Unclassified **NEA/CSNI/R(97)22**

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

OLIS : 03-Mar-1998
Dist. : 09-Mar-1998

English text only

NUCLEAR ENERGY AGENCY
COMMITTEE ON THE SAFETY OF NUCLEAR INSTALLATIONS

STATE-OF-THE-ART REPORT ON THE CURRENT STATUS OF METHODOLOGIES FOR SEISMIC PSA

62330

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

ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT

Pursuant to Article I of the Convention signed in Paris on 14th December 1960, and which came into force on 30th September 1961, the Organisation for Economic Co-operation and Development (OECD) shall promote policies designed:

- to achieve the highest sustainable economic growth and employment and a rising standard of living in Member countries, while maintaining financial stability, and thus to contribute to the development of the world economy;
- to contribute to sound economic expansion in Member as well as non-member countries in the process of economic development; and
- to contribute to the expansion of world trade on a multilateral. non-discriminatory basis in accordance with international obligations.

The original Member countries of the OECD are Austria, Belgium, Canada, Denmark, France, Germany, Greece, Iceland, Ireland, Italy, Luxembourg, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, Turkey, the United Kingdom and the United States. The following countries became Members subsequently through accession at the dates indicated hereafter: Japan (28th April 1964), Finland (28th January 1969), Australia (7th June 1971), New Zealand (29th May 1973), Mexico (18th May 1994), the Czech Republic (21st December 1995), Hungary (7th May 1996), Poland (22nd November 1996) and the Republic of Korea (12th December 1996). The Commission of the European Communities takes part in the work of the OECD (Article 13 of the OECD Convention).

NUCLEAR ENERGY AGENCY

The OECD Nuclear Energy Agency (NEA) was established on 1st February 1958 under the name of the OEEC European Nuclear Energy Agency. It received its present designation on 20th April 1972, when Japan became its first non-European full Member. NEA membership today consists of all OECD Member countries except New Zealant and Poland. The Commission of the European Communities takes part in the work of the Agency.

The primary objective of the NEA is to promote co-operation among the governments of its participating countries in furthering the development of nuclear power as a safe, environmentally acceptable and economic energy source.

This is achieved by:

- *encouraging harmonization of national regulatory policies and practices, with particular reference to the safety of nuclear installations, protection of man against ionising radiation and preservation of the environment, radioactive waste management, and nuclear third party liability and insurance;*
- *assessing the contribution of nuclear power to the overall energy supply by keeping under review the technical and economic aspects of nuclear power growth and forecasting demand and supply for the different phases of the nuclear fuel cycle;*
- *developing exchanges of scientific and technical information particularly through participation in common services;*
- *setting up international research and development programmes and joint undertakings.*

In these and related tasks, the NEA works in close collaboration with the International Atomic Energy Agency in Vienna, with which it has concluded a Co-operation Agreement, as well as with other international organisations in the nuclear field.

COMMITTEE ON THE SAFETY OF NUCLEAR INSTALLATIONS

The Committee on the Safety of Nuclear Installations (CSNI) of the OECD Nuclear Energy Agency (NEA) is an international committee made up of senior scientists and engineers. It was set up in 1973 to develop, and co-ordinate the activities of the Nuclear Energy Agency concerning the technical aspects of the design, construction and operation of nuclear installations insofar as they affect the safety of such installations. The Committee's purpose is to foster international co-operation in nuclear safety among the OECD Member countries.

The CSNI constitutes a forum for the exchange of technical information and for collaboration between organisations which can contribute, from their respective backgrounds in research, development, engineering or regulation, to these activities and to the definition of the programme of work. It also reviews the state of knowledge on selected topics on nuclear safety technology and safety assessment, including operating experience. It initiates and conducts programmes identified by these reviews and assessments in order to overcome discrepancies, develop improvements and reach international consensus on technical issues of common interest. It promotes the co-ordination of work in different Member countries including the establishment of co-operative research projects and assists in the feedback of the results to participating organisations. Full use is also made of traditional methods of co-operation, such as information exchanges, establishment of working groups, and organisation of conferences and specialist meetings.

The greater part of the CSNI's current programme is concerned with the technology of water reactors. The principal areas covered are operating experience and the human factor, reactor coolant system behaviour, various aspects of reactor component integrity, the phenomenology of radioactive releases in reactor accidents and their confinement, containment performance, risk assessment, and severe accidents. The Committee also studies the safety of the nuclear fuel cycle, conducts periodic surveys of the reactor safety research programmes and operates an international mechanism for exchanging reports on safety related nuclear power plant accidents.

In implementing its programme, the CSNI establishes co-operative mechanisms with NEA's Committee on Nuclear Regulatory Activities (CNRA), responsible for the activities of the Agency concerning the regulation, licensing and inspection of nuclear installations with regard to safety. It also cooperates with NEA's Committee on Radiation Protection and Public Health and NEA's Radioactive Waste Management Committee on matters of common interest.

* * * * * * * * * * * *

Requests for additional copies of this report should be addressed to:

Nuclear Safety Division
OECD Nuclear Energy Agency
Le Seine St-Germain
12 blvd. des Iles
92130 Issy-les-Moulineaux
France

ABSTRACT

This report is an up-to-date review of the state-of-the-art of the methodologies for conducting a seismic-PSA at a nuclear power station, including the six sub-methodologies that comprise the overall methodology. The review concentrates on evaluating the extent to which today's seismic-PSA methodology produces reliable and useful results and insights. The evaluation covers six sub-methodologies that must be combined to produce an overall full-scope seismic PSA: the seismic-hazard methodology; the local-ground-motion and building-motion methodology; the walkdown methodology; the failure-mode and fragility methodology; the systems-analysis methodology; and the consequence/release methodology. The report finds that all of these sub-methodologies are both reliable and useful, and that when combined together the overall seismic-PSA methodology can provide important engineering insights about how nuclear power-plants respond to earthquakes. This is true even though the numerical uncertainties in the bottom-line results can be large (plus-or-minus more than one order of magnitude or more is common.) However, a number of areas within the various sub-methodologies can be applied properly only if special expertise is available.

The report describes the technical issues in detail, and outlines the approaches that have proven to be most successful, based on world-wide experience with about a hundred seismic PSA studies.

FOREWORD

This project has been supported by the Nuclear Energy Agency of the Organisation for Economic Co-operation and Development in Paris. It is intended to support the work of Principal Working Group 5 under the NEA's Committee on the Safety of Nuclear Installations.

The author has based this work in important ways on an earlier review, now eight years old, on the same subject. This work is cited as (Budnitz and Lambert, 1989) in the reference list. Many portions of the descriptive part of this text were borrowed or adapted from that earlier review, although the substantive parts, the evaluations, were done anew, subarea-by-subarea.

While the citations and most of the specific experience that support the evaluations herein are from the United States, the experience of the author extends to a number of seismic PSA studies accomplished for nuclear power stations elsewhere. Specifically, to supplement his familiarity with over forty U.S. seismic PSAs, the author has relied, in part, on his personal familiarity with the seismic PSAs for the Kozloduy 5-6 units in Bulgaria, the Ulchin 3-4 units in Korea, the Krsko unit in Slovenia, and the Sizewell-B unit in Great Britain; and familiarity with the seismic designs of the four Paks units in Hungary, the Medzamor-2 unit in Armenia, the Ignalina 1-2 units in Lithuania, the Chernavoda-1 unit in Romania, and the Bohunice 1-2 units in Slovenia. He is also familiar with the seismic-safety aspects of the advanced LWR designs from ABB-Combustion Engineering and General Electric Company that were recently certified by the U.S. NRC.

The author wishes to offer special thanks to M.K. Ravindra and R.C. Campbell of EQE Engineering, R.C. Murray of Lawrence Livermore National Laboratory, N.C. Chokshi of the U.S. Nuclear Regulatory Commission, and R.P. Kennedy for their many years of innovative contributions to the discipline of seismic PSA, which contributions have been a source of continuing inspiration for his own PSA work.

This work is solely that of the author (R.J.B.) -- none of those thanked just above has reviewed it or has any responsibility for it.

Additionally, PWG5 Members and the NEA Secretariat wish to provide acknowledgement of the work performed by the author, Dr. Robert J. Budnitz, who provided the in-depth technical analysis as well as overall coordination in editing and compiling this report.

TABLE OF CONTENTS

1. INTRODUCTION

1.1 Introduction

There has never been an earthquake sufficiently damaging to any operating U.S. nuclear power station to cause safety concerns. In the past decade, there have been two very destructive earthquakes that were near enough to nuclear power stations to have caused at least temporary world-wide concern: the large Armenian earthquake of November, 1988, and the large Japanese earthquake near Kobe in January, 1995. However, in each case, these earthquakes produced only minor ground motion at the operating nuclear stations located close enough to merit some examination: in each case, inspections revealed essentially no damage, because the actual ground motions at the nuclear-plant sites were in both cases quite small and well within the design basis.

Nor has there ever been any other earthquake known to have damaged a nuclear power station. Therefore, the published historical record is not adequate for the analysis discussed here. The frequency of potential earthquake-initiated core-damage accidents at a given nuclear-power station can only be known from calculations, using a combination of real-earthquake data, test data, models of various phenomena, and systems analysis.

Despite the lack of actual earthquake experience at nuclear-power stations, almost every full-scope probabilistic safety analysis (PSA) that has examined earthquake-initiated accidents at nuclear power stations has found that this category represents one of the important initiator groups. Occasionally, one of the earthquake-initiated sequences is among the few largest contributors to calculated core-damage frequency and/or off-site risk. Of course, it is important to keep in mind that PSAs calculate a residual risk, which is judged to be acceptable by the regulatory authorities, and that the seismic contribution to this residual risk is what is being discussed here.

Usually the potential accident sequences identified in the PSAs are very plant-specific in character, such that the specific vulnerability is unlikely to exist at any other plant, even another similar plant. Sometimes the issue is site-related, and sometimes it is design-related. Also, even in areas where earthquakes are very uncommon phenomena, these types of accident sequences often appear as important contributors to the residual risk, typically because in such areas the attention given to designing nuclear stations against earthquakes is much less than in earthquake-prone areas.

Given this background, it is obvious that no full-scope PSA can be considered complete without an examination of earthquakes.

This report is a review of the methodology for conducting a seismic-PSA at a nuclear power station. The objective of this review is as follows:

> To provide an up-to-date review of the state-of-the-art of the various sub-methodologies that comprise the overall seismic-PSA methodology for addressing the safety of nuclear power stations, plus an overview of the whole methodological picture.

In preparing this review, the author has had in mind several categories of readers and users: policy-level decision-makers (such as managers of nuclear power stations and regulators of nuclear safety), seismic-PSA practitioners, and PSA practitioners more broadly.

