γ.

CONF-840647--17

UCRL- 90514
PREPRINT

UCRL--90514

DE84 011051

Observations on Seismic Probabilistic Risk Assessments (PRAs)

Paul D. Smith

MASTER

ASME Technical Session on Seismic Event
Probabilistic Risk Assessment (PRA)
San Antonio, TX

June 1984

Lawrence Livermore National Laboratory

DISCLAIMER

OBSERVATIONS ON SEISMIC PROBABILISTIC RISK ASSESSMENTS (PRAs)*

Paul D. Smith
Lawrence Livermore National Laboratory
Livermore, California

INTRODUCTION

The question formerly asked: "Should we perform probabilistic analyses of the seismic behavior of nuclear power plants?" is now moot. These analyses are now performed; in various ways; by utilities, regulators and researchers; for both generic and specific purposes; at large and small levels of effort; and in a number of different countries. More important, some of the studies suggest that seismic-initiated accidents dominate the public risk. Most important, there is a growing awareness that these analyses provide safety insights that are not easily obtained otherwise and that this is so in spite of weaknesses in the state-of-the-art. Fortunately research on seismic PRA methodology has been underway for some time. In our paper we present observations on a number of issues related to seismic PRAs including:

- Strengths and weaknesses of seismic PRAs.
- Uncertainty, sensitivity and variability.
- Common misconceptions.
- Possible improvements in seismic safety acceptance criteria.
- Recommended modifications to the NRC approach to safety goals.
- Major problems.

Some specific examples are provided.

STRENGTHS OF SEISMIC PRAs

First, because of the nature of the information sought for inputs to seismic PRAs, there is a focus on the following question: "How could this structure, equipment or component fail?" This is contrasted with the traditional engineering question: "Given these loads, how do we design against the failure of this structure, equipment or component?" This is a strength because we must: (1) identify possible failure mechanisms and (2) estimate failure levels. At a minimum this exercise provides insight on plants and seismic safety acceptance criteria that normally would not be obtained and this is true even if the analysis is not carried further to the system aspects of seismic PRAs.

Second, seismic PRA needs force a comprehensive, systematic and disciplined search for accident scenarios, or sequences of failures of structures, equipment, components as well as impact on evacuation procedures that could place the public at risk. This is a strength since we estimate the probable real plant performance where: (1) the conservatism used in safety evaluations are removed from the analysis where this is possible and (2) constraints used in safety evaluations are removed. For example, for (2) we typically will assess the effect of earthquakes that are larger than the earthquakes used in the safety evaluation.

Third, seismic PRA needs enhance communications between the parties who contribute to the design and safety evaluation. This includes: seismologists, geologists, and geotechnical, structural, mechanical, electrical and system engineers as well as owners, operators, maintenance personnel and architect and engineers. This is a strength because good communication of this type has probably not occurred before in the traditional design or safety evaluation process. This is also one of the most difficult areas to do well in performing a PRA or its review because of its difficulty, high cost and lack of a tangible product.

Fourth, a seismic PRA provides an overall measure of probable plant performance in terms of a single number, for example, probability of core melt or man-rems. This is a strength because such a capability provides the opportunity to simply summarize the conflicting factors that may affect the safety evaluation, backfit options or possible changes in safety criteria.

In spite of the deficiencies in the state-of-the-art of seismic PRAs (and they are manifold) it is clear that they are a powerful new tool that can be used effectively. A non-technical strength is that PRAs offer a path to insert more reason into the traditional adversary interactions between the regulator and the regulated and possibly reduce the degree of negative interactions. Utilities can also use the probability of core melt to determine whether the resultant financial risk is acceptable - and

*This work was supported by the United States Nuclear Regulatory Commission under a Memorandum of Understanding with the United States Department of Energy.

determine which backfits are economically prudent separate from any safety requirement.

WEAKNESSES OF SEISMIC PRAs

All PRAs have a weakness in the completeness issue. That is: "Are we sure that the PRA is complete in that all important failure modes, accident initiators, accident sequences, etc. have been included in the analysis?"

All PRAs are deficient in that the absolute value of the results (for example the probability of core melt) cannot be relied upon as the sole PRA result or relied upon uncritically. This complicates the most straightforward use of PRA results - comparison with a goal to make safety decisions. This leads us to state our uncertainty on our results but this step does not fully resolve this issue.

This is not a complete list of weaknesses or a full exposition of the ramifications of these two. Other problem areas will be described in the sequel.

