Transactions*, SMiRT 21, 6-11 November, 2011, New Delhi, India** ***Div-VII: Paper ID# 189

OPPORTUNITIES FOR ADVANCING TECHNOLOGY-NEUTRAL AND PERFORMANCE-BASED DESIGN METHODS FOR THE SEISMIC DESIGN AND REGULATION OF SSCs AT NUCLEAR POWER PLANTS

Robert J. Budnitz
Lawrence Berkeley National Laboratory, University of California, Berkeley CA 94720 USA
E-mail: RJBudnitz@lbl.gov

ABSTRACT

The current NRC regulations for design and analysis of nuclear power plants to resist large earthquakes use a framework that is partially risk-informed, in the sense that a target performance goal of 10^{-5} per year is the design target for the design of every individual SSC (structure, system, or component) that contributes significantly to the safety performance of the plant. However, the current framework does not admit design-specific or plant-specific PSA (probabilistic safety assessment) information directly as a part of the technical basis used to determine whether an SSC should be approved. Instead, the design rules and analysis provisions have used information from the body of seismic PSAs already in the literature to inform how the design process and analysis provisions themselves are framed. This is "risk informed" but not fully so, and it is also not fully performance-based because although the target is framed in probabilistic terms, most of the design rules are prescriptive, rather than allowing the designer to choose his/her own design approach. This paper will discuss a group of several proposals, any one of which could advance the situation significantly toward a more fully performance-based and risk-informed framework. This paper will discuss the technical basis for each of the several proposals, what valid reasons stand in the way of their early implementation, and what research could be undertaken to help move the seismic design and approval process along toward a more nearly risk-informed and performance-based framework.

INTRODUCTION

The current approach to assure that a new nuclear power plant (NPP) is properly designed to resist large earthquakes is embodied in the American Society of Civil Engineers consensus standard ASCE-43-05 [1]. Parts of this standard have been endorsed by the U.S. Nuclear Regulatory Commission (NRC). The standard uses a framework that in several important aspects is significantly different than the time-honored approach under which all operating NPPs were designed and analyzed. This new framework is partially risk-informed, but not fully so. In this paper the framework will be described, and opportunities will be discussed for advancing toward a framework with even more risk-informed features. This is a summary --- a more detailed paper [2] is available that provides a more extensive description of the subject summarized here.

In general, if a performance-based approach is contemplated, and if using explicit risk-type information is accepted as part of the intellectual basis for regulating seismic safety of NPPs, the "framework" for regulation would comprise a unified set of *performance criteria, acceptance criteria, prescriptions for analytical methods*, and *design criteria*. For each of these, the developer of a standard or the regulatory decision-maker has various options at his/her disposal. The main thrust of this paper is three fold. It will briefly discuss these options, then it will cover the choices that are embedded in the current approach, and finally it will discuss opportunities presented by the possibility of different choices being made in the future.

ELEMENTS OF THE FRAMEWORK

Performance of what?

In the context of this discussion, we are considering either the "performance" of an individual SSC (structure, system, or component), or alternatively the performance of a large complex system as a whole. Here "performance" generally means performance of whatever function the SSC or the overall complex system is designed to do or expected to do. The obverse or flip-side is that there is a "risk" of a "failure to perform." In this context, the "end point" of any discussion of risk contemplates just such a failure to perform. Such a failure is termed "unacceptable performance."

Performance Criterion: The most nearly "risk-informed" performance criterion would be requiring that the system under consideration not exceed an *annual frequency of unacceptable performance*. This risk-type criterion would be an integrated risk that represents an integration over the probabilistic seismic hazard and the probabilistic seismic fragility curve to yield an overall frequency of failure. An example could be a performance criterion not to exceed a frequency of failure of 10^{-5} per year. A second option would be some *surrogate* for the above. An example could be that the probability of unacceptable performance must be less than X% for a prescribed earthquake that would occur with frequency F per year (e.g., a 1% probability of "failure" for a 10^{-4}/year earthquake).