The review will concentrate on evaluating the extent to which today's seismic-PSA methodology produces reliable and useful results and insights, at its current state-of-the-art level, for assessing nuclear-power-station safety.

Also, this review paper will deal exclusively with seismic-PSA for addressing nuclear-power-station safety. Because the author is based in the U.S., it is natural that this review will contain more emphasis on U.S. experience than on experience in other countries. However, significant experience elsewhere is a major part of the basis for this evaluation.

1.2 Background

In recent years, the discipline of probabilistic safety analysis (PSA) for analysing the safety of nuclear power stations has become quite mature. It is now being used routinely to assist nuclear-safety decision-making throughout the world. This includes decision-making that affects every aspect of nuclear safety: design, construction, operation, maintenance, and regulation.

However, not all sub-areas within the larger discipline of PSA are equally "mature", and therefore the many different types of engineering insights from PSA are not all equally reliable. In particular, the sub-discipline of seismic-PSA analysis is in some areas thought to be relatively less mature and less reliable than the corresponding internal-initiators analysis. It is this author's conviction that this is not the case: whereas the methodology for seismic-PSA suffers from certain problems to be discussed below, these problems are not necessarily more several than the problems with the internal-initiators PSA methodology.

The problems that exist with the seismic-PSA methodology are due predominantly to three causes:

- First, seismic-PSA analysis began in a serious way only about 1980, more than a half-decade after the first internal-initiators analysis of WASH-1400. Also, compared to internal-initiator analysis the number of practitioners is fewer, and the number of full-scope PSAs that include earthquake initiators also fewer, so there is less opportunity for a broader community to have digested and re-digested the methods, models, data, and results. This is true for both the fragility-analysis methods and the hazard-analysis methods.

- Second, technical problems continue to exist with some components of the methodologies used. These problems occur in every sub-area of analysis.

- Third, the technical problems in seismic-PSA analysis lead to significant numerical uncertainties in the "bottom-line" core-damage and risk results.

Even today, these relative weaknesses in external-initiator PSA make many decision-makers, among both industry and regulatory bodies, reluctant to use external-initiator PSA results on a par with those from internal-initiator PSA. This is true despite numerous papers and reports pointing out that the methodology is reasonably mature. Some of these references go back a long way, while others are relatively recent (Ref. NRC, 1983; Budnitz, 1984; NRC, 1984; Budnitz, 1987; Bohn and Lambright, 1988; Budnitz and Lambert, 1989; NRC/1602, 1997; EPRI/Margins, 1991; Reed and Kennedy, 1994; ERI, 1997).

The fact is that seismic-PSA analyses, like analyses involving other external initiators such as high winds and external flooding, have developed a "bad reputation" in some quarters --- they are still considered too uncertain, or too conservative, or supported by too little solid data to be of important use. As mentioned just above, it is this author’s conviction that this bad reputation is undeserved.

Fortunately, this picture has begun to change recently. Among the signs of this change are:

- Around the world, almost all new full-scope PSAs accomplished with utility support in recent years have included earthquake initiators, as well as other external initiators, as an integral part.

- In the U.S., even though the IPEs (Individual Plant Evaluations) for external initiators were only required at a later stage than the internal-initiator evaluations (Ref NRC/IPEEE, 1991), the NRC’s approach to resolving the "severe accident" issue for existing plants recognises that external initiators must have an equal footing with internal initiators (Ref. NRC/SECY-97-077, 1997), and this equal footing is now recognised within the NRC for all applications of PSA within the agency (Ref. NRC/1602, 1997; NRC/DG-1061, 1997).

- The U.S. NRC has recently embarked on a regulatory initiative to use PSA-based risk information in a "risk-informed" manner. In its draft guidance on how to us PSAs for plant-specific changes to the licensing basis (Ref. NRC/DG-1061, 1997), the necessity of including earthquake-initiated PSA information along with information from internal-initiator PSA is emphasised.

- The U.S. NRC has recently issued draft guidance (Ref. NRC/1602, 1997) on how a PSA must be performed and documented to support a regulatory decision. The importance of seismic PSA is given equal weight with that of internal-initiators PSA in this guidance report.

1.3 Technical Approach of this Study

It is important to note that the thrust of this report is to describe and evaluate the principal aspects of the current state-of-the-art of seismic-PSA methodology, what aspects are more robust and therefore provide the most reliable insights, what aspects are less robust and therefore provide less reliable insights, and why. The evaluation concentrates on the sub-methodologies (for earthquakes these are the hazard methodology, the response methodology, and the systems methodology), and on how these sub-methodologies are combined together to provide overall seismic-PSA results and insights.

Although the various sub-methodologies are all being used today to perform PSA analyses, the evaluation here reveals important limitations in some aspects of them. Some of those limitations can be reduced or eliminated in any given application of the methodology by performing trial analyses and sensitivity studies to gain further understanding. In other cases, reducing the limitations will require physical experimentation, extensive data-gathering, the building and testing of complex phenomenological models, and so on.

1.4 Definitions of Terms

The terms "reliability", "usefulness", and "uncertainty" are used often in this report. These are all different, as follows:

The reliability of a PSA result describes how robust it is in the face of methodological approximations and incomplete underlying data. The concept is that a decision-maker can "rely" on the validity of the result if it is robust despite the shortcomings.

The usefulness of a PSA result describes how much use a safety decision-maker can make of it. In plain English, some results are simply more useful than others. Thus, it might be only moderately useful to identify a particular vulnerability per se, but much more useful to identify an easy remedy within the PSA analysis --- for example, a remedy involving a minor procedural change that, through PSA methods, can be shown to reduce one component of the core-damage frequency by several orders of magnitude. In this sense, intermediate PSA results (such as the results of a site-specific earthquake hazard analysis) tend to be less useful than final or bottom-line results, or of identified vulnerabilities in components or system configurations.

The uncertainty in a PSA result usually describes the numerical uncertainty in the results, but it could also describe a modelling uncertainty (such as an "either-or" uncertainty about whether a phenomenon actually occurs, or uncertainty as to whether a certain approximation is or is not valid) or an applicability uncertainty (such as whether the underlying data used actually apply to the case being studied). One example of applicability uncertainty would be the common situation that site-specific data for extremely large earthquake ground motions often do not exist, and the applicability of similar data derived from a different site may be suspect.

1.5 Guidance documents for the seismic-PSA methodology

There is an extensive literature that provides guidance on seismic-PSA and related methods such as the seismic-margin methodology. (Some of the guidance on margins methods is directly applicable to PSA analysis.) The discussion here relies on that literature to a major extent.

It is noteworthy that all of this literature originates in the United States -- this certainly reveals that the U.S. has been at the center of developing and applying seismic-PSA methodology from its inception; but it may also reveal the parochial viewpoint and interests of the author.

Among the key references are the following:

Seismic-PSA methodology:

SSMRP, 1981; NRC, 1983; Budnitz, 1984; Ravindra, 1984; NRC, 1984; Ravindra, 1985; Shieh, 1985; Brookhaven, 1985; Reed, 1985; SSMRP, 1986; RMIEP, 1987; Budnitz, 1987; Bohn & Lambright, 1988; Budnitz and Lambert, 1989; NRC-IPEEE-1407, 1991; Reed and Kennedy, 1994; NRC/1602, 1997

Seismic-margin methodology:

Budnitz/Margins, 1985; Prassinos/Margins, 1986; EPRI/Margins, 1991; SQUG-GIP, 1992; Reed and Kennedy, 1994; NRC-IPEEE-1407, 1991; Budnitz, Murray, and Ravindra, 1992

2. EVALUATION OF THE SEISMIC-PSA METHODOLOGY

2.1 Description of the Methodology

There have been dozens of full-scope PSAs that have studied potential earthquake-initiated accidents at nuclear power stations all around the world. The methodology has been exercised by several different groups of practitioners and is considered mature. Nevertheless, and despite continuing research work to develop and improve the various parts of the methodology, some aspects introduce considerable numerical uncertainties into the bottom-line results.

The overall seismic-PSA methodology consists of six sub-methodologies, which are combined together. (Of course, the division into these six sub-methodologies is quite arbitrary. Some analysts use a different division.) The six sub-methodologies to be evaluated here are:

- the seismic hazard methodology for calculating the frequency of earthquakes of various "sizes" at a given site and characterising the motion parametrically

- the seismic local-ground-motion and building-motion methodology for working out the motion at a given location on the site or within buildings, given the incoming earthquake motion

- the walkdown methodology that guides the essential plant walkdown that is at the heart of seismic PSA

- the seismic failure mode and fragility methodology for calculating the capacity of individual components and structures, and from that capacity the "seismic fragility curve" for each item and the correlations among these

- the seismic-PSA systems analysis methodology

- the seismic-PSA methodology for analysing plant response and off-site releases and consequences.

In the next sub-sections, we will discuss and evaluate each of the six sub-methodologies in turn.

2.2 Evaluation of the six sub-methodologies

2.2.1 Evaluation of the seismic hazard methodology

The term "seismic-hazard methodology" is used for the methodology that analyses for the frequency of earthquakes of various "sizes" at a given site, and the spectral shapes of the motion from these earthquakes. This methodology is sometimes called "probabilistic seismic hazard analysis" to emphasise that its results are intrinsically probabilistic in nature.

In most regions in the world, very large earthquakes have never been experienced. In those regions, it is not possible to rely on actual large-earthquake data as the primary basis for the hazard curves: rather, it is necessary to develop the hazard curves based on analysis of inferences from the other types of data that actually do exist. Sometimes these inferences are difficult to make and therefore controversial, even among the scientists who are most expert in understanding the relevant issues. (Even in regions with extensive large-earthquake experience, such as coastal California or the seismically active regions in Japan, the inferences are difficult to make.)

Description of the Methodology: The methodology for developing the seismic hazard for a given site is well developed in principle, although there is still much uncertainty in the detailed hazard results, as the discussion below will reveal. The outline of the four-step approach is shown in Figure 2-1, taken from the "PRA Procedures Guide" (Ref. NRC, 1983). Because the literature is so extensive, including a very clear description in the Brookhaven procedures guide (Ref. Brookhaven, 1985), we will not discuss the four steps in detail here.

Recently, an important research project has been completed that provides extensive guidance on how to perform a probabilistic seismic hazard analysis (Ref. Budnitz et al., 1997). This project, which was co-sponsored by the US Department of energy, the US Nuclear Regulatory Commission, and the Electric Power Research Institute, demonstrates that the methodology is mature, and gives explicit step-by-step guidance on all of the various sub-parts of the methodology. The discussion on seismic-hazard analysis here relies heavily on this recent project.

The methodology consists of four steps.

This following brief summary description of the four-step hazard methodology will not mention all of the technical issues, because the summary is intended mainly to introduce the various broad sources and types of information needed for seismic hazard analysis.

