

# USING SPARC TO EXTRACT RISK INSIGHTS FROM PERFORMANCE ASSESSMENTS FOR HIGH-LEVEL WASTE REPOSITORIES

S. Tina Ghosh / U.S. Nuclear Regulatory  
Commission

George E. Apostolakis / Massachusetts  
Institute of Technology

**Disclaimer:** One of the authors of this paper is an employee of the U.S. Nuclear Regulatory Commission (NRC). This work was completed at Massachusetts Institute of Technology before she joined the NRC. The NRC staff views expressed herein are preliminary and do not constitute a final judgment or determination of the matters addressed or of the acceptability of a license application for a geologic repository at Yucca Mountain.

## SUMMARY/ABSTRACT

Assessing risk from potential geologic repositories for high-level radioactive waste (HLW) poses a unique challenge because of the large temporal and spatial scales involved. The repository systems typically involve multiple engineered barriers and emplacement in a remote geologic location. It is not feasible to conduct experiments on the scale of a repository system. As a result, scientists rely on inference from shorter-term, controlled laboratory and field experiments, as well as historical data from natural analogs to model the system. This leads to uncertainties in modeling the risk from the repository.

The importance of these uncertainties depends on the context for conducting and using the risk assessments. The U.S. National Research Council has pointed out that the proper role of risk assessments is to inform societal decisions and facilitate deliberation among stakeholders. In the case of HLW repositories, the most important sources of risk information for these decisions are the results of integrated repository system modeling compiled in performance assessments (PA). These PAs inform critical decisions for national HLW repository programs. The important uncertainties in the PA are those to which decisions may be sensitive.

To build confidence given potentially important uncertainties, a typical safety case for a proposed HLW repository is comprised of PA results coupled with various defense-in-depth elements, such as the multiple-barrier requirement for repository design, and insights from supplementary analyses. This paper proposes an additional supplementary analysis, the Strategic Partitioning of Assumption-Ranges and Consequences (SPARC), whose goal is to construct a specific explanation of the PA results of interest, to aid risk-informed decision-making and risk-informed stakeholder deliberation. The method seeks to provide explanations for undesired system behavior that are similar to those provided by a reactor probabilistic safety assessment (PSA); i.e., if we want to know the reasons for possible reactor core damage (a result of interest), we can trace the responsible sequence of events through the event trees, and similarly trace the ways each event can occur through the underlying fault trees in a PSA. Since the system nature of repositories versus that of reactors is very different (continuous slow degradation in passive systems versus largely binary random failures in active systems), the scenarios of reasons for undesired behavior in a repository system would be defined on different terms.

The SPARC method extracts risk information from existing PAs and supporting databases to explain how the repository system may produce undesired behavior, specifically by uncovering what sets of model parameter values taken together could exceed a particular criterion set for the repository (e.g., an instance where a performance measure such as projected dose exceeds a particular goal.) The results can be displayed in SPARC trees. The method could be used: (1) in a safety case to help build confidence in a repository system; (2) to risk-inform decisions on how to allocate resources for future research; and (3) to risk-inform stakeholder deliberation. As a demonstrative example, the SPARC method is applied to a potential HLW repository at Yucca Mountain.

## INTRODUCTION

Performance Assessment (PA) results are important sources of information for the high-level radioactive waste programs. For example, in the U.S. the developer of a potential repository at Yucca Mountain, the U.S. Department of Energy (DOE), uses PA to help guide design and programmatic choices. The regulator, the NRC, uses PA results in its ongoing regulatory activities, such as prioritizing key technical issues in the pre-licensing phase, or focusing the review of a license application. Ideally, a total-system PA would explicitly include all potentially significant uncertainties. Unfortunately some uncertainties, such as many structural model uncertainties, may not be propagated explicitly through PAs [1]. Instead, analysts may choose a single 'best' model to represent a particular repository component or subsystem, and use engineering judgment to account for model uncertainty, often through sensitivity analyses targeted to that single component or subsystem, and a 'conservative' choice in modeling. The potential problems with applying this approach to uncertainty in the sub-models within a complex PA include: (1) It is not easy to determine whether a choice is conservative a priori, given the non-linearity of system response and coupling among sub-models. (2) With hundreds of uncertain variables and a multi-barrier system, rarely can one variable alone influence the system results significantly. (3) Without some kind of propagation of different model assumptions through the PA, or joint sensitivity analyses, the analysts can not find those sets of assumptions and parameter values that trip the decision threshold(s), which is of great interest to inform the decision problems. At the same time, it would require a tremendous amount of resources to propagate all possible model uncertainties through the PA, so including *all* uncertainties in the PA is precluded by practical limitations. It would be useful, therefore, to have some other way to systematically identify *important uncertainties*, i.e., those that may affect the decision that the PA is informing. This paper offers one approach to uncover these uncertainties by considering their significance in the decision context.