The selection of the above criterion is a policy choice, not a technical choice, involving how much protection is being sought. In this sense, it is imposed on the design, and on the designer, by an external decision-maker, such as a regulator, a consensus code committee, or the facility's owner.

Confidence level: The *required degree of confidence* must also be prescribed, usually in terms of a level of confidence that the criterion is met (median confidence? 95% confidence? etc.) This is another policy choice, not a technical choice. Crucially, this choice is closely related to the uncertainty in the analysis, which itself is closely related to how much "margin" a decision-maker (a consensus code committee, a regulator, or the owner of the facility) requires in order to conclude that the design will be "safe enough."

Performance – "success" vs. "failure": Is the "performance" in the probabilistic performance criterion a realistic definition of "success" vs. "failure" to perform the desired function, or is a conservative end-point chosen? (An example of the latter is to select the onset of modest inelastic deformation as the conservative definition of seismic-induced "failure" for a shear wall.)

End point for "failure": One SSC at a time, or failure of an entire complex system: For a nuclear reactor, given that the actual end-point of concern is preventing or mitigating "core damage", is the end-point of the design criterion the performance of a *single SSC* (one at a time), or is the end-point the "failure" of a *full accident scenario* to be kept within the target? At least three different options exist here: one SSC at a time? or one such scenario at a time? or by an integrated accounting for the full ensemble of accident scenarios of the whole complex system?

Analysis –realistic or requiring conservatisms: Is the analysis that is used to demonstrate compliance required to be *"realistic"*, or can it be a *demonstrably bounding* analysis, or is the analysis method specified in terms of certain *conservative analytical methods* or the use of specified *conservative data*?

Design-side options

Here, the options are very broad, ranging from true performance-based design to design using entirely prescriptive design rules.

Design-side Option A: True performance based design: In this option, the designer need not follow any prescriptive design rules, but is free to execute the design however he/she wants to or needs to, with the constraint that of course the design must meet all of the requirements set down for demonstrating compliance. This Option A places the heaviest burden not on the design but on the compliance analysis. Indeed, accounting for uncertainties and establishing the requisite degree of confidence is a major task, fraught with analysis difficulties, both with demonstrating the completeness of the analysis scope, and with documentation. A major advantage of Option A is that it imposes few constraints to stifle innovation in design, except that the design must be analyzable.

Design-side Option B: Specified deterministic design rules with a performance criterion: In this option, the designer is constrained by specified design rules, embedded in either a consensus industry code or a regulation. However, the design must be analyzed to show compliance with a performance criterion as discussed above.

Some of the advantages of Option A's flexibility are lost with Option B, because some innovative design ideas simply cannot be accommodated by the established design rules in the consensus codes. By their nature, the codes always lag behind the latest innovations.

Design-side Option C: Specified deterministic design rules without an explicit performance criterion: In this option, as with Option B, the designer is constrained by specified design rules, embedded either in a consensus industry code or a government regulation. However, the designer need not perform an analysis to demonstrate compliance with a performance criterion. Instead, *compliance is achieved (and implicitly demonstrated) simply by having followed the design rules.*

Summary discussion on the options for the "framework"

It is important to note that a regulator or a consensus code committee needs to select among all of these Options taking the whole picture into account – that is, it is crucial to account for the ramifications of all of these Options as a *set of decisions* is being made.

In practice, a consensus code committee or a regulator is free to adapt any of these options, or to mix-and-match. As we shall see (below), the actual codes and regulations in place today can best be characterized that way. Specifically, although there has been an evolution over time, only one combination of these options is endorsed today by the NRC to guide NPP seismic design and to govern how the analyses must be performed.

Next, we will describe the approach used in the existing design standards, analysis standards, and regulatory standards for U.S. nuclear power plants, and we will see how they fit together as a "framework." Then we will identify which of the several options for such a framework is represented by today's approach.

