3.3.1 Review of NEI 18-04

The NEI guidance is the main product of all the activities around the LMP [62] which involved several industry stakeholders for the development and deployment of Advanced Reactor Technologies. The guidance is around three themes:

1) The selection of LBEs (NEI 18-04, Section 3)
2) The safety classification and performance criteria for SSCs (NEI 18-04 Section 4)
3) The evaluation of DiD adequacy (NEI 18-04 Section 5)

NEI 18-04 was endorsed in the Regulatory Guide 1.233 [71] in June 2020 which forms the basis on the 10 CFR Part 53 rulemaking [72].

It is important to familiarize with the glossary and terminology of the guide which introduce sometimes new concepts and definitions not used before in the PRA community. Therefore, the review of the literature was quite cumbersome and revealed several inconsistencies or at least ambiguities. The resolution of such ambiguities was a necessary initial tedious step in this project. The motivation was the creation of a clearer ontology that could better serve the development of a suitable methodology for handling uncertainties in risk-informed applications.

This section highlights the resolution of some of these issues. Particular attention is given here at the discussion of uncertainties which is a primary objective of this project.

3.3.2 Selection of Licensing Basis

3.3.2.1 Event Sequence Families (ESF) Grouping

The NEI 18-04 guide proposes to organize the events starting from identifying a list of initiating events (IE) and event sequences (ESs) that follow such initiating events. It is important to note that IE and ES has special meaning in NEI 18-04 and there is no universal acceptance in the industry of such terminology.

One concept expressed in the guide is the grouping of ES in ESF. The ESFs are then classified into an LBE category (Event Type) depending on the frequency of occurrence. The listed event types are: AOOs, DBEs, BDBEs, and DBAs. The reader can refer to the guide for the definition.

The concept of grouping is one of those ambiguities mentioned before. While different approaches or possible views on the grouping are cited in different locations in the guide, the reader is left with the impression that such grouping hinges on "similarity" among the ES, an attribute that appears highly subjective.

The power reactor innovative small module (PRISM) LMP report [73] grouped the ES's as unique combination of the IE, plant response, and end state. This resulted in 591 ESFs.

The eVinci LMP report [74] describes the grouping being performed by varying combinations of assumed release paths and activity releases were included in the evaluation such that the radiological consequences associated with the range of IEs and combinations of available mitigation could be determined. They appear to have then determined a set of representative dose assumptions, which resulted in 15 different possible consequence simulations, then categorized the ESs as fitting into those dose cases. They also described the common threads between the sequences.

The Kairos Power LMP [75] report cited the ASME/ANS Non-LWR PRA standard [76] and ESFs are defined in the risk assessment to group event sequences having similar plant response characteristics. From the LMP report is possible to infer the rational. ES within an ESF are similar in that there is some "reactor shutdown" and some "decay heat removal" that works, but also they appear to be limiting the grouping to DBEs and do not include AOOs. This is done similarly for the BDBEs. Since the consequence estimates are done with a method that just sets a target based on frequency, there is no real detail on the implications of grouping on the simulations, but we suspect they would just model the ESF with the most limiting transient. Overall, this appears a reasonable approach, but there is the possibility that by omitting the AOO sequences from the group there would be a tendency of biasing the frequency if the 'representative' event in the ESF toward a lower value.

The Sandia report [77] states that there is no requirement as part of the LMP for a reactor designer to develop a PRA from design inception. There are alternate ways to develop the safety design philosophy at very initial stages of a design to ensure that the design will initiate from a robust basis for risk management. This statement is in line with a graded approach that will be discussed. The Sandia report stresses that a fundamental challenge for micro-reactors is developing this event classification in a manner that accounts for significant uncertainties in initiating event frequencies. The analysis in the Sandia report does not provide specific event trees were made, and instead, they provide high-level event sequences. In that context, ESFs are created using vague terms of functionality rather than explicitly listing event sequences that qualify within those ESFs.

In conclusion, regarding the ESF grouping, there is room for interpretation which make application of the process flexible but cumbersome as details and solutions needed to be developed by the applicants.

3.3.2.2 RISE Methodology Decision

For the RISE methodology a conscious decision was made to enable this degree of flexibility while stressing a degree of consistency in its implementation. Specifically, the grouping is decided by the user according to engineering judgment. Individual ES are first identified and described. The frequency of the end state of an ES is either input by the user (with its uncertainty if available) or computed through PRA analysis (tools like SAPHIRE). Then, the user groups ESs in an ESF with user-defined criteria and assign a 'representative' ES within the ESF. The representative ES will carry its own frequency and distribution (uncertainty) which will represent the ESF. The uncertainty considered should also consider the PRA completeness uncertainty is probably large when the design is in preliminary stage.

More sophisticated and mathematically based grouping could be conceived in the future, especially if aided by automation. For example, the analysis may begin with very limited PRA, especially in the early stage of reactor system design. This will lead to crude estimate of ESF frequency. As SSC reliability and other PRA analysis inputs become available, the determination of such frequency and distribution can be updated, almost following a Bayesian approach that consider the ES possibilities within an ESF. This

approach will de-facto reduce the completeness uncertainty over time, or over the iteration cycles of the design.

The guidance is unclear to which extent every ES needs to be analyzed for radiological release consequences. It is our interpretation that only the representative ES would be used to determine the consequences. The user may elect to bound the frequency of the representative ES that characterize the ESF such that a conservative estimate for the analysis of the cumulative QHOs can be provided.

3.3.2.3 ESF Classification Using Mean or Values with Uncertainty

The ESF are classified in AOO, DBE or BDBE depending on the frequency of the event, but it must consider the event frequency comes with its uncertainty. The question is: which frequency? The mean? The minimum? The 5th percentile? The 95th percentile? The maximum?

Again, the NEI 18-04, provides inconclusive guidance on this subject despite suggesting that the uncertainty must be accounted for. The PRISM LMP [73] suggests the use of the mean. The eVinci LMP [74] indicates that an uncertainty analysis on frequency is performed. The analysis suggests that they carried forward upper bound frequencies as classifier. It is not clear from the discussion if these upper bounds are 95th percentile. The review of the Kairos Power LMP [75] also suggests they use the mean of the frequency as classifier.

The Sandia report [77], as stated before, is more in-line with a graded approach interpretation of the risk-informed. However, they do stress that the evaluation should consider the impact of uncertainties in event frequency on meeting the F-C Target. In situations where uncertainties lead to an event frequency uncertainty band that overlaps two LBE categories, it is generally required to evaluate the event against F-C Targets for both categories. This is intended to ensure that uncertainties in frequency estimates do not impact whether a design can meet the F-C Target.

In conclusion the guidance in NEI-18-04 appears not fully clear on this topic and there is a range of interpretation among the users.

To provide some degree of consistency, the RISE methodology will adopt the 5th/95th results for LBE classification that, in our opinion, is the best and most consistent interpretation of NEI 18-04. Using both bounds, it is possible that a single ESF could be classified as different event types when the uncertainty crosses the frequency boundaries among those.

3.3.3 Meeting the F-C Target

In a risk-informed analysis, meeting specific risks target, defined in the frequency-consequence (F-C) map provides one useful and practical demonstration of how the design fulfills the NRC's expectations for enhanced safety. However, in the NEI 18-04 there are few language ambiguities than must be resolved.

For example, the guide states: "[…] The F-C Target values shown in the figure should not be considered as a demarcation of acceptable and unacceptable results […]". This is interpreted in the context that a guideline is indeed a guideline and not a rule.

Assessing if a design meets the F-C Target is useful and a critical part of assessing which safety functions (SFs) are required,that SSCs are safety related or risk-significant, etc. The risk-informed analysis exercise is really an assessment on how events moved in the F-C chart as a means to characterize margins and robustness of a design. As stated in the guide

> "[…] the results of the PRA which have been organized into LBEs will be evaluated against an F-C Target […] The figure does not define specific acceptance criteria for the analysis of LBEs but rather serves as a tool to focus the attention of the designer and those reviewing the design and

related operational programs to the most significant events and possible means to address those events".

The NEI 18-04 is not overly prescriptive when it comes to the details on how the analysis should benchmark design decisions to F-C Targets. This may lead to some confusion. However, in practice it is up the analyst to define those details and a degree of flexibility in the guidelines is reasonable. For instance, an analyst may decide that as assessment relative to the mean value of frequency and consequence is sufficient. Therefore, a more rigorous approach would consider the uncertainties considering the upper 95th quantile for both uncertainty and frequency for each LBE.

Regarding the definition of risk-significant LBEs, the guide says: "Risk-significant LBEs are those with frequencies within 1% of the F-C Target with site boundary doses exceeding 2.5 mrem. To consider the effects of uncertainties, the upper 95th percentile estimates of both frequency and dose should be used." Regarding the risk-significant SSCs, the guide states: "The LBE is considered within the F-C Target when a point defined by the upper 95th percentile uncertainty of the LBE frequency and dose estimates is within the F-C Target."

The F-C Targets play a role in the DID analysis. As stated in the guideline "An important input in evaluating DID adequacy is establishment of adequate margins between the risks of each LBE and the F-C Target.". However, what margin is adequate? Again, the guideline allows some degree of flexibility, and a specific methodology should reflect that possibility.

Compliance with F-C Target is probably purposely not a black and white determination of acceptability. Overall, the F-C Target seemingly has four explicit purposes, with varying treatment of uncertainty (e.g., using the mean or the 5th/95th F-C intervals).

When it comes to the interpretations in some of the LMP reports, the PRISM LMP [73] states a design objective of PRISM is to keep the LBEs well within the F-C Target such that the resulting margins can support the eventual demonstration of DID adequacy. As result, considering that there are large margins between the LBE mean-value points and the target line, no uncertainty upper and lower bars are presented for any of the LBEs and no uncertainty (frequency or consequence) analysis was performed during that LMP demonstration. The PRISM PRA included uncertainty analysis, but these were not resolved for individual ESFs and LBEs. Based on the large margin between the points and the target line and previous integrated uncertainty analysis performed in the PRA, these LBEs including uncertainties are not expected to challenge the target line.

The eVinci LMP report [74] also states as design objective to keep the LBEs well within the F-C Target with sufficient margins that can support the eventual demonstration of DID adequacy. The eVinci analysts avoid all the questions associated with the F-C Target not being a requirement by just making it one for themselves. Also, based on the report, they believe that uncertainty needs to be considered when comparing against the F-C Target, and they did so explicitly for frequency and dose. This resulted in frequency bands, but constant dose, because the dose cases were set taking epistemic uncertainties into account, which they expect to be dominant.

Kairos Power LMP report [75] states again as design objective of the KP-FHR is to keep the LBEs well within the F-C Target such that the resulting margins can reflect the KP-FHR safety design approach and support the demonstration of DID adequacy. When dealing with uncertainty they go further by stating that, in the context of risk-significant SSCs:

> "An LBE is considered within the F-C Target when a point defined by the upper 95th percentile uncertainty on both the LBE frequency and dose is within the F-C Target." Moreover, they go on by saying "Licensing basis events are represented on the frequency-consequence chart using six quantities from the PRA: Mean frequency of the event sequence occurring; Uncertainty in the frequency calculation characterized by 5th and 95th percentiles; Mean 30-day dose at the exclusion area boundary; Uncertainty in the dose calculation characterized by 5th and 95th percentiles."

The Kairos Power LMP work avoids all the questions associated with the F-C Target not being a requirement. However, for all practical purposes, uncertainty bands appear to be considered for the most part.

The Sandia report [77] stresses the importance of dealing with the uncertainties in many areas, from the determination of RSFs, to the evaluation of LBEs against the F-C Target to the DID analysis where they state: "*the magnitude of uncertainty in risk quantification is also an important consideration when focusing attention.*" However, the report does not reveal how these uncertainties are evaluated beside suggesting that these uncertainties could be rather large.

While it is inconclusive what NEI-18-04 requires for uncertainty treatment in the determination of RSFs, it seems to be a good strategy to assume uncertainty should be considered when comparing to the F-C Target. This is consistent with all LMP demonstrations. There appears to be universal agreement that uncertainty is needed for comparing to the F-C Target for risk-significance and DID adequacy.

The comparisons to the target line in RISE include the 5th to 95th space. This approach reinforces the DID argument in RISE by displaying uncertainty in comparisons to the line, regardless of if the user just wants to use mean.

3.3.4 Risk-Significance in LBEs

The NEI 18-04 guide states that for modular reactor designs, the ESs modeled in the PRA should also include scenarios involving single or multiple reactor modules or radionuclide sources. This approach provides useful risk insights into the design to ensure that ESs involving multiple reactor modules are not risk-significant. More in general, the primary purpose of comparing the frequencies and consequences of LBEs against the F-C Target is to evaluate the risk significance of individual LBEs. The question is: what determines an event to be risk-significant?

Figure 3-1 (Figure 3-4 from [62]) illustrates graphically the classification of risk-significant LBEs.

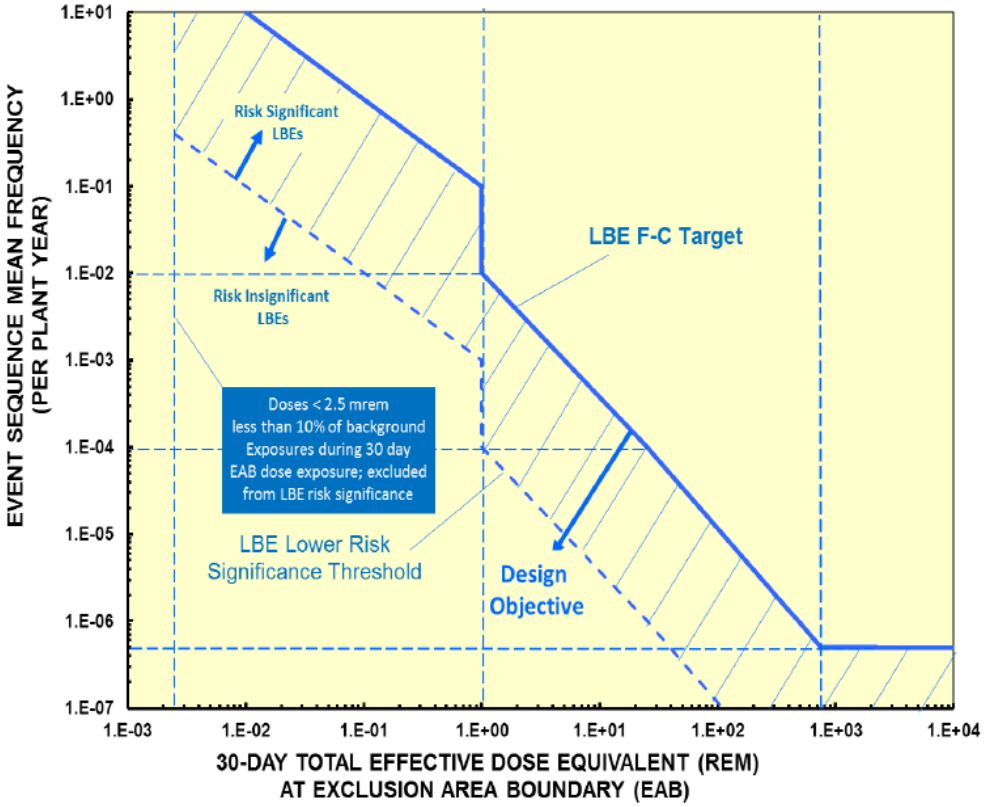

Figure 3-1 Graphical classification of Risk-Significant LBEs according to NEI 18-04.

Risk-significant LBEs are those with frequencies above 1% of the F-C Target with site boundary doses exceeding 2.5 mrem. NEI 18-04 also recommends considering the effects of uncertainties, defined as the upper 95th percentile estimates of both frequency and dose. The guides justify the use of the 1% metric as being consistent with the approach to defining risk-significant ESs in the PRA standards. It is also noted that the 2.5 mrem cut-off is selected as this is approximately 10% of the dose that an average person at the site boundary would receive in 30 days due to background radiation.

When LBEs are classified as risk-significant, the guide suggests several steps in the review of such events that is in-line with a risk-informed approach. The DID analysis also treats the risk-significant events with special attention.

The PRISM LMP report [73] shows that no LBEs are risk significant for PRISM IEAP scope. This situation is quite common in advanced reactor designs. For example, in the eVinci LMP analysis [74] there is no discussion of risk-significant LBEs. A similar discussion, in avoiding accepting events in the risk-significant region is presented in the Kairos Power LMP [75].

In conclusions, the NEI guidelines provides an explicit numeric definition of a risk-significant LBE, and it contains DID-relevant questions to such classification of the events.

For RISE, the designation of an LBE as risk-significant is automated consistent with NEI 18-04 recommendations.

3.3.5 Structures, Systems and Components (SSC) Classification

3.3.5.1 Definition of Safety Function

The primary objective of a risk-informed analysis is really the ability to properly classify the SSCs and characterize the safety parameters of a design. The first concept that requires a clarification is a proper

definition of SF. Moreover, the LMP introduces the term of PRA SF to specifically link the role of SFs in the PRA analysis.

The glossary at the end of NEI 18-04 define PSF as follows: “Reactor design specific SSC functions modeled in a PRA that serve to prevent and/or mitigate a release of radioactive material or to protect one or more barriers to release. In ASME/ANS-Ra-S-1.4-2013 these are referred to as "safety functions." The modifier PRA is used in the LMP GD to avoid confusion with safety functions performed by Safety-Related SSCs.”

We found that definition a little confusing or not particularly clarifying. Probably a most clear definition is provided in the text of guide: “*The PRA models the response of each SSC in the plant that performs a function to prevent or mitigate a release of radioactive material from any radionuclide source within the scope of the PRA. These SSC functions are defined in the LMP methodology as PRA Safety Functions (PSFs).*”

Given this clarification, in the context of risk-informed analysis, at least as implemented in RISE, there will be no distinction between PSF, and SF and the two terms will be used interchangeably. More specifically in this context the term SF represents a PSF.

The SFs are organized hierarchically starting with the main[e] SF at the top of the tree. The main SFs are defined in the International Atomic Energy Agency (IAEA) Safety Standard No. SSG-30 [78] as follows:

- Control of radioactivity
- Removal of heat from the reactor core and the fuel store
- Confinement of radioactive material, shielding against radiation and control of planned radioactive releases, as well as limitation of accidental radioactive releases

In short, as stated in the NEI 18-04 [62]: “i) Control heat generation; ii) Control heat removal; iii) Retain radionuclides.”

All reactors are designed to meet certain Main Safety Functions (MSF). However, application of the reactor-specific safety design approach leads to a set of reactor-specific PSFs that achieve the MSFs. The designer confirms the allocation of these PSFs to both passive and active SSCs. LBEs are defined as a series of successes and failures of SSCs to perform SFs.

All the LMP demonstration [73, 74, 75] appears to conform with this logic. The SFs tree is kept relatively high level in these demonstrations. SSCs are typically designed to perform SFs at the bottom of the SF tree. However, the details in the specification of the PSFs is up to the designer.

In RISE, SFs can be infinitely chained down from MSFs to all derivative PSFs that support the MSFs in a hierarchical tree. ESs are defined as success, or failure, or N/A of a SF/SSC pair. Future work should consider the possibility that an event in an event sequence is represented by an SSC failing or succeeding performing more than one SF. For now, a unique pair of SSC and PSF is what used to define an event in an event sequence or tree. Another possible extension in the future is ability to distinguish between passive and active SFs.

3.3.6 Determining Required Safety Functions

According to the NEI guide, some of the PSFs should be further classified as Required Safety Functions (RSFs) if they are necessary to ensure that all the DBEs and high-consequence BDBEs have

[e] The term ‘main’ safety function is preferred to ‘fundamental’, typically used in the NEI 18-04. This terminology is more consistent with the most recent IEAE Safety Standards No. SSG-30.

doses that fall within the F-C Target and to ensure that the doses for the DBAs meet the requirements of 10 CFR 50.34 using conservative assumptions. High-consequence BDBEs are those with consequences that exceed 10 CFR 50.34 dose criteria. For the DBEs these PSFs, when fulfilled, are responsible for mitigating the consequences within the F-C Target. RSFs for any high-consequence BDBEs are responsible for preventing them from increasing in frequency into the DBE region and outside the F-C Target by exhibiting sufficient reliability performance to keep the BDBE frequency sufficiently low.

The link between RSF and safety-related (SR) SSCs is provided in the guide as follows: “Once those RSFs are defined, SSCs that are available to support those functions on all the DBEs are identified. In addition, SSCs whose reliability needs to be assured to prevent any high-consequence BDBEs from migrating up into the DBE region are also identified. From these sets of SSCs, the designer selects a set of SR SSCs to perform each RSF.”

The process for identifying the RSFs typically involves sensitivity or challenge studies. The goal is to identify which of those PSFs, if not fulfilled, would likely increase the consequences of any of the DBEs beyond the F-C Target.

A possible implementation of such sensitivity analyses is with a systematic removal of PSFs that mitigates the consequences of each DBE followed by a re-evaluation of the consequences under such challenged scenario. From the RSFs, a top-down logical development is used to define the functional requirements that must be fulfilled for the reactor design to meet each RSF. These sensitivities would be performed for each DBEs and each high-consequence BDBEs to determine which set of PSF is necessary and sufficient to meet the F-C Target. That PSF set will be classified as RSFs.

In general, there may be two or more different sets of SSCs that could provide these RSFs. The design team selects which of the available SSCs can support the RSFs for all the DBEs and high-consequence BDBEs. Those SSCs are designated as SR.

It is noted that while the text in the NEI guide repeatedly references changing the consequences, in theory, both frequency and consequence may change when performing those challenge sensitivity studies.

The implementation method presented in the PRISM LMP analysis [73] appears consistent with the definition of PSFs, as well as the NEI RSF guidance. The LMP demonstration for the eVinci [74] appears to classify all the PSFs as RSFs, possibly due to the simplicity of its design. In the Kairos Power LMP demonstration [75], it appears that they use the term RSF to refer to high level PSFs derived from MSFs, and detailed pairs of a PSF and SSC. They also do not appear to use any sensitivity studies but do seem to acknowledge them as part of the process. The report by Sandia suggests that a comparative study may be performed in which different sets of SFs are credited against the set of DBEs and BDBEs. Those sets that meet the F-C Targets for all LBEs are eligible to be the declared RSFs. Both the PRISM and HTGC-PBR demonstrations determined that core heat removal and reactivity control were sufficient as RSFs and they are expected to be sufficient RSFs for the LBEs for both micro-reactor concepts which corresponds to Step 5a in Figure 3-1 of the NEI guide.

In RISE, for each DBE and BDBE, two options are supported: removing the credited PSFs entirely in the studies, and re-calculating the frequency of the resulting ESF, or simply maintaining the frequency of the ESF crediting the PSF. In either case, the consequences will be taken from a simulation which does not credit the PSF. Note that the consequences may correspond to an existing consequence evaluation (since every combination of SSC/function availability results in an existing event sequence; however, not every event sequence is expected to be simulated since some fall into the residual risk realm). Another observation is with sufficiently vague PSFs - and those PSFs having multiple SSCs to perform them - this may result in multiple branches being trimmed in the event tree.

3.3.7 Selecting SR SSCs

As stated in the NEI guide: "[…] for each of these RSFs identified in Task 5a, a decision is made on which set of SSCs is selected to perform these RSFs among those found to be available on each DBE. As a result of this selection, each DBEs is protected by a set of SR SSCs to perform each RSF. Structures and physical barriers that are necessary to protect any SR SSCs in performing their RSFs in response to any design basis external event are also classified as SR. SR SSCs are also selected for any RSF associated with any high-consequence BDBEs in which the reliability of the SSC is necessary to keep the event in the BDBE frequency region."

Once RSFs are determined, it is relatively straightforward to select a SR SSC for each RSF. Note that this assumes multiple SSCs can provide every RSF. If an RSF is provided by only one SSC, it is assumed that it would have to be SR. One complication that emerges from the review of the NEI guide is that the classification is focused on specific functions, rather than SSCs as a whole. This is likely relevant to requirements derivation.

The PRISM LMP [73] aligns well with this approach. The process outlines three critical tasks after the required functions are identified. The first task is to determine which SSCs are selected to meet the required SF for each DBE. These SSCs are then classified as safety-related. The next two tasks involve determining if the SSC should be classified as SR, Non-Safety-Related with Special Treatment (NSRST), or Non-Safety-Related with No Special Treatment (NST). Available sets of SSCs which could perform the three required SFs are identified. The LMP process does not attempt to determine which of the options should be chosen, as this is a designer's choice. The designer may consider many different parameters when selecting the safety-related SSCs, such as economic cost, regulatory uncertainty, and difficulty of performance requirements. Under LBE Selection Task 5b, "one specific combination of available SSCs" was selected that keeps all LBEs/ESFs within the F-C Target. Therefore, this set is classified as SR for the purposes of this demonstration.

Similarly, the eVinci LMP demonstration [74] appears to align well with the methodology described in the NEI 18-04, where it is stated: *"Reactivity control can be performed by three systems: Control Drum Subsystem; Emergency Shutdown Subsystem; The passive release of hydrogen from the moderator. The ESS is a Passive Category B system […] Due to the reliability of its passive features, it is selected as the safety-related system to perform the RSF of reactivity control. Therefore, it must be demonstrated in the PRA results that the ESS is capable of performing this RSF for all DBEs."*

This adds a stipulation that one SSC must perform the RSF for all DBEs and BDBEs. It is unclear if this is always the case, and that it wouldn't be acceptable to have multiple SSCs be SR for the same RSF. That said, obviously classifying the minimum number of SSCs as SR is probably good for operations.

The Kairos LMP demonstration [75] aligns with our interpretation of NEI-18-04 and also appears to confirm that a set of SR SSCs may be needed to cover all DBEs and BDBEs.

Once RSFs are determined, it is relatively straightforward to select a SR SSC for each RSF. Note that this assumes multiple SSCs can provide every RSF. If an RSF is provided by only one SSC, it is assumed that it would have to be SR. The RSF must be covered by a SR SSC or group of SR SSCs for all DBEs and BDBEs. The RISE data model was designed to support these conclusions, i.e., an SSC can be classified as SR, NSRST, or NST, but requirements of any type can be added.

3.3.8 Designating Non-Safety Related with Special Treatment SSCs

As stated earlier, SSCs are classified as risk-significant if the SSC function is necessary to keep any LBEs inside the F-C Target, or if the total frequency of LBEs with the SSCs failed is within 1% of any of the three cumulative risk targets identified in Task 7b. This information is used to provide risk insights, to identify safety-significant SSCs, and to support the RIPB evaluation of DID in Task 7e.

NEI 18-04 provides clear guidelines on determining if SSCs are risk-significant. NSRST is composed of risk-significant SSCs - those that are needed to keep all LBEs inside the F-C Target or significant to the QHOs - or those needed for DID adequacy. Those added to NSRST for DID adequacy are not risk-significant but are safety-significant. As we understand it, the difference between this study on impact to LBEs and the study on impact on DBEs for determining RSFs is that the study for RSFs requires removing the entire PSF family of top events, while these LBE risk-significance studies will only remove single top events corresponding to an SSC/function pair.

The PRISM and eVinci LMP demonstration analyses [73] [74] as well as the Sandia report [77] are consistent with our interpretation of NEI-18-04. The Kairos Power LMP demonstration is also mostly consistent with our interpretation. They introduce the concept of "hybrid point" shown in Figure 3-2 (taken from Figure 12 of [75]) which is basically equivalent to our "Challenge ESF" concept.

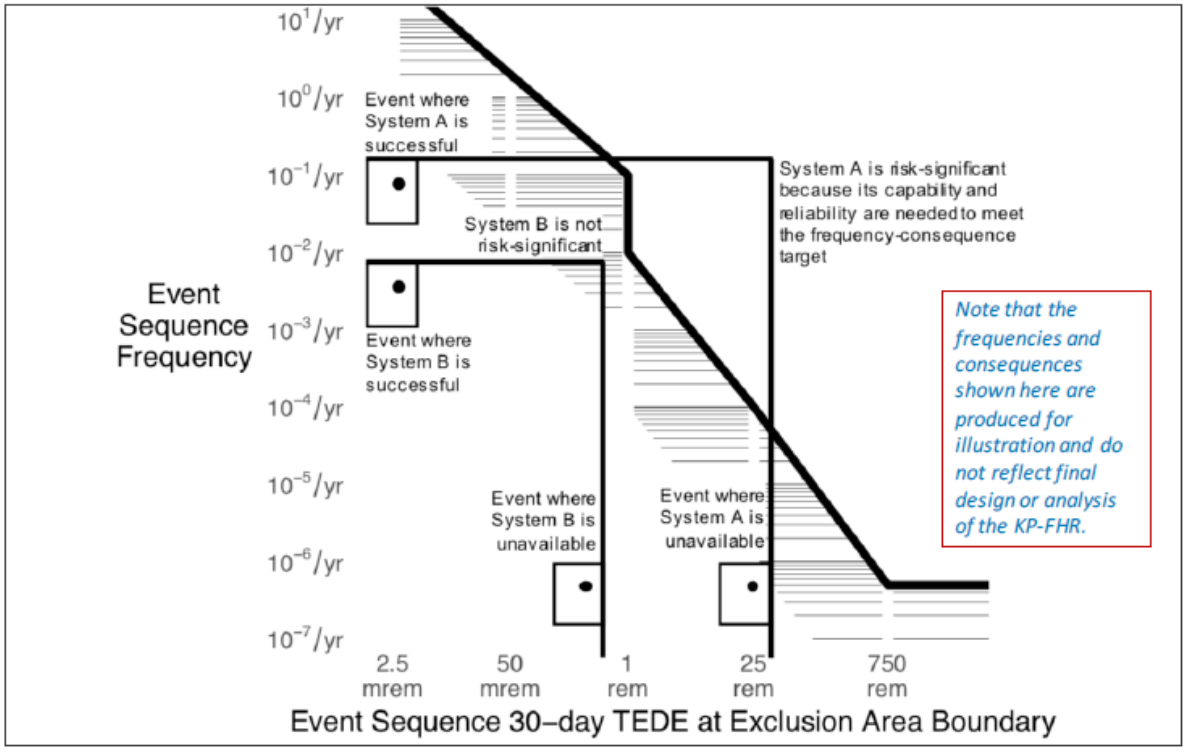

Figure 3-2 Kairos Power LMP method to determine risk-significance

The RISE data model supports these conclusions, in that an SSC can be classified as SR, NSRST, or NST, but requirements of any type can be added. The SF studies table in RISE also supports our interpretation or the hybrid point interpretation of risk-significance.