In the U.S., commercial nuclear operations and regulation are becoming increasingly more risk-informed. The NRC announced its intentions to use more risk information in its regulatory activities in 1995, and published a white paper in 1998 outlining what information risk assessments should provide. For nuclear power reactors, level-1 probabilistic risk assessments (PRAs) provide scenarios for the pre-defined undesirable end-state of reactor core damage and large early release of radioactivity. The risk/failure scenarios analyzed are those that lead to core damage. PAs for waste repositories are different in that they model the evolution of the repository system and display the resulting *performance*, as the name indicates. The results show probabilistic performance for all levels, e.g., even those that are several orders of magnitude below the performance goals. From a risk perspective, it is not as important to know all the different ways that the system works, except to help establish the probabilities for system success versus undesirable performance, and understand how challenges to the system are mitigated. In a reactor PRA, if we want to know the reasons for possible system failure, we can trace the sequence of events through the event trees, and similarly trace the ways each event can occur through the underlying fault trees. We seek to provide a similar explanation for PA results by finding undesirable scenarios for the repository system. These scenarios could then serve as the basis for further uncertainty studies.

For waste repositories, there is no analogous end state to core-damage for reactors because the nature of the systems are very different. We can, however, define an undesirable *scenario* of interest as one that leads to repository performance that crosses a reference point individual dose that is salient to the decision problem.

## APPROACH

We can adopt the common definition of a scenario as one collection of possible repository features, events, and processes (FEPs) that will be realized in the repository's future. Furthermore, in the context of gathering risk information, FEPs should be defined strategically to be meaningful from a risk perspective. We will define scenarios for HLW repositories more specifically in terms of a series of strategic partitions. The first strategic partition is for the end-states, the consequences of interest. Nuclear repositories do not *fail* in the way that nuclear reactors do, and hence the concept of *failure* must be redefined. We build on the Generalized Sensitivity Analysis (GSA) tool introduced by Hornberger and Spear [2] for environmental systems. These authors partitioned Monte Carlo Simulation (MCS) outcomes based on pre-defined system behavior of interest (in that case, the behavior of interest consisted of key measures of eutrophication in a lake, e.g., dissolved phosphorous concentration, exceeding specified thresholds). MCS realizations producing this behavior were placed in one bin, and those producing non-behavior were placed in another bin. This is an appropriate way to think of partitioning HLW repository PA results as well. For our purpose, the system behavior of interest is delivery of a specified dose (or greater) to any individual living near the repository sometime in the future. We will call this behavior *Substantially Increased Dose (SID)*. SID does *not* imply a violation of regulatory criteria, since the regulatory goal is based on expected dose only, i.e., the mean of all the dose projections produced by the PA should lie below a particular value.

Next, we have to partition FEPs strategically to create scenarios of risk. In the example application, the PA parameters represented both parameters and assumptions in the models. For example, some of the PA parameters are parameters in sub-models, such as one of the dominant parameters discussed later, the pre-exponential coefficient (PSFDM1) in the default spent fuel dissolution model. There are other parameters that represent a choice of alternative model structures; for example, there are two spent fuel dissolution models available in the PA code used in the application: (1) the default model is the bathtub model where spent fuel is immersed in water inside the waste package slowly dissolves; (2) the alternative model is the flow-through model where water contacting spent fuel as it flows inside the waste package is able to dissolve the fuel as it goes. The alternative spent fuel model can be invoked in the PA by toggling a 0/1 parameter in the PA. Yet other PA parameters are *lumped* parameters that can represent different physical processes and modeling assumptions as well. An example of this is the "Subarea wet fraction" parameter that describes what percentage of the soil above the repository is wet. It is a lumped parameter because there are numerous physical processes and modeling assumptions that could lead to different percentages of wetness.

Thus, our scenarios, which are collections of FEPs that lead to SIDs from the repository, will actually be defined in terms of PA parameter distribution intervals, since the parameter values capture *both* assumptions about the true values of model parameters *and* assumptions about different processes that may be at work in the repository.