OVERVIEW OF THE APPROACH USED IN EXISTING EARTHQUAKE DESIGN AND ANALYSIS CODES AND STANDARDS FOR NPPs

As mentioned above, today the basic procedures for seismic design of nuclear facility SSCs are provided in the Standard ASCE / SEI 43-05 [1], an important professional consensus standard. The intent of the ASCE 43-05 Standard is to produce engineered designs that achieve an acceptable *target seismic risk goal*, defined in terms of an *annual probability of seismically induced unacceptable performance*. This is accomplished by meeting established annual probability of unacceptable performance *target performance goals*, while simultaneously not exceeding *specified limit states*. The way a design is accomplished differs depending on the type of facility, and facilities are categorized into different *seismic design categories.* [ASCE 43-05 covers the design of a wide range of nuclear facilities besides nuclear power plants, using a graded approach, but the discussion that follows will be about nuclear power plants.]

Seismic Design Categories, Limit States, and Target Performance Goals:

Seismic Design Categories and Limit States: The graded approach in ASCE 43-05 introduces different levels of conservatism depending on the facility; this approach is intended to ensure that the levels of conservatism are consistent with the functionality and human hazards associated with a particular facility. To that end, every nuclear facility is assigned one of five "Seismic Design Categories" (SDCs), using guidance and following the SDC descriptors found in the American Nuclear Society Standard 2.26 [3]. In addition, a functionality descriptor, called a "Limit State", is assigned to each facility being designed, again using guidance in ANS 2.26. The choice of the SDC and LS for a given facility would normally be the prerogative of the facility's owner, but for vital SSCs in nuclear power plants a particular SDC and LS assignment has been made by the US NRC.

There are 5 Seismic Design Categories and 4 Limit States described in ASCE 43-05. The design approach embodied in ASCE 43-05 draws upon, and closely follows, the procedures from the U.S. Department of Energy Standard 1020 [4], developed by DOE for high consequence facilities. A very useful document that explains the basis for DOE STD 1020 was published by Kennedy and Short along with the original standard [5]. The U.S. Nuclear Regulatory Commission's Standard Review Plan NUREG-0800 [6] also cites ASCE 43-05, although the NRC does not endorse all of it for use by licensees and applicants.

In the ASCE 43-05 procedures, the *Seismic Design Category* is used to establish the design earthquake level motions. The *Limit State* is used to select the appropriate design procedures, analysis methodology, and acceptance criteria.

Target Performance Goals: ASCE 43-05 also provides Target Performance Goals for each SDC. These goals are described in terms of *mean annual probability of exceedance*.

SSCs in nuclear power plants that perform vital safety functions are assigned in ASCE 43-05 to SDC 5 and Limit State D. The US NRC has adopted this assignment. The Target Performance Goal P_F for SDC 5 is 1 x 10^{-5} per year. *It is this set (SDC 5, LS D, and* P_F *goal of 1 x* 10^{-5}*/year) that will be the focus of our discussion in what follows.* Also notice that the approach treats each individual SSC separately – that is, each SSC that falls under the coverage of the standard must individually meet the goal.

Design guidance in ASCE 43-05

A designer whose task is the design of an SSC to achieve adequate seismic performance needs a specified reference seismic design input to work with. This is usually called a "*design basis earthquake*", although other terms are also used. In laymen's terms, this is the "earthquake motion" that the designer must assure can be withstood. As discussed above, the goal of ASCE 43-05 is to meet a specified target performance goal. While in principle a designer could select his own reference seismic design input and then execute his design using it, and then analyze the design to demonstrate compliance with the target performance goal, that is not the approach used today. Instead, the methodology in ASCE 43-05 is prescriptive. The details of the prescriptive approach are beyond our scope here.

Design Basis Earthquake definition – ensuring adequately conservative design rules and acceptance criteria to achieve P_F. In the ASCE 43-05 provisions, it is assumed that the *Design Basis Earthquake* is based upon a site-specific Probabilistic Seismic Hazard Assessment (PSHA) [7, 8] that produces seismic hazard curves and Uniform Hazard Response Spectra (UHRS) associated with several hazard exceedance frequencies.