3.3.9 Design Criteria

NEI 18-04 makes a differentiation between Required Functional Design Criteria (RFDC) and Safety-Related Design Criteria (SRDC). The guide states the following: "*For SSCs classified as SR, the design criteria are referred to as Safety-Related Design Criteria (SRDC). These are derived from the Required Functional Design Criteria (RFDC) that are in turn developed from the RSFs determined in the LBE selection process as discussed in Section 3 of this guidance. RSFs are those safety functions that must be fulfilled to keep the DBEs within the F-C Target. RFDCs are taken down to a lower level and form a transition to SSC-level criteria. RFDCs are defined to capture design-specific criteria that may be used to supplement or modify the applicable General Design Criteria or Advanced Reactor Design Criteria in the*

formulation of Principal Design Criteria. RSFs and RFDCs are technology- and design-specific and are framed at the function level. After SR SSCs have been selected to perform the RSFs, the SRDCs are defined at the SSC level in a manner that assures meeting the RFDCs and the RSFs for the specific SSC selected to perform the RSFs."

The RFDCs are viewed as functional criteria that are defined in the context of the specific reactor design features that are necessary and sufficient to meet the RSF. The corresponding SRDCs are then developed from the RFDCs. DID attributes such as redundancy, diversity, and independence, and the use of passive and inherent means of fulfilling RSFs are used in the formulation of RFDCs.

Note that in the LMP exercises [73, 74, 75], the definition and evaluation of Design Criteria appears to be outside the scope of those demonstrations.

In the current state, the RISE methodology and data model does not dwell into RFDC. However, the model provides sufficient flexibility in the way the requirements are associated to SSCs and classified. The method will be eventually refined in the future as needed.

3.3.10 Special Treatment Requirements

The specification of Special Treatment Requirements is another quality attribute associated with the SSCs. Probably the most relevant characterization of this concept can be found in the NEI guidance in the following statement: "*NSRST SSCs are not directly associated with RFDC but are subject to special treatment as determined by the integrated decision-making process for evaluation of DID and for meeting the reliability and capability targets set in Task 6. The RFDC, SRDC, the reliability and capability targets for SR and NSRST SSCs, and special treatment requirements for SR and NSRST SSCs define safety-significant aspects of the descriptions of SSCs that should be included in safety analysis reports. The term "special treatment" is used in a manner consistent with NRC regulations and Nuclear Energy Institute (NEI) guidelines in the implementation of 10 CFR 50.69. In Regulatory Guide 1.201, the following definition of special treatment is provided: […] special treatment refers to those requirements that provide increased assurance beyond normal industrial practices that structures, systems, and components (SSCs) perform their design-basis functions.*"

The Table 4-1 and Table 5-7 in the NEI guidelines provide a list of special treatment and other requirements. Note that the applicability of special treatment to the SSC safety categories identified in Table 4-1 is provided for general guidance only, and it is not prescriptive.

Note that as stated in the NEI guide: "all safety-significant SSCs that are distributed between SR and NSRST are subject to special treatment requirements. These requirements always include specific performance requirements to provide adequate assurance that the SSCs will be capable of performing their PSFs with significant margins and with appropriate degrees of reliability. These include numerical targets for SSC reliability and availability, design margins for performance of the PSFs, and monitoring of performance against these targets with appropriate corrective actions when targets are not fully realized. Another consideration in the setting of SSC performance requirements is the need to assure that the results of the plant capability DID evaluation in Task 12 are achieved not just in the design, but in the as-built and as-operated and maintained plant throughout the life of the plant. The SSC performance targets are set during the design phase that is responsible for establishing the adequacy of DID. In addition to these performance targets, further special treatments may be identified."

The development of special treatment requirements was outside the scope of the PRISM and eVinci LMP exercises. The Kairos Power LMP and Sandia reports [75, 77] re-iterate the content in the NEI-18-04 guidance, but do not provide explicit examples on the implementation in the design process.

Special treatment requirements are needed for all SR and NSRST SSCs. They are meant to ensure SSCs perform their design-basis functions. Some guidance is given in Table 4-1 of the NEI guide for setting special treatment requirements, but this is not required. Table 5-7 of the NEI guide also has some

special treatment considerations for programmatic DID. Special treatment includes numerical targets for SSC reliability and availability, design margins for performance of the PSFs, and monitoring of performance against these targets with appropriate corrective actions when targets are not fully realized. Special treatment requirements need to be maintained through the life of the plant. RISE has the capability to make special treatment requirements and associate them to SSCs.

3.3.11 Reliability and Capability Targets

It is also possible that all SR SSCs need to have SRDCs for specific RSFs, but then also have special treatment requirements based on the overall SSC classification. In addition, specific SFs have special treatment requirements, and the needed special treatments will differ per function based on whether that function is an RSF or not. Another relevant term is what NEI 18-04 defines as Reliability and Capability Targets.

Information from the PRA is used as input to the selection of reliability targets and performance requirements for SSCs that set the stage for the selection of special treatment requirements. A good description in the NEI guide is the following: "*for those internal events caused by an equipment failure, the IE frequency is related to the unreliability of the SSC, i.e., SSCs with higher reliability serve to prevent the IE. Thus, higher levels of reliability result in a lower frequency of IEs. For SSCs that successfully mitigate the consequences of the IE, their capabilities and safety margins to respond to the IE are the focus of the safety classification process and resulting special treatment. For those SSCs that fail to respond along the LBE, their reliabilities, which serve to prevent the LBE by reducing its frequency, are the focus of the reliability targets derived from the classification and treatment process. The output of this task is the identification of the SSC prevention and mitigation functions for all the LBEs*."

Form the guide, all safety-significant SSCs, including those in the SR and NSRST categories, should be included in a Reliability Assurance Program (RAP) similar to that described in SRP 17.4. The reliability and availability targets established in the RAP are used to focus the selection of special treatments that are necessary and sufficient to achieve these targets and to assure they are maintained for the life of the plant.

The definition of reliability targets for passive systems brings some additional complexity and should be properly characterized. None of the LMP reports provide a clear example and this may be an area that requires further investigation.

In RISE, the reliability targets are a specific type of requirements, and the intent of NEI 18-04 is captured in its methodology.

3.3.12 Prevention versus Mitigation

The NEI guide often make a distinction between prevention and mitigation. In principle, in the view of F-C, prevention is an action that impacts the frequency of an event whereas mitigation is an action that impacts the consequences. Here is an excerpt from the NEI-18-04: "*[…] for those internal events caused by an equipment failure, the IE frequency is related to the unreliability of the SSC, i.e., SSCs with higher reliability serve to prevent the IE. Thus, higher levels of reliability result in a lower frequency of IEs. For SSCs that successfully mitigate the consequences of the IE, their capabilities and safety margins to respond to the IE are the focus of the safety classification process and resulting special treatment. For those SSCs that fail to respond along the LBE, their reliabilities, which serve to prevent the LBE by reducing its frequency, are the focus of the reliability targets derived from the classification and treatment process. The output of this task is the identification of the SSC prevention and mitigation functions for all the LBEs.*"

Interesting to note that all the LMP reports [73, 74, 75] do not have much discussion or emphasis on the distinction of mitigating and preventative. Figure 3-3, reproduced from Figure 3-7 from the Sandia report [77] provide an illustration which is consistent with our interpretation.

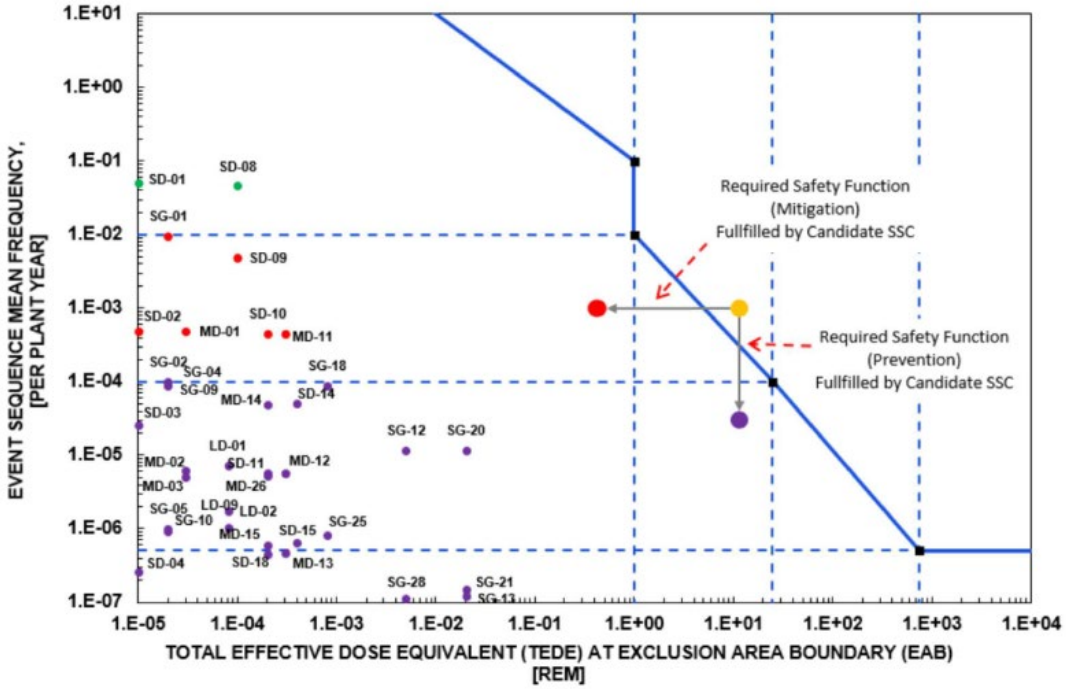

Figure 3-3 Illustration of prevention and mitigation from Figure 3-7 of [77].

The conclusion from the analysis of the NEI 18-04 guidance is that the main purpose of differentiating between mitigative and preventative is in the attempt to strike a balance for plant capability to support the DID.

It is possible that the underneath principle in striking a good balance between prevention and mitigations is to ensure there are redundant means to mitigate events that are going to happen frequently while preventing events that happens infrequently. However, this statement requires further corroboration.

RISE has the capability of providing a distinction between preventative and mitigating SF. The methodology will be refined as practical use cases will be implemented. One aspect we are considering is to update the data model by defining the preventative functions and mitigative functions lists in the ESF as SF/SSC pairs rather than just SF.

3.3.13 Safety-Significance of SSCs

NEI-18-04 states the following: “Safety-significant SSCs include all those SSCs classified as SR or NSRST. None of the NST SSCs are classified as safety-significant.” It also goes on by saying: “Safety-significant SSCs include those that perform risk-significant functions and those that perform functions that are necessary to meet DID criteria.”

As stated in the guidance the purpose of this task is to establish the specific design requirements for SSCs which include design criteria for SR classified SSCs, regulatory design and special treatment requirements for each of the safety-significant SSCs classified as SR or NSRST, and owner design requirements for NST-classified SSCs. Note that all SSC functions classified as either SR or NSRST are regarded as safety significant. All NST SSCs are not safety significant.

The relevance of this classification is because safety significant SSCs must have regulatory design and special treatment requirements and should be included in a RAP similar to that described in SRP 17.4. Adequacy of the programmatic measures for DID is driven by the selection of performance requirements for the safety significant SSCs.

A good summary of the SSC classification process is shown in Figure 3-4.

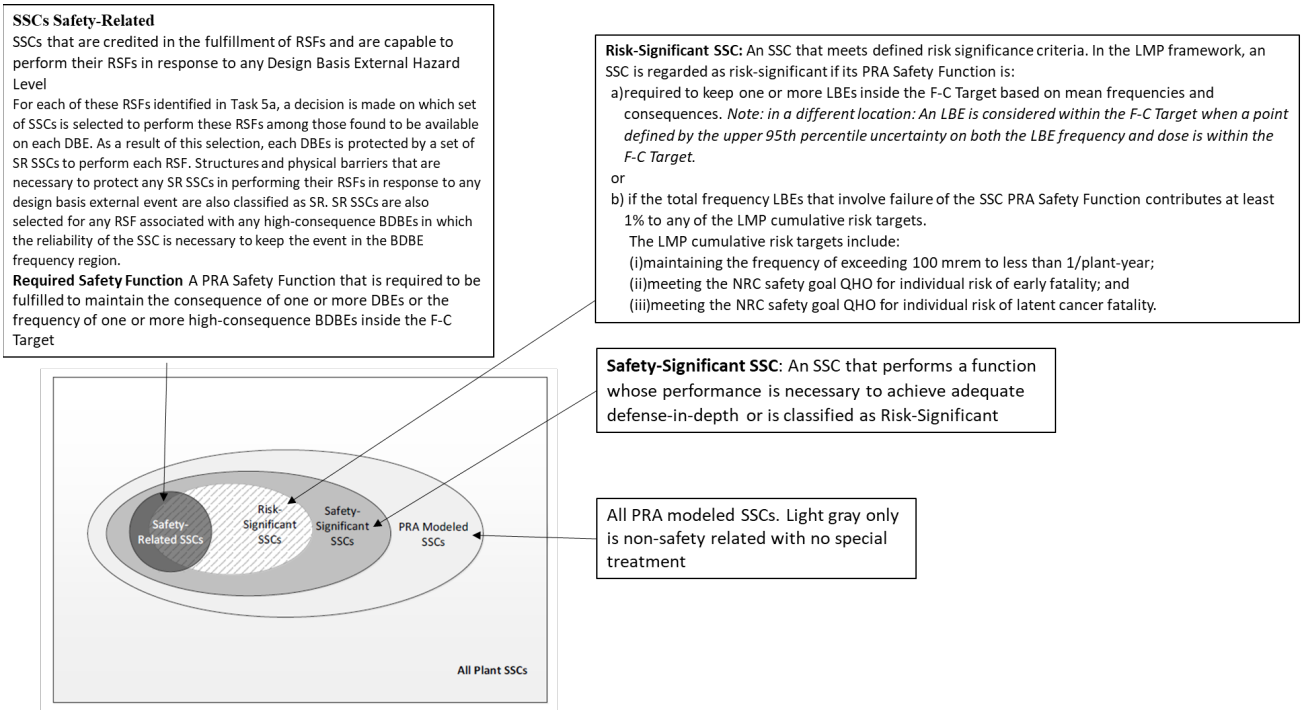

Figure 3-4 Summary of SSC Classification

The PRISM LMP [73] provides an actual application of this aspect of the guidance. On the other hand, the concept of safety-significant SSCs is not discussed in the eVinci LMP demo [74]. The Kairos Power LMP demonstration [75] is consistent with our interpretation of NEI 18-04. They also provide some examples of safety significant for DID, and mention that numeric performance targets are needed for safety-significant SSCs. The Sandia report [77] is also in agreement with our interpretation of NEI-18-04.

Currently RISE does not have a designation as safety-significant. The implementation of something like the Table 4-1 of Kairo Power LMP report is a considered improvement.

3.4 Evaluation of Defense-in-Depth (DID) Adequacy

According to NEI 18-04, the concept of using the layers of defense for performing the RIPB evaluation of plant capabilities and programs has been adapted from the IAEA “levels of defense” approach and shown in Figure 3-5.

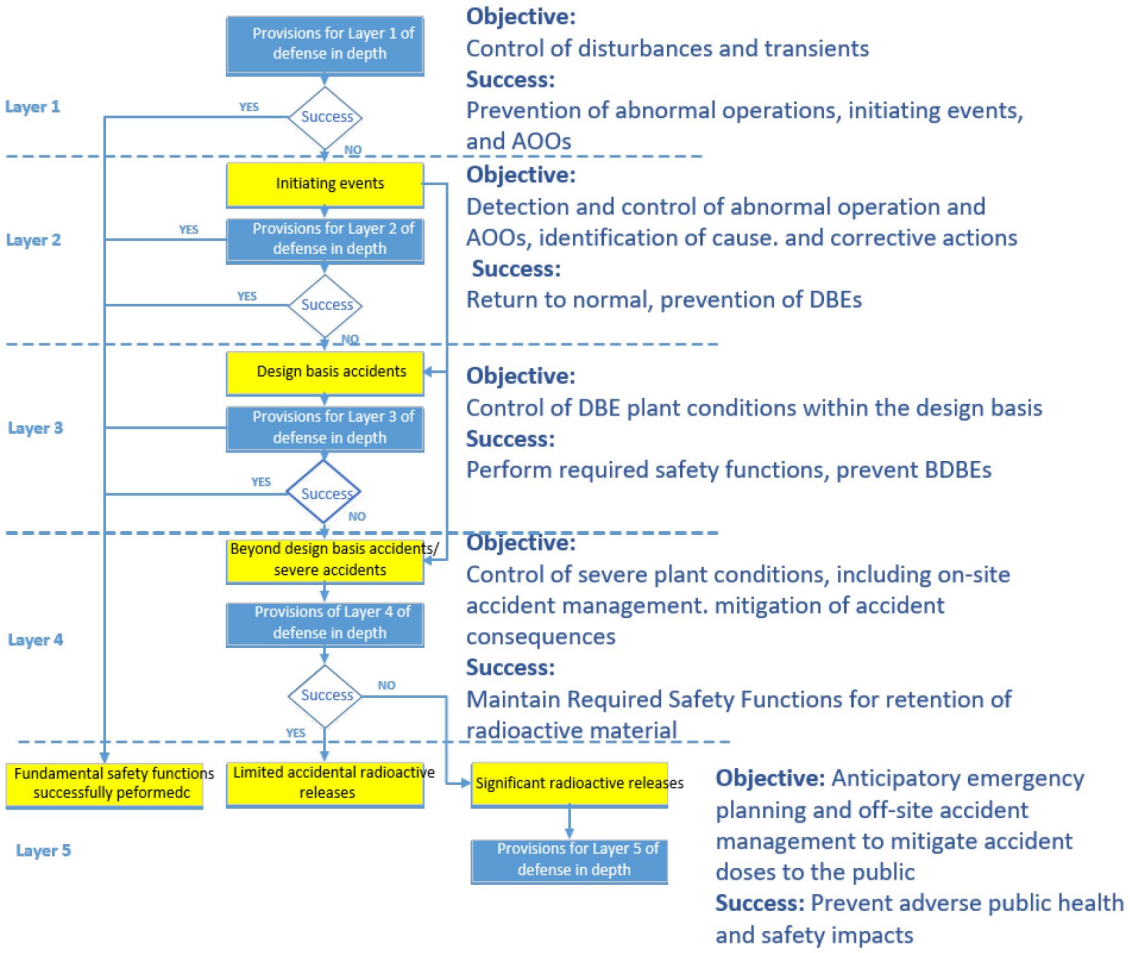

Figure 3-5 DID analysis framework (from Fig. 5-3 of NEI 18-04)

The layers of defense process are used in this task to evaluate each LBE with more attention given to risk significant LBEs to identify and evaluate the DID attributes to support the capabilities in each layer and to minimize dependencies among the layers.

Probably the most exhaustive demonstration of the DID analysis can be found in the Kairos Power LMP report [75]. The demonstration examines an entire IE (loss of flow), then defines the layers by timeline. The work appears in principle consistent with NEI 18-04.

However, a detailed review of the DID was considered outside the scope of this project.

Currently RISE has only partially implemented the workflow for the DID analysis and more work is needed to refine the details of this part of the process.

3.5 Notes on Uncertainty Treatment from the Review of NEI 18-04

Within the scope of [62], while many of the discussions of uncertainty are somewhat vague, there is some prescriptive guidance throughout. That said, much of the prescriptive guidance is unclear and sometimes even seemingly contradictory. In this section, specific guidance from [62] on the relevance of uncertainties is collected under the applicable step in the integrated process diagram (Figure 5-4) of [62] (copied herein as Figure 3-6). Different interpretations are discussed, and a position is taken for the remainder of the document.

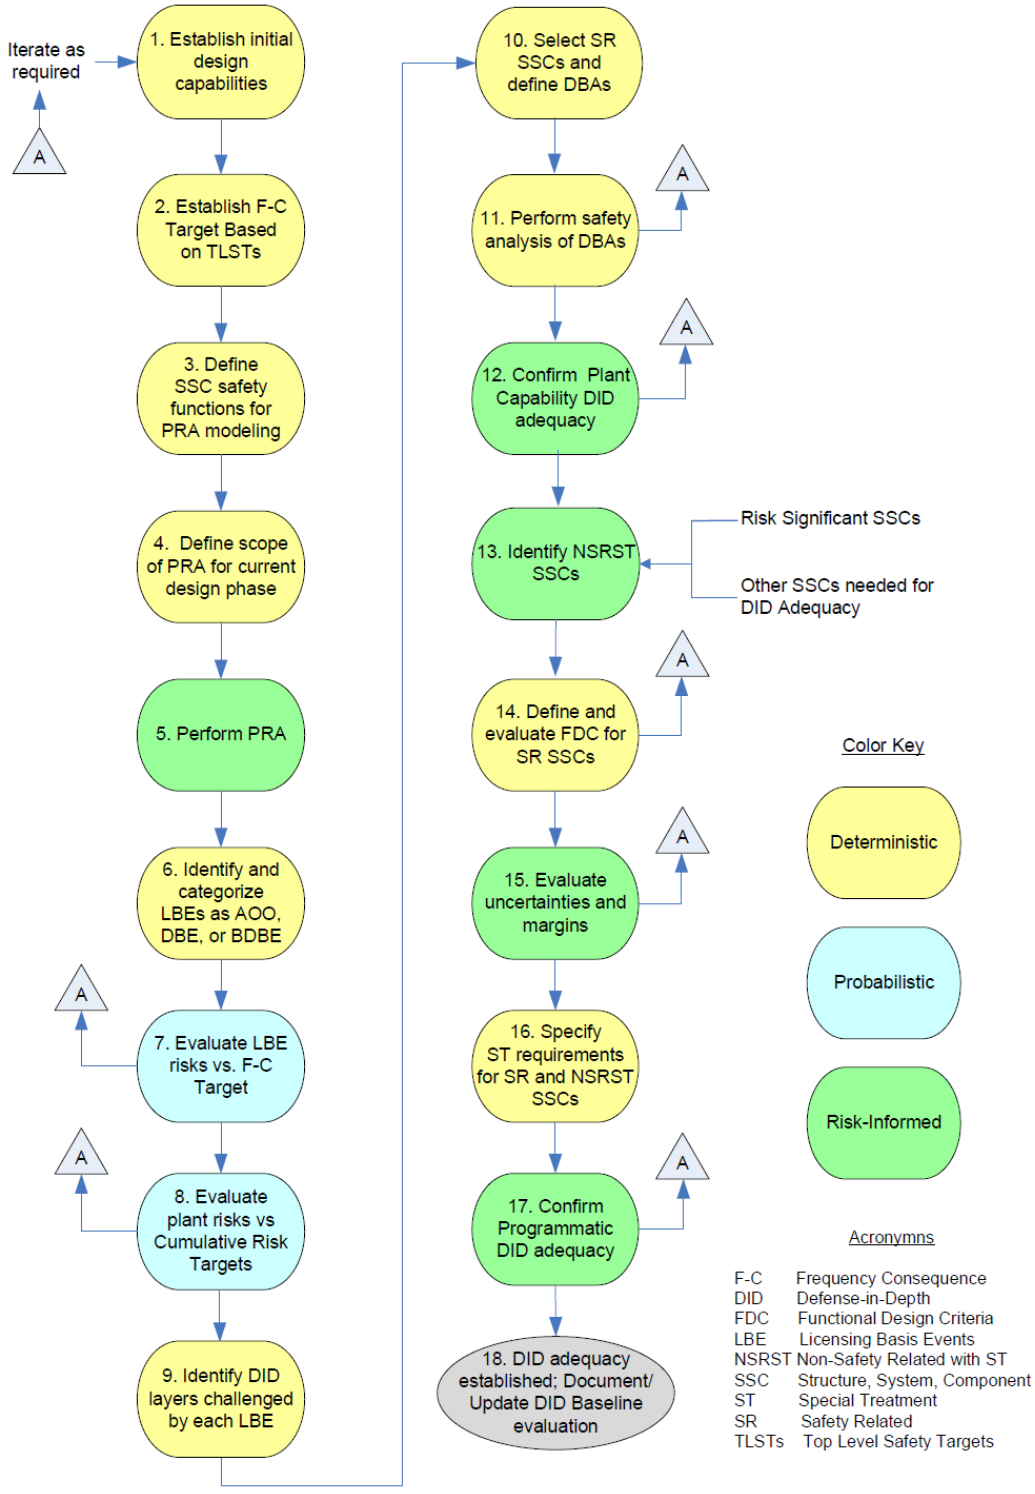

Figure 3-6 Figure 5-4 from [62].

1. Step 5 – Perform PRA
    - From the third paragraph of Section 3.2.2, Task 3 of [62]: "[…] The PRA process exposes sources of uncertainty encountered and provides estimates of the frequencies and doses for each LBE, including a quantification of the impacts of uncertainties using quantitative uncertainty analyses and supported by sensitivity analyses."
    - In a footnote of Section 3.2.2 of [62]: "[…] It is recognized that the PRA may not fully resolve the impacts of all sources of uncertainty, such as modeling uncertainty. The LMP approach to PRA recommends following the guidance in NUREG-1855 to address uncertainties. Uncertainties not quantified in the PRA are important inputs to the evaluation of defense-in-depth adequacy…"

Based on many locations in the discussions for the downstream steps, the results of the PRA with uncertainties are used for downstream activities (in particular, the distributions of frequency results). As such, Step 5 should be performed taking uncertainties into account, using guidance from [63].

2. Step 6 – Identify and categorize LBEs as AOO, DBE, or BDBE
    - In Table 3-1, for AOOs, DBEs, and BDBEs, mean frequency is specifically stated to be used for classification, for example: *"Event sequences with mean frequencies of 1×10-2/plant-year and greater are classified as AOOs."*
    - In the first and second paragraphs of Section 3.2.2, Task 4 (Identify/Revise List of AOOs, DBEs and BDBEs), it is stated that: "Each of these families is assigned to an LBE category based on mean event sequence frequency of occurrence per plant-year summed over all the ESs in the LBE family[…] AOOs are off-normal events that are expected to occur in the life of the plant with frequencies exceeding 10^{-2}/plant-year, where a plant may be comprised of multiple reactor modules. DBEs are less frequent events that may occur in a plant with frequencies between 10^{-4} to 10^{-2}/plant-year. BDBEs are rare events with frequencies less than 10^{-4}/plant-year but with upper bound frequencies greater than 5×10^{-7}/plant-year… ESs with upper 95th percentile frequencies less than 5×10^{-7}/plant-year are retained in the PRA results and used to confirm that there are no cliff-edge effects."
    - In the first paragraph of Section 3.2.2, Task 7a (Evaluate LBEs Against F-C Target), it is stated that: "The mean values of the frequencies are used to classify the LBEs into AOOs, DBEs, and BDBE categories. However, when the uncertainty bands defined by the 5th percentile and 95th percentile of the frequency estimates straddle a frequency boundary, the LBE is evaluated in both LBE categories. An LBE with mean frequency above 10^{-2}/plant-year and 5th percentile less than 10^{-2}/plant-year is evaluated as an AOO and DBE. An LBE with a mean frequency less than 10^{-4}/plant-year with a 95th percentile above 10^{-4}/plant-year is evaluated as a BDBE and a DBE. An event sequence family with a mean frequency less than 5×10^{-7}/plant year but with a 95%tile frequency estimate above 5×10^{-7}/plant-year is evaluated as a BDBE.

This appears to be a potential contradiction in the guidance. The word "Mean" is explicitly and exclusively mentioned in numerous places throughout the document, including several mentioned above. However, the last excerpt from above goes on to say that any category touched by the 5th to 95th percentile frequency range should also be evaluated. There may be a difference based on language used. For mean frequency, it appears to generally use "categorize" or "assign," while the passage about uncertainty bands says that it must be "evaluated" as whichever category. While there is this slight difference in language, it is not at all clear that this indicates any kind of practical difference. The purpose of categorization/classification appears to be to determine requirements for LBE evaluation, so it is unclear

how an LBE categorized by mean would differ from an LBE evaluated based on the 5th to 95th percentile frequency range. The position is taken herein to use the 5th to 95th percentile frequency range.

3. Step 7 – Evaluate LBE risks versus F-C Target
    - Section 3.2.1, third bullet point: "The F-C Target values shown in the figure should not be considered as a demarcation of acceptable and unacceptable results."
    - Section 3.2.2, Task 7a, first paragraph: "In this task, the results of the PRA which have been organized into LBEs will be evaluated against an F-C Target as shown in Figure 3-1. The figure does not define specific acceptance criteria for the analysis of LBEs but rather serves as a tool to focus the attention of the designer and those reviewing the design and related operational programs to the most significant events and possible means to address those events… the F-C Target provides one useful and practical demonstration of how the design fulfills the NRC's expectations for enhanced safety."
    - Section 3.2.2, Task 7a, first through third paragraphs: "Uncertainties about the mean values are used to help evaluate the results against the frequency-consequence criteria and to identify the margins against the criteria […] DBE doses are evaluated against the F-C Target based on the mean estimates of consequence [...] The primary purpose of comparing the frequencies and consequences of LBEs against the F-C Target is to evaluate the risk significance of individual LBEs. The objective for this activity is that uncertainties in the risk assessments are evaluated and included in discussions of design features and operational programs related to the most significant events and possible compensatory measures to address those events."
    - Section 3.3.5, Risk Significance Evaluations: "The first type is an evaluation of the frequencies and consequence of each LBE, expressed in the form of mean values and uncertainty (at the 5th and 95th percentiles), against the F-C Target. In this evaluation, the frequencies and consequences of individual LBEs are compared against an F-C Target derived from regulatory requirements and NRC safety goal policy. The objective is to keep the LBE frequencies and consequences within the F-C Targets."
    - Section 3.3.5, last paragraph: "Risk-significant LBEs are those with frequencies within 1% of the F-C Target with site boundary doses exceeding 2.5 mrem (see crosshatched region of Figure 3-4). To consider the effects of uncertainties, the upper 95th percentile estimates of both frequency and dose should be used."
    - Section 5.3, Task 7: "An important input in evaluating DID adequacy is establishment of adequate margins between the risks of each LBE and the F-C Target."
    - Section 5.7.1, all paragraphs: "Margins are established between the frequencies and consequences of individual LBEs and the F-C Target used to evaluate the risk significance of LBEs[…] Margins are developed in two forms. The margins to the F-C Target are measured based on mean values of the LBE frequencies and doses. In each case, margin is expressed as a ratio of the event's mean value (frequency and dose) to the corresponding F-C Target value (frequency and dose). These values are the best measure of the margins because traditionally in the PRA community, mean values are compared to targets such as design objectives for core damage frequency or large early release frequency and the NRC safety goal QHOs. A more conservative evaluation of margins, similar to the first form mentioned above, uses the 95th percentile upper bound values for both LBE frequency and dose to calculate the margins. This process is repeated for each individual LBE, grouped by LBE category as part of the DID evaluation during the design development."