In this paper, we present the SPARC method [3], which can help provide risk information for decisions about HLW repositories. This analysis has a two-prong focus: (1) finding SID scenarios and (2) finding 'savior' attributes. The SID scenarios are those combinations of possible FEPs, which are often embodied in parameter value intervals (as explained above), that can lead to individual doses that cross the reference point of interest. These scenarios tend to be tail scenarios; i.e., if we look at the distribution of outcomes from a typical PA, the very small number (if any) of the individual realizations that cross the decision threshold are on the tail of the distribution. The savior attributes are those combinations of possible FEPs that make it virtually impossible for the repository to lead to SIDs, even when challenged severely.

Note that the U.S. regulatory limit is defined for the mean of the PA results ("reasonable expectation"<sup>1</sup>); the PA projected-mean dose regulatory criterion for the base case is <15 mrem/yr for the

---

<sup>1</sup> 10CFR§63.311 states, "DOE must demonstrate, using performance assessment, that there is a reasonable expectation that, for 10,000 years following disposal, the reasonably maximally exposed individual receives no more than an annual dose of 0.15 mSv (15 mrem) from releases from the undisturbed Yucca Mountain disposal system." [10 CFR 63]. "Reasonable

first 10,000 yrs, as set by the U.S. Environmental Protection Agency in [4]; disruptive scenarios are considered separately. The purpose of the SPARC analysis is not to demonstrate compliance; rather, it can serve a supplementary purpose, to help explain PA results, and focus further analyses. The SPARC method consists of the following steps:

- (1) Identify important repository attributes.
- (2) Find (a) SID scenarios and (b) “savior” attributes.
- (3) Explain the reason(s) for SIDs in the SID scenarios, and display results in SPARC trees.

The next steps in an overall risk-informed analysis would be: present supporting evidence for assessing very low probabilities of SID scenarios and probe possible incompleteness in existing analyses; and identify important uncertainties.

PAs consist of hundreds of uncertain parameters and processes, not all of which contribute significantly to important variations in system-level performance. It is imperative, therefore, to identify the important uncertain attributes so that the scenario explanations can be defined based on these. There are several well established methods that can be used to find the important attributes; these methods are various kinds of sensitivity analyses (see [5] and [6] for more extensive description of these methods). Since the PAs typically use Latin Hypercube Sampling (LHS) in conjunction with a MCS, the first order estimate of scenario probability can be derived from the LHS size and associated frequency of the scenarios in the MCS, and the first order estimate of attribute probabilities can be taken from the parameter distributions. The probability estimates can be further refined by scrutinizing the underlying technical evidence supporting the scenarios. The results of these analyses should indicate which uncertainties and repository attributes are the most important from a risk perspective.

## APPLICATION

We illustrate the SPARC approach using version TPA 4.1j of the NRC’s total-system performance assessment code [7] (the dose projections presented here are based on a somewhat out-dated database and a code created as a tool to improve understanding of a potential Repository system at Yucca Mountain (YMR) system, *not* as a compliance demonstration tool; we use dose from drinking water only to the 10km receptor.) The actual numbers here are not important; rather, we are trying to demonstrate how the approach might be applied. The NRC’s PA is flexible in its capability to test different assumptions about a YMR, and because of its relative simplicity, has a short run time. In addition, it contains lumped parameters that could represent different physical processes, which lends flexibility in testing model sensitivities. It is a good basis for application of the SPARC method.

### **Step 1. Identify important repository attributes**

We take advantage of the fact that there are few radionuclides that contribute to the peak dose within the regulatory compliance period; among these, Np-237 (half-life of  $2.1 \times 10^6$  years) is the most important based on preliminary studies (e.g., [8]). Np-237 is the radionuclide with the best potential to produce the reference dose of 15 mrem/yr within 10,000 years. The NRC has performed sensitivity analyses on the parameters in the TPA 4.1, and has ranked the parameters according to numerous methods [9]. We use these findings to create the preliminary list of 12 potentially important repository attributes.

We tried to confirm the importance (for our purposes of identifying specific scenarios of risk) of the parameters by employing generalized sensitivity analysis (GSA)<sup>2</sup> (described in [2]) running the TPA base case multiple times with a Latin-hypercube size of 500 realizations; of these, 21 produced SIDs<sup>3</sup>. Specifically, the GSA can indicate (1) how important the attribute is, based on the largest vertical

---

*expectation*” is taken to mean the average of all PA realizations. Note that both 10 CFR 197 and 10 CFR 63 is currently under revision following the Federal Circuit Court of Appeals ruling NEI v. EPA [11].

<sup>2</sup> We used this as an independent check, since the USNRC had not used this as one of their methods (they used several others) to compile the list above.

<sup>3</sup> We ran the TPA base case multiple times with LHS size 500, with consistent results.