(i.) The first step is seismic source characterisation and assessment, to identify and characterise the seismic sources in the vicinity of the site. The sources are typically either identified faults, or point sources, or areas called source zones in which it is assumed that the occurrence of earthquakes is spatially uniform. (If spatial uniformity were not thought to be true, a zone should be sub-divided until each smaller zone is thought to be uniform in its seismicity.) The objective is to determine the frequency of earthquakes of different "sizes" from each identified source or source zone. Here "size" is a loose word that includes the distribution of magnitudes, the depth and physical location and extent of the source, various other physical parameters such as the distribution of stress drops during earthquakes, the annual frequencies of occurrence of measured smaller earthquakes from the recent instrumental record, and so on.

The assessment involves evaluating data about each of the various known nearby sources. Data could include geological and geotechnical information, surface topographic evidence, micro-seismicity records, and so on. Paleoseismic information is often of importance. Usually, unless specific knowledge indicates otherwise, the different zones are assumed to be independent of each other, in the sense that seismicity in one zone is uncorrelated with seismicity in adjacent zones.

The major problem with this type of assessment is that unfortunately, except in highly active regions where well-characterised faulting dominates, there is usually no clear understanding of the processes that give rise to earthquakes. Therefore, the analyst must postulate one or more models. The analyst's model(s) can then be applied to the set of seismic sources for use in the subsequent analysis. Models range from the simple to the complex, and can incorporate factors such as possible interactions among sources, time dependence or independence of earthquake occurrence due to stress build-up, and so on. A typical site-specific model might be founded on a continent-sized tectonic model, coupled with other regional and local features. A key element in the assessment, which requires expert judgement, is drawing inferences from various similarities with other regions.

One of the key problems with seismic-source characterisation is that in regions without active faulting or other well-identified sources, different experts often provide very different "maps" that characterise somewhat different zonation schemes, based on their deferring interpretations of the meager data that exist. Reconciling these differences in interpretation is usually not easy.

(ii.) The second step is determining the earthquake <u>recurrence</u> relationship. This is typically expressed in terms of an annual frequency as a function of magnitude (as shown in stylised form in Figure II), for each source or source zone. Among the issues that are accounted for are the historical seismic activity rate, the distribution of earthquake magnitudes, the lowest magnitude of concern for the given source, the upper-bound magnitude that the analyst believes is possible, the distribution of depths for the source, the spatial distribution of energy release (point, short plane, extended plane), and so on.

Models for the recurrence relationship by different experts can range from the simple to the complex, as is the case for the seismic-source assessment. Much judgement is necessary even in cases where the historical earthquake record is extensive; for most areas where this is not the case, and because except for very recent events the few important earthquakes usually have not been properly measured by good instruments, the judgements can be controversial.

Part of the problem of understanding large earthquakes in the distant past is that our knowledge is often limited to the spatial distribution of the damage that they caused. This is not easily transformed into a more scientific parameter like magnitude (which, itself, is only a rough scheme for categorising earthquake "size").

Another issue is that certain models require an upper-bound cut-off on the magnitude, because otherwise their mathematical form would allow at least a finite chance of earthquakes of essentially infinite energy release. Typically, physical arguments are used, based on a variety of types of evidence, to determine this cut-off, but the evidence can be difficult to interpret except in highly active areas. Thus the upper-bound magnitude cut-off is usually uncertain.

(iii.) The third step (see Figure 2-1) is working out the ground-motion attenuation relationship, which means associating a motion vs. distance relationship with each magnitude. Typically, either local spectral acceleration, peak ground acceleration, or spectral velocity is chosen as the motion parameter; sometimes more than one is chosen. Of course, all of these parameters are imperfect measures --- in fact, no one or two parameters can be other than an imperfect measure.

Two principal issues arise here: selecting an attenuation model and a ground response spectral shape. These two aspects are sometimes combined in a model that directly attenuates different frequencies differently. For some active seismic areas like parts of the western U.S., the strong-motion earthquake records are extensive enough to provide actual data for attenuation modelling. For less active seismic areas like the eastern U.S., the strong-motion information is usually absent, so theoretical models are often used, although based in part on data from high-seismic areas (like western-U.S. data) and on insights that may or may not always be applicable. Issues to be considered include the effect of local transmission paths, fault rupture characteristics, and frequency dispersion.

There is also uncertainty in selecting a ground response spectral shape. In PSAs for lower-seismicity areas like the eastern U.S., a standard broad-band spectral shape normalised to the zero-period acceleration value has usually been used in place of working out the spectral shape in a combined way with the attenuation model. However, if there is considerable soil amplification, a site-specific spectrum should be developed.

(iv.) As shown in Figure 2-1, the fourth and final step of the hazard assessment is producing the "hazard curves" themselves. These curves are usually expressed in terms of the annual frequency of exceedance vs. a ground-motion parameter like peak ground acceleration or a spectral acceleration.

It is important to note that a specific spectral shape used in the earlier parts of the analysis is implicitly embedded in the hazard-curve presentation --- the results of the hazard methodology include both the hazard curves and the spectral shape(s).

Use of Experts: The difficulties in performing a probabilistic seismic hazard analysis are fundamentally due to the lack of an adequate technical data base: there have not been enough earthquakes with well-instrumented records available for study, and those well-instrumented records that are available are not always easy to interpret, even for sites near the earthquakes studied. Other data, such as paleoseismic data and geological information, can be used indirectly, but mainly in an inferential rather than a direct way. The problem is exacerbated for sites where few if any strong-motion events have occurred in recorded history.

This problem has led to controversy among experts concerning how to interpret the sparse data, and this in turn has led in many seismic-hazard studies to the use of multiple experts. The problem of how to use the disparate insights from a number of experts, when the scientific underpinning of the enterprise is recognised by all to be weak, is not unique to seismic-hazard analysis, of course, but some of the most important recent work involving multiple experts has been done in this field. Two classic examples of multiple-expert seismic-hazard studies are the late-1980s studies of seismicity in the eastern United States sponsored by the Electric Power Research Institute (Ref. EPRI/Hazard, 1989) and the US Nuclear Regulatory Commission at Lawrence Livermore National Laboratory (Ref. LLNL, 1989). These two studies were the source of considerable controversy and consternation for several years after their completion because at many identical sites in the eastern U.S. their seismic-hazard results differed considerably.[1]

This difference between the EPRI and LLNL results led to the sponsorship by DOE, NRC, and EPRI of the project mentioned above (Ref. Budnitz et al., 1997) that developed the most modern methodological guidance on seismic-hazard analysis. The principals in that project, in turn, found that the principal pitfalls in performing such a study have historically been in how experts are used rather than in the technological earth-sciences aspects themselves, even though the latter are difficult enough to execute well. This led to a concentration in the project and its report on detailed guidance on using experts.

Without delving into the details of that guidance here -- guidance on how to use experts effectively is not a principal topic of this review -- it is important to note that serious problems arise unless the use-of-experts aspect is structured properly and implemented carefully.

Since the publication of the guidance mentioned above, the use-of-experts methodology has been successfully applied in a seismic-hazard project for the proposed Yucca Mountain high-level-radioactive-waste repository in Nevada in the western U.S. (Ref. Geomatrix, 1996). This study, similar in scope and depth to a seismic-hazard evaluation for a nuclear-power-plant site, is an excellent recent example of how to perform a seismic-hazard analysis well.

Evaluation: The methodology, as embedded in the above four seismic-hazard-analysis steps, is mature in the sense that it is well-developed and widely practiced. Nevertheless, problems with the methodology generally lead to major uncertainties in the hazard results. The central reason is that experts approaching this analysis problem from different perspectives have developed a large number of reasonable models, all of which are consistent with the available earthquake data, but which have important and different implications for the hazard results. The Brookhaven PSA guide (Ref. Brookhaven, 1985) states the situation succinctly: "The development [of a seismic hazard model] is a product of scientific interpretation of uncertain and incomplete physical evidence on geological structures, tectonic processes, and seismicity." In sum, reasonable experts differ in their assessments, and selecting which expert is correct is not possible today.

Summary: Suffice it to say that, from the perspective of a decision-maker, a legitimate and wide divergence of opinion among experts cannot but be taken on its face as genuine "uncertainty" in the best meaning of that term. The fact that different models can lead to PSA core-damage-frequency calculations differing by more than a factor of 10 is simply a manifestation of the current state-of-the-art in this discipline.

2.2.2 Evaluation of the seismic local-ground-motion and building-motion methodology

Discussion of the Methodology: The objective of this aspect of the methodology is to work out the local motion at the location of each significant item (each structure or component) necessary for the safety of the power station.

This aspect of the methodology is generally quite well developed, although when specific situations are being analysed there do remain uncertainties due to random variability and incomplete knowledge.

To begin this part of the analysis, the analyst usually starts with a family of earthquake motions that are postulated to arrive at the local site from afar (or, of course, perhaps from directly below the site). These

---

[1] The LLNL probabilistic seismic-hazard results were updated in 1993 (Ref. NRC, 1993).

motions can be in the form of either real or artificial time histories or another characterisation. Typically, several different earthquake "sizes" are calculated, by scaling the time histories up or down in amplitude anchored to different zero-period accelerations.

Floor spectra: Some items are located on the ground at grade level, while others are at different elevations, either above or below grade. For these latter, it is necessary to develop what are known as floor spectra, for each elevation in each important building, to represent the seismic input at the base (the "floor") of each component or structure.

There are several individual issues here, each involving its own methodology. Because sites differ so much, not all of the issues will be relevant to every site. It is beyond the scope of this discussion to cover the details of each aspect of the methodology: extensive discussion of the technical issues can be found in the literature (Ref. Budnitz et al., 1997; ASCE-4, 1997).

The first step in this part of the analysis is determining the ground response frequency spectrum at the site, which is a function of distance from the earthquake source, the size of the earthquake, and local subsurface (especially soil) conditions (Ref. Newmark and Hall, 1978; Brookhaven, 1985; Budnitz et al., 1997). Usually, generic broad-band spectra have been used in PSAs, and this should be acceptable, provided they are applicable to local soil conditions. If the local soil produces much amplification, the analyst should develop a site-specific response spectrum (Ref. Bernreuter, 1987; ASCE-4, 1997).

For those structural foundations that are on rock or stiff soil, the base motions should be more-or-less the same as for the free field. On soft soil, the soil-structure coupling can change both the frequencies and the amplitudes of the motion entering the building. For example, it is necessary to account for such factors as soil shear modulus and damping. Soil-structure interaction models developed over the years are quite reliable if all of the relevant site factors have been considered (Ref. Johnson, Schewe, & Maslenikov, 1984; Shieh, 1985; Budnitz et al., 1997). As a warning, it is especially important for the analyst not to take as necessarily correct the models used in the design; these often contain conservatisms or other unrealistic assumptions which cannot serve as a realistic representation of behaviour in an actual earthquake.