- Section 5.8.1, first bullet point: “Assuring that adequate margins exist between the assessed LBE risks relative to the F-C Target including quantified uncertainties.”

Based on the first two excerpts above, the guidance attempts to make it clear that the F-C Target is not a strict requirement or acceptance criterion. So, the description of Step 7 (Evaluate LBE risks verus F-C Target) must be considered in the context of the actual purposes of the comparison of each LBE to the F-C Target. This step is interpreted to represent two main separate activities:

1. Determining risk-significance of LBEs: There is some mention of using both mean and 95th percentile values for this determination. However, the excerpt from Section 3.3.5 explicitly states to use 95th percentile of frequency and dose for categorization. This is interpreted to mean that 95th percentile should be used for strict definition as risk-significant or not, but in the requirements derived as a result, the mean value may be relevant, in addition to the uncertainty (e.g., the risk-significance is a result of very large uncertainties rather than a more confident prediction of risk).
2. Determining margin to the F-C Target for DID adequacy: Both the mean and the 95th percentile appear to be explicitly required. This is likely because DID adequacy is somewhat subjective, and thus a more complete examination of the risk space is prudent.

There is an observed issue that the F-C Target line is explicitly stated to not be an “acceptability” line, despite the document giving no guidance if any points are above the line. In addition, most of the guidance appears to be written omitting the possibility of an LBE being above the F-C Target to begin with, for example:

- The studies for RSFs and risk-significance of SSCs explicitly and repeatedly discuss functions and SSCs being necessary and sufficient to keep the event within the target. It cannot be kept within the target if it is not there to begin with.
- The DID adequacy guidance specifically mentions that “margin” to the F-C Target must be considered in DID adequacy. There is an implication in the guidance that there has to be positive margin (everything under the target), but since it is explicitly stated that the target it is not an acceptance criterion, this cannot be the case. It is unclear how DID can be adequate with negative margin to the F-C Target.

4. Step 8 – Evaluate plant risks vs Cumulative Risk Targets
    - In the three bullet points of Section 3.2.2, Task 7b, it is stated that: “The total mean frequency of exceeding a site boundary dose of 100 mrem from all LBEs should not exceed 1/plant-year[…] The average individual risk of early fatality within 1 mile of the EAB from all LBEs based on mean estimates of frequencies and consequences shall not exceed 5×10-7/plant-year […] The average individual risk of latent cancer fatalities within 10 miles of the EAB from all LBEs based on mean estimates of frequencies and consequences shall not exceed 2×10-6 /plant-year.”

These indicate that in the evaluation of cumulative risk targets, uncertainty analyses must be performed for frequency and consequence estimates, and that 95th percentile of frequency and dose estimates should be used.

5. Step 10 – Select SR SSCs and define DBAs
    - Section 4, first bullet point: “Safety-Related… SSCs selected by the designer from the SSCs that are available to perform the RSFs to mitigate the consequences of DBEs to within the LBE F-C Target.”
    - Section 4.1, Task 3, second paragraph: “As explained previously, there are some SFs classified as RSFs that must be fulfilled to meet the F-C Target for the DBEs using

realistic assumptions and 10 CFR 50.34 dose requirements for the DBAs using conservative assumptions."

- Section 3.2.2, Task 7a, last paragraph: "The upper bound consequences for each DBA, defined as the 95th percentile of the uncertainty distribution, shall meet the 10 CFR 50.34 dose limit at the EAB."

All references to the F-C Target for determining RSFs are vague and without explicit discussion of uncertainties. However, per the excerpt from Section 4.1, Task 3, it says to compare to the target for DBEs using "realistic assumptions", as distinguished from "conservative assumptions" used for DBAs. Per the excerpt from Section 3.2.2, Task 7a, DBAs are evaluated with 95th percentile results. Thus, these two passages appear to, in aggregate, equate 95th percentile with "conservative." As a result, "realistic" may indicate mean. However, many "generic" discussions about meeting the F-C Target throughout the document and for other purposes say uncertainties are used. With a lack of specific context, it could be concluded that this should be applied to determining RSFs. As such, it is inconclusive whether uncertainties should be used to determine if F-C results meet the target for determining RSFs. Because comparison to the F-C Target and dose limits generally considers uncertainties within this methodology, it is judged that, in the selection of RSFs, uncertainties should be considered.

6. Step 11 – Perform safety analysis of DBAs
    - In the 6th paragraph of Section 3.2.2, Task 6 (Select Deterministic DBAs and Design Basis External Hazard Levels), it is stated that: *[…] the codes and models used in DBA analysis are expected to satisfy Regulatory Guide 1.203 requirements for evaluation models.*
    - In the fifth paragraph of Section 3.2.2, Task 7a (Evaluate LBEs Against F-C Target), it is stated that: […] The upper bound consequences for each DBA, defined as the 95th percentile of the uncertainty distribution, shall meet the 10 CFR 50.34 dose limit at the EAB. Sources of uncertainty in both frequencies and consequences of LBEs are identified and addressed in the LMP approach to DID.

These indicate that in DBA analyses, a 95th percentile result is required for comparison to 10 CFR 50.34 dose limits, and that the evaluation models used to perform DBA analyses should be performed consistent with the guidance in [64].

7. Step 13, part 1 – Identify NSRST SSCs, Risk-Significant SSCs
    - In the first bullet of Section 4.2.2 of [62], it is stated that an SSC is risk-significant if: […] A prevention or mitigation function of the SSC is necessary to meet the design objective of keeping all LBEs within the F-C Target. An LBE is considered within the F-C Target when a point defined by the upper 95th percentile uncertainty on both the LBE frequency and dose is within the F-C Target.
    - In the second bullet of Section 4.2.2 of [62], it is stated that an SSC is risk-significant if: […] The SSC makes a significant contribution to one of the cumulative risk metrics used for evaluating the risk significance of LBEs.
    - Section 3.3.6, last set of bullets: "The SSC makes a significant contribution to one of the cumulative risk metrics used for evaluating the risk significance of LBEs. A significant contribution to each cumulative risk metric limit is satisfied when the total frequency of all LBEs with failure of the SSC exceeds 1% of the cumulative risk metric limit based on the mean estimates of frequencies and consequences."

These indicate that 95th percentile results considering the PRA uncertainty analysis are to be considered in risk-significant SSC classification when comparing to the F-C Target, but the mean values should be used for determination of risk-significance of SSCs based on comparison to the QHOs.

8. Task 15 – Defense-in-Depth Evaluation
    - In the first paragraph of Section 3.2.2, Task 7e (Risk-Informed, Performance-Based Evaluation of Defense-in-Depth), it is stated that: *[…] In this task, the definition and evaluation of LBEs should be used to support a RIPB evaluation of DID. This task involves the identification of risk-significant sources of uncertainty in both the frequency and consequence estimates, and evaluation against DID criteria.*
    - In the first paragraph of Section 5.3, Task 15 (Evaluate Uncertainties and Margins), it is stated that: […] One of the primary motivations of employing DID attributes is to address uncertainties, including those that are reflected in the PRA estimates of frequencies and consequence as well as other uncertainties which are not sufficiently characterized for uncertainty quantification nor amenable to sensitivity analyses. The plant capability DID includes design margins that protect against uncertainties. The layers of defense within a design, including offsite response, are used to compensate for residual unknowns.

These indicate that outcomes used for DID evaluation should consider uncertainties, and that sources of uncertainty for which there is insufficient understanding for other treatment should be treated with additional DID layers.

3.5.1 Dose Calculation Uncertainty

There is minimal guidance in [62] related to the calculation of doses, but there are three dose calculation outputs explicitly required in the process:

- Dose at EAB for 30 days: Used for F-C Target comparisons, as well as the first cumulative risk target (total mean frequency of exceeding a site boundary dose of 100 mrem)
- Dose within 1 mile of the EAB: Used for the second cumulative risk target (The average individual risk of early fatality within 1 mile of the EAB from all LBEs based on mean estimates of frequencies and consequences)
- Dose within 10 miles of the EAB: Used for the third cumulative risk target (The average individual risk of latent cancer fatalities within 10 miles of the EAB from all LBEs based on mean estimates of frequencies and consequences)

The first dose output is used directly, while the other two are combined with other assumptions to produce the final figure of merit. In any case, it is evident that some form of calculation or argumentation is required to essentially produce Level 3 PRA output. This is a bit different than the traditional PRA of the current operating fleet of LWRs, where core damage frequency and large early release frequency can be generically used as surrogates for more complicated calculations.

In addition, with the standardized scenarios, assumptions, and tools used for the operating fleet of LWRs, the calculations of source terms and releases is drastically simplified relative to new reactor concepts. In such new concepts, different fuel, reactor designs, site designs, and other factors lend themselves to a need for ground-up construction of a dose calculation method.

When combining these factors, it becomes evident that dose calculation can become an extremely burdensome and high uncertainty exercise, especially early in the design cycle of advanced reactors. As a result, it is reasonable to assume that a nominal or conservative calculation may be used. This may take the form of creating some form of surrogate criteria which is like using core damage frequency, or simply imposing conservative assumptions in a more explicit calculation.

For the demonstrations herein, no attempt is made to perform a detailed dose calculation. Rather, surrogate criteria and conservative assumptions are used to provide numeric values for consequences to demonstrate on the F-C Target.

3.5.2 Graded Approach

A risk-informed approach is currently considered by the developers of advanced reactors [65]. However, the industry recognizes the challenges in reliance on a maximal PRA role in their design and licensing efforts. Rather, a degree of flexibility is deemed necessary, and the conversation is around a more pragmatic graded approach.

The unified industry position to the Part 53 was recently in a letter to the NRC coordinated by NEI [65]. The letter elaborates on key elements to be considered in the rulemaking. Specifically, the view of the industry stakeholders is for the new rule to be "used and useful", "efficient", "technology inclusive", "risk-informed", "to recognize confidence in licensee control" and consider "urgency" in the finalization of the rule. Regarding the use of PRA techniques in the risk-informed approach, the industry is seeking some degree of optionality through alternative requirements, a graded approach. An applicant may choose a PRA "leading approach", as articulated in [62]. Another may opt for a "confirmatory/supporting" role, more in line with the previous Part 52. The choice is based on the specificity of a particular technology aiming to the most efficient definition of the "safety case". The argument is that, for very simple designs, PRA may not provide any practical benefit over alternative methods considered for the definition of the safety case. However, for simple designs, the PRA should be correspondingly simple.

For the new rule to "used and useful", it must not encourage advance reactor developers to fall back in adopting Part 52 which remains an available licensing strategy. These new reactors tend to be relatively simple and inherently safe. A very bounding deterministic approach is often the most efficient method to develop the safety cases. The other complicating factor is that especially in the early stages of the design cycle, data is missing and not sufficient for the development of a sophisticated PRA model. In those circumstances the development of surrogate bounding scenarios, figures of merit, surrogate acceptance criteria, consideration of single failures, etc., are effective means in the developing a defensible safety case. Instead, as the design evolves, as reliability data become available, the PRA analysis may provide useful insights. Therefore, only a graded approach allows the industry to exploit the benefit of a risk-informed approach.

The common theme in the industry debate is the recognition that establishing an efficient and practical regulatory framework that will minimize unnecessary friction for a safe and rapid deployment of new nuclear reactor technologies is a complex exercise. This also reveals the need and opportunity for the development of instruments in the digital age that can facilitate and streamline the engineering processes involved.

3.6 The Probabilistic Digital Twin Concept

The vision of the probabilistic digital twin is to connect traditional digital twins to risk models that capture the contributions of various levels of uncertainties. The probabilistic digital twin is seen as an extension to the traditional digital twin. The concept is illustrated in the 'risk-informed decision" pyramid of Figure 3-7. The physical world, made of the environment and the assets, resides at the bottom of the pyramid. The digital twin in a traditional sense is the virtual representation of the physical system using physics models and data. The probabilistic digital twin is an extension of the digital twin and combines surrogate models, plausible events and scenarios, logic trees (fault trees and event trees), failure and degradation models, all framed within a probabilistic model.

Probabilistic digital twins are the basis for risk-informed decisions. The long-term vision is for the probabilistic digital twin apparatus to be continuously automatically updated to reflect new knowledge and information. Note that this view is also consistent with a graded approach where the level of sophistication of the analysis, being either the evaluation of the consequences for a postulated accident scenario or the development of the PRA model, evolves with the design. On the same token, the developer of a very safe, high margin design could opt for a bounding, deterministic approach for its

safety case. A bounding deterministic approach is simply an alternative method for managing uncertainties.

Regardless, the probabilistic digital twin concept is intended to form the basis for building the safety case of the plant within a risk-informed paradigm. The ability to continuously feed the probabilistic digital twin with new information will enable managing dynamic risks in operation and the development of 'smart' operational and emergency procedures. Figure 3-7 is used as a guide for the discussions throughout the rest of the section.

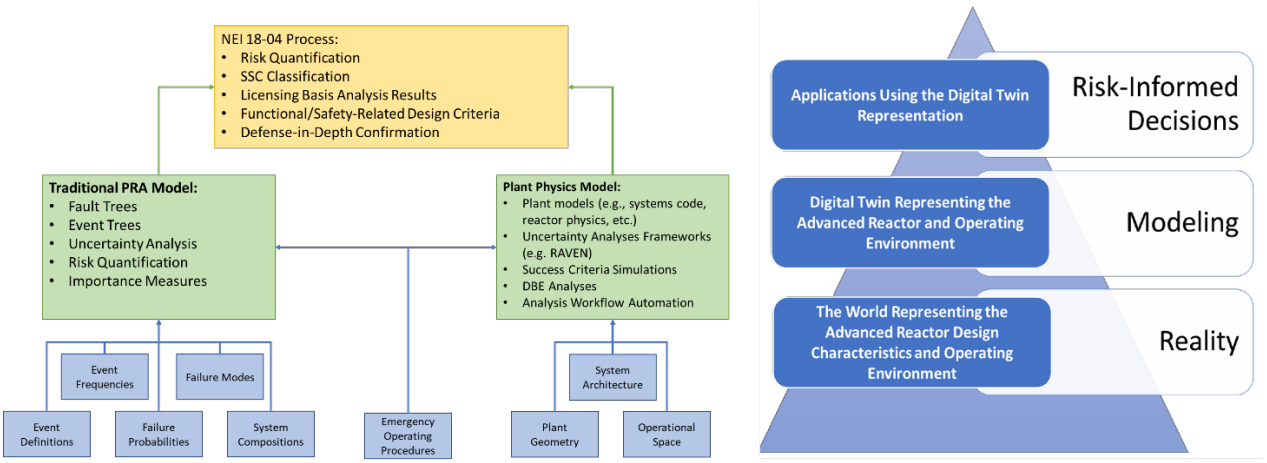

Figure 3-7 Risk-Informed Decisions Pyramid.

3.6.1 Risk-Informed Decisions in the Probabilistic Digital Twin

As stated above, a pragmatic methodology needs to consider a graded approach from a more traditional (10 CFR Part 52) deterministic safety case for a design, verified with a PRA analysis to full scope risk-informed design as outlined in the NEI 18-04 with maximal use of PRA artifacts since the early stages of the design cycle. In this section the focus is on the latter.

Within the context of this project, the process described in [62] is used as the basis for risk-informed decision making. The following outputs are produced in the process from [62]:

1. Risk Quantification: Part of the NEI 18-04 process is to quantify the frequency of ESF, as well as their consequences in terms of dose. Figure 3-8 illustrates the F-C Target concept from the NEI 18-04 process. This allows the risk of ESFs to be visualized along with their acceptance criteria.

In addition to the F-C Target for individual event sequence families, there are also the following cumulative risk targets:

a. The total mean frequency of exceeding a site boundary dose of 100 mrem from all LBEs should not exceed 1/plant-year.

b. The average individual risk of early fatality within 1 mile of the EAB from all LBEs based on mean estimates of frequencies and consequences shall not exceed 5×10^{-7}/plant-year.

c. The average individual risk of latent cancer fatalities within 10 miles of the EAB from all LBEs based on mean estimates of frequencies and consequences shall not exceed 2×10^{-6} /plant-year.

Taken together, the F-C Target and cumulative risk targets give the plant operator a number of quantitative margins to assess the true risk impact of a decision. In addition, they provide rigid requirements that allow quantification of margins.

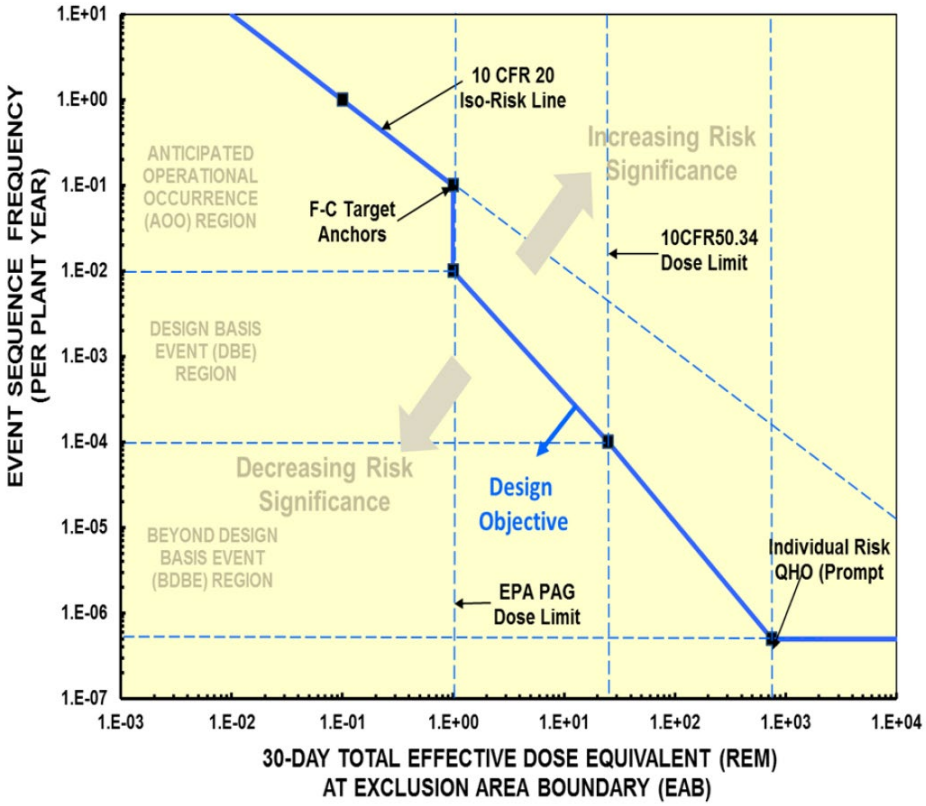

Figure 3-8 The Frequency-Consequence Target from the NEI 18-04 process.

2. SSC Classification: Part of the NEI 18-04 process is the designation of systems, structures, and components as being SR, NSRST, or NST. These classifications may result in requirements that provide increased assurance beyond normal industrial practices that SSCs perform their design-basis functions. As such, re-classification will generally result in changes to operating costs. Therefore, impact on SSC classification is an important concern for decision making.
3. Licensing Basis Analysis Results: As a result of the NEI 18-04 process, some subset of ESFs will be designated as DBAs and thus used to set design criteria and performance objectives for safety-related SSCs. Thus, impact on the analytical results of DBAs are of particular interest in risk-informed decision making.
4. Required Functional/Safety-Related Design Criteria and Special Treatment Requirements: The NEI 18-04 process establishes RSFs and the related RFDCs and SRDCs, as well as special treatment requirements. The RFDCs are essentially the specific requirements of a SF, while SRDCs are requirements for the SSCs that provide an RSF. As more design criteria are added or existing criteria become stricter, operation costs will increase; thus, the number and nature of RFDCs, SRDCs, and special treatment requirements is important to operators.
5. DID Confirmation: The NEI 18-04 process has some very explicit requirements for DID, such as having SFs for all layers and meeting risk targets, but also broadly covers a more subjective array

of questions such as whether appropriate conservatisms have been used. However, the reactor designer ends up confirming DID, its ongoing adequacy is a major concern in decision making.

Within this framework, decisions for the design and operation of an advanced reactor can be judged based on their impact on these outputs, and the acceptance criteria. For example, a decision may be considered ideal if it:

- Results in a negligible increase or a decrease to the overall risk profile
- Results in no changes to SSC classification, or relaxes some special treatment requirements
- Does not reduce the margins of licensing basis analyses to their acceptance criteria or increases the margin.
- Does not create additional design criteria or require more demanding design criteria.
- Does not challenge DID

If one or more of these bullets is not true, then the effects can be weighed against each other to determine the ideal path. This framework provides a robust set of information to perform such a comparison.

3.6.2 Modeling in the Probabilistic Digital Twin

For a given plant or design, plausible IEs and ESs are identified to form the set of LBEs. The analysis of frequency and consequences of the LBEs lead to their classification in AOOs, DBEs, BDBEs and DBAs.

The model requirements in the probabilistic digital twin flow down from the risk-informed decision-making layer. Within the NEI 18-04 process, there are two categories of overlapping analyses:

1. The traditional PRA model: This forms the skeleton of the probabilistic digital twin modeling. In essence, the PRA model defines all explicitly considered sequences of events and system availabilities for the plant design, as well as the frequencies of ESs. These sequences are meant to capture all risk-significant, credible scenarios and failure modes. The important aspects of the traditional PRA model are:
    a. Fault Trees: The fault tree is defined in the NUREG-0492 fault tree handbook as:
        i. An analytical technique, whereby an undesired state of the system is specified (usually a state that is critical from a safety standpoint), and the system is then analyzed in the context of its environment and operation to find all credible ways in which the undesired event can occur.
        ii. In practice, this generally manifests as a tool which is used to solve for the probability of branching conditions of a top event in an event tree. The fault tree is generally composed of basic events (such as the component failures, human errors, phenomenological events, etc.) and is used to express the top event in terms of combinations of necessary basic events.
    b. Event Trees: Event trees model the plant response to an initiating event. They identify ESs which result in some outcome of interest and consist of an initiating event, top events (system availability), branching, and end states (which in the context of the NEI 18-04 process is a dose estimation). Accident sequences are obtained by moving across the event tree from left to right, keeping track of successes and failures for each system top event until a specific end state is reached. In practice, the probabilities of branching conditions of top events are informed by the fault trees.

c. Uncertainty Analysis: The uncertainties in a PRA analysis are vast and the impact can be quite large. More detailed information on uncertainty analysis for PRA is provided in Section 3.2. For the purposes of the discussion in this section, it is important to understand that the PRA analysis and plant response simulations in the context of the NEI 18-04 process provide two-dimensional spaces on the F-C Target curve, generally reflecting 5^{th} to 95^{th} percentile frequency and consequence results, reflecting the numerous sources of uncertainty. Note that in the methodology herein, the PRA parameter uncertainty will generally be reserved to parameters that affect the frequency results.

d. Risk Quantification: As discussed in Item 1 of Section 2.1, a primary output of a PRA analysis is quantified risk (in terms of dose-frequency). This quantified risk is essential to all five outputs of the risk-informed decision-making framework in the NEI 18-04 process.

e. Importance Measures: These provide quantitative perspective on dominant contributors to risk and sensitivity of risk to changes in input values. In traditional PRA analyses, they are usually calculated at core damage frequency level, and some common importance measures include Fussell-Vesely, Risk Reduction, Risk Increase or Risk Achievement, and Birnbaum. Ultimately, these measures are meant to give operators a sense of safety-importance of SSCs, which is important to risk-informed decision making.

2. The plant response simulations: These models are used to perform simulations of ESs determined by the PRA analysis, to estimate consequences, and to provide licensing basis analyses. The important aspects of the plant response simulations are:

    a. Simulation Code Models: This is a potentially sprawling group of calculational codes and input models which forms a simulation that potentially goes from core modeling through offsite doses. With the intention of creating integrated PRA response simulations and DBA analyses, it is the intention of this project to create a model stack which may be composed of the following areas (and likely others):

        i. Reactor Physics – Code(s) or module(s) will be required to do reactor kinetics, fuel depletion, fuel management, neutronics, decay heat calculations, etc. This step is crucial to calculating power distribution and fission product inventories for downstream analyses.

        ii. Fuel Analysis – Code(s) or module(s) will be required to determine thermal hydraulic design, fuel temperature distributions, fuel housing integrity, etc. This step is crucial for determining some initial conditions for accident scenarios and for calculating potential fuel damage resulting in the release of fission products.

        iii. System Response – Codes(s) or module(s) are required to model the reactor system through the power generation system. The scope of this set of models will vary significantly between reactor designs, but should include the primary coolant system, and intermediate heat transfer systems, and the system which generates power. All relevant safety systems and accident mitigation systems should be modeled.

        iv. Dose Analysis – Codes(s) or module(s) which model pathways from releases through to doses to plant workers, operators, emergency personnel, and the public. This model group may consist of containment models, radionuclide transport models, and dose estimation (among others).

b. PRA Event Sequence Simulations: Using the ESs from the PRA analysis, a series of simulations will model the event sequence with the end goal of estimating the consequences (i.e., the resulting doses). It is envisioned that not every ES must be simulated, and ESFs can be represented by a bounding simulation. These simulations will use the physics models from part ‘a’ to calculate the event sequence from nuclear reaction to doses.

c. DBA Simulations: These are postulated ESs that are used to set design criteria and performance objectives for the design of SR SSCs. DBAs are derived from DBEs based on the capabilities and reliabilities of SR SSCs needed to mitigate and prevent ESs. DBAs are derived from the DBEs by prescriptively assuming that only safety elated SSCs are available to mitigate postulated event sequence consequences to within the 10 CFR 50.34 dose limits. These DBAs are required parts of a plant’s licensing basis.

d. Uncertainty Analyses Frameworks: The uncertainties in the plant response simulations are vast and can be quite large. More detailed information on uncertainty analysis for plant response simulations is provided in Section 3.3. There is some overlap of the uncertainties with the PRA analysis domain, but for the purposes of the discussion in this section, it is important to understand that the upper bound consequences for each DBA are used, defined as the 95th percentile of the uncertainty distribution. Thus, robust uncertainty analysis methods are necessary.

e. Analysis Workflow Automation: Between the potentially enormous number of ESs and simulations required, and the potential for that number to increase by orders of magnitude through uncertainty analyses, the automation and connection of these workflows is essential to the viability of a probabilistic digital twin. The significant upfront investment in defining a robust architecture for the automation framework provides in return the expedience and agility to perform a complex network of calculations automatically. This, in turn, enables quick updates of the analysis to reflect the effects of changes to input data.

3.6.3 Reality in the Probabilistic Digital Twin

In essence, reality is reflected in the probabilistic digital twin via input interfaces. In other words, the crucial problem to solve is the discretization of reality into usable inputs for the PRA and physics simulation models. The inputs are broadly grouped as follows:

1. Traditional PRA Model Inputs:

    a. Initiating Events: Perturbations to the plant during a plant operating state that challenge plant control and safety systems whose failure could potentially lead to an undesirable end state and/or radioactive material release. An IE could degrade the reliability of a normally operating system, cause a standby mitigating system to be challenged, or require that the plant operators respond to mitigate the event or to limit the extent of plant damage caused by the IE. These events include human-caused perturbations and failure of equipment from either internal plant causes (such as hardware faults, floods, or fires) or external plant causes (such as earthquakes or high winds). An IE is defined in terms of the change in plant status that results in a condition requiring shutdown or a reactor trip (e.g., loss of main feedwater system, small reactor coolant pressure boundary breach) when the plant is at power, or the loss of a key SF (e.g., decay heat removal system) for non-power modes of operation. A specific type of IE may be identified as originating from a specific cause as defined in terms such as “flood-induced transient” or “seismically induced reactor coolant pressure breach.”

b. Event Frequencies: For each IE, a frequency distribution must also be determined, as it is essential to obtaining the frequency portion of the F-C space. For some IEs which have actually happened, or are derived from actual occurrences, a narrower uncertainty distribution may be possible due to the availability of predictive data. However, as many of these IEs are only postulated, the frequency of occurrence is expected to often be a fairly broad distribution. An accurate accounting for uncertainty is necessary, as frequency is an essential component to risk.

c. Plant Responses: This is a broad category that involves preventative and mitigating SFs, automated instrumentation and controls, operating procedures, severe accident guidelines, and phenomenological events. To determine the ESs starting from an IE, all relevant plant responses must be considered to create branching conditions on the event tree.

d. System Compositions: For every top event in the event tree, every logical combination of basic events to reach that top event must be determined, so that a probability can be calculated. As such, system compositions refer to all the components of the system on which a top event is based. These components may be mechanical components (e.g., a valve), electrical components (e.g., a circuit), instrumentation and controls (e.g., a setpoint and its resulting signal), human actions (e.g., flipping a switch), or phenomenological events (e.g., leak rate exceeding a threshold).

e. Failure Modes: The failure modes of the system components discussed above must be determined, as these are the basis for basic events.

f. Failure Probabilities: With the failure modes determined, an estimate of the probability of those failures is needed. This estimate is most accurately expressed as a probability distribution of some kind.

2. Plant Response Simulations Inputs:

a. Plant SSCs, Geometries and Materials: This information is required to build all the models in the simulation code stack. This information includes anything essential to calculation of releases, which could include fuel structure, reactor coolant system, turbine, containment, or anything else essential to simulation of the plant.

b. Physics Models: This represents the mathematical quantification of the physics involved in the simulation of the plant. While it is recognized that these correlations are in fact models, in the context of this probabilistic digital twin concept, they are the nearest form of reality which can be expressed as input to a simulation. As such, they are included in the reality layer of the probabilistic digital twin. These physics models would be expected to include reactor physics, thermal-hydraulics, mechanical design, and anything needed to model phenomena anticipated in the reactor design.

c. Operational Space: The operating limits of a plant determine many of the initial conditions, boundary conditions, and timing aspects of plant response models. For example, limits on the core power and power distribution will strongly influence the prediction of core damage. The operational space is largely defined by the Technical Specification (TS), which presents a unique challenge. While the TS defines the operating limits of the plant, and thus that entire space should be considered, from a probabilistic standpoint, a plant likely spends most of its time in a smaller band of that operational space. One goal of this work is to develop a solution for this problem.