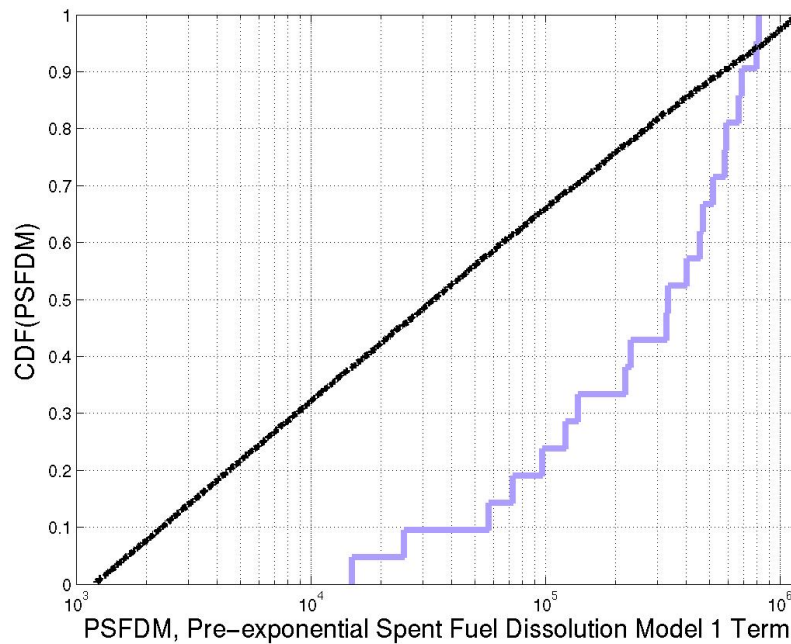
distance between the cumulative distribution functions for values in the two bins [2], AND (2) whether there are threshold effects, which is of particular importance for our goals.

We found that no one parameter *alone* can make the repository produce SIDs (an expected result for a multi-barrier system); we did find, however, that there are some parameters that seem to be able to prevent repository SID on their own. The PSFDM1 (SFD<sup>4</sup>) term is the most dramatic example of this. Figure 1 shows the partitioned CDFs for this parameter; the solid blue line shows the CDF for the SFD values in the 21 SID-producing realizations, and the dashed black line shows the CDF for the 479 non-SID realizations. Based on these realizations, where the lowest value of SFD resulting in SID is  $\sim 1.5 \times 10^4$ , it seems that SIDs would be very unlikely if the true value of the SFD is less than  $10^4$ , which corresponds to roughly the first third of its distribution. We confirmed that most of the remaining parameters were indeed important based on the GSA results. As a control, we checked several parameters that should not show up as important; one example is the AA\_1\_1 parameter which is the key WP general corrosion parameter and one that is very important for repository performance at larger time scales (40,000-100,000 years). As expected, GSA showed no importance for AA\_1\_1 for the 10,000 year time frame.

Starting with this list of potential explanations for repository SID/success, we can build SID scenarios, step 2 of the analysis.