NRC SAFETY GOAL AND SEISMIC RISK

There is a fundamental factor that distinguishes seismic and indeed all natural and external hazards as possible accident initiators from accident initiators that arise from causes within the plant:

> The public is at risk from earthquakes whether or not the nuclear plant exists.

We call the public risk that exists assuming the nuclear plant does not exist the background seismic risk. The additional* risk due to the nuclear plant is called the incremental seismic risk. As an alternative or supplement to the present thinking on the safety goal the acceptable incremental seismic risk could be specified in terms of some "small" percentage of the background seismic risk. This approach has a number of advantages:

- Engineers familiar with the seismic design of nuclear and conventional facilities can give many examples of how the design of nuclear power plants is much more conservative than the design of conventional facilities. This fact often leads to an enhancement of the negative interactions between the regulator and the regulated. A comparison of the background and incremental seismic risk would explicitly quantify the degree of this conservatism. This comparison would be of great assistance in developing informed professional judgment as to whether or not regulatory concerns are well founded.

- A comparison of background and incremental seismic risk would provide the public, the NRC and utilities with precisely the information needed for decisions on whether or not reducing the incremental seismic risk is very effective in reducing the total public seismic risk (or where earthquake hazard reduction money should best be spent).

*We assume that the conventional facilities that are replaced by the nuclear ones are risk-free.

- A destructive earthquake is a large-scale event, particularly if it occurs in the Eastern United States. This may mitigate against the public being adverse to a nuclear accident, at least relative to the non-nuclear aspects of a large destructive earthquake, especially if the background and incremental risks are explicitly described in detail.

- The comparison between background and incremental seismic risk will be more credible than a comparison of nuclear seismic risk with a numerical goal. This is because one of the largest contributors to uncertainty in estimates of seismic risk, that due to uncertainty in estimates of the hazard, will tend to be common for both background and incremental seismic risk. The error in our estimate of the hazard will tend to be biased in the same direction for the background and incremental risk and as discussed in the sequel this improves the quality of the comparison.

- Both background and incremental seismic risks are involuntary.

If we follow the existing NRC safety goal philosophy we might infer that the incremental seismic risk should be 0.1% or less of the background seismic risk. However, if we consider that the above advantages provide a means to increase the usefulness or credibility of seismic PRAs for safety goal purposes compared to the same use of PRAs from internal initiators, we might on the other hand infer that the acceptable incremental seismic risk is greater than 0.1%.

There are two fundamental problems associated with seismic PRAs.

First, while there has been much concern and discussion on the large size of the uncertainty in seismic PRA results, the implications of the nature of the sources of uncertainty are equally if not more important and are rarely discussed. The complications introduced by consideration of the nature of the sources of uncertainty are pervasive and far-reaching, but only one specific example will be given.

One goal (not a safety goal) that has been identified as important in the past is to determine the "seismic contribution to reactor risk*" (Refs. 1 and 2). Indeed, the conclusion in some commercial PRAs has been that seismic initiators dominates the public risk.

This goal is unattainable and such a conclusion is speculative rather than rational.

To see why this is so first consider the common presentation of PRA results shown in Fig. 1. Here we have shown the annual probability of core melt due to seismic and internal initiators in terms of the median values and the 10% and 90% confidence levels.

*In these terms the concept of background and incremental seismic risk can be stated as "reactor contribution to seismic risk."

This type of presentation easily leads to misinterpretation. For example, one conclusion that might be drawn from Fig. 1 is that the seismic initiator is not a significant contributor to the probability of core melt. It is not rational to draw this conclusion from these results. More important this conclusion may be completely incorrect. The seismic initiator may in reality completely dominate the probability of core melt at the plant for which the results in Fig. 1 were obtained.

The justification for this assertion will become clear if we first describe the results in Fig. 1 more carefully. For both the seismic and internal initiators the most robust meaning of the results in Fig. 1 is:

> "We do not know what the true value of the annual probability of core melt is but its median is as shown in Fig. 1 and we have 80% confidence that the true value is somewhere in the interval shown in Fig. 1."

Our true state of knowledge is even weaker than suggested by this statement but an exposition of this point is not germane here.

Our inability to describe the annual probability of core melt precisely or the reason that intervals are shown in Fig. 1 rather than a single value is because of the existence of what we have called modeling uncertainty (Ref. 3) but there are also many other terms used to describe this uncertainty.

For seismic-initiated core melt sources of modeling uncertainty include:

- Uncertainty in our models of the hazard (Fig. 2).
- Uncertainty in our models of the dynamic response of the site and plant due to an earthquake.
- Uncertainty in our models of human behavior.
- Uncertainty in our models of failure scenarios or accident sequences (event-tree/fault-tree models).