It is noted that many aspects of reality mentioned above are speculative but a necessary element in the safety analysis. Many of the events are only postulated, the physics models are imperfect representations

of reality, there are failure modes that are not anticipated or believed to be incredible, and much more. This results in questions of the validity of the representations of reality. While the imperfect nature is acknowledged, most of these inputs are designed around using probability distributions. As such, uncertainty on their validity can be captured as broad distributions, and the effects of these broad distributions will be apparent in the uncertainty analysis. As a result, confidence in the validity of these representations of reality can be quantified.

Ultimately the vision for the digital-twin is also the necessity to create a feedback loop between reality (measurements) and digital representation of the plant throughout its life. This will provide the bridge between activity such as condition monitoring, sensing during operations, etc. and the model calibration.

3.6.4 Integrated Platform to Facilitate a Probabilistic Digital Twin

FPoli developed a digital solution for managing the challenges in adopting a risk-informed approach early in design. The application is called RISE and introduced in this Section. The platform was architected with a list of key quality attributes that enable the user/analyst to:

- Create a collaborative environment for engineering teams and stakeholders within their organization as they build the 'safety case' for their plant
- Digest large and complex data structures needed to characterize the engineered safety features and relationships with scenarios and events
- Optimize design to satisfy safety and economics goals
- Guide analysts through complex workflows of simulations, data processing and qualification, analyses, and documentation
- Maximize the value of enterprise technical data with enhanced security and process automation
- Automate the creation of documentation and smart procedures for quality, transparency and expedited regulatory review
- Provide a platform for maintaining the safety case throughout the life of the plant
- Fit seamlessly within established processes of the organization

These quality attributes were the motivation behind the development of the RISE platform. The need was to construct a framework that can intelligently guide and organize critical engineering data and decisions toward the creation of a robust, defensible, and traceable safety case of new reactor system and the ability to maintain the safety case throughout the entire product cycle including the conceptual design, the detailed design, testing, deployment, and operation.

RISE is one of the applications powered by the FPoli Agile Application Platform, FPoliAAP [66]. The generic platform uses modern data management and simulation management tactics to orchestrate complex workflows with a highly and rapidly customizable UI/UX. Several use cases can be supported from test data management (FPoliDON) [67], document management (FPoliDOX) and simulation management application (FPoliSIM). The simulation manager, FPoliSIM, leverages RAVEN technology as workflow engine. RAVEN provides a vast library of optimization and machine learning algorithms which can be invoked as needed. RISE was recently added to the suite as the service to automate and facilitate the workflow associated with NEI-18-04 as shown in Figure 3-9. All artifacts are stored in a central relational database which represents a single point-of-truth accessible simultaneously from multiple sites and multiple users within an organization. The data is shared among the subscribed applications.

RISE was chosen as the instrument to illustrate the methodology developed in this project. The analysis presented in this report was casted into the RISE application as a mean of illustrating the benefit

of performing complex workflows, such as the one shown in Figure 3-9, in a powerful digital infrastructure.

Plant design data is loaded and warehoused from a variety of sources, reviewed and qualified. When the data is ingested and prepared for use, the powerful RISE workflow orchestrator guides the engineer through the intricacies of building a risk-informed ‘safety case’ of the plant design without impeding, but rather augmenting, the creative aspects and value of the engineering exercise, in a collaborative framework typical of an enterprise solution.

The system engineer starts by collecting reactor design data and metadata. APIs are available to interface with user system engineering tools. Relationships to systems design requirements, engineered SFs and parameters are established. The data set includes owner design requirements, economics targets and safety constraints. Once the available data is entered, this forms the current snapshot of the plant design with the characterization of considered SSCs. At the same time, the designer identifies and describes a list of plausible IEs and ESs which are then organized into a database of LBEs, the safety basis events.

The safety analyst enters safety evaluation criteria and metrics for the safety assessment. The user may elect to enter criteria described in NEI 18-04 (Part 53), but they are not limited to those criteria. If desired, PRA artifacts and results are entered to inform the classification of LBEs. LBEs are evaluated leveraging the powerful simulation manager service, FPoliSIM. The simulation manager guides the analyst in developing the evaluation model, performing simulations – the digital twin representation of the plant – and postprocessing the results following the established metrics and criteria.

Complex workflows that include a multi-physics representation of the plant are handled by the simulation manager FPoliSIM to ultimately lead to a realistic estimate of the event consequences. The level of sophistication of the analysis or physics tools of choice is up to the analyst and commensurate to the analysis goals. The analyst may choose to adhere closely to NEI 18-04 or follow alternate paths as deemed necessary to build the body of evidence that forms the safety case following principles of DID. Finally, the analysis results are collected and synthetized in digital reports.

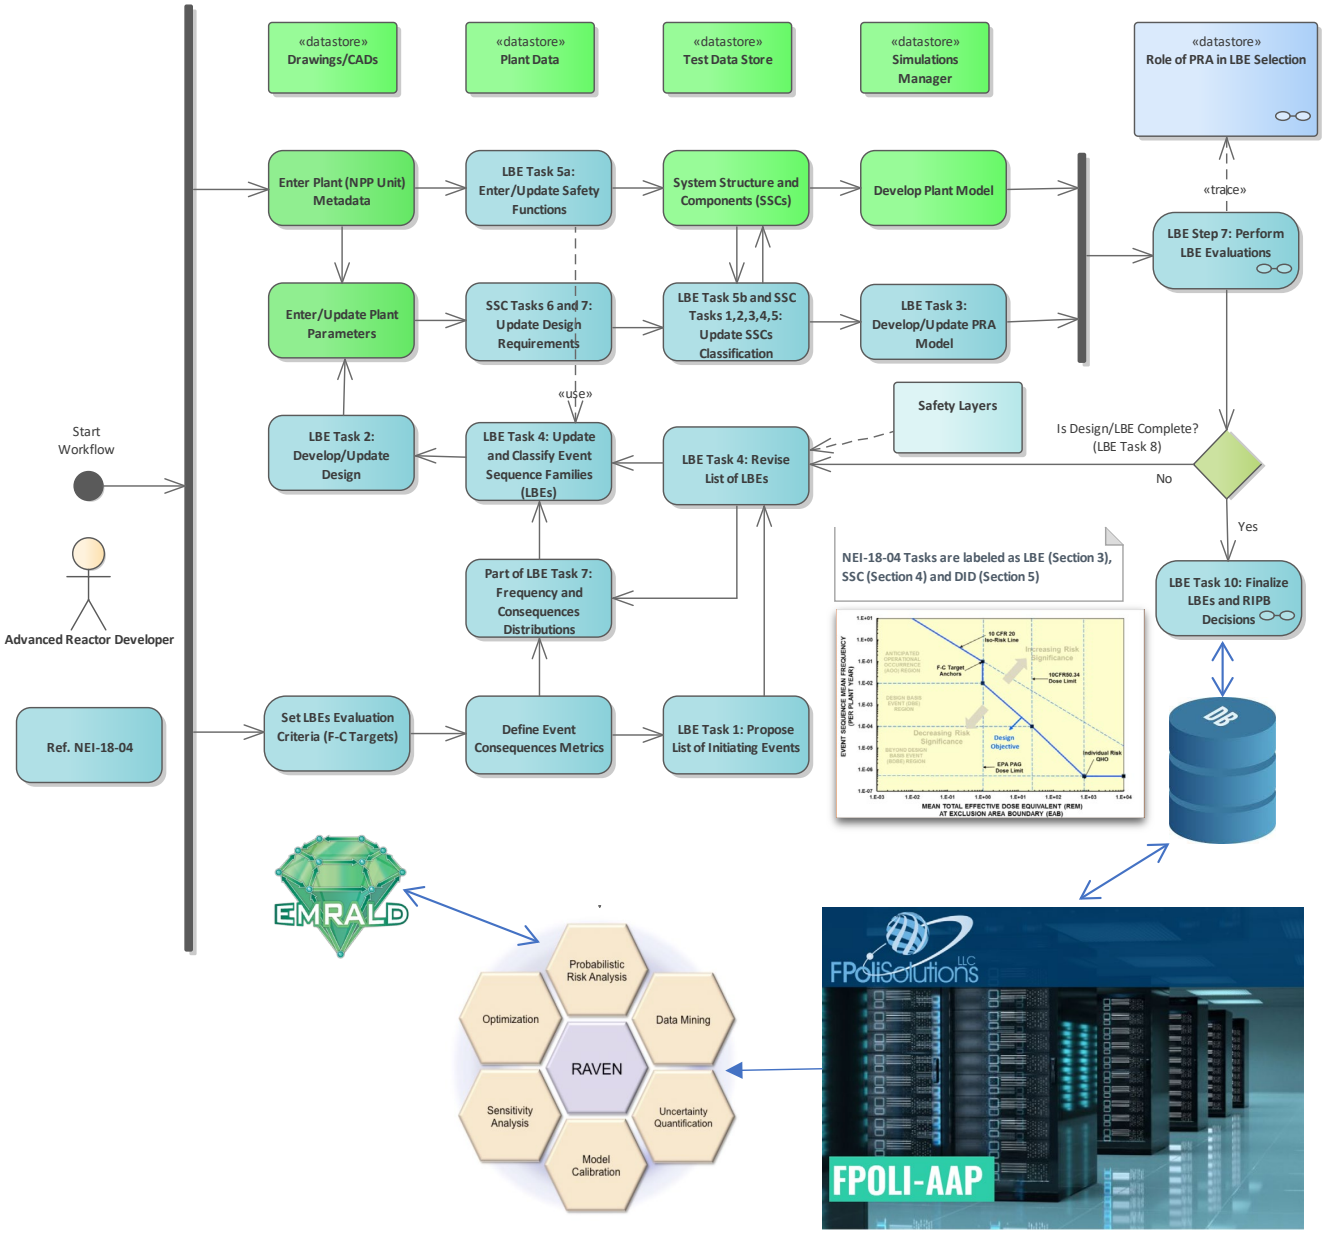

Figure 3-9 Risk-Informed System Engineering (RISE) application.

The Safety Case Manager has a comprehensive synthesis of the results through the RISE dashboard which acts as a wizard to trace, document and communicate such a safety case to regulators and other stakeholders. The platform is architected to streamline reviews before and during the licensing process by leveraging the scrutability and transparency of a browsable digital media rather than traditional flat reports. The evolution of a design can be easily tracked and maintained in the system. The RISE dashboard (Figure 3-10) is where all the actions are coordinated. The dashboard is a browsable site where a team can collaborate to visually construct the safety case, from the classification of the LBEs, the benchmark against risk metrics, SSCs classification and DID.

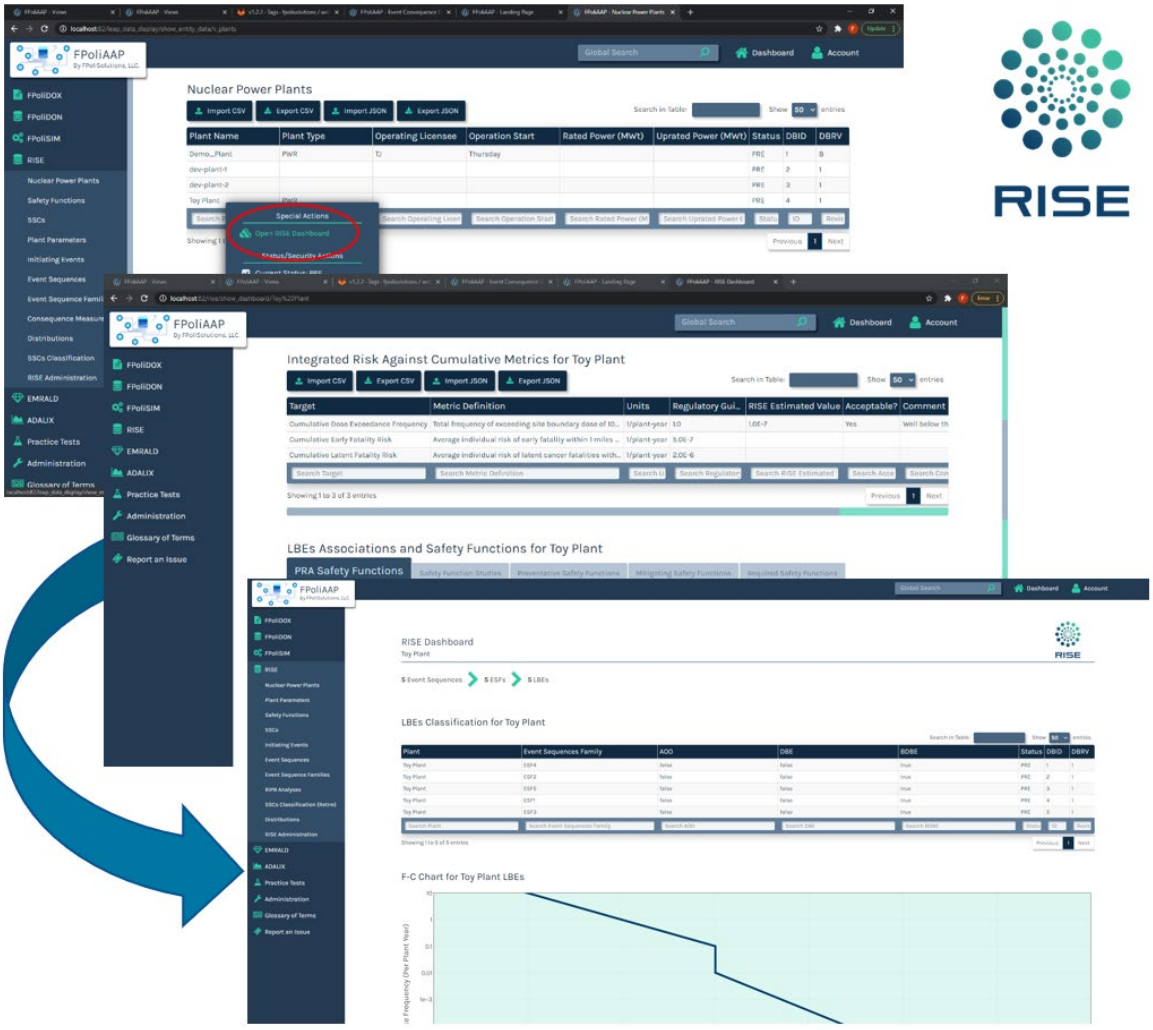

Figure 3-10 Risk-Informed System Engineering (RISE) application – the dashboard.

The RISE analyst relies on the powerful simulation manager, FPoliSIM, to manage the simulations and PRA analyses necessary to evaluate frequency-consequence profiles for the events. Uncertainties are tracked and handled along the workflow and/or deterministic approaches are adopted when needed. Once safety basis events are identified, the corresponding DBA are evaluated. The analysis of the DBAs should be consistent with RG 1.203, the Evaluation Model Development and Assessment Process (EMDAP). Part of the platform infrastructure with FPoliDOX, FPoliDON and FPoliSIM was architected to provide a digital representation of the EMDAP roadmap. The vision was presented in [67]. Finally, the equivalent of a SRP Chapter 15 analysis report can be generated to form the licensing basis for the safety analysis. Figure 3-11 is a synthesis of the architecture of the RISE application.

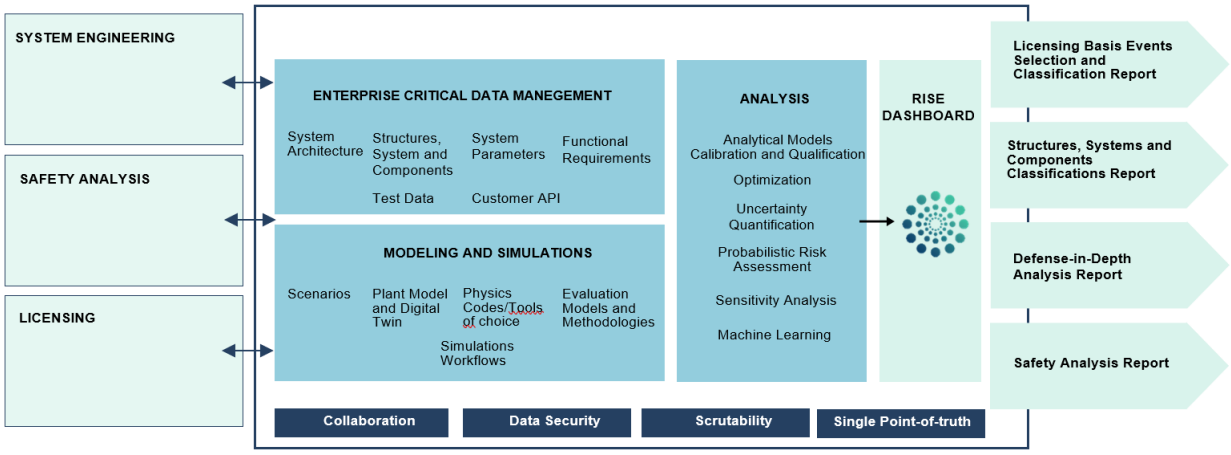

Figure 3-11 Risk-Informed System Engineering (RISE) application architecture.

3.6.5 Uncertainty Parameter Treatment Classification

The probabilistic digital twin of the plant needs to consider relevant uncertainty parameters and their influence on decision-making. Both PRA and plant response simulation model key parameters comes with uncertainties and a process must be established to handle their propagation throughout the analysis.

3.7 Categorization of Uncertainties

Based on modeling layer of the probabilistic digital twin, there are two primary areas of uncertainty analysis: The PRA analysis and the plant response simulations. Moreover, based on the NEI 18-04 process additional consideration for uncertainties is needed for the DID analysis. In summary, the following uncertainty considerations are:

1. PRA uncertainty analysis – This is focused on uncertainties that produce changes in event sequence frequency distributions, and it is guided by [63].
2. Plant response simulation uncertainty analysis – This is focused on uncertainties that produce changes in event sequence consequence distributions, and it is guided by [64].
3. DID evaluation of uncertainties – This is used to identify sources of uncertainty that are to be treated with additional DID layers.

These categorizations are discussed in the subsections below.

3.7.1 PRA Uncertainty Analysis

From [63], PRA analysis uncertainties are generally categorized as epistemic uncertainty or aleatory (also known as stochastic) uncertainties. PRA models explicitly address aleatory uncertainty which results from the randomness associated with the events in the model logic structure. The random occurrence of different IEs with subsequent failure of components to operate and human errors lead to many possible accident sequences that are accounted for in the event and fault trees used in a PRA model.

Epistemic uncertainties arise when making statistical inferences from data and, perhaps more significantly, from incompleteness in the collective state of knowledge about how to represent plant behavior in the PRA model. Following the earlier discussion in Section 2, epistemic uncertainties can be categorized into the following three types:

1. Completeness Uncertainty – Completeness uncertainty relates to uncertainty from risk contributors that are not accounted for in the PRA model. This type of uncertainty may further be categorized as either being known, but not included in the PRA model, or unknown.
2. Parameter Uncertainty – Parameter uncertainty relates to the uncertainty in the determination of the input parameter values used to quantify the frequencies and probabilities of the events in the PRA logic model. Examples of such parameters are initiating event frequencies, component failure rates and probabilities, and human error probabilities.
3. Model Uncertainty – Model uncertainty relates to the uncertainty associated with some aspect of a PRA model that can be represented by any one of several different modeling approaches, none of which is clearly more correct than another.

3.7.2 Plant Response Simulation Uncertainty

From [64], DBA analysis uncertainties are categorized less clearly. The sources of uncertainty are summarized in the Glossary as follows:

- The inaccuracy in experimentally derived data typically generated by the inaccuracy of measurement systems.
- The inaccuracy of calculating primary safety criteria or related figures of merit typically originating in the experimental data or assumptions used to develop the analytical tools
- The analytical inaccuracies related to approximations and uncertainties.

This list of uncertainties is largely related to uncertainties from evaluation model development. In application of an EM, there are other sources of uncertainty that must be considered. The following categories are used herein for plant response simulation uncertainties:

- Physics Model Uncertainties
- Operational Space Uncertainties

Special consideration must be given to methods to handle technical specifications. The definition of technical specifications which are intended to bound the operation envelope of an advanced reactor is an area that requires research, especially in the context of the risk-informed approach.

Similarly, to the PRA uncertainty parameters, a distinction between aleatory and epistemic warrant some considerations in developing the methodology to combine them.

3.7.3 DID Evaluation Uncertainty

These uncertainties will flow down from the PRA analysis and plant response simulation uncertainties.

3.8 PRA Analysis Method

An existing process for identifying and treating uncertainties in PRA analyses is described and endorsed within [63] (along with a series of accompanying EPRI reports). For this project, the guidance in [63 and the associated EPRI reports will be followed for the ground-up development of a process which conforms to [62], focusing on PRA analysis uncertainty treatment. Some additional goals of this exercise are to find potential gaps in the guidance, identify areas in the guidance where new methods can be developed, and add supplementary guidance where it is deemed helpful. In this section, the process from [63] is reviewed and adapted to the NEI 18-04 process. For this purpose, the discussion is broken into the stages from [63].

3.8.1 NUREG-1855 Process Stage A

Stage A in [63] provides guidance to both the licensee and the NRC staff on determining whether the approach for treating PRA uncertainties, as provided in [63], should be used for a risk-informed decision under consideration. The guidance involves determining whether the results from a PRA are used in the application and how the results are being used to support the decision. Because the NEI 18-04 process requires a PRA analysis, and this document is built around NEI 18-04, Stage A is largely irrelevant. As such, it will not be discussed further.

3.8.2 NUREG-1855 Process Stage B and C

The NUREG-1855 process is based on making changes to existing PRAs for existing plants. In this context, Stage B is used to determine if a change to a plant is already in the scope of the PRA analysis. If not, then this becomes part of the completeness uncertainty, which is to be evaluated in Stage C. The purpose of Stage C is to determine if a change not already covered in the PRA is a significant contributor to risk, and if so, the analysis may be refined to include it, or the change may be altered to be in the scope of the existing PRA analysis. In either case, both stages assume an existing PRA analysis for a single or small group of changes.

In the context of the NEI 18-04 process, Stage B and Stage C can be envisioned as a loop which is executed for every design feature of the advanced plant to build the PRA model from the ground up. Figure 3-12 shows Figure 4-1 from [5], which is an overview of Stage B. In essence, this stage is just determining if the existing PRA covers the scope of the feature.

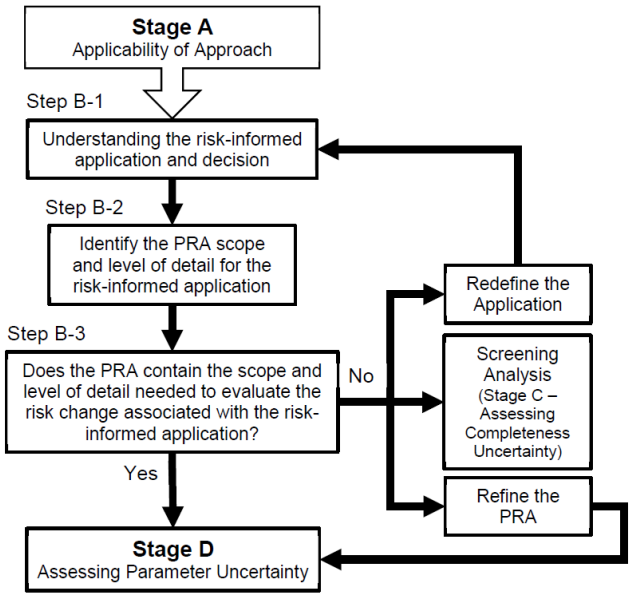

Figure 3-12 Figure 4-1 from [62] (overview of Stage B).

Figure 3-13 shows Figure 5-1 from [62], which is an overview of Stage C. This stage is basically determining if the feature is a significant contributor to risk, and then determining if the feature can be handled with additional parameter uncertainty or an upgrade to the scope of the PRA, or if the feature should be redesigned.

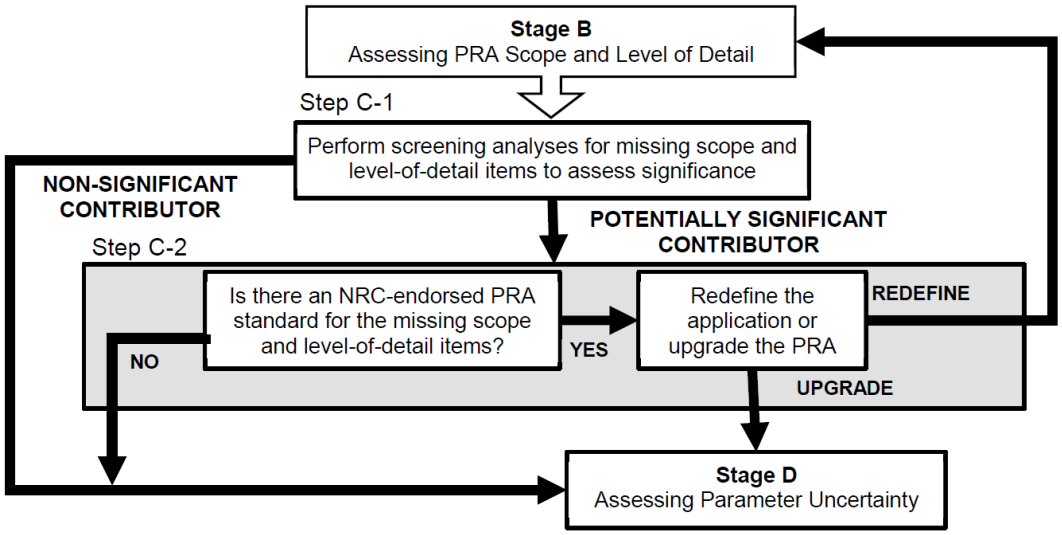

Figure 3-13 Figure 5-1 from [62] (overview of Stage C).

Repeating these stages for each design feature will result in building a PRA analysis which addresses the desired design features and provides an initial list of parameter uncertainties.

3.8.3 NUREG-1855 Process Stage D and E

Stages D and E are meant to quantify parameter uncertainty and model uncertainty. These steps can be performed largely as described in [62], but with the impacts limited to frequency distributions. Any uncertainties related to the consequence calculations should be pushed to the LBE simulations uncertainties.

3.8.4 NUREG-1855 Process Stage F

The purpose of Stage F is to help ensure that sufficient justification is provided for the acceptability of the risk-informed application. Further, the guidance for this stage helps ensure that the argument for acceptable justification is included in the documentation clearly and concisely. Much of this section remains applicable for application herein. The “compensatory measures” part of Stage F is comparable to the handling via DID adequacy. The “performance monitoring” part of Stage F may be transformed into requirements which must be developed into operating procedures in the eventual plant technical specifications.

3.8.5 NUREG-1855 Process Stage G

The purpose of Stage G section is to describe the process used by the staff for determining whether a licensee’s risk-informed application demonstrates an acceptable treatment of uncertainties and that the proposed application represents an acceptable risk impact to the plant. For the work herein, review under the NEI 18-04 process replaces the Stage G process; however, much of the guidance remains useful to consider.

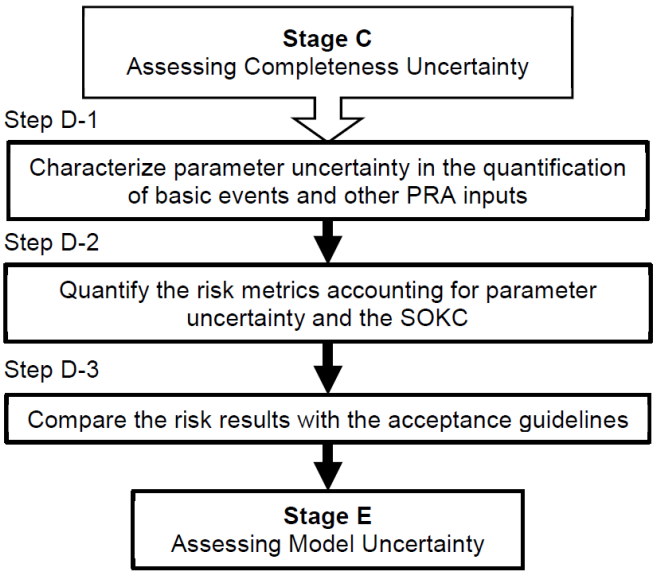

Figure 3-14 Figure 6-1 from [62] (overview of Stage D).

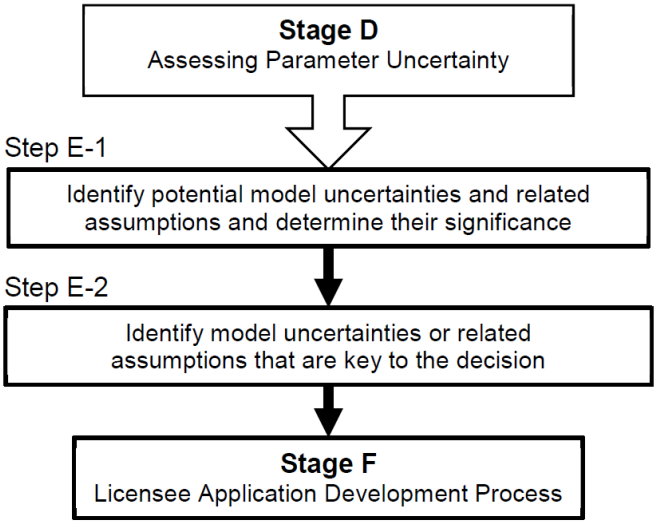

Figure 3-15 Figure 7-1 from [62] (overview of Stage E).

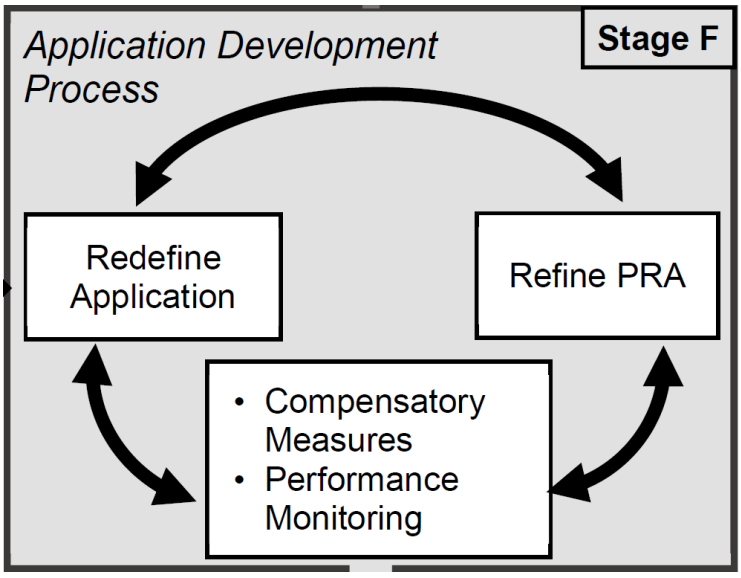

Figure 3-16 Figure 8-1 from [62] (overview of Stage F).