Achieving desired performance – FOSID for NPPs: For nuclear power plants, ASCE 43-05 also selects the targeted design performance conservatively. Specifically, Limit State D is defined so that, *"An SSC designed to this Limit State shall maintain its elastic behavior."* This has been interpreted more specifically to be the threshold point for the *"onset of significant inelastic deformation"*. The frequency for reaching this threshold has been given the moniker FOSID (frequency of the onset of significant inelastic deformation).

The NRC has specifically adopted this procedure for development of site-specific ground motions in Regulatory Guide 1.208 [10]. This NRC adoption is equivalent to specifying that the SSCs vital to safety in nuclear power plants are *Structural Design Category 5* SSCs for which only *Limit State D* outcomes are permitted under loads from the Design Basis Earthquake. Thus these SSCs must achieve ruggedness at least sufficient to assure that they achieve an annual probability of exceedance of 1×10^{-5} with respect to FOSID.

It is important to note here that the NRC endorsement of the ASCE 43-05 approach for defining the site-specific earthquake ground motion, although important, does not mean that the NRC has endorsed the rest of the procedures in ASCE 43-05. Only the part of ASCE 43-05 leading to a site-specific design response spectrum has been explicitly endorsed. The NRC, of course, has its own regulations and regulatory guidance, and although much of what the NRC requires or offers as guidance is similar to what is in ASCE 43-05, there are differences in detail, some of which are important although many others are not.

Achieving desired performance – design provisions: It is beyond our scope here to provide details about how the ASCE 43-05 design specifications and acceptance criteria achieve the desired performance. (As mentioned, much of the technical detail in ASCE 43-05 is based on DOE Standard 1020-2002 [4].) The technical issues that a designer must be attentive to include, among others, lateral force provisions; story drift/damage control provisions; detailing for ductility provisions; and quality assurance provisions. As DOE Standard 1020 states [4], "These provisions are comprised of the following four elements taken together: (1) seismic loading; (2) response evaluation methods, (3) permissible response levels; and (4) ductile detailing provisions."

Performance implications of the ASCE 43-05 design approach: As discussed above, the target performance goal adopted for nuclear power plant applications is a mean annual frequency of $1\text{x}10^{-5}$/year with respect to the earthquake-caused onset of significant inelastic deformation (FOSID). Crucially, ASCE 43-05 notes that the seismic demand and structural capacity evaluation criteria are aimed at providing sufficient conservatism to achieve two different outcomes for nuclear power plant applications, both of which are stated in *probabilistic terms.* Both of these should be true for any individual SSC that is designed and evaluated according to the ASCE 43-05 framework:

- Less than about a 1% probability of unacceptable performance for the Design Basis Earthquake Ground Motion (with the Design Spectra defined per ASCE 43-05), and
- Less than about a 10% probability of unacceptable performance for a ground motion equal to 150% of the Design Basis Earthquake Ground Motion.

The conclusion that these outcomes are achieved if the design rules and acceptance criteria in ASCE 43-05 are followed is based on considerable experience and judgment, backed up by extensive calculations found in Kennedy [5, 9].

Basis for the 1×10^{-5} target performance goal: PRA analysis of current plants

Kennedy notes [9] that the basis for selection of the 1×10^{-5} per year target performance goal is related to the calculated seismic contribution to annual core-damage frequency (CDF) in current U.S. nuclear power plants. For the U.S. plants that have completed full Seismic Probabilistic Risk Assessments (SPRAs), which includes about 25 U.S. plants, the average annual seismic contribution to CDF is $\sim 1 \times 10^{-5}$ per year. By selecting the FOSID target as the target performance goal at 1×10^{-5}/year, there is additional significant conservatism embedded with respect to the potential for core damage, because FOSID corresponds to a significantly lower response level than that required to reach any type of important damage for an individual SSC that might lead to overall reactor core damage.