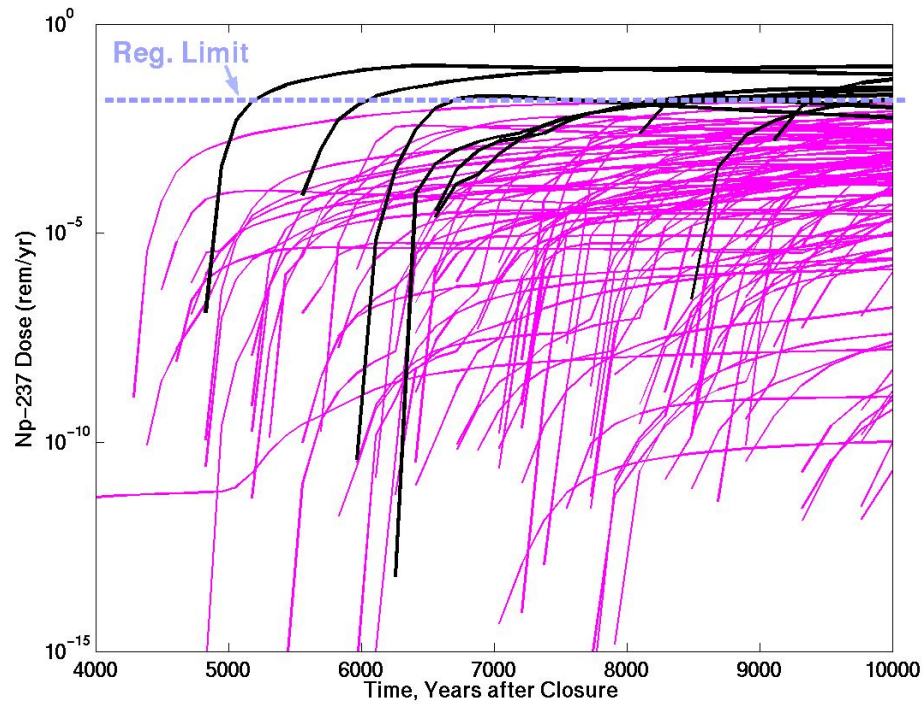
### Step 2a. Find SID scenarios

To reduce the dependence on particular LHS samples, we ran the TPA again with a LHS size of 200, calculating doses from the Cm-245/Am-241/Np-237 radionuclide chain only. Of these 200 realizations, 9 (4.5%) crossed the reference point of 15 mrem/yr. Figure 2 shows the curves for the 120 realizations that resulted in non-zero doses within the 10,000-yr time frame; the nine realizations that produced SIDs are shown in black. Figure 3 shows summary statistical measures for all the realizations. The mean curve (and percentile curves) in Figure 3 is not the outcome of any particular set of parameter values; rather, the mean is constructed by calculating for each time on the x-axis the average of all doses from the multiple realizations, and then connecting these points.

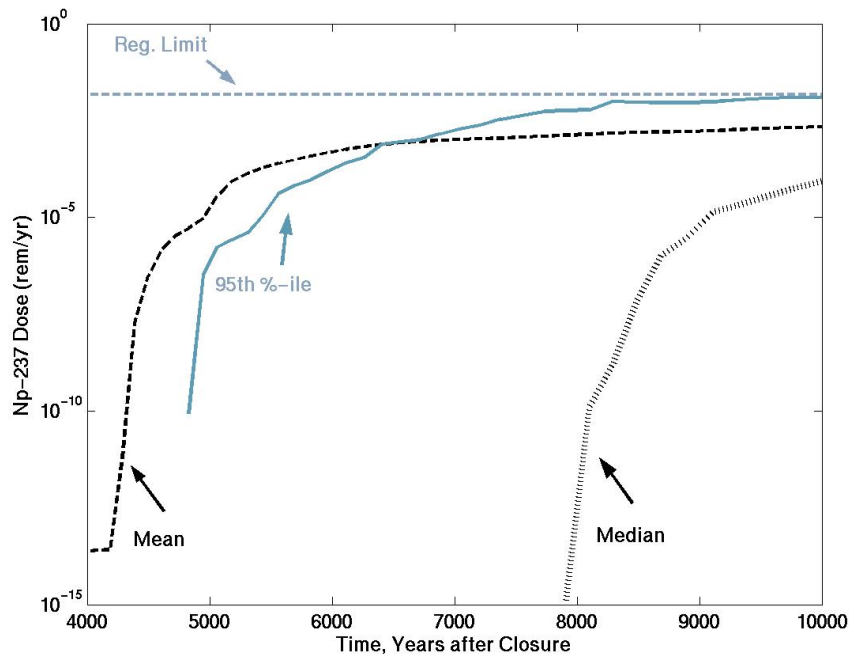


<sup>4</sup> A key term in the spent fuel dissolution model, PSFDM1 (SFD), is the pre-exponential spent fuel dissolution term in fuel dissolution model 1 (which is the default model in the PA).

**Figure 1. Generalized Sensitivity Analysis (GSA) based on Cumulative Distribution Function (CDF) partitioning for the Pre-exponential term in the Spent Fuel Dissolution Model 1. CDF for SIDs plotted in the blue solid line; CDF for success values is plotted in the dashed black line.**



**Figure 2. Np-237 dose (rem/yr) to 10 km-receptor using TPA 4.1j Code, 200-realization base case. SID realizations in black, 'successes' in magenta.**



**Figure 3. Np-237 dose (rem/yr) summary measures for 200-realization base case in Fig. 1.**

For the 9 SID realizations shown in Figure 2, we checked the sampled values for the key attributes identified in step 1, in order to find an explanation for the SIDs; these are shown in Table 1. On inspection, it appears that the WP flow factor (WPFF), Initial WP defect rate<sup>5</sup> (IWPD), and Np solubility<sup>6</sup> (SOL) have a significant potential to create large doses if high percentiles of these parameters are sampled together. At the same time, these factors alone do not determine whether there will be a high dose because there were plenty of realizations where high values of these parameters were sampled without resulting in a high dose. In order to build a class of scenarios based on this, we can perform conditional sensitivity analyses on the remaining attributes, given high percentiles sampled for these three parameters, to find repository features that may save us from the challenge posed by high WPFF, IWPD, and SOL. The purpose is to build an understanding of what poses a challenge to the repository system, and what repository attributes can mitigate these challenges, and try to quantify the probabilities of different possible scenarios.