3.9 Plant Response Simulation Method

The first step in the plant response simulation uncertainty parameter treatment is to develop the initial lists of uncertainties for the simulation models. Some of these may be inherited from the PRA uncertainty analysis process, but the analyst should also attempt to determine other relevant uncertainties. This may be through use of a Phenomena Identification and Ranking Table, through review of the physics models/codes, or something else.

There are realistically functionally unlimited uncertainties in the analysis of any complex system; however, for practical reasons, explicit treatment must be reserved for the most significant contributors to the outputs of interest. Figure 3-17 shows an example flowchart for a proposed method of treating the uncertainties for the plant response simulations. In essence, the flowchart seeks to first perform dimensionality reduction by finding parameters which are amenable to bounding or nominal treatment. From there, the remaining uncertainties are categorized as plant parameter or model uncertainties, and epistemic or stochastic uncertainties. When the final convolution of uncertainties is performed, these categorizations will be helpful for choosing a specific method. Figure 3-17 is the basis for the remaining discussions in this section.

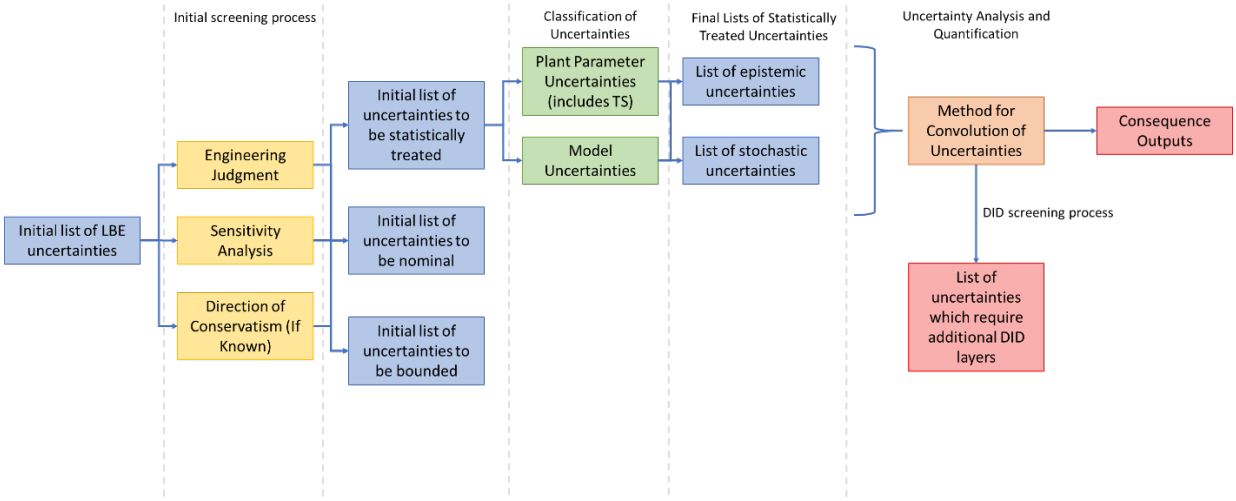

Figure 3-17 Flowchart of the uncertainty treatment for plant response simulations.

3.9.1 Initial Screening Process

An initial screening process is used to determine those parameters, which can be treated nominally or bounded, and which need to be carried forward for statistical treatment. This screening is performed with the following considerations:

- Direction of Conservatism – In order for a parameter to be bounded, by definition, the direction of conservatism must be known. If there are competing effects or interactions with other parameters, it may be impossible or difficult to predict which effect is more influential. In those circumstance assigning a bounding value could lead to misleading conclusions and should be avoided unless derived from sensitivities analysis, which capture all the potential interactions. In some situations, it may be more appropriate to determine the bounding value from engineering judgment.
- Engineering Judgment – This is a collection of techniques that allow an analyst to make categorizations and inference based on expertise and knowledge. Some considerations that may be used in engineering judgment are:
    - Magnitude of variation – If a parameter is known to vary within the resolution of a simulation code, this can be used to justify categorization. For example, in a systems code simulating the reactor coolant system of a PWR, the node volumes would technically be affected by manufacturing tolerances on the vessel components. However, these tolerances are miniscule compared to the volumes, so a nominal or bounded treatment can easily be justified.
    - Magnitude of effect – Similar to the magnitude of variation, a parameter that is known to have minimal effect, regardless of the size of variation, this can be used to justify categorization. For example, for many transient analyses in a PWR, there are only several operating conditions that strongly influence the calculation, and the rest are set to nominal values.
    - State of knowledge – If the effects of a parameter are not well-known or the magnitude of variation is unknown, then it can be very difficult to justify nominal or bounding treatment. As such, parameters with a low state of knowledge are likely to be treated statistically.

- Sensitivity Analysis – If there is insufficient knowledge to determine conservativism basis from engineering judgment, but this knowledge may be readily obtained from running simulations, then sensitivity analyses are a good approach to inform treatment.

3.9.2 Classification of Uncertainties

From the initial screening process, uncertainties that are to be statistically treated are then classified as follows:

- Operational Space Uncertainties – These are uncertainties related to the operating ranges utilized in the plant. Things such as core power, operating temperature, operating pressure, and many others. These parameters establish initial conditions and influence the timing of phenomena and setpoints in the transients. This category also includes measurement bias and uncertainty and any other tech spec type parameters
- Physics Model Uncertainties – These uncertainties may include:
    - The inaccuracy in experimentally derived data typically generated by the inaccuracy of measurement systems.
    - The inaccuracy of calculating primary safety criteria or related figures of merit typically originating in the experimental data or assumptions used to develop the analytical tools
        - The analytical inaccuracies related to approximations and uncertainties.
        - This can be adequately summarized as uncertainty resulting from inherent imperfection in the modeling, quantification, and validation of physical phenomena.

3.9.3 Final Lists of Statistically Treated Uncertainties

From the parameters classified in the last section, they are then divided into the following categories:

- Epistemic Uncertainties – This is the uncertainty related to the lack of knowledge about or confidence in the system or model and is also known as state-of-knowledge uncertainty. This uncertainty can be thought of as due to things one could in principle know but do not know, in practice.
- Stochastic (or Aleatory) Uncertainties – This uncertainty is based on the randomness of the nature of the events or phenomena and cannot be reduced by increasing the analyst's knowledge of the systems being modeled.

This classification is important for the uncertainty analysis. In general, it is acceptable to treat stochastic uncertainties with simpler methods, such as Monte Carlo sampling. However, for epistemic uncertainties, the treatment is generally more rigorous and may involve more imposed conservatism.

3.9.4 Uncertainty Analysis and Quantification

The final convolution of uncertainties is used to provide a probability distribution on the consequences of a given event sequence family. This will require explicit methods for the treatment of epistemic and stochastic uncertainties. As this method is finalized, it is envisioned that several specific methods for both epistemic and stochastic uncertainties will be presented.

3.10 Case Study 1: Simplified PWR Station Blackout

3.10.1 Scenario Description

The scenario of interest in this demonstrative study is a hypothetical (i.e., generic model not representing any actual power plant) PWR subject to a station blackout (SBO) initiating event. In this sample SBO, the PWR features SSCs engineered to perform critical functions that support the main SF of maintaining a sufficient coolant inventory for removing the decay heat from the core and ultimately preventing core melt through a feed and bleed operation. The SSCs considered in this analysis are the reactor system and the rest of components that form the Nuclear Steam Supply System (NSSS), and the following engineered safety features:

1. Low Pressure Safety Injection (LPSI) pumps powered by Diesel Generators (DGs)
2. Low Pressure Safety Injection (LPSI) powered by diverse and flexible coping strategies (FLEX) pumps
3. Refueling Water Storage Tank (RWST)
4. Sump recirculation

The LPSI is designed to maintain sufficient inventory of liquid in the reactor vessel while steam is generated from the removal of the decay heat and assumed vented to the containment. The possible scenarios are represented by the even tree depicted in Figure 3-18 which leads to five event sequences, ES1 through ES5.

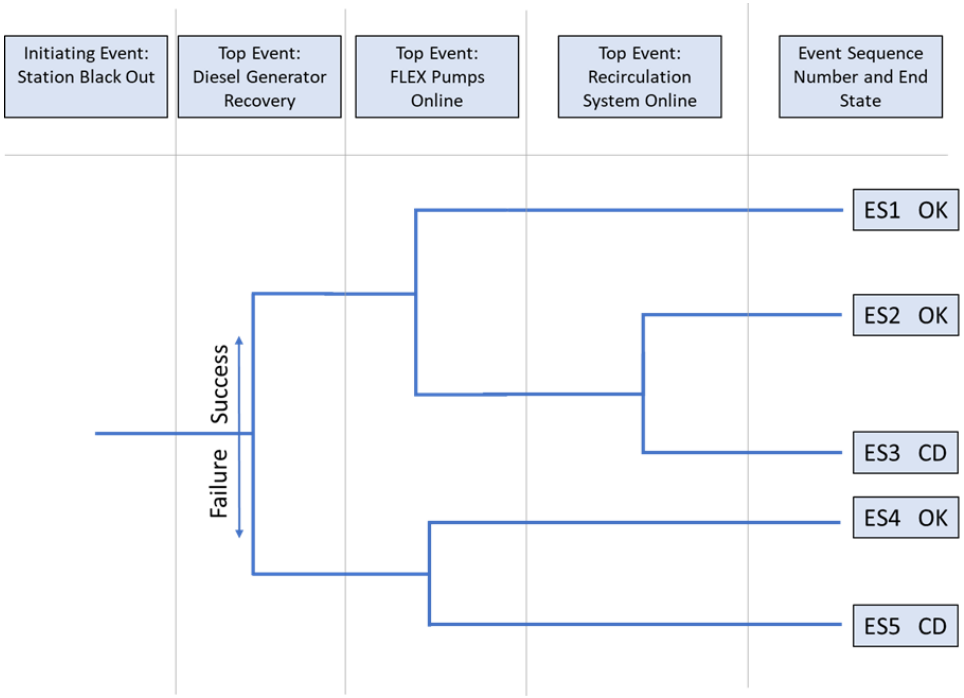

Figure 3-18 Event tree for the SBO case.

It is noted that two of the ESs end in core damage and have no additional mitigating safety features. However, they fall into the BDBE range. While it is unrealistic that there are no mitigating systems being used, but an acceptable dose is achieved, it is noted that this limited scope demo is assumed to have other SFs (potentially many) that are simply not explicitly considered herein.

For this demonstrative problem, a simple evaluation model was developed to represent the NSSS under the assumed SBO scenario. The evaluation model was built with RELAP5-3D and like the simple model presented in [69]. The entire NSSS is represented by a single node. The volume contains a heat source to describe the core. The heat structure produces nuclear power using the point kinetics model, which provides a realistic representation of decay heat during the transient. The safety injection is represented by a flow boundary conditions connected to the volume and the venting (bleed) is represented by a pressure boundary condition at the top of the volume. A combination of valves and control variables mimic the behavior of the safety systems (LPSI Pumps, DGs, FLEX, etc.). As the fuel rod model is extremely rudimentary, a simple success criterion was set to be the mixture level covering the core. A first-principles level swell model was included to predict the location of the mixture level without reliance on detailed modeling of the vessel, this was done to keep the model generic and not representative of any actual facility.

A basic set of simulations is performed, where the safety injection flow is assumed to begin at 5 hours (when it succeeds), the FLEX flow is assumed to begin at 5.5 hours (when it succeeds), and switchover to recirculation is assumed to begin at 10 hours and last for 1 hour (when it succeeds). The liquid fraction results of the five event sequences are shown in Figure 3-19, which shows that the ES-1, ES-2, and ES-4 succeed, while ES-3 and ES-5 cause core damage.

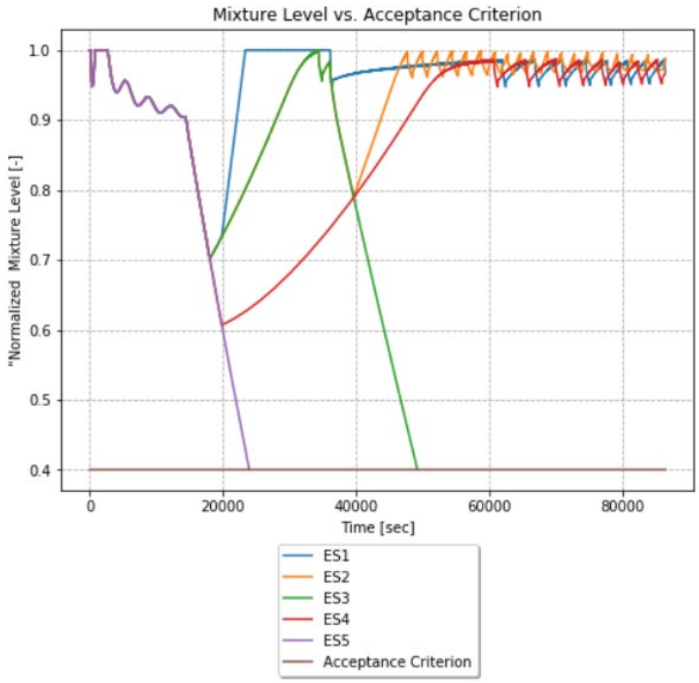

Figure 3-19 Event sequence simulation results.

3.10.2 Uncertainty Analysis

There are two uncertainty analyses performed for this scenario. The first relates to the PRA and the second relates to the plant response simulations. For the PRA uncertainty analysis, the fault and event tree models in SAPHIRE are provided with parameter uncertainty ranges, and a Monte Carlo analysis is performed, which produces frequency distributions for each of the ESs. For the uncertainty analysis of the plant response simulations, a Monte Carlo analysis of the RELAP5-3D simulations is performed, and the

results are post-processed into a consequence distribution. A simplified method of dose calculations is performed which assumes a low dose rate prior to core damage, and then a significantly increased dose rate after core damage. As such, the dose for each sequence will be a function of time of core damage. As a result of running many cases in a Monte Carlo analysis, a probability distribution of dose can be obtained.

3.10.3 PRA Uncertainty

There are three main categories of PRA uncertainty:

- Completeness Uncertainty – This uncertainty step is represented by the process of developing the PRA analysis from the ground up. So, when the PRA analysis begins, it has no SFs modeled, and thus in our example, repair of a DG, deployment of FLEX pumps, and initiation of recirculation could be thought of as completeness uncertainty. However, once the PRA analysis is built, this completeness uncertainty is gone, but likely replaced with some combination of parameter and model uncertainty. In this sense, completeness uncertainty can be thought of in the context of NEI-18-04 as a means to track intended PRA scope that has not been implemented yet.
- PRA Parameter Uncertainty – Every parameter in the SAPHIRE model has some parameter uncertainty assigned to it. As an example, the DG has a failure to start basic event, with a nominal failure probability, and a beta distribution defining the uncertainty range.
- PRA Model Uncertainty – Model uncertainty will largely be represented in the plant response simulations. An example (not necessarily considered in this case) of a PRA model uncertainty is one where failure modes are not completely understood or are unknown, such as in the failure of digital instrumentation and controls.

With the various parameter uncertainties, a Monte Carlo simulation was run in SAPHIRE, and the resulting event sequence frequency distributions are shown on the F-C plot in Figure 3-22. Note that all parameters affecting the analysis, both the PRA and ESs evaluations are databased, including the uncertainty attributes. An example on how such data is represented in the RISE database is shown in Figure 3-20.

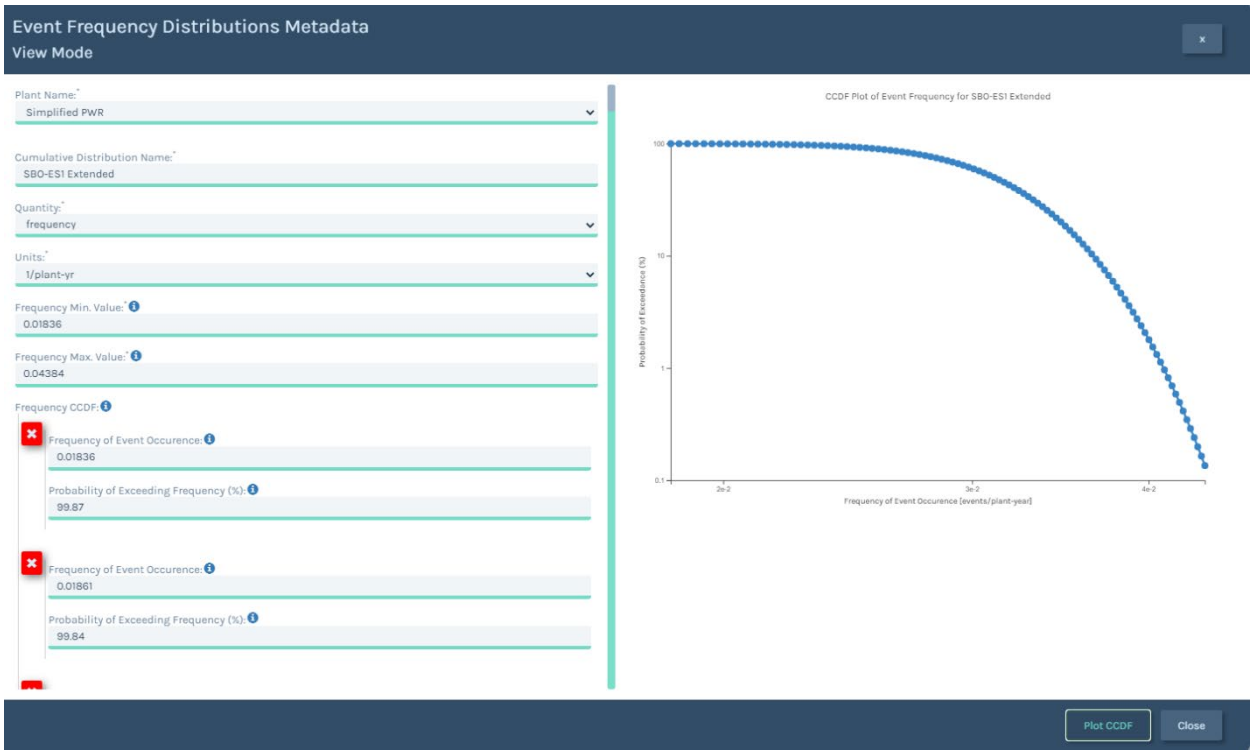

Figure 3-20 Example of frequency CCDF for an ES.

3.10.4 Plant Response Simulations Uncertainty

The following classes of uncertainty are sampled in this analysis:

- Bounded Operational Parameter Uncertainty – The core power is chosen as a bounded operational uncertainty. This is because it is a high impact uncertainty with a known direction of conservatism. Thus, it is set to 102% of the nominal power.
- Nominal Operational Parameter Uncertainty – The pressure of the RCS is set to a nominal value. This is because it is a relatively low impact value with minimal variation.
- Stochastic Operational Parameter Uncertainty – The SI/FLEX flows are chosen as a stochastic operational uncertainty, because it has a relatively large uncertainty band, but is also high impact, as the flow rates determine if there is enough liquid to keep the core covered. Therefore, it is set to sample nominal +/- 20%.
- Stochastic Operational Parameter Uncertainty – The Temperatures of the SI/FLEX flows are chosen as a stochastic operational uncertainty, because it has a relatively large uncertainty band, but is also relatively high impact, as the subcooling affects how much energy is removed with a given amount of liquid injected. Therefore, it is set to sample nominal +/- 20%.
- Epistemic Physics Model Uncertainty – The decay heat model is considered an epistemic physics model uncertainty in this case. It is chosen to be nominally treated with the ANS 1979 decay heat model in RELAP5-3D.

With the various parameter uncertainties, a Monte Carlo simulation was run in FPoliSIM, and representative event sequence consequence distributions are calculated. It is noted that while an uncertainty analysis was run, the consequence results were tweaked to provide a better demonstration. An example on how such data is represented in the RISE database is shown in Figure 3-21.

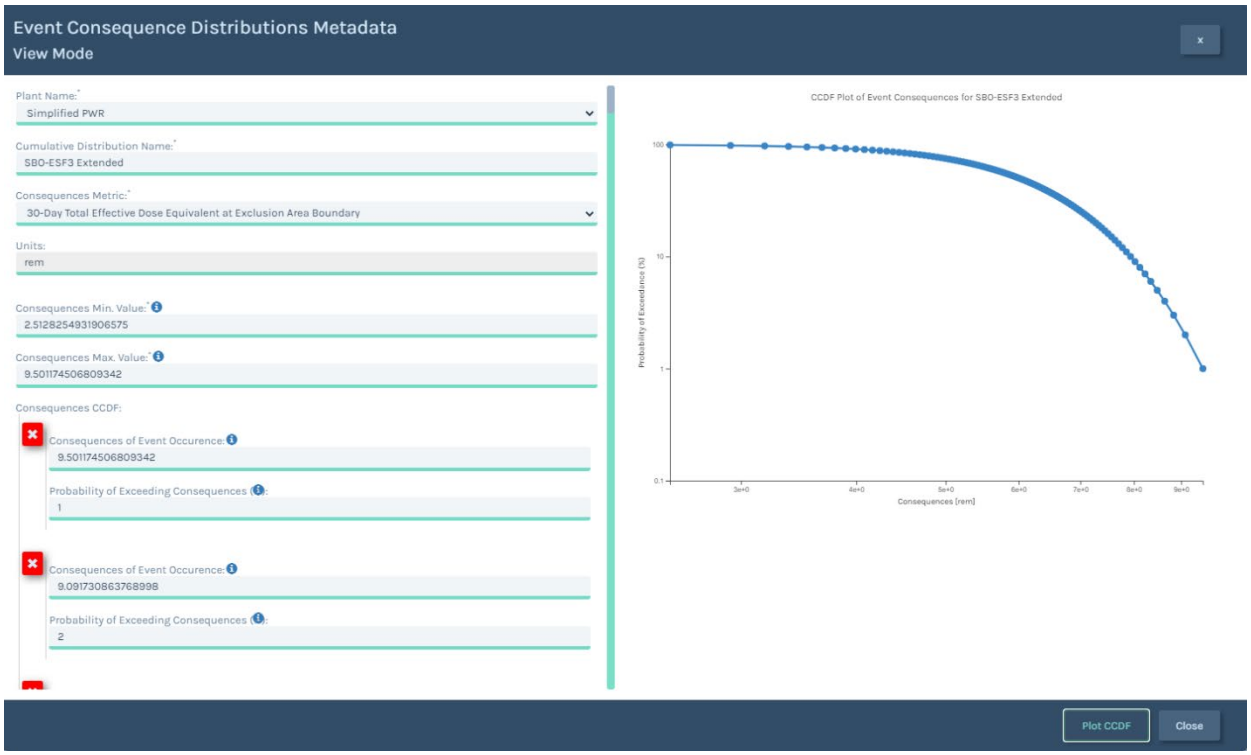

Figure 3-21 Example of Consequence CCDF for an ESF.

3.10.5 Entry into RISE

The analysis is orchestrated through the RISE dashboard (Figure 3-22). The dashboard is where the Safety Case Manager administers the activity associated with the definition of the safety case. The RISE dashboard was architected to provide consistency with the draft 10 CFR Part 53, whose roadmap is described in the NEI 18-04. However, the workflow is suitable to a graded approach in which the role of the PRA analysis can range from design validation step at the end of the design cycle to a fully coupled PRA embedded in design process itself.

From the frequency distributions, the system automatically identifies the LBEs and classifies them in AOO, DBE and BDBE. In this simple demo case, all five ESFs make it into the final LBEs selection. However, for a typical application, there could hundreds ESFs that funnel down to 20-30 LBEs. From the consequences, the LBEs can be displayed in the Frequencies-Consequences (F-C) chart. In this example the acceptance criteria are based on NEI 18-04.

For those results the integrated risk against cumulative metrics defined by the user are computed and displayed in the table below the F-C chart. Next, a multi-tab table displays the LBEs associations to the SF.

The system first filters and lists the applicable PSF. In the following two tabs, the user can view which ones were identified as “Preventative SFs” for at least one ESF and which one were set to “Mitigating SFs” for at least one ESF. The system then queries the associations in the database and lists which LBE was prevented or mitigated by the SF.

For the demonstration analysis, two PSFs are specified: “RCS Inventory Control” and “Provide Water after Depletion of the RWST.” Both functions are derived from the Fundamental Safety Function (FSF) of “Removal of heat from the reactor and from the fuel store.”

Figure 3-22 RISE Dashboard.

In the tab ‘Safety Function Studies’ the user can organize studies to identify which SFs are required and which SSCs are risk-significant. This is typically performed via sensitivity studies (simulations) as described in [61]. From that information the analyst can determine which SFs are required and which SSCs are risk-significant.

To demonstrate the selection of RSF, a DBE must be selected to perform the sensitivity study. ESF 2 is used for this purpose. In this study, the RCS Inventory Control is removed from consideration. This functionally results in an F-C space with frequency distribution based on ESF 2, but consequences based on ESF 5. This is because in ESF 2, RCS inventory control is powered by the DG. Changing this assumption results in a simulation consistent with ESF 5. As such, the consequence results of ESF 5 can be used directly. The results of this study are shown in Figure 3-23. From this figure, it is observed that ESF 2 now violates the F-C Target. This results in the definition of RCS Inventory Control as a RSF. A similar study results in the classification of Provide Water after Depletion of the RWST as a RSF as well (in this study, ESF 2 is again used as the base, and the consequences now come from ESF 3).

The next table organizes the SSC classification. The SSCs that perform SFs are classified as SR, NSRST and NST. From the RSFs, one SSC that provides the SF must be declared a SR SSC. For this exercise, the DGs and SI pumps are chosen as the SSCs which provide RCS inventory control, and the recirculation system is chosen as the SSC, which provides water after depletion of the RWST.

Once the SR SSCs are chosen, the remaining SSCs must be checked for risk significance. In this case, the only remaining SSC is the FLEX pumps. The risk-significance is determined from a sensitivity study which is like the study considered for determining the RSFs. It is noted that there are two key differences in these studies. The first is that it the analysis is not limited to DBEs but must be extended to all. The second is that this study is removing a specific SSC providing the SF, not the entire SF (and thus all SSCs which provide that SF). In practice, in this case, these differences are not of particular importance, but they are noted for clarity. In this study, ESF 4 is used as the base, and the consequences now come from ESF 5. The resulting F-C space is over the target line, and thus the FLEX pumps are designated as NSRST.

In this demonstration, no examples of NST SSCs are included, but they would simply be SSCs whose SF removal does not result in violation of the F-C Target.

At this point the analyst has all the information required to determine the set of DBAs. According the [64] the DBAs are derived from the DBEs, but are performed crediting only the SR SSCs needed. A DBA is essentially a variation of the associated DBE where only SR SSCs are available to mitigate the postulated ESFs to within the 10 CFR 50.34 dose limits. The list of DBA is presented in the last table of the dashboard.

The next step in the process is requirements derivation. Based on the PRA analysis and simulations, as well as the SSC classification, RFDCs, SRDCs, and Special Treatment Requirements must be derived. Many of these requirements, such as maintenance, testing, and quality assurance programs, are beyond the scope of this demonstration. However, there are some requirements that are relevant to the performance and availability of SSCs and are shown below.

The following are examples of SRDCs:

- SI Pump Minimum Flow
- DG Maximum Repair Time
- DG Minimum Probability to Start after Repair
- Sump Recirculation Switchover Maximum Time

The following are examples of Special Treatment Requirements:

- FLEX Pump Minimum Flow
- FLEX Maximum Deployment Time

The RISE[f] Dashboard has not yet been completed for the scope of DID adequacy.

Finally, the safety case manager can automatically generate the key reports that define the safety case.

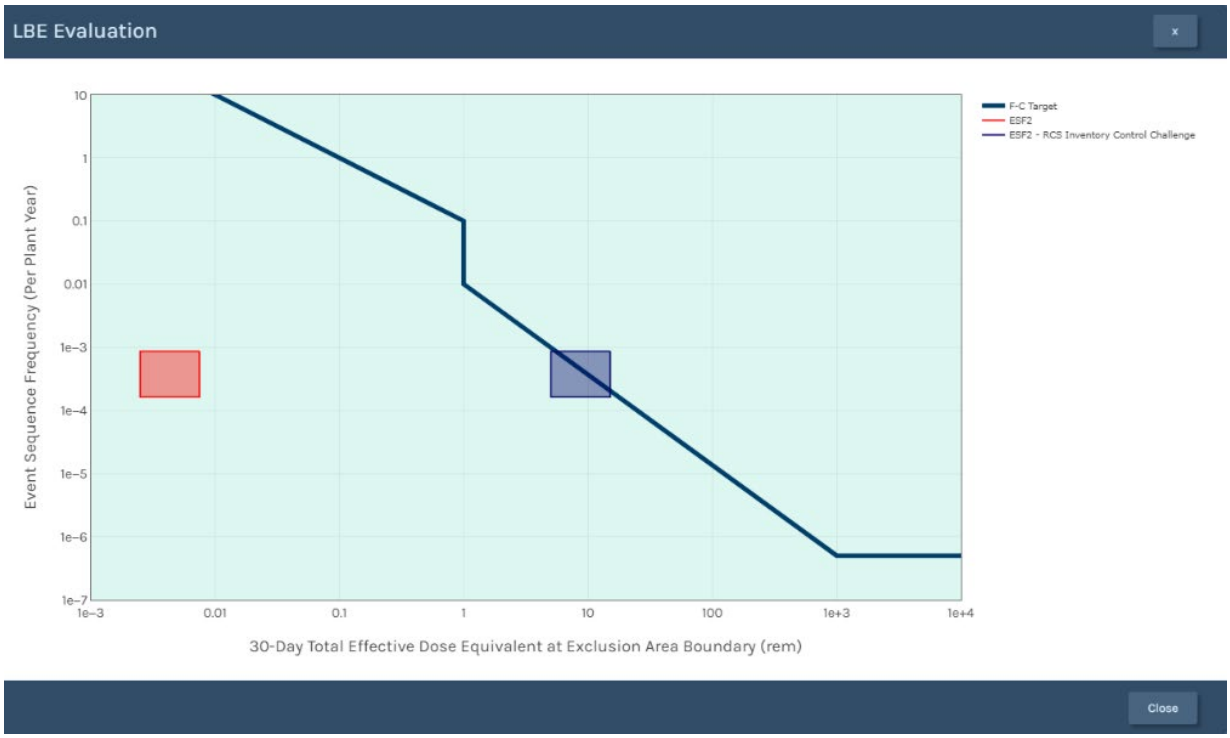

Figure 3-23 F-C results for the sensitivity studies.