Within the structure the input seismic motion is transmitted, and this transmission must be determined, from the foundation to any given elevation and location. The analyst needs to develop a structural model for the building, unless a model developed earlier, such as in the original design or for the safety analysis report, can be relied on. As elsewhere, it would not be correct for the analyst to use uncritically the floor response spectra found in the design analysis or safety report, because these will in all likelihood be highly conservative or contain other unrealistic assumptions.

In developing realistic floor spectra, it is typical to use linear dynamic analysis for the structure, and then to account for non-linear effects by estimating the inelastic energy absorption capacity of each component, so that the response for the equipment item represents the floor spectrum modified to account for how each equipment item responds in frequency space. The modifications account for several factors specific to each item such as damping and modal response combination --- all of which have variability which must be included in the analysis.

The variabilities in the earthquake source(s) themselves are usually accounted for by using several time histories, each of which captures the correlations properly for itself; a multiple-time-history set is used to capture, as an ensemble, the variability from earthquake to earthquake. Guidance on carrying out this aspect can be found in several references (Ref. Bohn, 1984; Brookhaven, 1985; Kennedy, 1981; Budnitz et al., 1997), and discussion of a computer code specially developed for this analysis can be found in a Lawrence Livermore report that was part of their SSMRP project (Ref. Johnson/SMACS, 1981).

Evaluation: While uncertainties certainly exist in this aspect of the seismic-PSA analysis, arising from both variabilities and modelling approximations, the analytical approaches for the several topics are all generally well-developed and robust in the hands of experienced analysts. Therefore, a summary evaluation of the state-of-the-art of the seismic local-ground-motion and building-motion methodology would label it as mature.

2.2.3 Evaluation of the walkdown methodology

Discussion of walkdown issues: There is a broad consensus among seismic-PSA analysts that the plant walkdown is the most crucial aspect of the entire process. By using a well-planned and effectively executed walkdown, the analysis team can develop vital information about the plant configuration, specific spatial relationships, anchorages, and other features that cannot be found any other way. Furthermore, if a good walkdown is not performed, neither the seismic-capacity analyst nor the systems analyst can properly perform the required work.

A walkdown team usually consists of expertise drawn from at least the following areas: seismic-fragility-analysis, systems-analysis, and plant operations/maintenance. Sometimes, the walkdown teams can consist of several representatives of each area, although having too many individuals on a walkdown can lead to a clumsy and inefficient evaluation.

For the seismic-capacity team, the crucial benefit of the walkdown is that they can determine, for each important item (structure or component), whether that item is "typical" of its generic category, or somehow atypical or even unique. If it is judged to be "typical", then information from the broad class in which the item fits can usually be used, eliminating the need for special analysis. If an unusual component or structure is identified, it can be given the special attention that it deserves. [Sometimes, these unusual components are termed "outliers", but this word implies that something may be "wrong" or "unsafe" about the component, whereas in fact all that is implied is that the component does not fit well into the broad generic category of similar components for the purposes of grouping. Therefore, the word "outlier" should be avoided.]

Interactions: Another important benefit of the walkdown is the interaction that occurs among the systems analysis team, the seismic-capacity analysis team, and the plant's operating and maintenance staff. These various groups should be working together throughout the seismic-PSA effort, but their interactions are most crucial during the walkdown, when each can assist the others in identifying the more important issues and screening out the less important. This is one of the major lessons learned in performing seismic PSAs over the years: the earliest seismic PSAs suffered because these interactions among analysts were insufficient, whereas today no seismic PSA would be considered competent without the significant analyst-to-analyst and analyst-to-plant-staff interactions that have become a central element of the walkdown.

For the systems-analysis team, another major reason for the walkdown is the opportunity to understand just how the operating crew has been trained to carry out its tasks, especially during emergencies. This understanding is crucial to the development of a correct system model (event trees and fault trees).

The literature now contains excellent guidance on how to plan and carry out a walkdown (Ref. Brookhaven, 1985; Budnitz/Margins, 1985; Prassinos/Margins, 1986; EPRI/Margins, 1991; SQUG-GIP, 1992; SSRAP, 1992; NRC-IPEEE-1407, 1991). As an example, an extensive table in the Brookhaven guide (Table 9.3.5 in Brookhaven, 1985) is especially useful, since it provides a list, for almost every category of equipment, of what to look for and why. The SQUG-GIP report (Ref. SQUG-GIP, 1992) and the EPRI/Margins report (Ref. EPRI/Margins, 1991) also contain walkdown guidance that includes specific and detailed check-list sheets for different types of equipment and structures.

The documentation of the walkdown's findings is an important aspect, not only for archival reasons, but more importantly because the documentation is needed by both the seismic-capacity and systems-analysis engineering teams. Today, extensive guidance on the documentation aspects is also available (Ref. SQUG-GIP, 1992; EPRI/Margins, 1991).

Evaluation: Because a large number of seismic-PSA walkdowns have been performed, and there exists excellent guidance on how to perform and document a walkdown, the methodology for seismic-PSA walkdowns should now be considered very mature. The guidance is sufficiently detailed, and the number of teams that have accomplished an excellent walkdown is large enough, that a new team should not have difficulty in learning how to perform a satisfactory walkdown.

2.2.4 Evaluation of the seismic failure-mode and fragility methodology

The seismic fragility methodology is the methodology for calculating the seismic capacity of individual structures and equipment components, and from that capacity the "fragility curve" for each item and the correlations among these. This methodology is intrinsically probabilistic in character, because it produces a probability of failure as a function of earthquake "size".

When analysing any specific structure or component, there are two different aspects of the analysis: the definition of "failure" and the determination of the fragility.

Determining "failure" modes: "Failure" must be defined before a seismic capacity can be determined.

For a structure, failure would usually be severe buckling or collapse that would compromise the safety equipment within the structure, or collapse in which the structure could fall onto and damage equipment. "Failure" usually does not include minor structural damage.

The decision about what constitutes "failure" must be made by the structural analyst on a case-by-case basis, with the advice of a competent systems analyst, and taking into account the specific safety equipment and safety functions that would be vulnerable. Sometimes more than one failure mode must be considered in the analysis. The walkdown is an essential part of the engineering determination of what "failure" means, because drawings often cannot properly capture the actual configuration of adjacent vulnerable items, nor reveal damage such as erosion that might affect a structure.

Based on considerable experience with failure-mode determinations for structures, this aspect of the methodology -- identifying structural failure modes -- is highly reliable and useful. This is especially true if a conservative assignment of "failure" is adequate.

For an item of equipment, "failure" means the inability to perform its safety function --- inability of a pump to move water, of a battery rack to provide DC power, of a valve to close or open, and so on. Sometimes "failure" can involve a short-term phenomenon involving no lasting damage, such as relay chatter that affects other equipment functions. The definitions are highly individualised for specific equipment items. As with structural failures, the decision about which failure mode(s) to consider must be made with the advice of a competent systems analyst.

Guidance on assigning failure modes is available in the various methodology guides (Ref. NRC, 1983; Brookhaven, 1985; SQUG-GIP, 1992; Reed and Kennedy, 1994).

Evaluation: Today, the determination of the appropriate failure modes for structures and components is generally a robust and reliable aspect of the seismic-PSA methodology.

The fragility analysis: The fragility of a component is defined as the conditional probability of its failure as a function of a response parameter, which in seismic PSAs is usually an acceleration parameter, such as peak ground acceleration or local spectral acceleration. Usually, a family of "fragility curves" is generated, as described well in the early procedures guides (Ref. NRC, 1983; Brookhaven, 1985). These fragility curves are typically characterised mathematically by lognormal expressions, anchored to median values and using various uncertainty parameters to capture both variability from randomness and uncertainty from lack of knowledge (Ref. Kennedy, 1980; Kennedy & Ravindra, 1984; Casciati and Faravelli, 1991; Reed and Kennedy, 1994).

It is beyond the scope here to present a thorough discussion covering either the standard mathematical formulation or its pitfalls. It is important to note, however, that the use of lognormal mathematics is known to be an erroneous approach in the tails of the lognormal distributions, even when the lognormal shape adequately describes the data in the main parts of the distribution, because the data do not fit a lognormal in the tails beyond a couple of log-standard-deviations. Despite these limitations, the lognormal is commonly used, principally for its calculational convenience (Ref. Casciati and Faravelli, 1991; Reed and Kennedy, 1994).

To develop a family of fragility curves for an item, there are three types of information that can be relied on: data from real earthquake experience, test data, and analysis. For a structure, analysis is almost always used, because structures are almost all designed individually and because they are more amenable to accurate calculation, once a determination is made of the important failure mode(s).

For equipment, reliance on test and experience data is the common approach. There are extensive data compilations in existence now, too numerous even to list here as references. Good recent compilations can be found in the two NRC seismic margins reports (Ref. Budnitz/Margins, 1985; Prassinos/Margins, 1986); the EPRI margins report (Ref. EPRI/Margins, 1991); the compilation from Livermore (Ref. Cover et al., 1995); and the SQUG report and its support (Ref. SQUG-GIP, 1992; EQE, 1991).

For many years, earthquake experience data have been used to supplement test data (Ref. SQUG-GIP, 1992; SSRAP, 1992; EQE, 1991; Merz, 1991). These data do not involve reactor equipment in earthquakes, because no nuclear power reactor has ever experienced a large earthquake, but they do involve similar or identical equipment from other installations (conventional steam power plants, refineries, etc.) that have experienced earthquakes. The experience data have improved the ability of analysts to anchor their fragility analyses to real-world experience. Also, there are enough practitioners doing this kind of analysis today that a variety of independent viewpoints are being brought to the analysis of equipment fragility.

One key outcome of the multi-year effort to compile and understand earthquake-experience data is that some important categories of equipment are now known to be generically quite rugged. This knowledge was first developed as part of the work of the Seismic Qualifications Utility Group (Ref. SQUG-GIP, 1992; SSRAP, 1992) and is also embedded in a set of screening tables for seismic capacity, that can be found in the NRC and EPRI seismic margin reports (Ref. Prassinos/Margins, 1986; EPRI/Margins, 1991). Using these screening tables and the SQUG insights, fragility analysts can screen out certain items as rugged provided that various conditions are met for each individual item so that it qualifies as a member of the ensemble represented.

An important recent report on fragility analysis is by Reed and Kennedy (Ref. Reed and Kennedy, 1994); it provides specific methods for specific types of equipment, as well as an excellent introduction to the overall methodology.

Although the discipline of seismic-fragility analysis is very mature, some uncertainties remain for many items of equipment. Specifically, for many categories there are still key unknowns, or differences among approaches, or different ways to interpret the underlying data. This means that different analysts will produce different capacities and fragility curves for the identical equipment item.