CHARACTERIZATION OF TODAY'S APPROACH: WHICH "OPTIONS" ARE REPRESENTED

Above, a set of "options" was presented for how a "framework" could be constructed, covering the *performance criteria*, the *acceptance criteria* (including a confidence level), *analytical methods*, and *design criteria*. Here we will characterize today's "framework", by which we mean the framework that follows ASCE 43-05 and the NRC's corresponding guidance for nuclear power plant safety applications, vis-à-vis these several "options." The next several paragraphs explain how today's framework can be characterized vis-à-vis these options.

Performance criterion: As discussed above, today's framework uses an *annual frequency of unacceptable performance* as its performance criterion. For SSCs vital to the safety of nuclear power plants, this is linked to Seismic Design Category 5, and selected to be a target frequency of 1 x 10^{-5}/year.

Confidence level: The approach in today's framework explicitly seeks *mean confidence* that the performance criterion is met. This mean confidence is typically about one standard deviation above the median, and sometimes higher.

Performance figure-of-merit: The *onset of significant inelastic deformation* as the threshold for "failure" in Limit State D is definitely not a realistic characterization of the earthquake-caused "failure" of the SSCs under consideration. *FOSID is a conservative characterization of failure,* and indeed has been explicitly chosen to be so. As noted above, the onset of inelastic deformation is conservatively a long way short of the point where a given SSC will fail to perform its safety function.

End point – a single SSC: The undesired end-point of concern in the seismic "framework" under discussion here is the *seismic-caused failure of a single SSC*. The next paragraph explains why for nuclear power-plant applications this is conservative, and indeed almost always highly so.

The undesired end-point in a nuclear power plant is damage to the core, as analyzed in a PRA. In a seismic PRA (SPRA), the scenarios of concern are all characterized by the initiating earthquake causing damage to one or more SSCs, leading to core damage. Sometimes, depending on the scenario, one or more non-seismic failures or human errors must occur too. In all of the few dozen SPRAs to date that have studied large NPPs, no scenario has been found to be important in which core damage is caused by the seismic failure of a single SSC – it always requires more than one, or an SSC failure plus a non-seismic failure or a human error. This of course is largely because the systems design of the plants, with its redundancy, diversity, and defense in depth, explicitly assures that such a "singleton" scenario cannot by itself lead directly to core damage. Thus, choosing a single SSC as the end point for the "framework" is highly conservative, vis-à-vis the true end-point of safety concern, core damage.

Realistic or conservative analysis: The analysis that the current "framework" requires to show compliance, as described above, is a *combination of realistic and conservative analysis.* Certain features of the analysis are explicitly conservative, while others are more nearly realistic.

Design requirements: The design requirements are *detailed and prescriptive*. These requirements aim toward an *explicit performance target.* Experience from the seismic PRA literature for large LWRs demonstrates that this design approach is successful, in that the performance target is met by the current fleet of plants.

Summary: As can be seen, the current framework is effective, in that it will produce an NPP design that will meet the performance target with large margins. These margins are embedded explicitly in some of the aspects of the framework and implicitly in others. The framework is also one that can be used effectively by engineers at the bench who do the design and analysis work, and it can be efficiently and effectively peer-reviewed. All of these features are attractive, and one of the most attractive of them is the explicit numerical performance target.

Shortcomings: There are, however, shortcomings that make the framework less than ideal. None of these gets in the way of the effective implementation of today's framework, but some of them, if overcome, could provide a much sounder foundation for the future. Some possible paths forward are discussed next.

POSSIBLE PATHS FORWARD BEYOND TODAY'S "FRAMEWORK"

The current "framework" for seismic SSC design and analysis for nuclear power-plant applications has many strengths, among the most important of which is that it already has demonstrated that it is technically sound and useful in the hands of both the routine designer and the sophisticated analyst. A major positive attribute is that it is anchored in a technically achievable *and analyzable* performance goal that is considered "safe enough." Another major attribute is its strong endorsement by both the practitioners and the regulators. Therefore, any suggestions for "improvements" must be made with a certain amount of humility, plus a large dose of "look before you leap" skepticism.