For internal-initiated core melt, sources of modeling uncertainty include:

- Uncertainty in our models of accident initiators.
- Uncertainty in our models of failure rates for the components of a plant.
- Uncertainty in our models of human behavior.
- Uncertainty in our models of accident sequences.

These are not complete lists of categories of sources of modeling uncertainty.

There are similarities in the categories in each list. Both lists include uncertainty in models of accident sequences for example. We have little confidence however that our bias or error will be the same for the two different accident initiators. For example we have no reason to believe that if the true level at which a component will fail due to earthquake loading is actually higher than we have estimated that the true random failure rate of this same component is also lower than we have estimated. Even if our errors are in the same direction we do not know that their magnitudes are the same or that their effect on estimates of the probability of core melt will be similar. Many other examples could be given. The point is that the uncertainties in the two PRAs (seismic- and internal-initiated) are probably widely unrelated.

This is the fundamental problem: The two uncertainties (the one on seismic- and the other on internal-initiated core melt) shown in Fig. 1 are probably widely unrelated. This means that we cannot draw any conclusions on order or on which initiator dominates the annual probability of core melt*.

This fundamental problem cannot be solved to the extent that it would not be a regulatory problem. This does not mean that it is an unpenetrable barrier to effective regulation. All that is required is to change the objective through decoupling. By this we mean that the determination of the seismic contribution to reactor risk should not be set as a necessary input to regulation (it is not now part of the stated goal). We should instead simply determine if the seismic risk satisfies some explicit safety acceptance criteria that are not related to risks from other initiators, except perhaps quantitatively.

Second, data on the observed performance of nuclear power plants during earthquakes is virtually nonexistent. Further, a destructive earthquake is a rare event and it is an even rarer event for one to occur at the site of a nuclear power plant. This means that we cannot expect to obtain sufficient data by observing the field performance of nuclear power plants to confidently state that we know what performance to expect. Most important, since observation of field performance is the only way to assess the design adequacy of large one-of-a-kind facilities, this means that there is an unremovable residual uncertainty on how nuclear power plants would perform during and after an earthquake.

This is contrasted with the ever-increasing data base that is being obtained on the operational performance of nuclear power plants. This includes failure rates, precursor accident sequences, operator performance, design, construction and maintenance errors, and so on. This is precisely the kind of information that is needed for internal-initiated PRAs. While we will always desire more and better such data we at least have some confidence because we are gathering, assessing and using it in our analyses.

This is also contrasted with the state-of-the-art of the design of conventional facilities against the effects of earthquakes. Numerous large earthquakes have led to observations of the performance of numerous conventional facilities. These data have led to many revisions in design codes. There is probably not a single example of where the observed performance of a nuclear power plant in an earthquake led to a revision of nuclear design practice. Lastly, in spite of the large data base on the observed performance of

*Of course, the same conclusion may also be reached considering risks from any natural or external hazard and their comparison with each other or with internal initiators.

conventional facilities, significant experimental research continues.

These observations lead to the following conclusions:

- In the present budget era, the NRC will not be able to afford the experimental seismic research that is required to overcome these deficiencies.
- There should be greater emphasis on coordinated seismic research, both within the United States and internationally.
- More emphasis is required to insure that the various research efforts will fit together in a coherent and efficient way.

In the above we have discussed some fundamental problems in applying the NRC safety goal for seismic-initiated accidents and suggested perhaps an impractical way of overcoming some of the problems through the use of background and incremental seismic risk. We will now flesh out one of the problems and give an example of how seismic risk-related safety acceptance criteria can be devised to overcome regulatory problems.

A more detailed presentation of uncertainty in estimates of annual probability of core melt is shown in Fig. 3 (Ref. 4). As shown in this figure the uncertainty in the annual probability of core melt is about three orders of magnitude between the 10% and 90% confidence levels. The single most important contributor to this uncertainty is the uncertainty in the seismic hazard shown in Fig. 2 by the multiple curves.

The importance of the uncertainty in the seismic hazard leads to the following problems:

- Wide uncertainties on seismic-initiated risk make it difficult to conclude that a numerical goal has been met.
- The importance of uncertainty in the seismic hazard as a contributor to uncertainty in estimates of the annual probability of core melt leads to the uncomfortable appearance that safety decisions based on seismic PRAs are good or poor decisions depending on how capricious the unpredictable natural phenomena of an earthquake is rather than on sound decisions on plant design, operation and maintenance.