**Table 1. Distribution percentiles for parameter values sampled for six key uncertain parameters in four worst realizations in 200-realization base case for Np-237 Dose to 10-km receptor**

Np-237 Dose (mrem/yr)	Infiltration @start (Infil)	WP Flow Factor (WPFF)	Sub-Area Wet% (SAW%)	Initial WP Defects (IWPD)	Np-237 Solubility (SOL)	SF Dissolution Term (SFD)
101	19%	99%	36%	98%	97%	90%
97	8%	97%	43%	80%	96%	94%
48	97%	95%	99%	69%	56%	38%

<sup>5</sup> The initial waste package defect rate, WPDef% (IWPD), is the percentage of waste packages that are initially defective (assumed because of weld manufacturing defects), and hence this percentage of waste packages have the only inventory that is available to become the source term before general corrosion begins to fail waste packages tens of thousands of years into the future.

<sup>6</sup> The solubility of Np-237, SolbNp (SOL), is an important parameter because in the default spent fuel dissolution model, the solubility limits how much Np-237 can enter water contacting the spent-fuel waste form.

30	35%	94%	94%	96%	89%	45%
----	-----	-----	-----	-----	-----	-----

### **Step 2b. Find savior attributes**

We wish to find what set of attributes can save us from the challenge of high IWPDP, WPF, and SOL. As a basis for our conditional sensitivity analyses, we ran the TPA again, with a LHS size 50 starting with the condition that the WPF, IWPDP, and SOL values lie in the top 5% of their pdfs (i.e.,  $F(x) > 0.95$ ), and that the drip shield fails before 10,000 years (the first 82% of the DSFailTi distribution), since the dose is zero in the base case if the drip shield is intact for the first 10,000 years. We used GSA for single parameters, both step-wise and on the entire set of realizations. For the step-wise case, we first found which parameter could explain best the successful realizations out of the whole set, and removed the realizations 'saved' by this parameter; then found which parameter could best explain the successes in the remaining set of realizations, and removed those realizations explained by it; and so on, until single parameters could no longer explain the success/SID divide in the remaining realizations. Using GSA for single parameters yielded the following insights: (1) the drip shield, which is a barrier (like an umbrella) to water reaching the waste package, obviously must fail before the 10,000-year compliance period. The dose is zero for the 18% of realizations where the drip shield failure time (DSFT) is beyond 10,000 years. And, with a high degree of confidence, we find that the peak dose will not exceed 15 mrem if; (2) SFD, a key term in the spent fuel dissolution model lies within the first 26% interval of its distribution; (3) SubAreaWet% (SAW%) is less than 10%; and (4) SFWt% < 10%. But this still leaves a lot unexplained. So we looked at how sets of parameters might explain the results further.

Scatter plots of SIDs and successes for two parameters at a time can help identify threshold effects, similar to those identified by GSA for single parameters. It makes sense that some parameters working in concert will become more important than either one alone – for example, while the rain infiltration parameter (AAMI@s, or Infil) alone did not yield any threshold effects, the infiltration considered along with the percent of the subarea that is wet (SAW%) might. Using scatter plots two parameters at a time, we found that the rain infiltration parameter combined with either the subarea wet % or condensate moving towards the repository (FOCTR) parameters, *did* exhibit threshold effects. The scatter plot of AAMI@s (Infil) and SAW% shows the absence of SIDs in the lower left corner of the sample space could indicate a threshold effect for combinations of low values for these two parameters. This adds the insight that, with a high degree of confidence, we find that the peak dose will not exceed 15 mrem/yr if; (5) Infil. < 8.5 mm/yr and SAW% < 45%; or (6) Infil. < 8.5 mm/yr and FOCTR < 25%. Two of the 50 realizations could not be explained by these savior attributes. Most likely we would have to seek out higher order explanations for these, e.g., numerous parameters working in concert to prevent repository SIDs (which of course is what a multiple barrier repository system is designed to do.)