Nevertheless, it is easy to see where, at least in principle, some further advances in the framework are available. The problem is that closing the gap between "available" and "achievable" is, in this arena as in many others, both difficult and likely to take a long time.

For each of these, one could formulate a research program that could help define how far each of these ideas could realistically be pushed along.

Limitation: one SSC at a time

A major limitation is that the current framework forces each SSC in the reactor with the requisite safety significance to be designed, analyzed, and approved by the regulator *one at a time*. But given that there are no earthquake-caused accident sequences that involve only a single SSC, this means that the current approach is surely conservative, and sometimes significantly so.

At least in principle, the remedy for this limitation is to use a full PRA analysis to understand the specific role of each SSC in all of the accident sequences where it participates. With this understanding, it is likely that some SSCs may not require the full robustness (and cost) of meeting every requirement in the framework. For example, for a given SSC it might be that the only important accident sequences in which it participates are characterized by the SSC being in Boolean logical AND with another different SSC which is demonstrably very strong in earthquakes, much stronger than the framework demands. (What this Boolean AND concept means in plain English is that both SSCs would need to fail to cause the accident sequence.)

This approach, at least as expressed here, comes down to *using the overall CDF (core damage frequency) for the entire seismic PRA as the figure of merit.* This figure of merit would become a tool to use in modifying the design requirements for some SSCs, based on their individual roles in achieving overall reactor safety. However, in the end *a set of specific design rules, linked to performance targets like those in today's framework, would be necessary for each SSC.* What would be different would be that *the design rules could be tailored to the individual SSC through a tailored performance target*.

A modification to this approach of using the entire PRA could be to examine the role of a given SSC in only those accident sequences to which it contributes, and to study these sequences individually, one by one, not necessarily in the context of the CDF from the entire PRA as a figure-of-merit but in relative isolation – albeit with the context of the entire PRA also kept in mind. A relaxation in design and performance requirements for that SSC would be permitted when it would not materially affect just those sequences.

Limitation: FOSID is not failure

A second limitation of the current "framework" is that FOSID, the onset of significant inelastic deformation, is acknowledged not to represent the failure of most SSCs to perform their safety functions. The technical issue that must be wrestled with is that for some SSCs the earthquake "size" (however defined) that actually does compromise the safety performance is not much larger than where FOSID occurs, whereas for other types of SSCs there is quite a large margin. This leads to the obvious suggestion that perhaps advantage can be taken of this fact at least for the latter class of SSCs, while leaving the current framework in place for the former class.

How to differentiate? Clearly this must start with some sort of knowledge, derived in part from data and in part from analysis, that can support careful deliberations on how to use the knowledge to suggest a possible modification, case by case, of today's framework.

One area where such a change could provide immediate benefit is for the class of SSCs whose behavior in earthquakes is *to "go inelastic" by design* for a certain "size" earthquake but to retain adequate safety performance even while "going inelastic" unless a much "larger" earthquake were to occur. Again, if adopted by the code committee(s) or the regulators, a philosophy that would recognize that this type of SSC can have an important role in nuclear power plants could "open a door" – open the door to engineering innovations, cost savings, safety improvements, better understanding and/or analyzability, larger safety margins, or some combination. All of these advantages are closed off for this class of SSCs to the extent that the FOSID threshold continues to be the definition of "failure" for all nuclear power plant SSCs, regardless of the true "failure" behavior of the SSC.

The ideal: true (pure) performance-based design

In a true (pure) performance-based design framework, there would be no constraints at all on the designer: no code rules, no code allowables that must be used, no restrictions on how the design could be developed. The only constraint, and it is a crucial one, is that *an analysis, presumably a robust realistic analysis, would be required to demonstrate that the required performance is accomplished*.

This ideal world should probably never be realized in any real-world application, of course! The reasons are several, perhaps the most important of which is that discarding all of the code restrictions and rules also discards a century of design experience developed by the broader engineering community and distilled into the consensus codes and regulations.