[f] FPoli deployed a demonstration of RISE application and other FPoli-AAP services on INL HPC enclave. Please contact FPoliSolutions for further information on its usage.

3.11 Case Study 2: Generic Small Modular Reactor Scenario

One important aspect of new reactors is an enhanced interest or attention to physical security scenarios. As these scenarios are composed of actors intending to inflict harm in unpredictable and reactive ways, traditional, static analysis methods are not necessarily well-positioned to quantify the associated risks. In other words, the uncertainties associated with a physical security scenario and its effects on plant states are large and require the consideration of novel approaches.

In this section, a hypothetical approach to the analysis of a physical security scenario is performed based around Event Modeling Risk Assessment using Linked Diagrams (EMRALD), Neutrino, RAVEN, and RELAP5-3D. The workflow is built to be managed with the FPoliAAP technology to promote ease of use. The plant model used was created by the team and does not represent any actual facility or attack target.

3.11.1 Scenario Description

The generic scenario is that there is a representative small modular reactor (SMR) plant with events that result in a water jet or flooding which could damage or disable other systems. The loss of these systems, combined with an existing event or another failure, could result in an unsafe plant state. This scenario was created to demonstrate the analysis platform and does not represent any real scenario.

The considered technology framework is:

- EMRALD as the workflow manager and uncertainty analysis engine. EMRALD can interact directly with Scribe-3D or Avert, Neutrino, and RAVEN/RELAP for 2-way coupling.
- Scribe-3D or Avert are used to simulate force on force scenarios.
- Neutrino is used to simulate pipe break flows and flooding of rooms or compartments.
- RAVEN/RELAP are used to simulate the RCS and determine if core damage is experienced.

In the eventual solution, all these tools would be connected via the FPoliAAP-EMRALD simulation framework, which would manage the simulations with 2-way coupling and get the most detailed responses. However, at this stage, a proof of concept was done with the following simplifications/assumptions:

1. No actual force-on-force simulation will be done, but instead, distributions will be developed to represent results that might be obtained from a simulation.
2. EMRALD and Neutrino run with two-way coupling to simulate scenarios using the probability distributions from part 1. These simulations are then used to determine probability and timing distributions of flooding safety systems.
3. The distributions from parts 1 and 2 are used as input into a dynamic event tree using RAVEN/RELAP.

The scenario is described as follows. An event occurs, with the objectives to damage offsite power source to begin the event. Potential damage may occur to other structures such as the hypothetical auxiliary building and the feedwater lines. It is assumed that the loss of the offsite power source alerts the site and allows security forces to begin responding (again, no detailed force-on-force modeling was used in this demonstration case). In addition, the feedwater lines were assumed to be accessible due to easier access that we created in the generic model.

Once the security forces are mobilized, they will engage the attackers to prevent them from reaching their objective. However, it is assumed for simplicity, we ran the example case assuming that the attackers eventually make it to their object, just that the time delay may be significant (again, not a realistic scenario).

Once the feedwater piping is damaged, there is another concern of flooding the computer systems that control safety injections systems. It is noted that it may not be realistic for the safety injection systems to be flooded by a break in the feedwater line; however, as no specific design is being considered, this is meant to be a demonstration of the technology stack, and not a demonstration of a purely realistic scenario. This facility used for the model does not actually exist and may not represent any actual working system.

With the feedwater piping damaged, the attackers hold position and interfere with plant recovery actions. Therefore, the security forces must secure routes through which operators can interact with the plant. With the destruction of the feedwater system and the safety injection pumping equipment, the credited actions are the transport and deployment of hypothetical FLEX pumps for reactor coolant system makeup. Once these pumps are deployed, flows to the reactor coolant system can resume, and the calculation will generally end in successful mitigation.

3.11.2 Modeling

3.11.2.1 Force-on-Force Modeling

At this early demonstration stage, there is no simulation of this, but in theory, the simulation would have the following:

- It is assumed that the attack is discovered when the offsite power source is destroyed.
- Once the offsite power source is destroyed, the attackers move towards the feedwater line in the auxiliary building (there is an arbitrary assumption that the reactor building would be better defended, so they go for this part of the facility).
- There is some amount of time until they successfully destroy the feedwater line, then they hold their position to prevent intervention, which includes prevention of accessing FLEX equipment or hooking it up.
- At some point, the security forces secure the site, and FLEX equipment can be used.

The output of this simulation is essentially a timing distribution when feedwater lines are successfully destroyed by the attackers, and a timing distribution for securing the site. The exact nature of the destruction feedwater lines is unknown, and thus it could vary from a split break to a full double-ended rupture.

3.11.3 EMRALD/Neutrino Run

Figure 3-24 and Figure 3-25 show the EMRALD model of this simulation. Note that because the current implementation of EMRALD in FPoliAAP enterprise system does not support two-way coupling, this simulation was performed in two separate simulations. First, an EMRALD/Neutrino simulation was performed to determine timing of just the flooding portion. This was then characterized as a normal distribution and entered back into EMRALD. Then a separate EMRALD simulation is executed as part of RISE workflow which queries the distribution that was stored in the database after EMRALD simulation.

The simulation begins with the offsite power being lost as an assumption. Next, the simulation determines the time required for the damage of the feedwater. In this initial demonstration, this time will be set to a normal distribution with a mean of 1200 seconds and a standard deviation of 400 seconds (note there is no technical basis for this time, it was just created as a hypothetical time to demonstrate the use of the software infrastructure). The destruction of the feedwater is tracked as a key state.

Once the feedwater piping is destroyed, A key state is entered for the purpose of tracking the time of the destruction of the feedwater. In addition, in the eventual solution, the Neutrino simulation would be dispatched. For this initial demonstration, the EMRALD/Neutrino simulation is performed with an assumed nominal pressure. This simulation samples the size of the break in the feedwater lines, which

strongly influences the time required for liquid accumulation to reach the height that is considered to disable the safety injection. Once Neutrino calculates that the safety injection systems are flooded, a key state is entered to track this timing. Additional details on the Neutrino and EMRALD two-way coupling model are given in Appendix A. The results of this simulation are then loaded back into a normal distribution in the main EMRALD simulation.

The Neutrino model used for this simulation is an auxiliary building for a representative SMR. This building is assumed to house the turbine, feedwater piping, and some sort of electrical boxes that control safety systems (among other miscellaneous components). The Neutrino model is shown in Figure 3-26.

Once the feedwater piping is destroyed, in addition to the start of the Neutrino simulation, there is a time distribution for securing the site. While in the eventual solution, this would be determined with a continuation of the coupled force-on-force simulation, in this initial demonstration, this will be set to a normal distribution with a mean of 2500 seconds and a standard deviation of 750 seconds (note there is no technical basis for this time, it was created as a hypothetical time to demonstrate the use of the software infrastructure). After the site is secured, there is another distribution for the time to deploy the FLEX equipment. This is assumed to vary significantly, as damage to the site that may impede actions may occur. While in the eventual solution, this would be determined with a more complex EMRALD simulation, in this initial demonstration, this will be set to a normal distribution with a mean of 1800 seconds and a standard deviation of 600 seconds (there is no technical basis for this time, it was created as a hypothetical time to demonstrate the use of the software infrastructure). The deployment of FLEX equipment is tracked as a key state.

Ultimately, the EMRALD simulation produces timing distributions for the three key states: feedwater destruction, the timing of loss of safety injection based on flooding, and the timing of FLEX being hooked up based on securing the site.

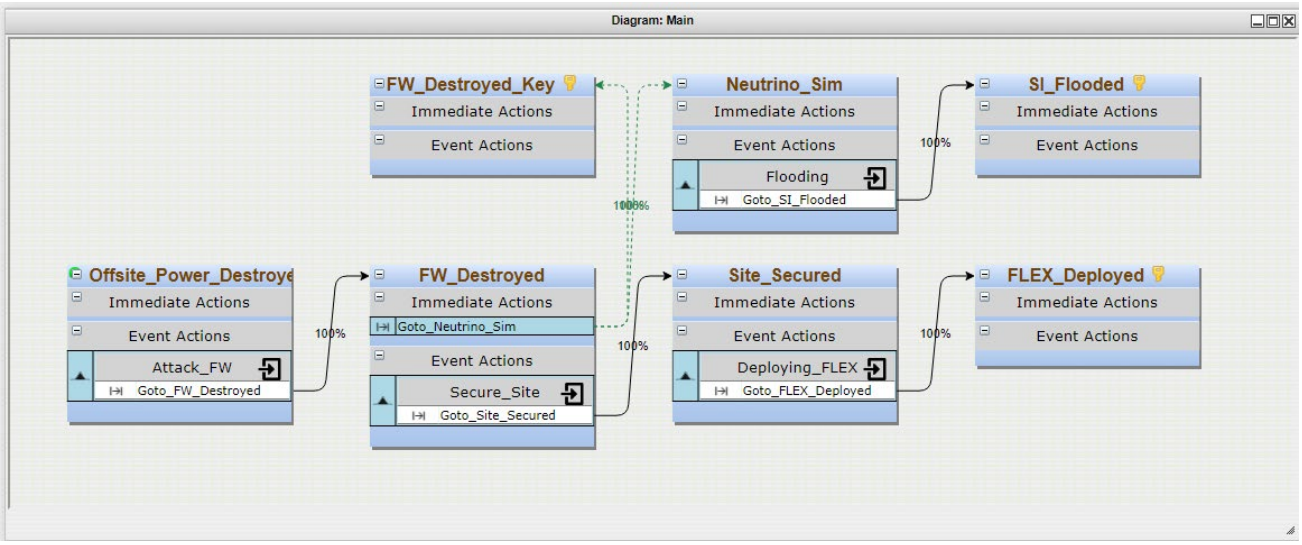

Figure 3-24 Main EMRALD diagram.

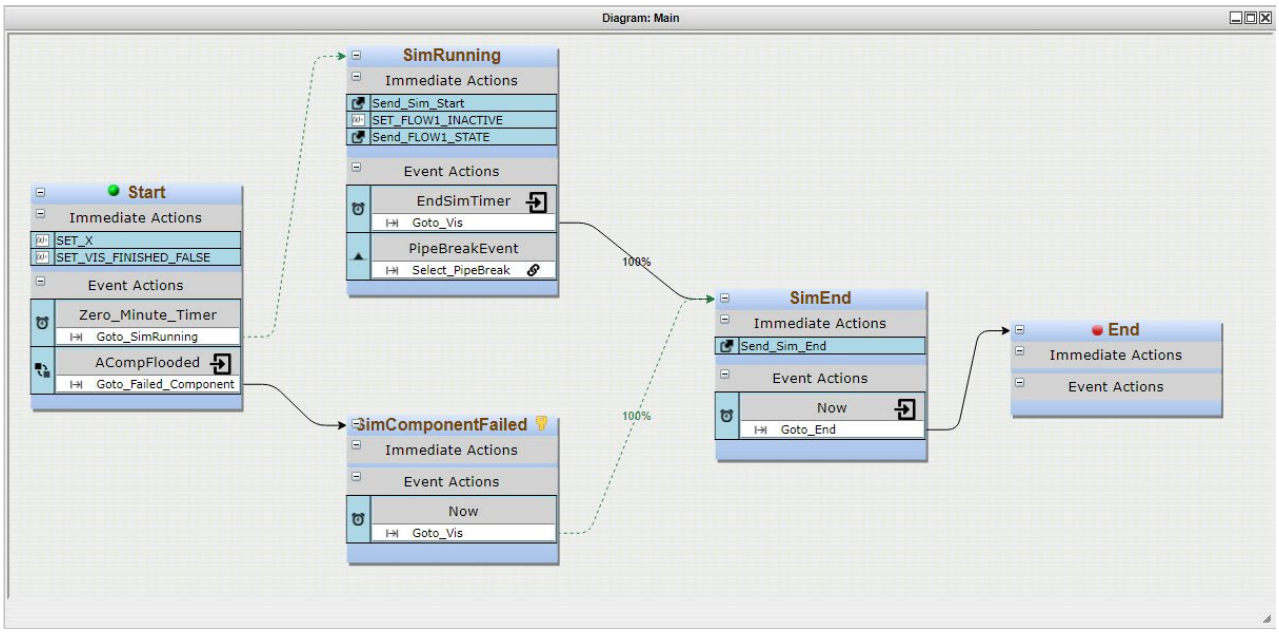

Figure 3-25 EMRALD diagram which runs Neutrino.

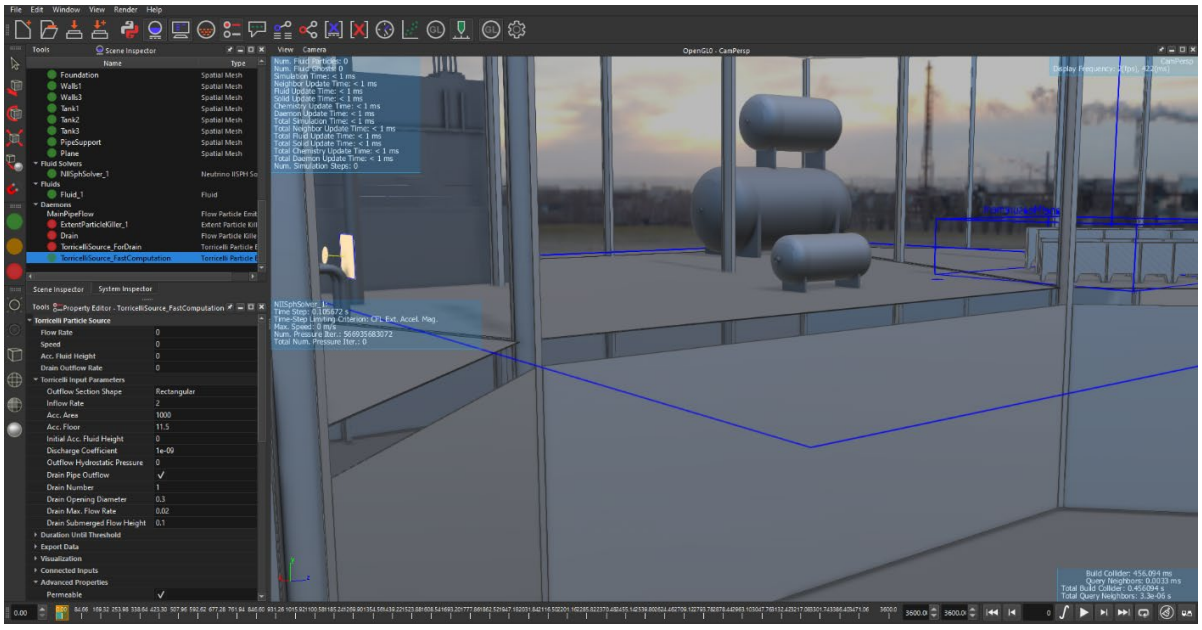

Figure 3-26 Neutrino model for the hypothetical SMR.

3.11.4 RAVEN/RELAP DET

The simulation manager, FPoliSIM, orchestrates the simulation of the dynamic event tree through RAVEN. RAVEN simulates the dynamic event tree with RELAP for this scenario. The RELAP model, shown in Figure 3-27, is a model of a hypothetical 350 MWt pressurized SMR. The simple model consists of a vessel containing a downcomer, lower plenum, core, and upper plenum; a hot leg with pressurizer; a steam generator with secondary side modeling from feedwater to steam valve; and a

crossover leg which splits into two pumps and cold legs. The model also has control logic to achieve a balanced steady state.

A transient is modeled to represent the scenario described in this section. The initial transient begins by modeling a loss of offsite power (LOOP). Starting from the LOOP initiating event, there are three main ESs that are driven by the dynamic event tree. The first is the destruction of the feedwater piping. This results in a loss of auxiliary feedwater in the RELAP simulation. After the loss of feedwater, it is assumed that feed and bleed is initiated as a backup cooling method, consistent with usual station blackout procedures. Feed-and-bleed usually consists of feeding the RCS with water via the safety injection system, and bleeding off pressure via the pressurizer relief valves.

The next major event is the flooding of the safety injection system. This results in a loss of safety injection flows in the RELAP simulation. At this point, the simulation will begin boiling off RCS inventory and eventually, core uncover will occur. At this point, core heatup will begin, and an acceptance criterion for core damage is set at 1500 K, based roughly on the limits used for loss of coolant accident analyses in 10 CFR 50.46.

The final event in the dynamic event tree is the initiation of flows from FLEX equipment. It is assumed that these flows are the same as those for the safety injection, for simplicity. Once the FLEX pumps begin injection, the heatup will quickly end, and the core will begin cooling again.

Ultimately, the dynamic event tree is meant to show that varying combinations of timing of these events can lead to many possible scenarios in which core damage is avoided. Like the Simplified PWR simulations in Section 5, a simplified dose modeling method is used, which basically relies on timing of core damage.

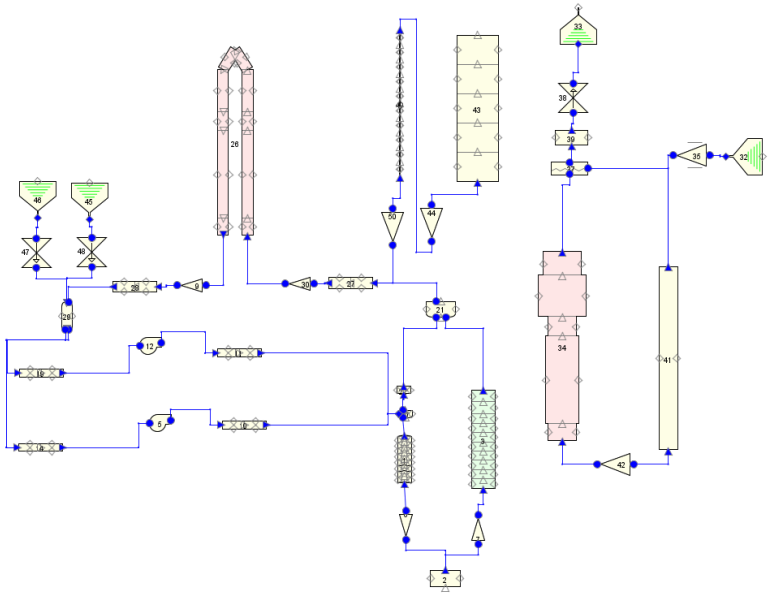

Figure 3-27 RELAP model of a representative SMR.

3.11.5 Results

The first step of the process was to run the coupled EMRALD and Neutrino simulation; this is shown in Figure 3-28. 100 runs were executed, as the computational time is significant, but this provides an adequate sample size to define a reasonable mean and standard deviation.

Downstream of this execution, a standalone representative Neutrino case was executed to create a rendered visualization, and a screen shot is shown in Figure 3-29. In the Neutrino simulation, fluid particles are emitted from a break in a piping line representing the feedwater line. The boxes in the background are chosen to represent components that control the safety injection systems. The simulation calculates the accumulated water level and terminates when the water level reaches a value assumed to be the bottom of crucial components.

The output of this coupled EMRALD and Neutrino simulation is 100 different values for the timing of the flooding of the safety injection systems. From these values, a mean and standard deviation were calculated by EMRALD. This mean and standard deviation were used for the “Flooding” event from Figure 3-24.

This EMRALD model was then executed from FPoliAAP with 100,000 runs to get a good statistical characterization. For each of the three key states, the 100,000 individual timing values were used to produce a 100-point cumulative distribution function (CDF) to characterize the timing. The produced CDFs are shown in Figure 3-30.

The CDFs produced by postprocessing the EMRALD results are used directly as custom 1D distributions in RAVEN for the dynamic event tree in FPoliSIM. For the purposes of simulation speed, each distribution is only simulated with four even subdivisions of the distribution (resulting in branching conditions at the 25^{th}, 50^{th}, 75^{th}, and 100^{th} percentiles). The execution of the dynamic event tree results in 32 unique histories.

The results of the dynamic event tree were post-processed in the FPoliAAP notebook API. The probability and end state for each unique history is shown in Figure 3-31. The clad temperature results for each unique history are shown in Figure 3-32. The core liquid level results for each unique history are shown in Figure 3-33. An initiating event frequency of 1.0E-3 /plant-yr was used to calculate the frequencies based on the branch probability for each unique history. A simplified consequence calculation based on the timing of core damage is calculated for each of the histories. The specifics of this calculation were set simply to provide demonstrative dose results, as the details of the dose calculation are beyond the scope of the work herein.

With frequencies and doses calculated for each individual history, the results are then plotted as a point cloud against the F-C Target in Figure 3-34. This shows that the simulations result in histories under the F-C Target. Although this shows compliance with the target, it is noted that taking each unique history as an LBE would be unwieldy.

It is noted that the transients in Figure 3-32 and Figure 3-33 show that several of the runs that did not encounter core damage would likely result in core damage if executed for a longer time. An attempt was made to run longer, and an error was encountered that could not be fixed in time in a timely manner. However, this really has no effect on the demonstration herein, as the demonstration is for the methodology rather than the results.

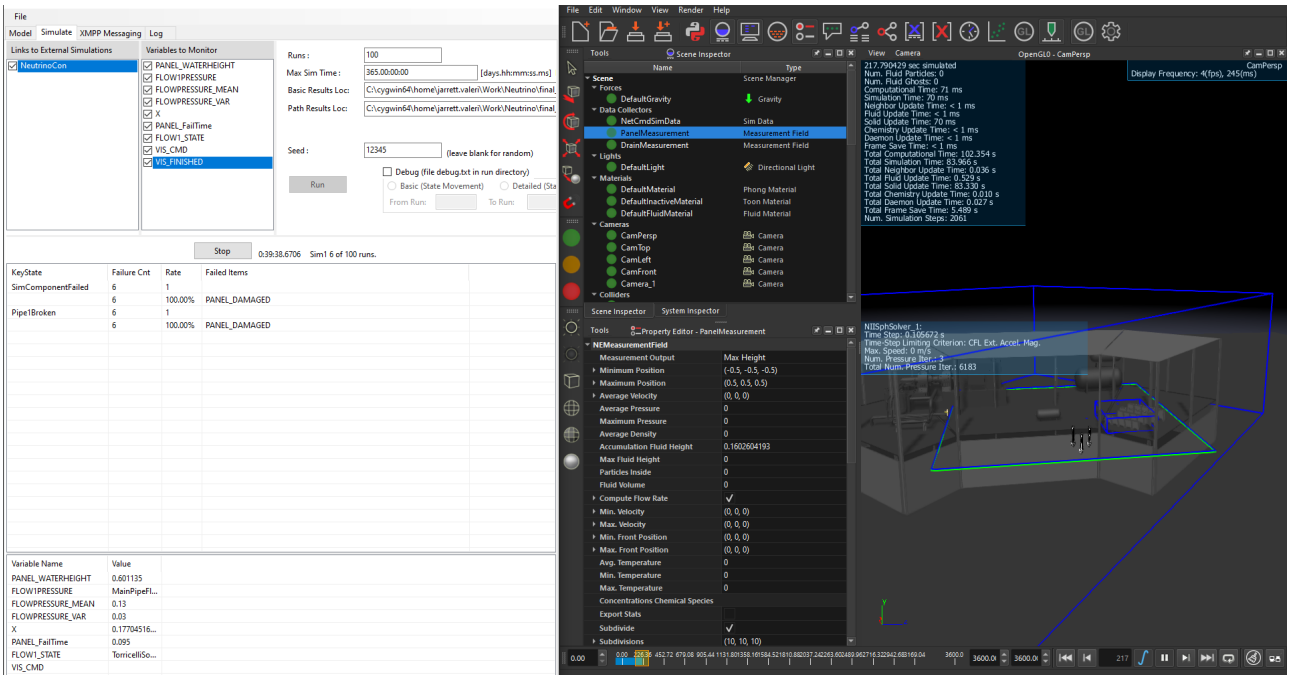

Figure 3-28 EMRALD running Neutrino.

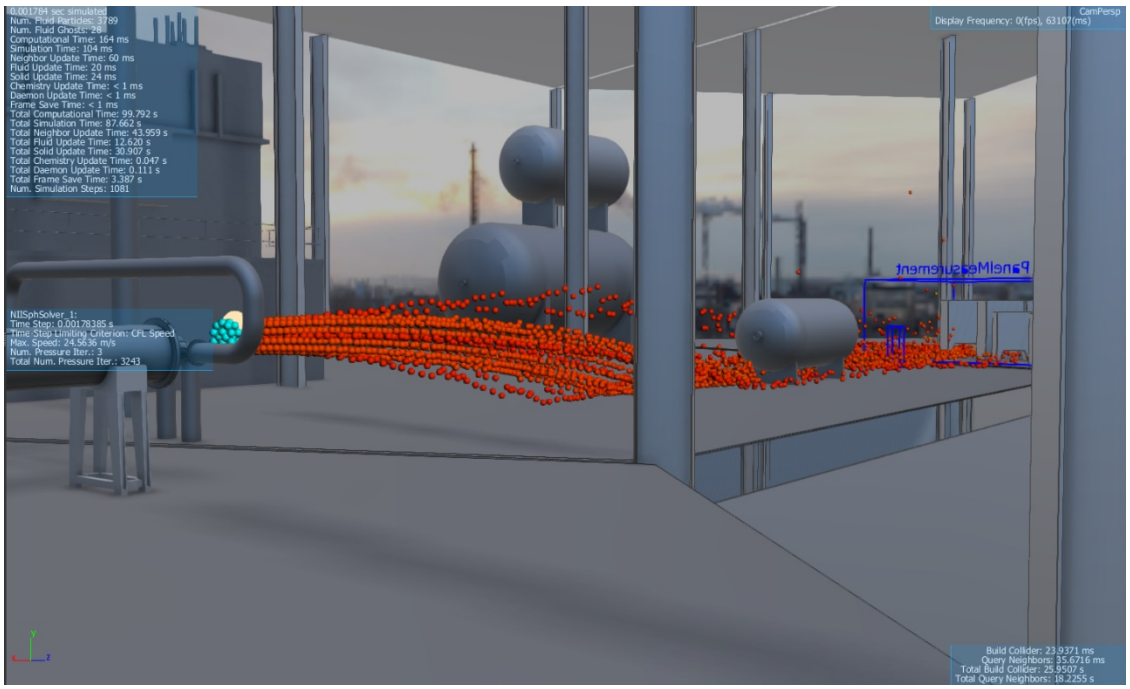

Figure 3-29 Rendered Neutrino simulation.

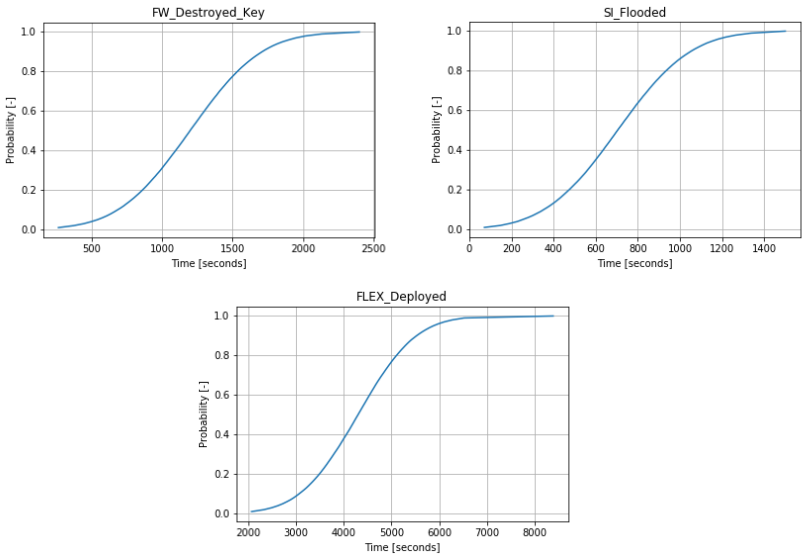

Figure 3-30 CDFs produced by the EMRALD simulation.

Run	Frequency	Result	CD Time	Dose	Run	Frequency	Result	CD Time	Dose
[-]	[/plant-yr]	[-]	[sec]	[rem]	[-]	[/plant-yr]	[-]	[sec]	[rem]
1	6.59E-06	OK	N/A	5.00E-03	17	2.34E-05	CD	8776.08	7.69E-01
2	2.20E-06	OK	N/A	5.00E-03	18	7.03E-05	CD	8812.83	7.46E-01
3	6.59E-06	OK	N/A	5.00E-03	19	1.98E-05	OK	N/A	5.00E-03
4	1.98E-05	OK	N/A	5.00E-03	20	5.93E-05	OK	N/A	5.00E-03
5	8.79E-06	OK	N/A	5.00E-03	21	4.69E-05	CD	8381.88	1.02E+00
6	2.64E-05	OK	N/A	5.00E-03	22	1.56E-05	CD	8383.33	1.01E+00
7	5.86E-06	OK	N/A	5.00E-03	23	7.91E-05	OK	N/A	5.00E-03
8	1.76E-05	OK	N/A	5.00E-03	24	2.64E-05	OK	N/A	5.00E-03
9	6.59E-06	OK	N/A	5.00E-03	25	7.91E-05	OK	N/A	5.00E-03
10	1.98E-05	OK	N/A	5.00E-03	26	2.64E-05	OK	N/A	5.00E-03
11	2.64E-05	OK	N/A	5.00E-03	27	1.05E-04	CD	9976.05	2.00E-02
12	8.79E-06	OK	N/A	5.00E-03	28	3.52E-05	OK	N/A	5.00E-03
13	5.27E-05	CD	9195	5.08E-01	29	1.76E-05	OK	N/A	5.00E-03
14	5.86E-06	OK	N/A	5.00E-03	30	5.27E-05	OK	N/A	5.00E-03
15	1.76E-05	CD	9194.1	5.08E-01	31	7.03E-05	CD	9707.63	1.88E-01
16	1.76E-05	OK	N/A	5.00E-03	32	2.34E-05	OK	N/A	5.00E-03

Figure 3-31 Probabilities and end states for each unique history in the DET.