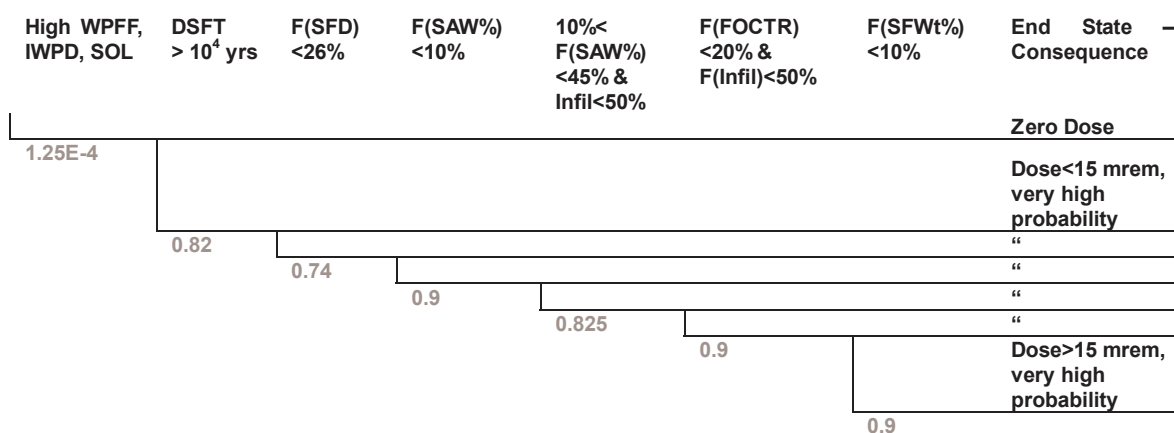
In order to confirm the above hypotheses about savior attributes, we compared them with GSA findings from the 500-realization base cases, and we tested some of them using smaller LHS sizes (~20); e.g., we forced the key spent fuel dissolution term (SFD) into the bottom 30% and forced the other parameters into the ranges of their respective distributions where SIDs were possible, and checked whether the postulated savior attribute did still prevent repository SIDs. While this is not enough to say with certainty that our hypotheses are indeed true for all cases (the entire response surface,) it does increase our confidence in the results.

Of course this is not the only class of scenarios, there may be many other ways for the repository to produce SID; this set is just an example. Using these results, we can move to step 3.

### **Step 3. Explain reasons for SIDs and display results in SPARC tree**

Based on the above analyses, we can construct the tree shown in Figure 3, for the scenario that begins with the challenge that IWPDP, WPF, and SOL are all in the top 5% intervals of their respective pdf's. The SPARC tree displays the results that, even if the WPF, IWPDP, and SOL were all in their worst 5% intervals of their respective pdfs (the initial postulated challenge) the repository produces SIDs *only if* DSFT > 10<sup>4</sup> years and F(SFD) > 0.26 and F(SAW%) > 0.1 and F(SFWt%) > 0.1 and the savior conditions for Infil with SAW% or FOCTR are not true, as shown on successive branches. The probability for each branch that can lead to SID is shown below the branch.

The preliminary probability estimates for the undesirable end state from this class of scenarios that starts with high WPFF, IWPDP, and SOL is on the order of  $(0.05)^3 \cdot 0.82 \cdot (0.6)^2 \cdot 0.5 \cong 2E-5$ . One may ask why we should bother with these analyses, since our preliminary assessment shows a low probability of exceeding the decision threshold. One answer is that we may be wrong about the low probabilities we assign to some of the challenges, because of incompleteness. And, if we make investments to strengthen the repository system, we should target risk-significant areas, including areas of large important uncertainties. An example of possible incompleteness is the potential concern raised by stakeholders about localized WP corrosion in the USDOE's high-temperature (HTOM) design [10]. The USDOE PA assumes that more than 1 early WP failure is very unlikely, but not all stakeholders are convinced that this will always hold true. We can use the early weld failure mode in the TPA 4.1j code as a proxy for small holes created by localized corrosion and test sensitivity of repository performance to changes in the IWPDP and the localized corrosion concern. The PA base case average dose can get to 15 mrem if we postulate early WP failures lie between 0.01% and 5.5%; whereas the current IWPDP distribution lies between 0.01% and 1%. But the localized corrosion concern also implies that the number of early WP failures and amount of water reaching the WP would no longer be independent, since water and humidity contribute to potential localized corrosion. This synergistic effect could increase the probability of early WP failure. Based on our preliminary scenarios above, we may be more confident in the repository's capability to mitigate these challenges if we improve our knowledge about the distribution of subarea wet fraction and the key parameter in the spent fuel dissolution model, since these have a large capacity to mitigate the flow focusing and high waste package defect challenges. A risk-informed research program could focus on these attributes for example.