To illustrate this, it is illuminating to review the comparison (Ref. Kennedy, 1989) among four internationally-renowned expert analysts of their fragility calculations for five specific items: a flat-bottom vertical water storage tank, an auxiliary contactor for a motor starter in an older motor control center, a starting air tank, a component-cooling heat exchanger, and a cantilevered reinforced block wall. Specific design details and failure mode assumptions were provided as input. The approach was for the experts to do independent calculations first, then to compare and review the results to identify sources of differences, and finally to revise the calculations as appropriate. After the second round, the calculated median capacities differed, from the highest to the lowest of the experts, by factors in the range of about 1.5 for most of the components[2]. If the median of the four experts is considered as an approximate "best estimate", this means that the highest and lowest calculations by the experts differ by about ± 20% to ± 25%. (For a typical eastern-U.S. site, the hazard-curve frequencies differ by factors of about 2 to 3 when the acceleration differs by factors of 1.3 to 1.5, so if a single component completely dominates a given sequence, core-damage frequency would vary by a factor of about 2 to 3.)

When the same standard-problem exercises were given to engineers less well qualified than the four in the main study, the variances among the fragility results were even greater.

The lesson from this comparison is that the determination of seismic fragility curves, even by the most qualified experts, will still result in non-negligible differences, which can only be considered, for the purposes of overall PSA analysis, to be "uncertainty" in the best meaning of that term.

---

[2] For the so-called "HCLPF (High-Confidence-of-Low-Probability-of-Failure) capacity" the ratio from highest to lowest among the experts was in the range of about 1.3 to 1.4. The HCLPF capacity is the capacity at the point on the fragility curve representing a 95% confidence of a 5% probability of failure, and is a figure-of-merit in seismic-margin reviews. Calculating HCLPF capacities is described in the seismic-margin-review literature (Ref. Budnitz/Margins, 1986; EPRI/Margins, 1991; Kennedy, 1989; Reed and Kennedy, 1994).

Evaluation: The most important overview comment is that, even though the discipline of fragility analysis is quite mature, with numerous practitioners around the world, there is still important variability in the calculation of fragility parameters for items of equipment such as those cited in the test comparison study. This variability usually propagates through to modest uncertainties (but sometimes to significant uncertainties, especially where test data are limited) in the bottom-line risk results such as core-damage frequencies.

Therefore, while we conclude that the fragility estimates are reasonably good for many purposes, such as identifying the few important risk contributors at a plant, it is important not to misinterpret the numerical fragility values as implying too much accuracy.

2.2.5 Evaluation of the seismic-PSA systems analysis methodology

Given which equipment would be damaged by the postulated earthquake (typically characterised by a fragility curve that embodies a probability distribution as a function of earthquake "size" as measured by an acceleration parameter), the systems analyst must determine which core-damage accident sequences may result, and the core-damage frequencies for each. This is the primary objective of the systems-analysis part of seismic PSA.

It is important to emphasise at the outset that seismic PSAs typically identify not only accident sequences involving one or more seismic-induced failures, but also sequences involving a combination of seismic failures, human errors, and non-seismic failures such as "random" failures or maintenance unavailabilities. It is often found that accident sequences of this latter type are as important overall as the sequences involving only seismic failures. This comment should serve as a notice that the systems-analysis aspect of the overall seismic-PSA methodology is every bit as important as the seismic-failure aspect. All seismic-PSA analysts and users of seismic-PSA results must understand this.

Discussion of the Methodology: The seismic-PSA systems-analysis work is broadly similar to traditional PSA systems analysis for internal initiators. It uses the same tools and types of data, and the same way of setting up the analysis and solving it numerically. The following paragraphs will point out a few special considerations.

Logically, the systems analyst should begin with the results of the seismic fragility analysis, which will have determined which structures and equipment have been damaged by the various postulated earthquakes (as a function of earthquake "size" in terms of, say, peak ground acceleration, frequency, etc.). The systems analyst must then take into account issues such as the random (non-earthquake-caused) likelihood that other vital equipment might be out-of-service due to testing, maintenance, operator error, or failure; possible correlations among failures; and the procedures used by the operators, including their ability to recover certain earthquake-damaged or failed equipment, or to substitute other equipment, or to perform the needed safety function another way.

Of course, in practice the process is not quite as linear as described in the above paragraph, but rather it will be more iterative: the systems analysts and the seismic-capacity analysts should have been working together from the start to screen out certain potential issues, develop input information on others, and help each other to focus on the issues deemed important. There will have been several iterations in any well-executed seismic-PSA study.

At the center of the systems analysis work is developing one or more accident sequence event trees, that include the various functions or systems needed for safe shutdown, possible operator prevention and recovery actions, and the like. The success-or-failure numerical values on the event-tree branch points are then worked out using either data or inputs from fault trees. If we assume that the analyst has access to a completed internal-initiators PSA (which should almost always be the situation), then direct use can be made of such vital information as the random failure data, the operating crew's procedures, and the support-system matrix. (Support systems such as AC power, instrument air, service water, and so on support the vital front-line safety equipment.) Otherwise, the analyst must develop this information anew.

If fault trees from an internal-initiator PSA analysis are used, they must be modified somewhat to account for location correlations and to introduce different seismic failure modes. Also, seismic failures of certain passive items such as internal block walls or piping will not have been modelled in the internal-initiators PSA, and the effects of such failures will need to be incorporated into the fault trees.

The outcome of the systems analysis is the numerical value of core-damage frequency (actually, a density function that captures uncertainties) for each of several (usually discrete) earthquake sizes.

Four special issues need to be discussed here, because the methodologies for them are distinct from other methodologies: correlations among failures, relay chatter, design and construction errors, and post-earthquake operator response.

Correlations among failures: It can sometimes be difficult to analyse correlations among earthquake-induced failures.

The most logical assumption, which seems at first to be universally appropriate, is that the earthquake motion coming into the site will affect all buildings and components in a fully correlated way. However, at different locations in a building, and certainly in different buildings, this correlation is diluted by several intervening factors.

Typically, the PSA analysis will assume complete correlation in the response for nearby and similar equipment that is subject to the same floor motion. However, different equipment types, even if located in close proximity, are usually assigned only minor (if any) response correlation. Furthermore, even high response correlation does not always imply high capacity correlation, which would arise most obviously when, for example, two valves come from the same manufacturer and the same assembly line, with adjacent model numbers.

The difficulty is that there is only very limited experimental information on correlations, from either testing or actual earthquakes, upon which to rely. Therefore, while the methodology for coping with correlations is well-developed (Ref. Ravindra, 1984; Reed, 1985), the underlying knowledge needed to perform the calculations is almost always inadequate.

To overcome the problem, the usual fallback approach is to perform a sensitivity analysis, for example assuming complete correlation and then complete independence and ascertaining what difference these two assumptions make. The difference is then understood as representing a measure of the uncertainty in the final results.

The analyst must take care about correlations not only in the central values but in the uncertainties. If neither of two parameters is known well, but what little is known comes from the same data set, the correlation in the uncertainty can be high.

Whenever the accident sequences of concern involve components for which correlation might or might not be large, and the sensitivity analysis shows that the bottom-line results are sensitive to correlations, this issue is one of the important sources of uncertainty in the overall analysis. (Conversely, if a key sequence is dominated by a single failure, or by two failures of very different kinds --- an example would be a sequence in which a large yard tank fails seismically simultaneously with a battery rack --- both response correlations and capacity correlations should be minor and the sensitivity of the results should also be minor.)

A summary evaluation of the correlation issue is that, while the methodology for analysing correlations certainly exists in an adequate form, the underlying data are usually inadequate, so that uncertainties in the final PSA results can sometimes be important due to the issue of correlation.

Relay chatter: In the first several seismic PSAs (early 1980s), the issue of relay chatter was not analysed at all. Instead, the assumption was made that all relay chatter was recoverable by the operating crew, which assumption is tantamount to assuming no relay chatter. Later, however, this issue received significant attention, especially after an NRC-sponsored study of chatter at two power plants (Ref. Budnitz, Lambert, & Hill, 1987) demonstrated that if there is no operator recovery the chattering of key relays could lead to core-damage accident sequences with high annual frequencies. This study also developed and applied a methodology for examining relay-chatter issues in the context of a full-scope PSA. Later, the Diablo Canyon PSA (Ref. PG&E, 1988) and the Hatch seismic-margin and PSA review (Ref. Moore, Wooten, and Kassawara, 1990; Orvis and Moieni, 1990) included thorough examinations of relay chatter. Furthermore, the test data base on seismic capacities for relay chatter has become more and more extensive (Ref. MPR, 1990; Merz, 1991a; SQUG-GIP, 1992), at least for relays installed in U.S. nuclear plants.

The relay-chatter seismic-PSA methodology usually employs a successive-screening approach (Ref. Hardy and Ravindra, 1990), in which the initial step is to compile a list of all relays that participate in safety-important circuits and functions. This long list is reduced by a combination of screening steps: many relays are eliminated by demonstrating that they do not participate in any important safety functions; others are screened out by reference to the test data base that demonstrates that they are very rugged; still others are screened out by a detailed circuit analysis to show that their chatter is benign; and so on. In U.S. nuclear-power plants, using such an iterative process, the original relay list that may involve hundreds of relays is usually reduced to a very few (typically less than ten, sometimes even none) of concern.

Typically, after the analysis has identified any relays whose chatter can be troublesome to important safety functions, the next step is to remedy the situation by either (i) replacing the particular relay; or (ii) changing the circuit to eliminate, for example, seal-in or reset problems; or (iii) instructing the operators to be alert to post-earthquake relay-chatter problems; or some combination. The analysts in the most recent U.S. seismic PSAs have almost never actually included relay-chatter failures in the PSA analysis itself using an appropriate seismic fragility curve. Rather, they have almost always identified a few potential problems -- sometimes none at all -- which have been "fixed".

When the U.S. NRC required all U.S. nuclear power stations to perform a seismic review beyond their original design basis (Ref. NRC/IPEEE, 1991; NRC-IPEEE-1407, 1991), about half of them selected a seismic PSA, and several dozen PSA-based relay-chatter reviews have been performed on U.S. plants in the past five years. The guidance for these analyses can be found in (Ref. Hardy and Ravindra, 1990; NRC/1602, 1997).

The successful execution of these numerous studies, which as mentioned usually identified only a few relays of concern, demonstrates that as of today (i) an acceptable methodology does exist for treating relay-chatter properly; and that (ii) the issue should not be ignored, because it certainly has a potential for contributing significantly to the overall seismic risk at a nuclear power plant.

Design and construction errors: Except in the rare case that a design or construction error may be identified during the walkdowns or the study of design drawings, the seismic-PSA methodology does not systematically take into account possible design and construction errors.