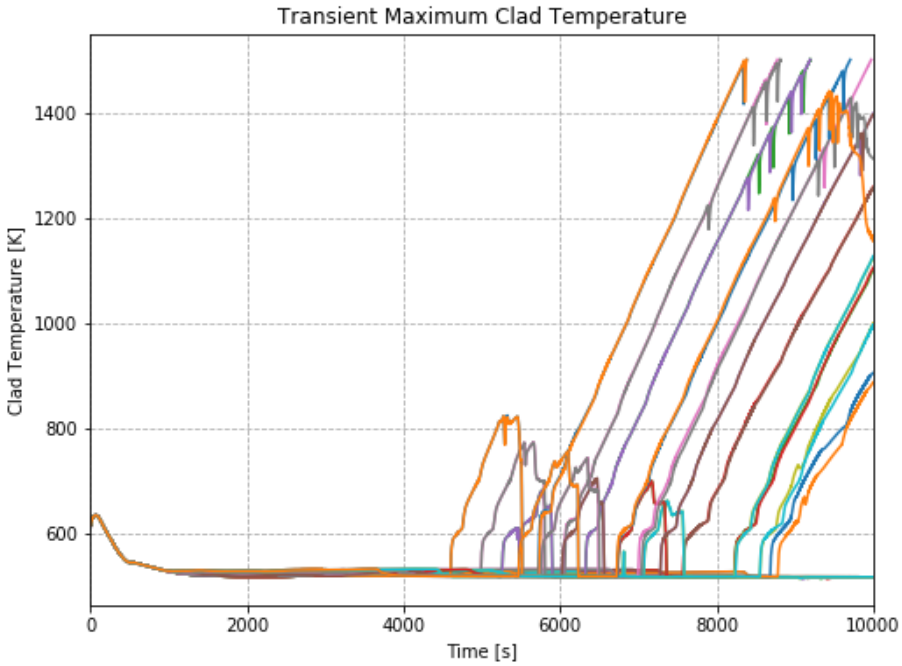

Figure 3-32 Unique clad temperature histories from the DET.

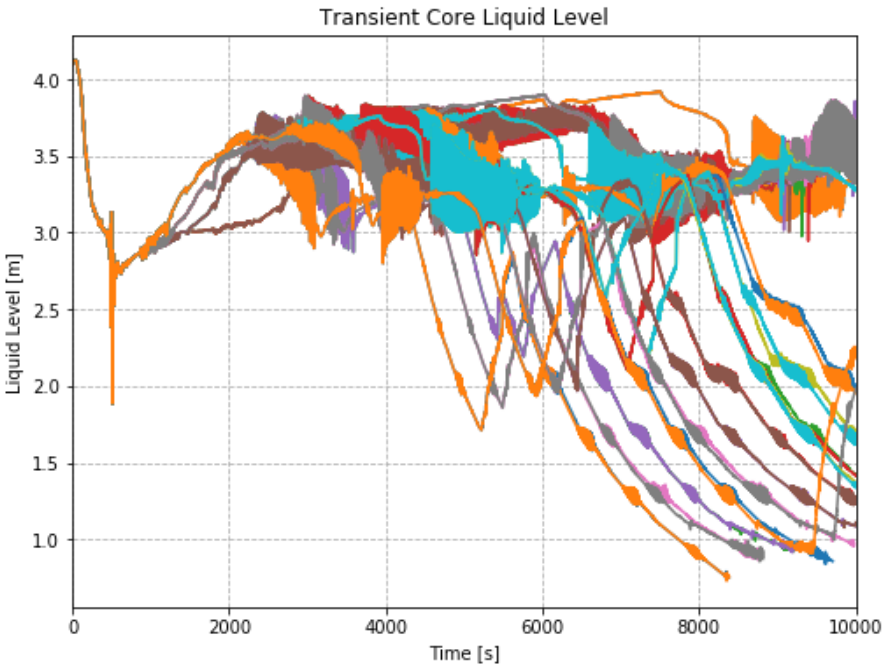

Figure 3-33 Unique core liquid level histories from the DET.

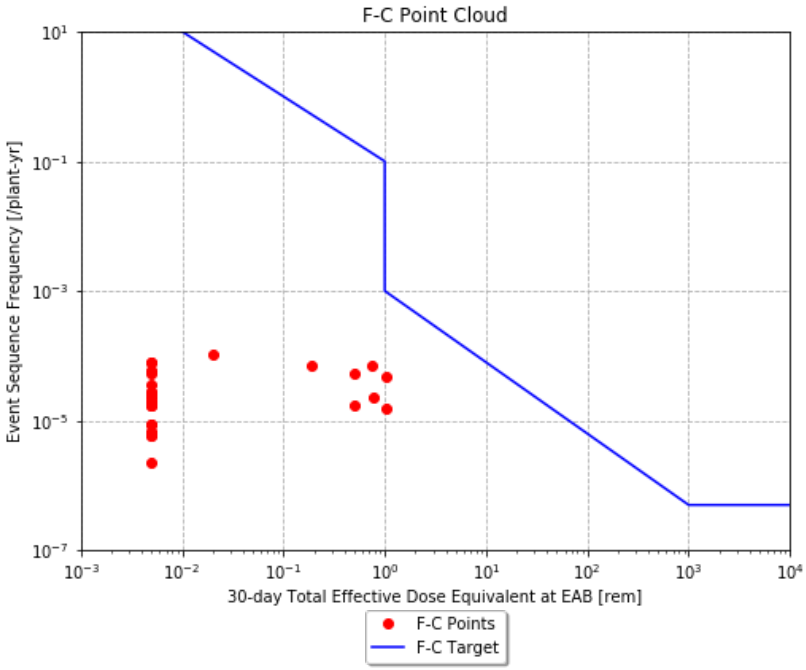

Figure 3-34 Frequency-consequence point cloud.

3.11.6 RISE Workflow

The current implementation of EMRALD in FPoliAAP does not support two-way coupling. As such, for this initial demonstration, some data had to generated externally and then imported back into the FPoliAAP data base. Once the two-way coupling development is completed, the user will be able to execute this complex workflow with more ease. The main EMRALD simulation was executed from the FPoliAAP UI as shown in Figure 3-35. The postprocessing of the simulation results into CDFs loaded in the database are automated when EMRALD is executed in the FPoliAAP framework.

Next, the dynamic event tree execution via RAVEN is set up in FPoliSIM, as shown in Figure 3-36. Currently, the implementation only supports even probability interval subdivisions for DETs, but future functionality could include customized probability intervals. Upon execution in FPoliSIM, the results are saved to the database, and can be post-processed from the browser in the FPoliAAP notebook API.

The results of the dynamic event tree analysis are shown as a point cloud. This was produced using a relatively coarse DET, and it still creates a somewhat unwieldy number of LBEs. As such, part of the incorporation of the dynamic PRA process herein into the RISE process involves grouping the dynamic event tree results into ESFs. For simplicity herein, the sequences are divided into two ESFs: sequences that end in core damage, and those that do not. This grouping is based largely on the simplicity of the dose calculations. If the modeled scenario had more potential complications to the dose calculations, more ESFs could be reasonable. With those ESFs established, the minimum and maximum frequencies and doses are used as bounding estimates of the 5th and 95th percentile results for each ESF. The results of this grouping are shown in the F-C chart of the RISE dashboard in Figure 3-37.

It is noted that the dynamic event tree simulation performed herein does not include the failure or success of the relevant SSCs, such as the emergency core cooling system, the auxiliary feedwater system, and the FLEX equipment itself. To fully take this into account, a hybrid dynamic event tree approach would be used to add the possibilities of outright failure, in addition to the timing that is tested here. As such, the remaining parts of the NEI-18-04 process are not explored in detail for this case study.

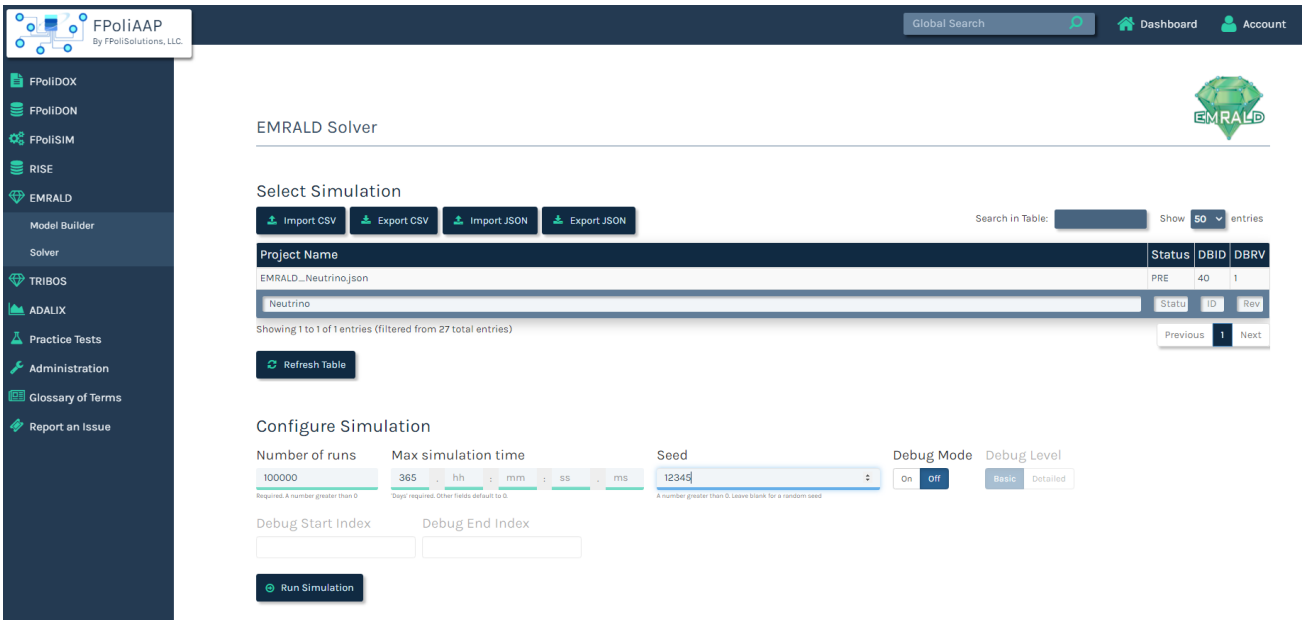

Figure 3-35 EMRALD solver in FPoliAAP.

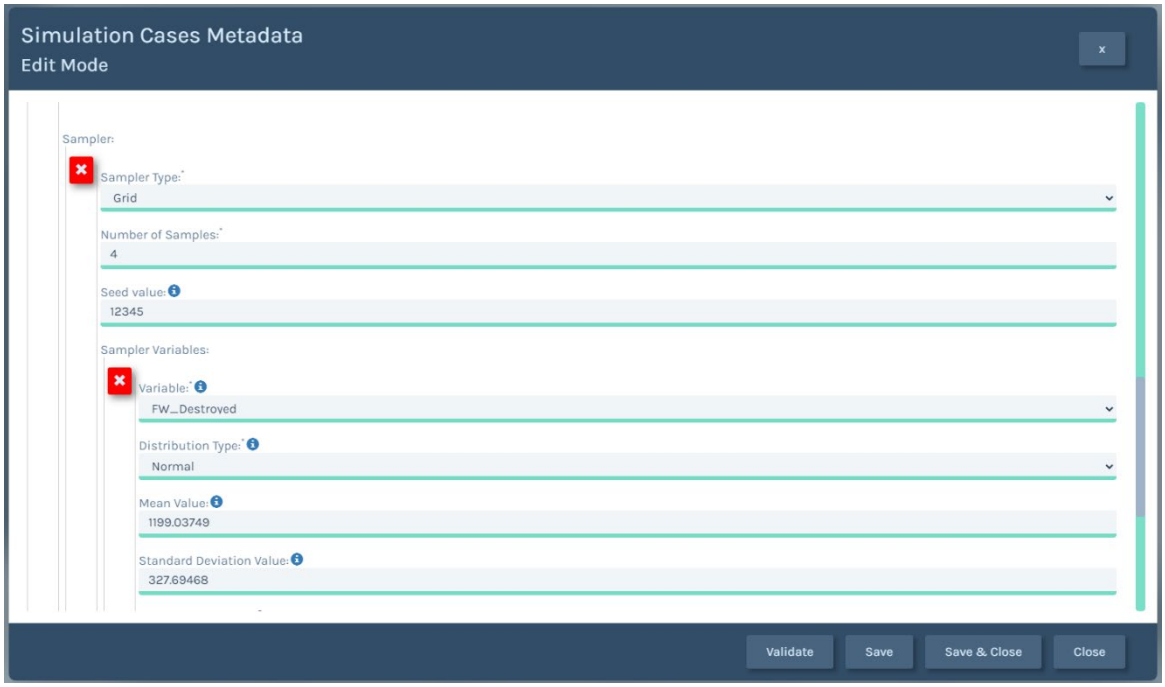

Figure 3-36 Dynamic event tree setup in FPoliSIM.

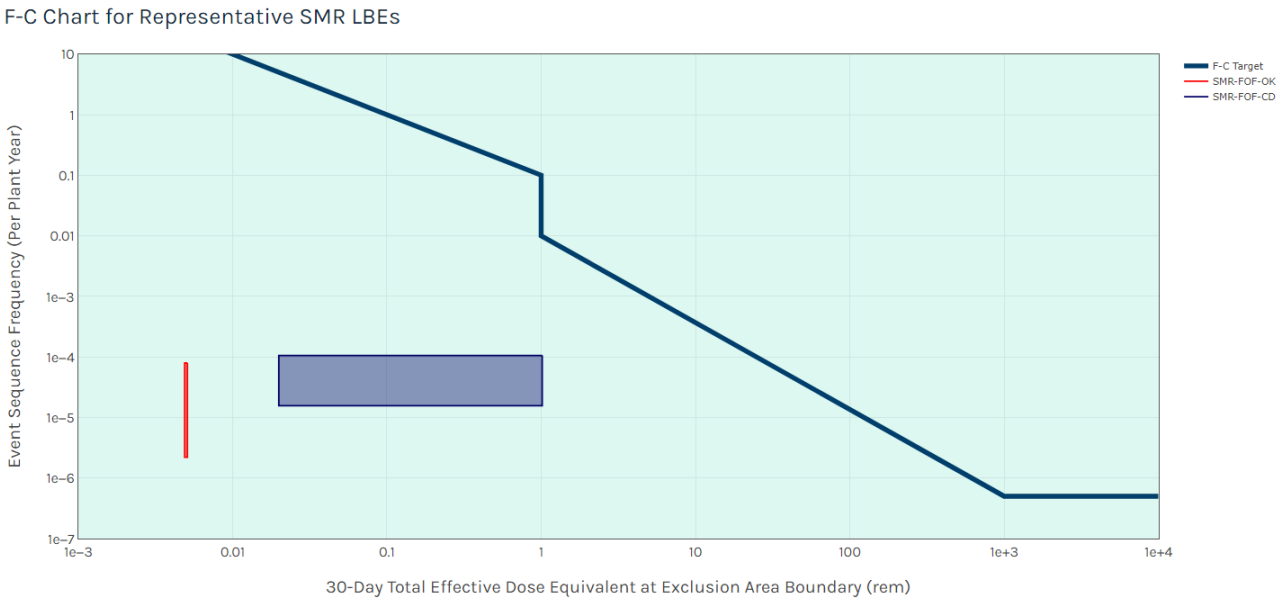

Figure 3-37 Dynamic event tree ESFs on the F-C chart.

3.12 Case Study 3: Graded Approach for a Hypothetical Molten Salt Reactor

For reactors concepts where components and safety systems are being designed from the ground up, a detailed PRA analysis may not be possible in the early stages of design and licensing. In such situations, the LMP process in NEI-18-04 may not be well suited for the initial development of the safety case. In addition, depending on the simplicity and inherent safety margins of the design, a bounding deterministic analysis may be the most straightforward route to a safety case. As these circumstances are likely to exist for many advanced reactors, an example of the graded approach is performed in this section. This approach tends to use bounding scenarios, figures of merit, surrogate acceptance criteria, consideration of single failures, etc.

3.12.1 Scenario Description

The scenario of interest in this demonstrative study is a hypothetical molten salt reactor subject to a Loss of Flow (LOF) initiating event. In this sample LOF, the reactor features SSCs engineered to perform critical SFs that support the FSFs of controlling heat generation ad controlling heat removal. The SSCs considered in this analysis are the following engineered safety features:

1. Normal reactivity control
2. Safety reactivity control
3. Normal decay heat removal
4. Safety decay heat removal

The exact design of these components is unspecified, but they are simply meant to reflect a redundancy in design of two key safety features. The possible scenarios are represented by the even tree depicted in Figure 3-38 which leads to five event sequences, ES1 through ES7.

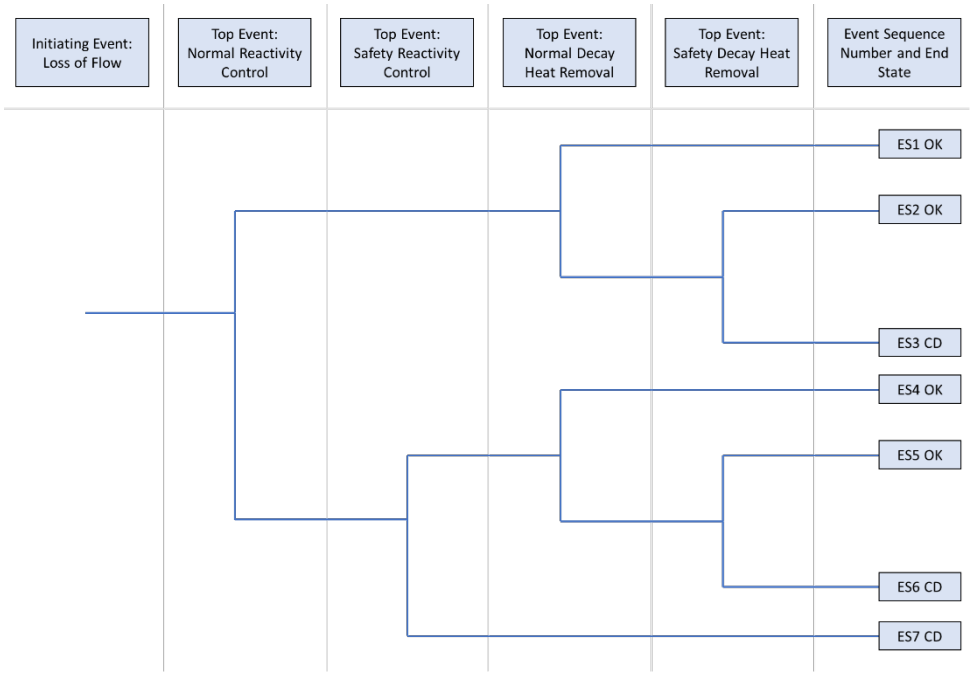

Figure 3-38 Event tree for the LOF case.

Starting from the event tree above, the sequences are individually examined to create representative scenarios. In particular, two ESs are chosen to be executed. ES1 is chosen to be the representative AOO. This is based on it being the successful operation of the two normal systems of event mitigation, which is philosophically consistent with the idea of an AOO. ES5 is chosen to be the representative DBE. This is based on it being the successful operation of the two safety grade systems of event mitigation, which is philosophically consistent with the idea of a DBE.

For this demonstrative problem, a RELAP5-3D model was built of a hypothetical molten salt reactor under the assumed SBO scenario (not that this model is generic and does not represent the details of any specific reactor design under consideration). This reactor is a pebble bed design, with a primary coolant loop using RELAP5-3D ms1 fluid properties, and a secondary heat removal loop using RELAP5-3D "ms2" fluid properties. Simulations are performed wherein the primary coolant pumps are shut off.

In lieu of calculating dose, consistent with the graded approach, a surrogate acceptance criterion of peak fuel temperature less than 1500 K is set. This limit is set with the intention of ensuring fuel integrity, and minimizing releases. The peak fuel temperature results of the two chosen ESs are shown in Figure 3-39, which shows that both sequences succeed.

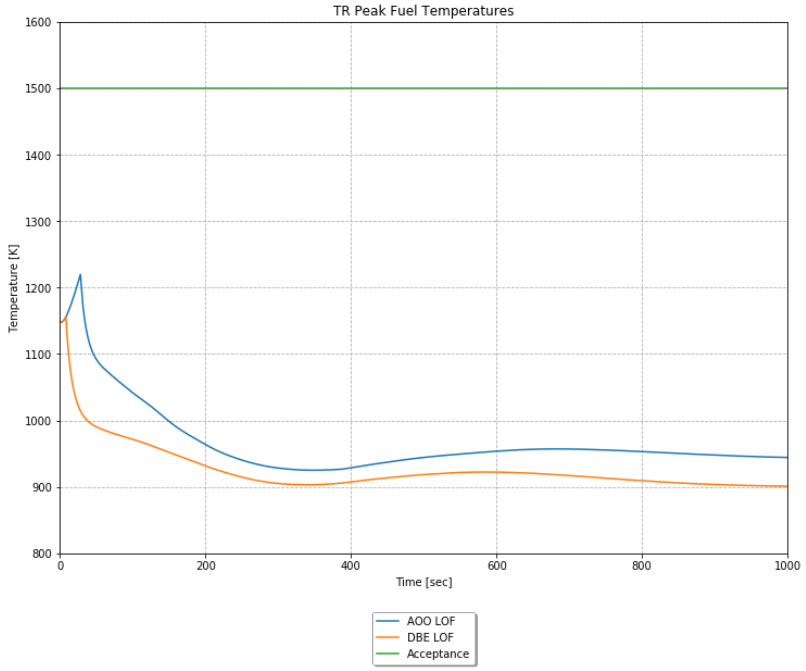

Figure 3-39 Event sequence simulation results.

3.12.2 Uncertainty Analysis

Uncertainties are typically very large in the early stages of the design cycle. Bounding or representative approaches is more suitable and consistent with the graded approach. As design evolves PRA may start playing more of a leading role, which would align better with NEI-18-04 methodology.

In the molten salt demonstration case, a PRA was developed only far enough to determine plausible ESs. It is assumed that the basic event probabilities have enough uncertainty that any calculated event sequence frequencies would be too uncertain to use in this simple hypothetical case. As a result, representative sequences were chosen in the last section. With these classifications, representative

frequencies are chosen for illustrative purposes (and to aid in setting reliability targets downstream). For the AOO, an event frequency of 1 /plant-yr is chosen. For the DBE, an event frequency of 1E-2 /plant-yr is chosen. Both frequencies will essentially be treated as point values in the F-C curve, as they are intended to be bounding/representative.

For the uncertainty analysis of the plant response simulations, bounding plant parameters are used, as well as representative physics models. With the peak fuel temperature acceptance criterion set, a conservatively high consequence of 0.01 rem at EAB is used. This value is chosen to be acceptable in the F-C Target, with some margin, and can be used to develop performance targets. This consequence result will essentially be treated as point values in the F-C curve, as it is intended to be bounding/representative.

3.12.3 RISE Workflow

This analysis is orchestrated through the RISE dashboard (Figure 3-40). In this section, usage of the dashboard for the graded approach is demonstrated. In this simple demo case, only two ESFs are entered (the representative AOO and the representative DBE), and they both make it into the final LBEs selection. It is expected that a similar process would be used to generate a much larger set of representative events in a typical graded approach. In this example the acceptance criteria are based on NEI 18-04; however, in the graded approach, it is likely that the designer may choose an alternative set of acceptance criteria.

For those results, the integrated risk against cumulative metrics defined by the user is computed and displayed in the table below the F-C chart. In this demonstration, an arbitrarily low result is calculated for these metrics based on the peak fuel temperature limit precluding any significant dose.

For the representative molten salt demonstration analysis, two PSFs are specified: “Decay Heat Removal” and “Reactivity Control.” Both functions are derived directly from the MSFs.

To demonstrate the selection of RSFs, a DBE is chosen for performing a sensitivity study. MS-LOF-DBE is used for this purpose. In this study, the Reactivity Control is removed from consideration. This functionally results in an F-C space with frequency distribution based on ES 5, but consequences based on ES 7. The results of this study are not explicitly plotted on the F-C Target, since the frequencies and consequences are rough approximations. Instead, it is simply assumed that ES 7 results in a violation of the F-C Target. This results in the definition of Reactivity Control as a RSF. A similar study results in the classification of Decay Heat Removal as a RSF as well (in this study, ES 5 is again used as the base, and the consequences now come from ES 6).

From the RSFs, one SSC that provides the SF must be declared a SR SSC. For this exercise, the safety reactivity control SSC is chosen as the SSC which provides Reactivity Control, and the safety decay heat removal is chosen as the SSC which provides Decay Heat Removal.

Once the SR SSCs are chosen, the remaining SSCs must be checked for risk significance. In this case, the remaining SSCs are the “normal” reactivity control and decay heat removal SSCs. A sensitivity study, like one used to check for the RSFs, is performed (note the differences with determining risk-significance and RSFs described earlier). In this study, ES 1 is used as the base, and the consequences now come from ES 5. This results in an F-C space that is not over the line. However, in the interest of defense in depth, the remaining systems will be determined as NSRST. In this demonstration, no examples of NST SSCs are included.

At this point the analyst has all the information required to determine the set of DBAs. The DBAs are derived from the DBEs, but are performed crediting only the SR SSCs needed. A DBA is essentially a variation of the associated DBE where only SR SSCs are available to mitigate the postulated ESFs to within the 10 CFR 50.34 dose limits. The list of DBA is presented in the last table of the dashboard.

The next step in the process is requirements derivation. Based on the PRA analysis and simulations, as well as the SSC classification, RFDCs, SRDCs, and Special Treatment Requirements must be derived. Many of these requirements, such as maintenance, testing, and quality assurance programs, are beyond

the scope of this demonstration. However, there are some requirements that are relevant to the performance and availability of SSCs and are shown below.

The following are examples of SRDCs:

- Safety decay heat removal rate
- Safety reactivity control time
- Safety reactivity control worth

The following are examples of Special Treatment Requirements:

- Normal decay heat removal rate
- Normal reactivity control time
- Normal reactivity control worth

The RISE Dashboard has not yet been completed for the scope of DID adequacy but will be in the future. In addition, the limited scope of this demonstration would obviously leave significant holes in defense in depth. This was considered acceptable for the scope of work herein.

Finally, the safety case manager can automatically generate the key reports that define the safety case.

3.12.4 Comments on Graded Approach

The RISE methodology was designed to accommodate a flexible graded approach to risk-informed. The analyst could start from a deterministic approach in the early stage of the design. In the long term, as more data become available and uncertainties are characterized, PRA may take more of a leading role in the analysis. The NEI 18-04 already allow some degree of flexibility to manage a graded approach.

For example, LBEs are Event Sequence Families (ESFs) with frequency higher than the residual risk events. In RISE (our interpretation of NEI 18-04), the user can really ‘decide’ how to define the ESF and there is a layer of engineering judgment. The ESF could be a representative event sequence or surrogate within the family. A sophisticated user may drill down the exercise to identify all possible ESs directly from PRA and then group them in mathematical coherent framework to determine the ESF and their frequencies and consequences. Another user may simply define ES to explore the landscape of possibilities and then combine those in ‘representative’ events that can be treated deterministically. NEI 18-04 provides great latitude to the analysts on how sophisticated the PRA model is needed to inform frequencies of LBEs.

Ultimately a reactor designer of a design with high margin to safety may simply choose to continue using the bounding/representative approach and use a detailed PRA analysis as a confirmatory step of their safety case.

This section really described a process to incorporate updated state of knowledge on an ongoing basis. In the future this could evolve in creating a Bayesian framework for this type of analysis. You could consider initial ESs, combine them in surrogate metrics (the ESFs) and update your prior as you gather data. This can in principle be framed in a mathematical coherent framework via automation. One caveat is that major design changes could disrupt this approach, as the past lessons-learn may be impacted by change. However, as design reaches a reasonable maturity level, this Bayesian approach is in principle feasible.

Figure 3-40 RISE dashboard.

4 Conclusions

The classical nuclear plant safety goals have absolute significance. They try to compare risks from nuclear plant operation against risks from competing technologies. In practice, they focus on nearby population affected by accidents (and not the broader issues associated with theft of nuclear material), and (again in practice) they focus on radiological consequences, generally without regard to costs of long-term evacuation.

The Commission's 1986 policy statement on safety goals explicitly put risks from sabotage and theft of nuclear material *outside* the purview of the safety goals. Though far from perfect, PRA for nuclear plants had some credibility at that point, so the safety goals could at least be interpreted and applied to ordinary safety issues. Unfortunately, our understanding of the risk of sabotage or theft of SNM is only partial; we can identify sabotage and/or theft scenarios and work against their occurrence, but we do not know attack likelihood to the same degree that we know initiating event likelihood, and if we did know it, the adversaries might change their minds.

The idea of plotting frequency versus consequences has been tremendously fruitful, not only (or even primarily) in nuclear safety. As illustrated originally by Farmer, and currently by the LMP, one can articulate meaningful policies based on such a plot (adequacy of a given design for a given site, as Farmer did) or reason about classification of SSCs, or generally about mitigating ESs at the design stage.

It would be useful to be able to do this for security, however "frequency" is a challenge in the security domain. But RIMES proposes something like a proxy for frequency: attack difficulty. This does not support comparison of security risks directly with safety risks, but it offers a way to compare security risks with each other. This, in turn, points the way to deciding which security scenarios to mitigate first.

The report also describes an approach to uncertainty treatment and provide an end-to-end demonstration of a possible deployment model of the technology leveraging FPoli's existing digital platform FPoliAAP. FPoliAAP is a knowledge management framework that was built on the INL codes RAVEN and EMRALD as result of DOE projects. This demonstration of uncertainty via a risk informed approach is quite relevant considering that the NRC and the industry are considering a new regulatory framework, 10 CFR Part 53, to facilitate the deployment of advanced reactor technologies. We also demonstrated a comprehensive analysis approach that capture a safety case and allows for treatment of uncertainties found within a complex risk analysis.