**Figure 4. SPARC tree for class of scenarios starting with WPFF, IWPDP, and SOL in the worst 5% intervals of their respective parameter distributions. The probability for strategic partitions for each repository attribute (represented by the appropriate parameter) is shown below the branch.**

## DISCUSSION

The SPARC method presented in this paper explicitly constructs scenarios of risk (i.e., SID-producing) for a HLW repository system by identifying the uncertain parameter-assumption space that has a very high probability of resulting in SIDs or preventing SIDs (savior attributes). While other sensitivity and uncertainty analysis methods identify parameters that are *generally* important at the total-system or sub-system level, the SPARC method identifies the *specific* combined distribution-intervals of these important parameters and assumptions that produce or prevent SIDs.

Findings from the SPARC method could be used as the first element in risk-informed integrated decision-making (RIIDM) for the HLW programs. The NRC has outlined a generic RIIDM approach for nuclear reactors [10] which could be useful for waste repositories. Since the supplementary analysis helps identify important uncertainties and repository attributes, it paves the way to defining how other RIIDM elements, such as various defense-in-depth approaches, could be used to compensate the potential adverse effects of these important uncertainties.

The supplementary analysis here could be extended to off-normal disruptive scenarios as well.

## ACKNOWLEDGMENTS

The authors would like to thank the numerous individuals who provided valuable feedback on early parts of this work, in particular B. John Garrick, George Hornberger, Milt Levenson, Michael Ryan, John Larkins, Timothy McCartin, Robert Budnitz, and Richard Lester. The views expressed in this paper and any potential errors are the authors' only.

## REFERENCES

- [1] Ghosh, S. T., and G. E. Apostolakis, 2002, "Assessing Structural Model Uncertainty in Performance Assessments," *6th International Conference on Probabilistic Safety Assessment and Management (PSAM 6)*, San Juan, PR.
- [2] Hornberger, G. M. and R. C. Spear, 1981, "An Approach to the Preliminary Analysis of Environmental Systems," *Journal of Environmental Management*, 12, pp. 7-18.
- [3] Ghosh, S. Tina and George E. Apostolakis, 2006, "Extracting Risk Insights from Performance Assessments for High-Level Radioactive Waste Repositories," *Nuclear Technology*, Vol. 153, pp. 70-88.
- [4] U.S. Environmental Protection Agency, "Public Health and Environmental Radiation Protection Standards for Yucca Mountain, Nevada," *U.S. Code of Federal Regulations*, Title 40, Chapter I, Part 197.
- [5] Codell, R. B., M. R. Byrne, T. J. McCartin, S. Mohanty, J. Weldy, M. Jarzemba, G. W. Wittmeyer, Y. Lu, and R. W. Rice, 2001, *System-Level Repository Sensitivity Analyses Using TPA Version 3.2 Code*, Center for Nuclear Waste Regulatory Analysis.
- [6] Helton and Davis, 2003] Helton, J. C. and F. J. Davis, 2003, "Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems," *Reliability Engineering and System Safety*, 81, pp. 23-69.
- [7] Center for Nuclear Waste Regulatory Analysis, 2002, *Total-System Performance Assessment (TPA) Version 4.0 code: Module Descriptions and User's Guide*.
- [8] McCartin, T.J., 2003, "Understanding Performance Assessment Results," *Presentation to the U.S. Nuclear Regulatory Commission's Advisory Committee on Nuclear Waste*, Washington, DC, 26 March.
- [9] Mohanty, S. R. Codell, J.M. Menchaca, R. Janetzke, M. Smith, P. LaPlante, M. Rahimi, and A. Lozano, 2002, *System-Level Performance Assessment of the Proposed Repository at Yucca Mountain Using the TPA Version 4.1 Code*, CNWRA 2002-05 Rev. 1, Center for Nuclear Waste Regulatory Analyses, San Antonio, TX.

[10] U.S. Nuclear Regulatory Commission, 2002, Regulatory Guide 1.174, *An Approach for Using Probabilistic Risk Assessment in Risk-informed Decisions on Plant-specific Changes to the Licensing Basis*, Revision 1, November.

[11] U.S. DC Circuit Court of Appeals, Judges Edwards, Henderson and Tatel, 2004, *NEI v. EPA*, Nuclear Energy Institute, Inc., State of Nevada, Natural Resources Defense Council, et al., Petitioners, versus Environmental Protection Agency, Department of Energy, Nuclear Regulatory Commission, Respondents; Petitions for Review of Orders of the Environmental Protection Agency, the Department of Energy, and the Nuclear Regulatory Commission, July 9.