This may seem like a serious flaw in the methodology. In actual fact, there is no way to know whether it is or not. It is important to note, however, that these omissions are directly parallel to possible omissions in the rest of PSA, such as in the analysis of internally-initiated accidents, where possible design errors affecting the configuration of systems are also not accounted for properly. This is not an excuse, but rather a generic weakness of all PSAs -- it is actually a weakness in our overall knowledge about how the plants are actually configured, and this weakness propagates through to PSA but is not intrinsically a PSA "methodological flaw". (Of course, a rigorous pre-operational testing program and design review should identify most of these errors.)

Post-earthquake operator response: During and after a strong-motion earthquake, it seems likely that the ability of control-room operators to perform their assigned tasks without error should be substantially degraded, because of high levels of stress and confusion. This issue has been examined (Ref. Budnitz, Lambert, & Hill, 1987), and a model has been proposed to account more effectively for possible high operator stress. Several other high-stress models are also in the literature (Ref. Spurgin and Moieni, 1991; Cooper et al., 1996; Hirschberg et al., 1996). Unfortunately, because good data are lacking, there is no way today to sort out with confidence which of the several models of degraded post-earthquake operator performance is best.

However, at least for U.S. nuclear power plants, this issue does not have as much effect on the results of seismic PSAs as might be thought at first, principally because the assumption is commonly made that no credit is allowed for operator control actions during the early minutes --- often for as long as a half-hour --- after a large earthquake. This assumption is justified by the existence in U.S. plants of automatic equipment that usually does not require any operator intervention in the first half-hour or so. By that time, things should have settled down (literally and figuratively), so that the normal (non-seismic) PSA methodology for analysing operator errors should apply. If this is true of other plants in other countries, then the same argument would apply.

Based on this, the conclusion is that in U.S. plants the operator-response aspect of the seismic-PSA methodology, while not as strongly based on knowledge as would ultimately be desirable, is reasonably mature, and it is as robust (more-or-less) as the approach for operator error analysis used in internal-initiators PSA studies.

Evaluation of the Overall Systems-Analysis Methodology: As mentioned above, the seismic systems-analysis methodology is, in its basic outline, a variant of the type of systems analysis that is now a well-developed, mature PSA discipline.

Although certain important issues require special attention and treatment, every aspect of the methodology, including correlations, relay chatter, and operator response, is fully within the routine capability of PSA systems analysts. Therefore, any competent PSA systems analyst can perform this work, with little special training and only the minimal guidance that is readily available and easily learned.

2.2.6 Evaluation of the seismic-PSA consequence/release methodology

Discussion of the Methodology: The objective of the seismic-PSA consequence/release methodology is to calculate, for various postulated earthquake "sizes" associated with various probabilities of core damage, the conditional probability that the postulated accident will evolve into a "radiological release" scenario.

This conditional probability differs from one postulated core-damage accident sequence to the next. Therefore, each sequence requires separate treatment, depending on which items of safety equipment have been damaged by the postulated earthquake, which operator actions have contributed to the damage or mitigated the situation, which equipment has failed from other (non-seismic) causes, and so on.

The size of the release also obviously depends on how phenomena develop both within the primary system, outside the vessel, and in the containment after core damage begins; how ex-plant radiological dispersion phenomena develop; and how sheltering and evacuation are accomplished.

It is important that the analysis team consider a few special issues, such as the possibility that the earthquake may affect containment integrity, either for the structure itself or, more likely, for the penetrations or other ways in which integrity can be compromised.

Also, if the earthquake has caused extensive damage off-site, such as to roads, buildings, and bridges, or widespread panic among the public, the effect of these issues on emergency evacuation must be assessed.

Evaluation of the Methodology: In its basic outline, the consequence/release methodology is a variant of the type of level-2 and level-3 analysis that is now a well-developed, mature PSA discipline. The methods and data used are similar or identical, including the use of containment event trees (or accident-progression event trees, as they are now often called) and off-site-consequence codes. While a few issues must be specially treated, we conclude that any competent PSA level-2/level-3 analysis team can perform this work, with no special training.

Because some of the special issues --- such as off-site seismic damage and panic and their effect on evacuation --- are difficult and highly uncertain to analyse, the reliability and usefulness of the results can be significantly compromised. This is not a fault of the methodology per se, but rather a potential problem if the analyst is not thorough in developing all of the issues fully.

2.3 Evaluation of the Reliability and Usefulness of the "Bottom-Line" Results for Core Damage Frequency and Off-site Risk, and the Key Risk Insights

The numerical uncertainties in the bottom-line core-damage-frequency and off-site-risk results of a seismic PSA can certainly be large (plus-or-minus more than one order of magnitude or more is common). This is due to several factors in the various sub-methodologies, but dominantly due to the uncertainty in the seismic-hazard evaluation. The uncertainties in the fragility estimates per se make smaller but important contributions to the overall uncertainty. Perhaps the other major source of possible uncertainty arises when several components must fail together to cause the accident sequence, and the correlations among them are not well understood --- the differences between assuming full correlation and zero correlation can also amount to about an order of magnitude difference in core-damage frequency in some cases.

Despite the numerically large uncertainties, these uncertainties should generally not invalidate the key insights concerning potential earthquake-related vulnerabilities. These insights include, among others:

- the identification of specific equipment and structural weaknesses, including weaknesses in components and systems that are not specifically designed or qualified against earthquakes;
- specific non-seismic-initiated failures and human errors that may contribute to a key sequence;
- the possible role of post-earthquake operator recovery actions;
- whether a given sequence would have major or only minor off-site-release consequences; and
- the places where support-system vulnerabilities can compromise different safety systems in subtle ways.

One of the major seismic-PSA lessons is that an integrated examination of the plant, by a team including both seismic-capacity engineers and systems engineers, can be of major benefit in identifying issues that neither type of expert could find alone.

Another major benefit is that an integrated examination of seismic issues in the context of the rest of the plant's safety functions and systems is crucial --- and PSA analysis can accomplish this integrated examination very well.

Finally, the benefits of a seismic walkdown by an integrated team cannot be overemphasised: only in this way can some of the important issues be identified.

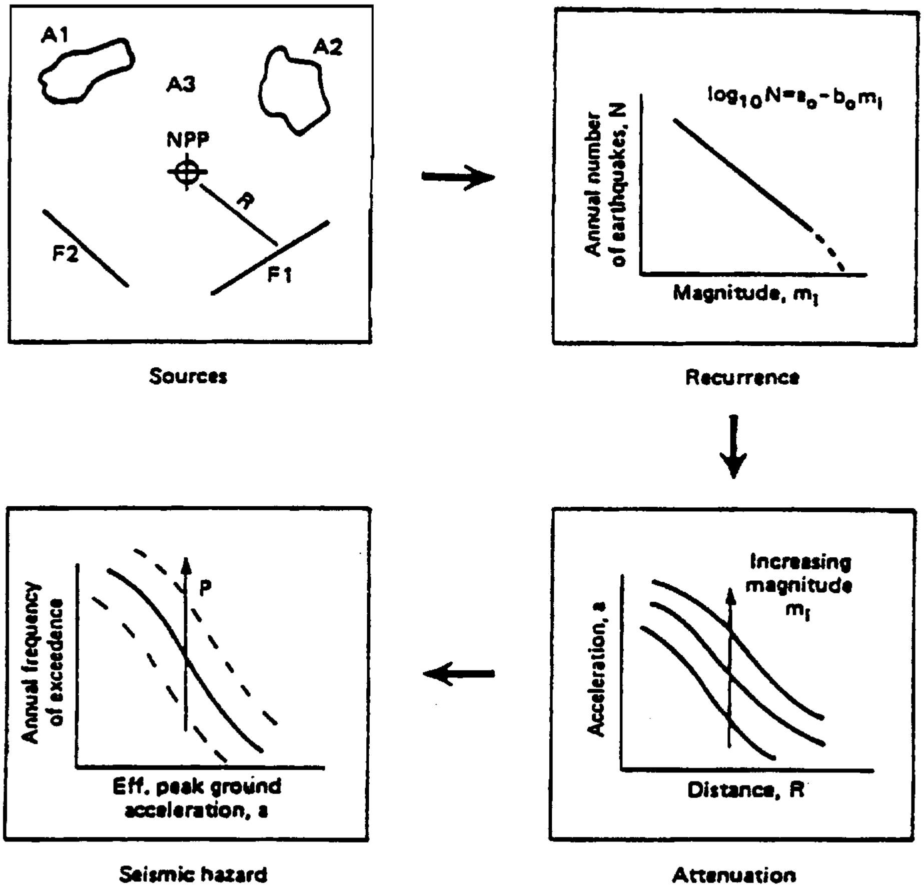

Steps to Evaluate the Frequency of Exceedance of Ground Motion from Earthquakes

*** taken from the "PRA Procedures Guide" (Ref. NRC, 1983)**

Figure II

3. SUMMARY EVALUATION OF THE SEISMIC-PSA METHODOLOGY

This chapter will provide a summary evaluation of the reliability and usefulness of the PSA methodology for studying earthquakes. The full text in Chapter 2 supports the summary statements here. The six sub-methodologies will be evaluated first, followed by an evaluation of the overall methodology.

3.1 Evaluation of the Six Sub-Methodologies

3.1.1 How reliable and useful is the seismic hazard methodology?

The "hazard methodology" analyses for the frequency of earthquakes of various sizes at a given site, and the spectral shapes of the motion from these earthquakes. The methodology has four steps (see Figure 2-1). A fair characterisation is that each of the steps is straightforward to describe, but difficult to implement.

The four-step approach begins with a seismicity assessment, to delineate and characterise the seismic sources. The second step involves determining the earthquake recurrence relationship, which is usually expressed in terms of an annual frequency as a function of magnitude for each source or source zone. The third step is associating a motion vs. distance relationship with each magnitude. Usually, acceleration (in terms of peak ground acceleration or local spectral acceleration) is chosen as the motion parameter, even though it is an imperfect measure --- in fact, no single parameter can be other than an imperfect measure. The fourth step produces the ultimate product of the hazard assessment, the "hazard curves" themselves, typically in terms of the annual frequency of exceedance vs. a motion parameter like peak ground acceleration.

Evaluation: The methodology, as embedded in the above four seismic-hazard-analysis steps, is mature in the sense that it is well-developed and widely practiced. Nevertheless, problems with the methodology generally lead to important uncertainties in the hazard results. The central reason is that experts approaching this analysis problem from different perspectives have developed a large number of reasonable models, all of which are consistent with the available earthquake data, but which have important and different implications for the hazard results. The Brookhaven PSA guide (Ref. Brookhaven, 1985) states the situation succinctly: "The development [of a seismic hazard model] is a product of scientific interpretation of uncertain and incomplete physical evidence on geological structures, tectonic processes, and seismicity." In sum, reasonable experts differ in their assessments, and selecting which expert is correct is not possible today.