5 References

1. Speech-97-26: [NRC] Chairman Shirley Ann Jackson Keynote Address to the Plant Life Management and Plant Life Extension International Conference and Exhibition. (Available at the NRC web site)
2. Safety Goals for the Operations of Nuclear Power Plants; Policy Statement; Republication, 51 FR 30028, published 8/21/86.
3. NRC White Paper on Risk-Informed and Performance-Based Regulation (https://www.nrc.gov/docs/ML1522/ML15223A685.pdf)
4. Reactor Safety Study—An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants, WASH-1400 (NUREG-75/014) (U.S. Nuclear Regulatory Commission, 1975).
5. NEI 18-04: Risk-Informed Performance-Based Guidance for Non-Light Water Reactor Licensing Basis Development, September 28, 2018 (Nuclear Energy Institute, 2018).
6. Miriam-Webster Dictionary, https://www.merriam-webster.com/dictionary/mW.
7. Oxford English Dictionary, https://www.oed.com.
8. Wikipedia, https://en.wikipedia.org/wiki/Risk.
9. NASA. NASA/SP-2011-3422, *NASA Risk Management Handbook*, Washington, DC. 2011.
10. NASA. NPR 8000.4B, Agency Risk Management Procedural Requirements, Washington, DC. 2017.NASA. NASA/SP-2010-576, NASA Risk-Informed Decision Making Handbook, Washington, DC. 2010.
11. NASA. NASA/SP-2016-6105, *NASA Systems Engineering Handbook Rev2*, Washington, DC. 2016.
12. NASA. NASA/SP-2014-612, *NASA System Safety Handbook Volume 2*, Washington, DC. 2014.
13. NASA. SMA-HQ-WBT-220, "Risk Leadership," SATERN course. Washington, DC.
14. NASA. NPR 8705.2C, Human-Rating Requirements for Space Systems, Washington, DC. 2017.
15. ITAA, GEIA-STD-0010, Standard Best Practices for System Safety Program Development and Execution, 2008.
16. NASA. NASA/SP-2010-580, *NASA System Safety Handbook Volume 1*, Washington, DC. 2011.
17. JPL, "2020 ARMW – Epsilon Briefing," 2020.
18. Reevaluating the Current U.S. Nuclear Regulatory Commission's Safety Goals, V. Mubayi and R. Youngblood, Nuclear Technology (American Nuclear Society, 2020). DOI: https://doi.org/10.1080/00295450.2020.1775452
19. U.S. NRC Regulatory Guide 1.174, An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis, Revision 2 (May 2011).
20. "Practical Elimination Applied to New Nuclear Power Plant Designs- Key Elements and Expectations, a RHWG [Reactor Harmonisation Working Group] Report for the Attention of WENRA," September 17, 2019, Western European Nuclear Regulators Association (2019).
21. "Deliberately small reactors and the second nuclear era," D. T. Ingersoll, Progress in Nuclear Energy 51 (2009) 589-603 (Elsevier, 2009).

22. “Siting Criteria – A New Approach,” F. R. Farmer, in “Containment and Siting of Nuclear Power Plants,” Proceedings of a Symposium, Vienna, 3-7 April, 1967 (International Atomic Energy Agency, Vienna, 1967).
23. Risk-Informed Management of Enterprise Security: Methodology and Applications for Nuclear Facilities, Felicia A. Duran, G. D. Wyss, S. E. Jordan, and B. B. Cipiti, SAND2013-5095C (Sandia National Laboratories, 2013).
24. Risk-Informed Management of Enterprise Security (RIMES), Gregory D. Wyss, Presentation to the Institute of Nuclear Materials Management Workshop on Quantification of the Likelihood of an Attack, November 2020.
25. Regulatory Analysis Guidelines of the U. S. Nuclear Regulatory Commission, NUREG/BR-0058.
26. STPA [System-Theoretic Process Analysis] Handbook, N. Leveson and J. P. Thomas, March 2018.
27. “HAZCADS / Hazards and Consequences Analysis for Digital Systems,” EPRI Report 3002012755, M. Gibson (EPRI, December 2018).
28. Hazard and Consequence Analysis for Digital Systems – A New Approach to Risk Analysis in the Digital Era for Nuclear Power Plants, Andrew J. Clark, Adam D. Williams, Alice Muna, and Matt Gibson, Transactions of the American Nuclear Society 119, November 11-15, 2018.
29. Youngblood, R. and L. F. Oliveira. 1989. "Application of an Allocation Methodology." In Proceedings of "PSA '89 / International Topical Meeting / Probability, Reliability, and Safety Assessment." Pittsburgh, Pennsylvania, April 2-7. La Grange Park, Illinois: American Nuclear Society, Inc.
30. Youngblood, R. W. and R. B. Worrell. 1995. "Top Event Prevention in Complex Systems." In Proceedings of the 1995 Joint ASME/JSME Pressure Vessels and Piping Conference. PVP-Vol. 296, SERA-Vol. 3. Risk and Safety Assessments: Where Is the Balance?” New York, NY: The American Society of Mechanical Engineers.
31. Multi-State Top Event Prevention Analysis, R. Youngblood, Proceedings of the 30th European Safety and Reliability Conference and the 15th Probabilistic Safety Assessment and Management Conference, Edited by Piero Baraldi, Francesco Di Maio and Enrico Zio, Research Publishing, Singapore. ISBN: 981-973-0000-00-0 :: doi: 10.3850/981-973-0000-00-0
32. Youngblood, R. W. 1998. “Applying Risk Models to Formulation of Safety Cases.” Risk Analysis 18, No. 4, p. 433. Society for Risk Analysis.
33. Risk-Informed Physical Security; Dynamic Allocation of Resources, R. B. Worrell, G. B. Varnado, and D. P. Blanchard, International Topical Meeting on Probabilistic Safety Assessment and Analysis 2005 (PSA '05), September 2005.
34. “Top Event Prevention Analysis to Eliminate Requirements Marginal to Safety,” Worrell, R. B. and D. P. Blanchard, Proceedings of the 1995 Joint ASME/JSME Pressure Vessels and Piping Conference (Honolulu, 1995).
35. Worrell, R. B. and D. P. Blanchard. 2002. “Top Event Prevention Analysis – A Method for Identification of Combinations of Events Important to Safety.” PSA ‘02. La Grange Park, Illinois: American Nuclear Society, Inc.
36. Theoretical Possibilities and Consequences of Major Accidents in Nuclear Power Plants / A Study of Possible Consequences if Certain Assumed Accidents, Theoretically Possible but Highly Improbable, Were to Occur in Large Nuclear Power Plants, WASH-740 (Atomic Energy Commission, March 1957).

37. “Survey of cyber risk analysis techniques for use in the nuclear industry,” Shannon Eggers and Katya Le Blanc, Progress in Nuclear Energy 140 (2021) 103908.

38. Global Statistics vs. PRA Results: Which Should We Use?, Presentation by Commissioner George Apostolakis, U.S. Nuclear Regulatory Commission, to American Nuclear Society, Eastern Carolina Section, North Carolina, March 27, 2014.

39. Baseline Risk Index for Initiating Events (BRIIE), NUREG/CR-6932 (NRC, 2007)

40. Fiscal Year 2011 Results of the Industry Trends Program For Operating Power Reactors, SECY 12-0056 (NRC, 2012).

41. Know the Risk / Learning from Errors and Accidents: Safety and Risk in Today’s Technology, Romney Beecher Duffey and John Walton Saull (Butterworth-Heinemann, 2003).

42. Value-Focused Thinking: A Path to Creative Decisionmaking, Ralph L. Keeney (Harvard University Press, 1992).

43. PRA Procedures Guide, A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants, Prepared under the auspices of: The American Nuclear and The Institute of Electrical and Electronics Engineers, NUREG/CR-2300 (Nuclear Regulatory Commission, 1983).

44. Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts, Prepared by Senior Seismic Hazard Assessment committee (R. J. Budnitz, G. Apostolakis, D. M. Boore, L. S. Cluff, K. J. Coppersmith, C. A. Cornell, and P. A. Morris), NUREG/CR-6372, U.S. Nuclear Regulatory Commission, 1997.

45. T. Bedford and R. Cooke, Probabilistic Risk Analysis, Cambridge University Press, UK, 2001.

46. “Summary from the epistemic uncertainty workshop: consensus amid diversity,” Scott Ferson, Cliff A. Joslyn, Jon C. Helton, William Oberkampf, and Kari Sentz, Reliability Engineering and System Safety 85, Numbers 1-3, July - September 2004. The entire issue consists of papers from a workshop on epistemic uncertainty.

47. Improving the Safeguardability of Nuclear Facilities, T. Bjornard, R. Bari, D. Hebditch, P. Peterson, M. Schanfein, INL/CON-09-15834, Institute of Nuclear Materials Management 50th Annual Meeting, July 2009.

48. Walker, W.E., Harremoës, P., Rotmans, J., Van Der Sluijs, J.P., Van Asselt, M.B.A., Janssen, P., & Krayer von Krauss, M.P. (2003). Defining uncertainty: A conceptual basis for uncertainty management in model-based decision support. *Integrated Assessment, 0*(0), 000-000.

49. van der Bles, A.M., van der Linden, S., Freeman, A.L.J., Mitchell, J., Galvao, A.B., Zaval, L., & Spiegelhalter, D.J. (2019). Communicating uncertainty about facts, numbers, and science. *R. Soc. Open sci., 6*, http://dx.doi.org/10.1098/rsos.181870

50. Kahneman, D., Sibony, O., & Sunstein, C.R. (2021). *Noise: A flaw in human judgment*. Little, Brown Spark.

51. Digdon, N. (2020). The Little Albert controversy: Intuition, confirmation bias, and logic. History of Psychology, 23(2), 122-131. https://doi.org/10.1037/hop0000055

52. Lee, Y., Dunbar, N.E., Miller, C.H., Lane, B.L., Jensen, M.L., Bessarabova, E., Burgoon, J.K., Adame, B.J., Valacich, J.J., Adame, E.A., Bostwick, E., Piercy, C.W., Elizondo, J., & Wilson, S.N. (2016). Representativeness bias mitigation through a digital game. *Simulation & Gaming, 47*(6), 751-779. https://doi.org/10.1177/1046878116662955

53. Gesser-Edelsburg, A. & Shir-Raz, Y. (2018). Communicating risk for issues that involve "uncertainty bias": What can the Israeli case of water fluoridation teach us? *Journal of Risk Research, 21*(4), 395-416. https://doi.org/10.1080/13669877.2016.1215343

54. Fischhoff, B. (2012). Communicating uncertainty: Fulfilling the duty to inform. *Issues in Science and Technology, 28*(4), 63-70. https://www.jstor.org/stable/43315647

55. National Aeronautics and Space Administration (2016). *Standard for models and simulations*. https://standards.nasa.gov/standard/nasa/nasa-std-7009

56. Nuclear Regulatory Commission (2004). *Effective risk communication: Guidelines for internal risk communication*. https://www.nrc.gov/reading-rm/doc-collections/nuregs/brochures/br0318/guidance/index.html

57. Nuclear Regulatory Commission (2004). *Effective risk communication: The Nuclear Regulatory Commission's guidelines for external risk communication*. https://www.nrc.gov/reading-rm/doc-collections/nuregs/brochures/br0308/index.html

58. Fischhoff, B. & Davis, A.L. (2014). Communicating scientific uncertainty. *PNAS, 111*, www.pnas.org/cgi/doi/10.1073/pnas.1317504111

59. Gribok, A., Wang, M., Wu, J., Wang, Q., Venketeswaran, A., Zhao, K., & Chen, K. (2020). Concept for integrated multi-modal online piping monitoring system along with data fusion and advanced data analytical algorithms using high-resolution fiber optics sensors. Idaho National Laboratory.

60. US NRC Policy Issue SECY-19-0117, "Technology-Inclusive, Risk-Informed, and Performance-Based Methodology to Inform the Licensing Basis and Content of Applications for Licenses, Certifications, and Approvals for Non-Light-Water Reactors," December 2019. (ADAMS Accession Number ML18312A253)

61. US NRC Draft Regulatory Guide DG-1353, "Guidance for a Technology-Inclusive, Risk-Informed, and Performance-Based Methodology to Inform the Licensing Basis and Content of Applications for Licenses, Certifications, and Approvals for Non-Light-Water Reactors," April 2019. (ADAMS Accession Number ML18312A242)

62. NEI Technical Report NEI-18-04, Revision 1, "Risk-Informed Performance-Based Technology Inclusive Guidance for Non-Light Water Reactor Licensing Basis Development," August 2019. (ADAMS Accession Number ML19241A472)

63. US NRC Report NUREG-1855, Revision 1, "Guidance on the Treatment of Uncertainties Associated with PRAs in Risk-Informed Decisionmaking," March 2017. (ADAMS Accession Number ML17062A466)

64. US NRC Regulatory Guide 1.203, "Transient and Accident Analysis Methods," December 2005. (ADAMS Accession Number ML053500170)

65. NEI Letter to Mr. John Tappert, "Unified Industry Position on the NRC's Rulemaking on 'Risk-Informed, Technology-Inclusive Regulatory Framework for Advanced Reactors' (RIN-3150-AK31; NRC-2019-0062)", July 14, 2021.

66. FPoliSolutions LLC Homepage, www.fpolisolutions.com (Accessed July 2021).

67. C. Frepoli, "Digitalization and Management of Thermal-Hydraulic Legacy Data Using FPoliDON Platform," Proceedings of The 18th International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-18), Portland, OR, August 18-23, 2019, Paper 28618.

68. Martin, Robert & Frepoli, Cesare, et. al. (2019). "Design-Basis Accident Analysis Methods for Light-Water Nuclear Power Plants."

69. NRC NUREG-0800, “Standard Review Plan for the Review of SafetyAnalysis Reports for Nuclear Power Plants: LWR Edition,” https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr0800/index.html (Accessed Jul 2021).
70. R.P. Martin and C. Frepoli, “Resurrection of FLASH for High-Tier Scale Verification of RELAP5-3D,” Proceedings of The 17th International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-17), September 10, 2017, Paper 20475.
71. US NRC Regulatory Guide 1.233, “Guidance for a Technology-Inclusive, Risk-Informed, and Performance-Based Methodology to Inform the Licensing Basis and Content of Applications for Licenses, Certifications, and Approvals for Non-Light-Water Reactors,” June 2020. (ADAMS Accession Number ML20091L698)
72. Title 10 in the Code of Federal Regulations Part 53, “Risk-Informed, Technology-Inclusive Regulatory Framework for Advanced Reactors,” Forthcoming.
73. Southern Company Document SC-29980-201, “Modernization of Technical Requirements for Licensing of Advanced Non-Light Water Reactors, PRISM Sodium Fast Reactor Licensing Modernization Project Demonstration,” December 2018. (ADAMS Accession Number ML19036A584)
74. Southern Company Document SC-29980-202, “Modernization of Technical Requirements for Licensing of Advanced Non-Light Water Reactors, Westinghouse eVinciTM Micro-Reactor Licensing Modernization Project Demonstration,” August 2019. (ADAMS Accession Number ML19227A322)
75. Southern Company Document SC-29980-203, “Modernization of Technical Requirements for Licensing of Advanced Non-Light Water Reactors, Fluoride-Cooled High Temperature Reactor Licensing Modernization Project Demonstration,” September 2019. (ADAMS Accession Number ML19247C198)
76. ASME/ANS RA-S-1.42013, “Probabilistic Risk Assessment Standard for Advanced Non-LWR Nuclear Power Plants,” December 2013.
77. Sandia National Laboratories Report SAND2020-4609, “Technical and Licensing Considerations for Micro-Reactors,” April 2020.
78. IAEA Safety Standard No. SSG-30, “Safety Classification of Structures, Systems and Components in Nuclear Power Plants,” 2014.
79. A. Alfonsi, C. Rabiti, D. Mandelli, J. Cogliati, & R. Kinoshita, “Hybrid Dynamic Event Tree Sampling Strategy in RAVEN Code,” ANS PSA 2015 International Topical Meeting on Probabilistic Safety Assessment and Analysis, January 2015.

Appendix A – Neutrino and EMRALD Model Details

A.1 Neutrino Setup

A.1.1 3D Model

The 3D Model was created a computer aided design (CAD) modeling package and imported into Neutrino as STL files. The model was loosely based on the Open 100 reactor design and the turbine building was chosen as a hypothetical area to be modeled.

The following figures depict the modeling process.

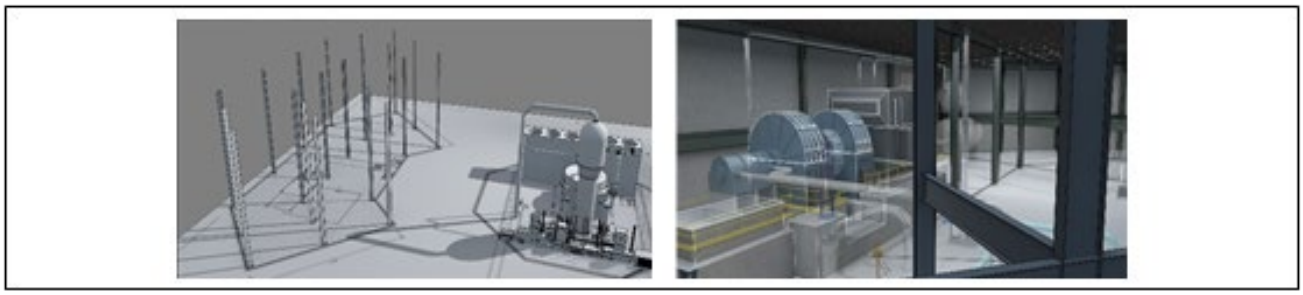

Figure A-1 - CAD Model and Layout of Reactor and Turbine Room

A.1.2 Simulation

There are two ways to simulate a pipe break and the subsequent flooding in Neutrino:

a) Representing the liquid as a set of interacting, moving particles in 3D, and solving the Navier-Stokes equation using smoothed-particle hydrodynamics (SPH). This approach makes it possible to account for the three-dimensional and dynamical aspects of the liquid motion and its interaction with solid structures. The main drawbacks are that the level of discretization of the vertical dimension depends on the particle size and that it may be computationally intensive, especially when the particle size is small and/or when a portion of the flow has a high speed, which may yield very small time-steps.

b) Representing the liquid accumulation within a room as an evolving height field based on a mass balance and an energy balance. It accounts for inflows and outflows occurring within the room (e.g., flow through a pipe break or a drain) or from/to adjacent rooms (e.g., flow through the bottom of a door). Flow rates between adjacent rooms can be estimated using Torricelli's law-based calculations, whereas specific simple drain models can be used to estimate the flow rates through drains. This approach makes it possible to achieve very fast simulation runs. The main drawback is that it cannot account for spatial information beyond connectivity characteristics between adjacent rooms.

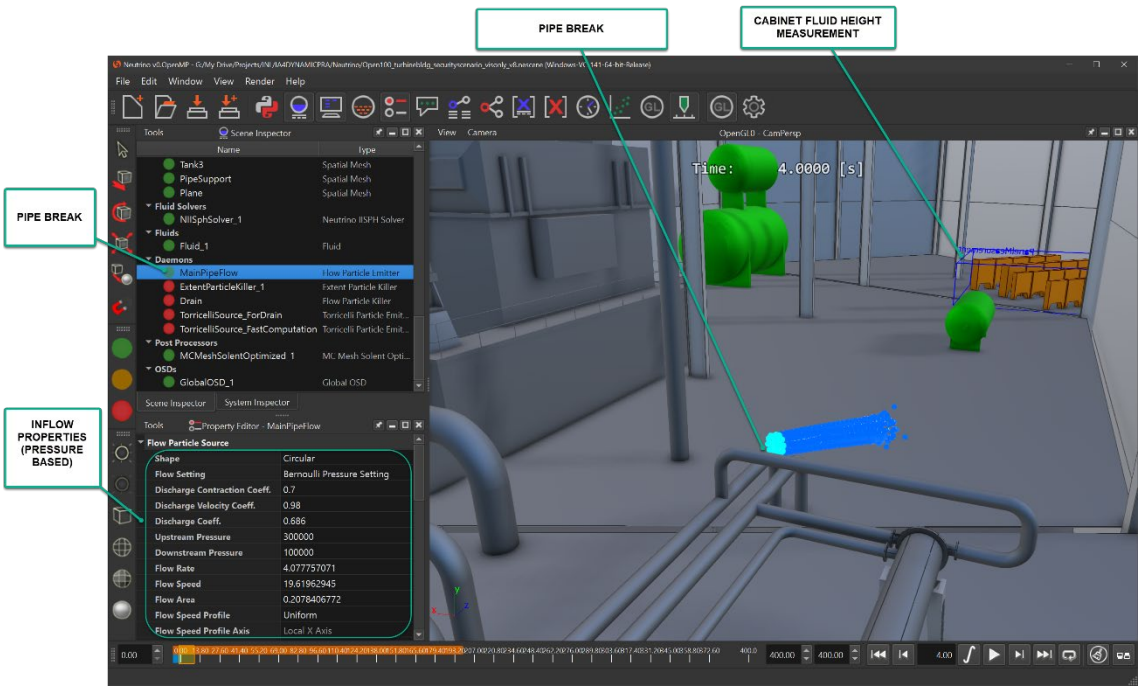

Figure A-2 - Neutrino Particle (SPH) CFD setup for the pipe break

The scene setup for approach (b) is significantly simpler, and the elements strictly necessary and relevant for approach (b) are indicated in the following with italics and underlined text.

Structures are modeled by importing CAD geometries. Those that are thought to possibly interact with the liquid are set active, and boundary particles are automatically created to sample the surface of those solids. The boundary particles are offset towards the interior or exterior of the solid they sample, depending on whether, respectively, the exterior or the interior is expected to be in contact with the liquid.

The pipe break is modeled by a circular flow source to which the position, orientation and radius are set in accordance with the location and dimension of the break. The flow rate is calculated using Bernoulli's equation based on the specified upstream-to-downstream pressure difference, specified discharge contraction and velocity coefficients, and the given radius. Acting as an inlet with prescribed velocity, the flow source adds new fluid particles into the simulation domain.

A Drain is modeled by a flow killer to which the position and scaling are set in accordance with the location and dimensions of the drain. Acting as an outlet, the flow killer removes fluid particles from the simulation domain. The deletion rate is dynamically adjusted based on the liquid height measured at this location and a simple drain model relating liquid height and outflow rate.

In approach (a), a Torricelli source is set to calculate the drain flow rate. The characteristics of the drain system are specified: opening diameter, maximum flow rate, and liquid height corresponding to a submerged state in which the flow rate is maximum.

In approach (b), a Torricelli source is set not only to calculate the drain flow rate but also to track and evolve the liquid accumulation. The floor elevation and the room area where the liquid can accumulate are specified. The inflow rate is dynamically taken as the flow rate of the pipe break. The characteristics of the drain system are specified: opening diameter, maximum flow rate, and liquid height corresponding to a submerged state in which the flow rate is maximum.

In approach (a), two measurement fields are set, one at the location of the drain and another at the location of the target, to measure the height of the liquid columns that accumulate there. These pieces of information are transmitted to, respectively, the Torricelli source and EMRALD. The position and scaling of each measurement field are set such that its bottom is at the same elevation as the floor.

In approach (b), a measurement field is set to access the liquid height evolved by the Torricelli source and to transmit it to EMRALD.

The frame rate is set to 1 Hz, implying that cache data are generated only every one second of simulation. *In approach (b), this also enforces a time step of one second for evolving the liquid accumulation with the Torricelli source.*

The same scene could be used for both simulation setups (a) or (b) by just enabling and disabling components in the Neutrino scene file. The following figure illustrates the setup in (b).

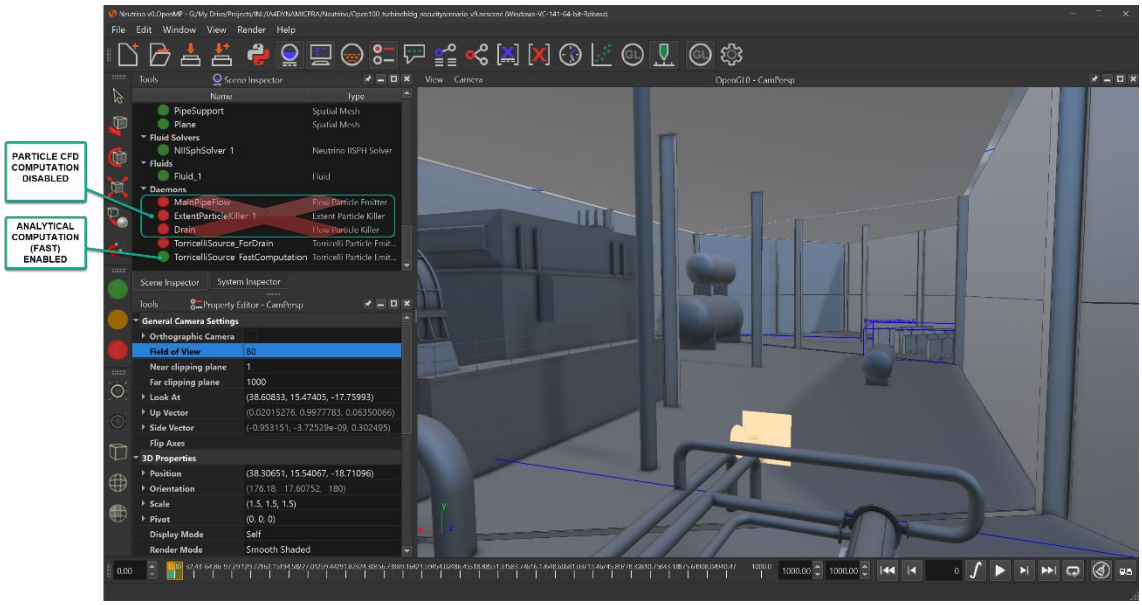

Figure A-3 - Neutrino analytical solver setup.

A.2 EMRALD Setup

In either of the approaches EMRALD acts as a front-end driver loading the appropriate Neutrino scene file and varying the inflow conditions according to the appropriate distribution and varying the size of the pipe break. Such a setup is indicated in the following figures.

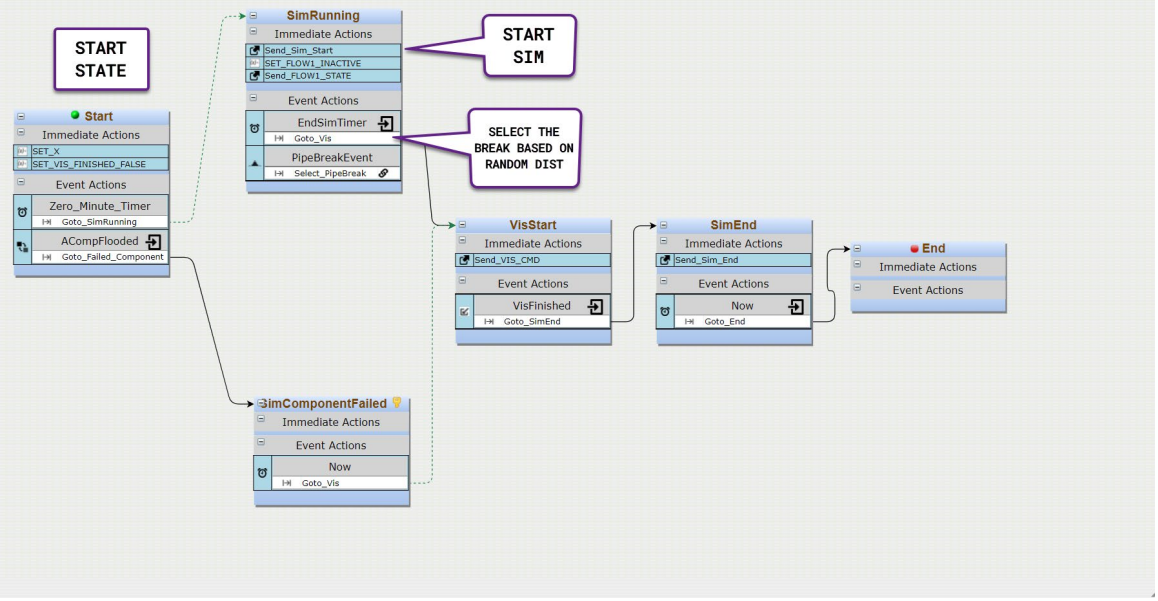

Figure A-4 - EMRALD main setup

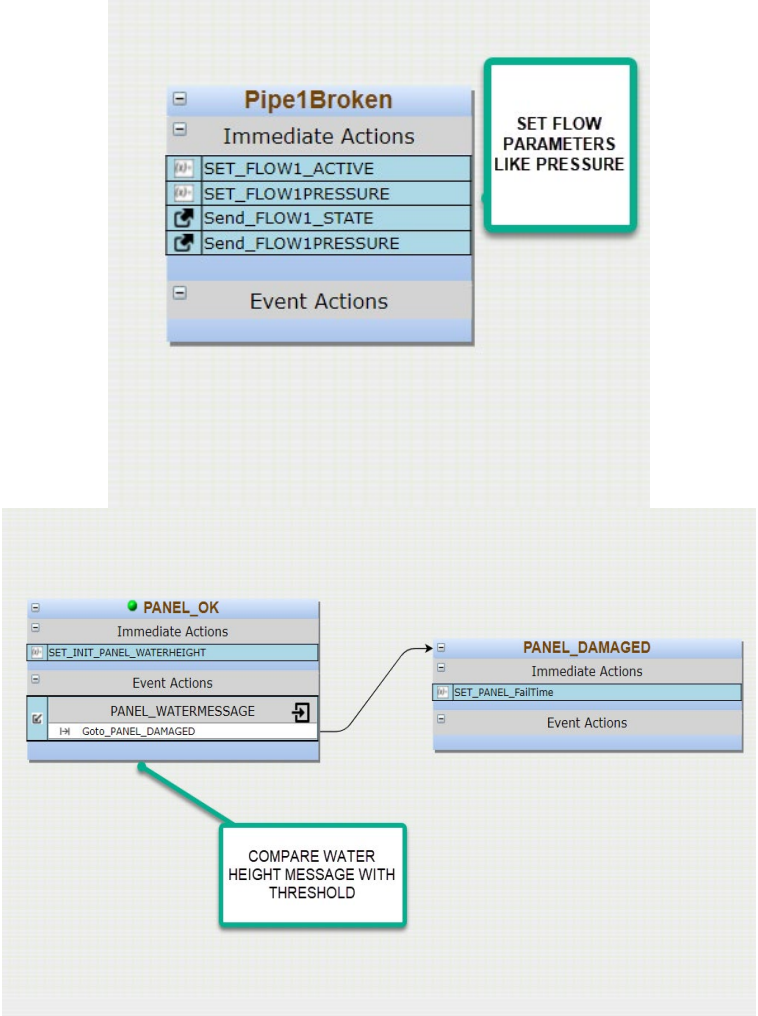

Figure A-5 - EMRALD support diagrams setup