The following criteria are offered as a model that to a great extent overcomes these two problems as well as others that are not discussed here. The proposed seismic risk safety acceptance criteria are stated only in terms of core melt to simplify the presentation. Similar criteria could be developed for other measures of risk such as prompt fatalities.

There are two criteria:

1. The mean annual probability of core melt shall be shown to be less than some numerical goal, presumably 1E-4.
2. At least 90% of the mean annual probability of core melt shall arise from accelerations greater than twice the SSE value. Conversely, no more than 10% of the mean annual probability of core melt shall arise from accelerations less than twice the SSE value.

In the first criteria the use of the mean accounts for the large uncertainties in seismic PRA results. This is because the use of the mean will provide a conservative bias in terms of probability of exceedance. For example, the SSMRP mean of 2E-4 has a probability of about 0.2 of being exceeded, see Fig. 3. This is the traditional way to deal with uncertainty, that is, introduce conservatism. The use of the mean will also result in a heavy weighting of the most conservative hazard curve (the top one in Fig. 2). This criteria is closely related to the present thinking on a safety goal. The present presumed goal is 1E-4 but our specification of the mean is not now part of the specified goal.

The second criteria is predominately related to engineered features of the plant and it insures that these features satisfy some risk-related criteria regardless of uncertainty in the seismic hazard. This criteria also provides a simple relationship between seismic risk and a traditional deterministic seismic parameter (the SSE acceleration) and thus provides a practical means of interpreting the meaning of achieving the first criteria. In a crude way this criteria also provides a means to relate the background and incremental seismic risk. This criteria is not now part of the safety goal.

Another interpretation of these two criteria is that the first provides a means to design for the threat and the second provides a practical means to design for uncertainty in the threat. This provides a good example of how (perceived) barriers to effective regulation can be overcome. Other examples could also be given.

A COMMON MISCONCEPTION

We will give a single example of a common misconception of seismic PRAs.

The following question is often asked:

"If the uncertainty in the seismic hazard is the single most important contributor to uncertainty in our estimates of the annual probability of core melt, why should we expend anything but a minimal effort in the fragility area? These efforts are expensive and will not significantly reduce our uncertainty."

The presumed answers are (1) on the part of the NRC a desire to insure that regulation is effective and (2) on the part of utilities a desire that backfit decisions or new design requirements be cost-effective.

To see why this is so consider the following example. Assume that one of the conclusions of the NRC review of a seismic PRA is that the seismic risk is too high and must be reduced. Assume further that we would prefer to be efficient in any modifications to the plant and that we would want to strengthen a (few) key component(s) rather than strengthen the entire plant or a significant portion of it. Uncertainty in our estimates of the seismic hazard has very little to do with a decision as to which component to strengthen. This is because uncertainty in the seismic hazard primarily affects our ability to describe the risk quantitatively. In addition, while we may not be able to describe the true hazard we do know it is the same for every component of the plant. The plant risk can be reduced only by changing some physical element, operational or maintenance procedure, etc. Uncertainty in estimates of fragility now become most important. This is because if we strengthen a component that we believe is the weak link but is really relatively strong we may not reduce the risk at all.

Presumably the NRC does not want the utility to backfit if it is ineffective in reducing risk and neither does the utility. The key question is: When is it more cost-effective to backfit a significant portion of the plant to insure that the risk is truly reduced and when is it more cost-effective to develop more accurate estimates of the fragilities so that we can justify backfitting a more limited portion of the plant?

CONCLUSIONS

It is clear that seismic PRA represent a powerful new tool that opens new and intriguing vistas for the management of risk, regulation and research. The task that lies before us is to develop logical and practical uses of seismic PRAs while at the same time we restrain a natural tendency to oversell their capabilities.

REFERENCES

1. January 11, 1983 letter from ACRS to NRC Chairman Palladino.

2. January 18, 1984 letter from ACRS to NRC Chairman Palladino.

3. P. D. Smith, R. G. Dong: Seismic Safety Margins Research Program (Phase I) Interim Definition of Terms, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53001, December 1980.

4. M. P. Bohn, L. C. Shieh, J. E. Wells, L. C. Cover, D. L. Bernreuter, J. C. Chen, J. J. Johnson, S. E. Bumpus, R. W. Mensing, W. J. O'Connell, D. A. Lappa: Application of the SSMRP Methodology to the Seismic Risk at the Zion Nuclear Power Plant, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53483 (also published by the U.S. Nuclear Regulatory Commission, NUREG/CR-3428) May 1983.