To summarise, suffice it to say that, from the perspective of a decision-maker, a legitimate and wide divergence of opinion among experts must be taken on its face as genuine "uncertainty" in the best meaning of that term. The fact that different models can lead to PSA core-damage-frequency calculations differing by more than a factor of 10 is simply a manifestation of the current state-of-the-art in this discipline.

3.1.2 How reliable and useful is the local-ground-motion and building-motion methodology?

This part of the overall methodology is generally very well developed, although uncertainties do remain in some site-specific situations.

The analyst usually starts with a family of earthquake motions, either time histories or another characterisation, that are postulated to arrive at the local site from the source. As a set, the time histories are intended to capture variability in the source. Usually, several different earthquake "sizes" are calculated. The objective is to work out the local motion at the location of each significant item (equipment items and structures) necessary for the safety of the power station.

For items located at different elevations, either above or below grade, it is necessary to develop what are known as floor spectra, using a structural model of the building. This part of the analysis begins with determining the ground response frequency spectrum at the site. Usually, generic broad-band spectra have been used in PSAs, and this should be acceptable, provided they are applicable to local soil conditions. If the local soil produces much amplification, the analyst should develop a site-specific response spectrum. If structural foundations are on rock or stiff soil, their motions should be the same as for the free field. On soft soil, the soil-structure coupling can change both the frequencies and the amplitudes of the motion entering the building.

Evaluation: While uncertainties certainly exist in this aspect of the seismic-PSA analysis, arising from both variabilities and modelling approximations, the analytical approaches for the several topics are all generally well-developed and robust in the hands of experienced analysts. Therefore, the state-of-the-art of the seismic-PSA local-ground-motion and building-motion methodology is mature, and this aspect should be very reliable and very useful.

3.1.3 How reliable and useful is the walkdown methodology?

Among seismic-PSA analysts, the plant walkdown is considered to be almost the most crucial aspect of the entire process.

One of the most important benefits of the walkdown is the interaction that occurs among the systems analysis team, the seismic-capacity analysis team, and the utility staff. This is one of the major lessons learned in past seismic PSA studies. Another crucial benefit of the walkdown is that the seismic-capacity team can determine, for each important item (structure or equipment), whether that item is "typical" of its generic category, or somehow atypical or even unique. Still another benefit is the opportunity for the systems-analysis team to understand just how the operating crew has been trained to carry out its tasks, especially during emergencies.

Evaluation: Because a large number of seismic-PSA walkdowns have been performed, and there exists excellent guidance on how to perform and document a walkdown, the methodology for seismic-PSA walkdowns should now be considered very mature. The guidance is sufficiently detailed, and the number of teams that have accomplished an excellent walkdown is large enough, that a new team should not have difficulty in learning how to perform a satisfactory walkdown. In summary, this aspect should be very reliable and very useful.

3.1.4 How reliable and useful is the failure-mode and fragility methodology for earthquakes?

The seismic fragility methodology calculates the capacity of individual structures and equipment items, and from that capacity the "fragility curve" for each item and the correlations among these.

3.1.4.1 Failure Modes:

Before capacity can be determined, "failure" must be defined, for both structures and equipment items. The definitions are highly individualised for specific equipment items, and the assignments must be made with the advice of a competent systems analyst. However, excellent guidance exists and there are numerous practitioners with experience in failure-mode identification.

Evaluation: Today, the determination of the appropriate failure modes for structures and components is generally a robust and reliable aspect of the seismic-PSA methodology.

3.1.4.2 Fragility

The fragility of a component is defined as the conditional probability of its failure as a function of a response parameter, usually an acceleration parameter, such as peak ground acceleration or local spectral acceleration. A family of "fragility curves" is generated, typically characterised mathematically by lognormal expressions, anchored to median values and using various uncertainty parameters to capture both variability from randomness and uncertainty from lack of knowledge.

To develop a family of fragility curves, the analyst can use test data, data from real earthquake experience, and/or analysis. For a structure, analysis is usually used, since structures are all so individualised and since they are more amenable to calculation given a determination of the important failure mode(s). For equipment, reliance on test and experience data is the common approach, because there are now extensive test data compilations in existence, as well as extensive earthquake-experience data. Some important items of equipment are now known to be generically quite rugged. This knowledge is embedded in a set of screening tables for seismic capacity that can be found in the literature.

Despite major progress in our understanding, some important uncertainties remain, especially for many items of equipment. Specifically, there are still many unknowns, or differences among approaches, or different ways to interpret the underlying data, so that different analysts will produce different capacities and fragility curves for the identical item. Based on experience with trial analyses, the conclusion is that there is still some important uncertainty in the calculation of seismic fragility curves.

Evaluation: The most important overview comment is that, even though the discipline of fragility analysis is quite mature, with numerous practitioners around the world, there is still important uncertainty in the calculation of fragility parameters, especially for items of equipment but also for structures. This uncertainty usually propagates through to modest uncertainties (but sometimes to significant uncertainties, especially where test data are limited) in the bottom-line risk results such as core-damage frequencies.

Therefore, while we conclude that the fragility estimates are reasonably good for many purposes, such as identifying the few important risk contributors at a plant, it is important not to misinterpret the numerical fragility values as implying too much accuracy.

3.1.5 How reliable and useful is the systems-analysis methodology?

The objective of the systems-analysis methodology, given which equipment would be damaged by the postulated earthquake, is to determine which core-damage accident sequences may result, and the core-damage frequencies for each.

The systems-analysis work is broadly similar to traditional PSA systems analysis for internal initiators, and is within the technical capability of any competent PSA systems analyst, with no special training. There are only a few special issues: correlations among failures, relay chatter, design and construction errors, and operator response.

3.1.5.1 Correlations

The problem of analysing correlations among earthquake-induced failures can sometimes be difficult, especially for co-located equipment. Typically, the assumption is made of complete correlation in the response for nearby and similar equipment subject to the same floor motion. However, different equipment types, even if located in close proximity, are usually assigned only minor if any response correlation. The problem for the analyst is that there is only very limited experimental information, from either testing or actual earthquakes, upon which to rely.

A summary evaluation of the correlation issue is that, while the methodology for analysing correlations certainly exists in an adequate form, the underlying data are usually inadequate, so that uncertainties in the final PSA results can sometimes be important due to the issue of correlation. Whenever the accident sequences of concern involve components for which correlation might or might not be large, this issue is one of the important sources of uncertainty in the overall analysis. The usual fallback approach is to perform a sensitivity analysis, to obtain a measure of the uncertainty in the final results.

3.1.5.2 Design and construction errors

The seismic-PSA methodology does not systematically take into account possible design and construction errors. This may seem like a serious flaw in the methodology. In actual fact, there is no way to know whether it is or not. It is important to note, however, that these omissions are directly parallel to possible omissions in the rest of PSA, such as in the analysis of internally-initiated accidents, where possible design and construction errors affecting the configuration of systems are also not accounted for properly.

3.1.5.3 Relay chatter

Recently, the relay-chatter issue has received significant attention. While the earliest seismic PSAs did not examine this issue, today an acceptable methodology does exist for treating it properly.

The relay-chatter seismic-PSA methodology usually employs a successive-screening approach, in which the initial step is to compile a list of all relays that participate in safety-important circuits and functions. This long list is reduced by a combination of screening steps: many relays are eliminated by demonstrating that they do not participate in any important safety functions; others are screened out by reference to the test data; still others are screened out by a detailed circuit analysis; and so on. In U.S. nuclear-power plants, using such an iterative process, the original relay list that may involve hundreds of relays is usually reduced to a very few (typically less than ten, sometimes even none) of concern.

Typically, after the analysis has identified any relays whose chatter can be troublesome to important safety functions, the next step is to remedy the situation. The analysts in the most recent U.S. seismic PSAs have almost never actually included relay-chatter failures in the PSA analysis itself using an appropriate seismic fragility curve. Rather, they have almost always identified a few potential problems -- sometimes none at all -- which have been "fixed". The successful execution of these numerous studies demonstrates that as of today an acceptable methodology does exist for treating relay-chatter properly.

3.1.5.4 Post-earthquake operator response

During and after a strong-motion earthquake, it seems likely that the ability of control-room operators to perform their assigned tasks without error should be substantially degraded, because of high levels of stress and confusion. This issue has been examined, and numerous models have been proposed. Unfortunately, because data are lacking, there is no way today to sort out with confidence which of the several models of degraded post-earthquake operator performance is best.

However, at least for U.S. nuclear power plants, this issue does not have as much effect on the results of seismic PSAs as might be thought at first, principally because the assumption is commonly made, based on plant operating instructions, that no credit is allowed for operator control actions during the early minutes after a large earthquake. By that time, the normal (non-seismic) PSA methodology for analysing operator errors should apply. If this is true of other plants in other countries, then the same argument would apply. Based on this, the conclusion is that in U.S. plants the post-earthquake operator-response methodology, while not as strongly based on knowledge as would ultimately be desirable, is reasonably mature, and it is as robust (more-or-less) as the approach for operator error analysis used in internal-initiators PSA studies.

3.1.5.5 Evaluation of the Overall Systems-Analysis Methodology

As mentioned above, the seismic systems-analysis methodology is, in its basic outline, a variant of the type of systems analysis that is now a well-developed, mature PSA discipline.

Although certain important issues require special attention and treatment, every aspect of the methodology, including correlations, relay chatter, and operator response, is fully within the routine capability of PSA systems analysts. Therefore, any competent PSA systems analyst can perform this work, with little special training and only the minimal guidance that is readily available and easily learned. This aspect should, therefore, be very reliable and useful.

3.1.6 How reliable and useful is the consequence/release methodology?

The objective of the seismic-PSA consequence/release methodology is to calculate, for various postulated earthquake "sizes" associated with various probabilities of core damage, the conditional probability that the postulated accident will evolve into a "radiological release" scenario. This conditional probability differs from one postulated core-damage accident sequence to the next. Therefore, each sequence requires separate treatment.

It is important that the analysis team consider a few special issues, such as the possibility that the earthquake may affect containment integrity, and the effect on emergency evacuation of possible extensive damage off-site, such as to roads, buildings, and bridges, or widespread panic among the public.

Evaluation of the Methodology: In its basic outline, the consequence/release methodology is a variant of the type of level-2 and level-3 analysis that is now a well-developed, mature PSA discipline. The methods and data used are similar or identical, including the use of containment event trees (or accident-progression event trees, as they are now often called) and off-site-consequence codes. Although a few issues must be specially treated -- and may be difficult and highly uncertain in some circumstances -- we conclude that any competent PSA level-2/level-3 analysis team can perform this work, with no special training.

3.2 Evaluation of the Overall Results of Seismic PSA

How reliable and useful are the "bottom-line numbers" for core-damage frequency and off-site risk, and the key risk insights?