Another major reason why this ideal is unapproachable at present is that, except for a very few simple and idealized systems, a robust, realistic, comprehensive analysis with only modest uncertainties and nearly complete confidence in its correctness is simply not feasible. The fact is that we are, as an engineering community, dealing with complex real systems for which one cannot at present get close enough to the above ideal to believe we could do away with our current *requirements* for margin, conservatism, and the like. And crucially, because even "the best of us" can make errors, complete trust in analysis without requiring extra margins is not likely to become a reality any time soon, for any system within the scope being discussed here.

Nevertheless, the advantages of a certain amount of movement in this direction are manifest, and easy to write down (if not to achieve in practice.) The advantages come in at least two categories, in each of which important innovation could be stimulated that is now partially stifled:

- innovation in the form of advanced design concepts, or advanced approaches to achieving better designs with today's concepts.
- innovation in the form of advanced analysis methods, including analysis using simulation and testing working together more effectively.

The stifling of innovation is never a good thing, one would think. True enough. A major effort to develop more steps in the right direction than are now being considered – perhaps only baby steps at first – might let loose a snowballing of innovative ideas not only in design space and analysis space, but also in code-committee/regulatory-philosophy space.

REFERENCES

[1] American Society of Civil Engineers, "Seismic design criteria for structures, systems and components in nuclear facilities," Standard ASCE/SEI 43-05, 2005

[2] Budnitz, R.J., "Opportunities for advancing technology-neutral and performance-based design methods for the seismic design and regulation of nuclear power plant structures, systems and components," Lawrence Berkeley National Laboratory, Report to U.S. Nuclear Regulatory Commission, 2010

[3] American Nuclear Society, "Categorization of nuclear facility structures, systems, and components for seismic design," Standard ANS-2.26-2004, 2004

[4] US Department of Energy, "Natural phenomena hazards design and evaluation criteria for department of energy facilities," Standard DOE-STD-1020-2002, 2002

[5] Kennedy, R.P. and Short, S.A., "Basis for seismic provisions of DOE-STD-1020," Lawrence Livermore National Laboratory Report UCRL-CR-111478, 1994

[6] US Nuclear Regulatory Commission, "Standard review plan for the review of safety analysis reports for nuclear power plants," NUREG-0800, 2007

[7] Budnitz, R.J., Apostolakis, G, Boore, D.M., Cluff, L.S., Coppersmith, K.L. Cornell, C.A. and Morris, P.A., "Recommendations for probabilistic seismic hazard analysis: guidance on uncertainty and use of experts," NUREG/CR-6372, Lawrence Livermore National Laboratory, sponsored by the US Nuclear Regulatory Commission, US Department of Energy, and Electric Power Research Institute, 1997

[8] Budnitz, R.J., Apostolakis, G, Boore, D.M., Cluff, L.S., Coppersmith, K.L. Cornell, C.A. and Morris, P.A., "Use of technical expert panels: applications to probabilistic seismic hazard analysis," *Journal of Risk Analysis*, Vol. 18, No. 4, p. 466, 1997

[9] Kennedy, R.P. "Performance-goal based (risk informed) approach for establishing the SSE site-specific response spectrum for future nuclear power plants," *Transactions of the 19th International Conference on Structural Mechanics in Reactor Technology* (SMiRT), Toronto, Canada, 2007

[10] US Nuclear Regulatory Commission, Regulatory Guide 1.208, "A performance-based approach to define the site-specific earthquake ground motion", March 2007

ACKNOWLEDGMENT AND DISCLAIMER

The author wishes to acknowledge the strong intellectual support received from Dr. Annie Kammerer (US Nuclear Regulatory Commission's Office of Nuclear Regulatory Research), who also served as the technical contact on the NRC staff.

This manuscript has been written by an author at the Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy (DOE). The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes. This project was supported by the U.S. Nuclear Regulatory Commission under a Federal Interagency Agreement with DOE.