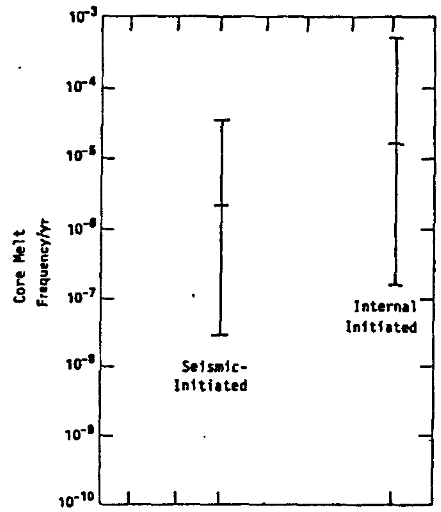

Figure 1.

Typical presentation of annual probability of core melt from seismic and internal initiators. The vertical band indicates the 10% and 90% confidence interval. The value of the median within this interval is indicated. As discussed in the text this type of presentation easily leads to misinterpretation. This is because the placing of these two results on the same diagram leads the observer to draw conclusions on the relative importance of the two initiators on core melt through a comparison of the relative vertical position of the results. This comparison is not valid because the true value of the annual probability of core melt is unknown and the sources of the uncertainties that lead to the confidence intervals are probably widely unrelated for the two initiators.

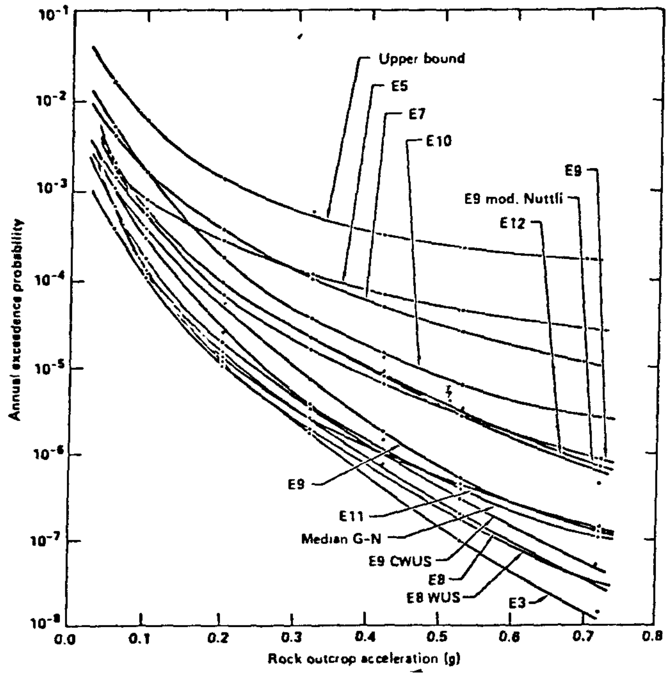

Figure 2.

Typical presentation of the results of a hazard analysis (from Ref. 4). Each of the curves displayed describe the randomness associated with estimating acceleration and is a complimentary cumulative distribution or a hazard curve. The existence of more than one hazard curve is a presentation of the results of an uncertainty analysis and these curves together describe what we have called modeling uncertainty (Ref. 3) for the hazard.

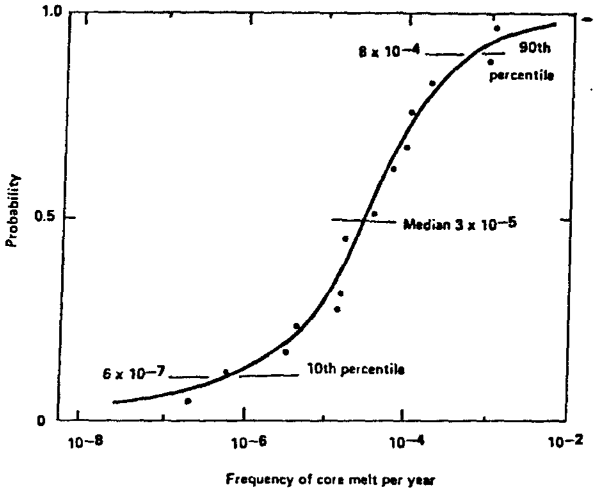

Figure 3.

SSMRP results for the annual probability of core melt at the Zion Unit 1 plant (from Ref. 4). The 10% and 90% confidence level values are shown as is the median. The mean value is 2E-4. These results were obtained from 14 complete risk analyses and these 14 results are also shown in the indicated data points. The curve shown is a lognormal cumulative distribution function fit to the 14 results. This curve is one way to describe the modeling uncertainty in our estimates of the annual probability (frequency) of core melt. The single most important contributor to this uncertainty is the modeling uncertainty in the hazard (see Fig. 2).