The numerical uncertainties in the bottom-line results can certainly be large (plus-or-minus more than one order of magnitude or more is common). This is due to several factors in the various sub-methodologies, but dominantly due to the uncertainty in the seismic-hazard evaluation. The uncertainties in the fragility estimates per se make smaller but important contributions to the overall uncertainty. Perhaps the other major source of possible uncertainty would arise when several components must fail together to cause the postulated accident sequence, and the correlations among them are not understood well --- the differences between assuming full correlation and zero correlation can also amount to about an order of magnitude difference in core-damage frequency in some cases.

Despite the numerically large uncertainties, these uncertainties should generally not invalidate the key engineering insights concerning potential earthquake-related vulnerabilities.

4. REFERENCES

1. ASCE-4, 1997: American Society of Civil Engineers, "Seismic Analysis of Safety-Related Nuclear Structures and Commentary - Final Draft" (1997)

2. Bernreuter, 1987: D.L. Bernreuter, J.C. Chen, and J.B. Savy, "Development of Site Specific Response Spectra", Report NUREG/CR-4861, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1987)

3. Bohn, 1984: M.P. Bohn et al., "Application of the SSMRP Methodology to the Seismic Risk at the Zion Nuclear Power Plant", Report NUREG/CR-3428, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1984)

4. Bohn and Lambright, 1988: M.P. Bohn and J.A. Lambright, "Recommended Procedures for Simplified External Event Risk Analyses", Report NUREG/CR-4840, Sandia National Laboratories and U.S. Nuclear Regulatory Commission (draft report, 1988)

5. Brookhaven, 1985: M. McCann, J. Reed, C. Ruger, K. Shiu, T. Teichmann, A. Unione, and R. Youngblood, "Probabilistic Safety Analysis Procedures Guide", Report NUREG/CR-2815, Volume 2, Brookhaven National Laboratory and U.S. Nuclear Regulatory Commission (1985)

6. Budnitz, 1984: R.J. Budnitz, "External Initiators in Probabilistic Reactor Analysis ---- Earthquakes, Fires, Floods, Winds", Risk Analysis 4, 323 (1984)

7. Budnitz/Margins, 1985: R.J. Budnitz, P.J. Amico, C.A. Cornell, W.J. Hall, R.P. Kennedy, J.W. Reed, and M. Shinozuka, "An Approach to the Quantification of Seismic Margins in Nuclear Power Plants", Report NUREG/CR-4334, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1985)

8. Budnitz, 1987: R.J. Budnitz, "Recent Developments in Methodology and Applications for Seismic PRA", paper published in proceedings of PSA'87 Symposium (see Ref. PSA'87/Zurich, 1987)

9. Budnitz, Lambert, and Hill, 1987: R.J. Budnitz, H.E. Lambert, and E.E. Hill, "Relay Chatter and Operator Response After a Large Earthquake: An Improved PRA Methodology with Case Studies", Report NUREG/CR-4910, Future Resources Associates, Berkeley, California and U.S. Nuclear Regulatory Commission (1987)

10. Budnitz and Lambert, 1989: R.J. Budnitz and H.E. Lambert, "An Evaluation of the Reliability and Usefulness of External-Initiator PRA Methodologies", Report NUREG/CR-5477, Future Resources Associates Inc. for the U.S. Nuclear Regulatory Commission (1989)

11. Budnitz, Murray, and Ravindra, 1992: R.J. Budnitz, R.C. Murray, and M.K. Ravindra, "NRC Seismic Margin Methodology: Updated and Supplementary Guidance", Report NUREG/CR-6033, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1992)

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

13. Casciati and Faravelli, 1991: F. Casciati and L. Faravelli, Fragility Analysis of Complex Structural Systems, John Wiley & Sons Inc., New York (1991)

14. Cooper et al., 1996: S.E. Cooper, A. Ramey-Smith, J. Wreathall, G.W. Parry, D.C. Bley, J.H. Taylor, and W. J. Luckas, "A Technique for Human Error Analysis (ATHEANA): Technical Basis and Methodology Description", Report NUREG/CR-6350, U.S. Nuclear Regulatory Commission (1996)

15. Cover et al., 1985: L.E. Cover et al., "Handbook of Nuclear Power Plant Seismic Fragilities", Report NUREG/CR-3558, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1985)

16. EPRI/Hazard, 1989: Electric Power Research Institute, "Probabilistic Seismic Hazard Evaluations at Nuclear Plant Sites in the Central and Eastern United States: Resolution of the Charleston Earthquake Issue", Report EPRI-NP-6395D (1989)

17. EPRI/Margins, 1991: NTS Engineering, RPK Structural Mechanics Consulting, Pickard Lowe & Garrick, Woodward Clyde Consultants, and Duke Power Company, "A Methodology for Assessment of Nuclear Power Plant Seismic Margin", Report EPRI-NP-6041, Electric Power Research Institute (1991)

18. EQE, 1991: EQE Engineering, Inc., "Summary of the Seismic Adequacy of Twenty Classes of Equipment Required for Safe Shutdown of Nuclear Plants", Report EPRI-NP-7149, Electric Power Research Institute (1991)

19. ERI, 1997: Energy Research Inc., "Individual Plant Examination of External Events (IPEEE) Program: Preliminary Insights on Results, Plant Improvements, and Implementation Strengths and Weaknesses", Report ERI/NRC-97-502, published for U.S. Nuclear Regulatory Commission by ERI, Rockville, MD (1997)

20. Geomatrix, 1996: Geomatrix, Inc. TRW, "Probabilistic Volcanic Hazard Analysis for Yucca Mountain, Nevada", Report BA-1717-2200-0082, prepared for U.S. Department of Energy Yucca Mountain Site Characterization Project (1996)

21. Hardy and Ravindra, 1990: G.S. Hardy and M.K. Ravindra, "Guidance on Relay Chatter Effects", Report NUREG/CR-5499, EQE International, Inc. and U.S. Nuclear Regulatory Commission (1990)

22. Hirschberg et al., 1996: S. Hirschberg, V.N. Dang, and L. Reiman, "A Survey of International Practices in Human Reliability Analysis", in "Proceedings of PSA'96, International Topical Meeting on Probabilistic Safety Assessment", Vol. 1, p. 78, published by American Nuclear Society (1996)

23. Johnson/SMACS, 1981: J.J. Johnson, G.L. Goudreau, S.E. Bumpus, and O.R. Maslenikov, "Seismic Safety Margins Research Program -- Phase I Final Report, SMACS [Seismic Methodology Analysis Chain with Statistics] (Project VIII)", Report NUREG/CR-2015 (Volume 9), Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1981)

24. Johnson, Schewe, and Maslenikov, 1984: J.J. Johnson, E.C. Schewe, and O.R. Maslenikov, "Soil Structure Interaction Response of a Typical Shear Wall Structure", Report UCID-20122, Vol. 1 and 2, Lawrence Livermore National Laboratory (1984)

25. Kennedy, 1980: R.P. Kennedy et al., "Probabilistic Seismic Safety Study of an Existing Nuclear Power Plant", Nuclear Engineering and Design, Vol. 59, No. 2, pp. 315-338 (1980)

26. Kennedy, 1981: R.P. Kennedy, R.D. Campbell, D.A. Wesley, H. Kamil, A. Gantayat, and R. Vasudevan, "Subsystem Response Review - Seismic Safety Margins Research Program", Report NUREG/CR-1706, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1981)

27. Kennedy, 1989: R.P. Kennedy, R.C. Murray, M.K. Ravindra, J.W. Reed, and J.D. Stevenson, "Assessment of Seismic Margin Calculation Methods", Report NUREG/CR-5270, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1989)

28. Kennedy and Ravindra, 1984: R.P. Kennedy and M.K. Ravindra, "Seismic Fragilities for Nuclear Power Plant Risk Studies", Nuclear Engineering and Design, Vol. 79, No. 1, pp. 347-68 (1984)

29. LLNL, 1989: D.L. Bernreuter, J.B. Savy, R.W. Mensing, and J.C. Chen, "Seismic Hazard Characterization of 69 Nuclear Plant Sites East of the Rocky Mountains", Report NUREG/CR-5250, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1989)

30. Merz, 1991: K. Merz, "Generic Seismic Ruggedness of Power Plant Equipment in Nuclear Power Plants", Report EPRI-NP-5223, Revision 1, prepared by ANCO Engineers for Electric Power Research Institute (1991)

31. Merz, 1991a: K. Merz, "Seismic Ruggedness of Relays", Report EPRI-NP-7147, prepared by ANCO Engineers Inc. for Electric Power Research Institute (1991)

32. Moore, Wooten, and Kassawara, 1990: D.P. Moore, K.D. Wooten, and R.P. Kassawara, "Results of the Seismic Margin Assessment of Hatch Nuclear Power Plant", paper presented at "Second Symposium on Current Issues Related to Nuclear Power Plant Structures, Equipment, and Piping", Orlando, Florida, sponsored by Electric Power Research Institute (1990)

33. MPR, 1990: MPR Associates Inc., "Procedure for Evaluating Nuclear Power Plant Relay Seismic Functionality", Re;port EPRI-NP-7148, prepared by MPR for Electric Power Research Institute (1990)

34. Newmark and Hall, 1978: N.W. Newmark and W. Hall, "Development of Criteria for Seismic Review of Selected Nuclear Power Plants", Report NUREG/CR-0098, U.S. Nuclear Regulatory Commission (1978)

35. NRC, 1983: J. Hickman et al., "PRA Procedures Guide: A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants", Report NUREG/CR-2300, American Nuclear Society, Institute of Electrical and Electronic Engineers, and U.S. Nuclear Regulatory Commission (1983)

36. NRC, 1984: U.S. Nuclear Regulatory Commission, "PRA Reference Document", Report NUREG-1050 (1984)

37. NRC/IPEEE, 1991: U.S. Nuclear Regulatory Commission, Generic Letter No. 88-20, Supplement 4, "Individual Plant Examination for External Events (IPEEE) for Severe Accident Vulnerabilities - 10 CFR 50.54(f)" (1991)

38. NRC-IPEEE-1407, 1991: U.S. Nuclear Regulatory Commission, "Procedural and Submittal Guidance for the Individual Plant Examination of External Events (IPEEE) for Several Accident Vulnerabilities", Report NUREG-1407 (1991)

39. NRC, 1993: U.S. Nuclear Regulatory Commission, "Revised Livermore Seismic Hazard Estimates for 69 Nuclear Power Plant Sites East of the Rocky Mountains", Report NUREG-1488 (1993)

40. NRC/SECY-97-077, 1997: U.S. Nuclear Regulatory Commission, Policy Paper, SECY-88-147, L. Callan (EDO) to Commissioners, "Draft Regulatory Guides, Standard Review Plans and NUREG Document in Support of Risk Informed Regulation of Power Reactors" (April 8, 1997)

41. NRC/1602, 1997: U.S. Nuclear Regulatory Commission, "The Use of PRA in Risk-Informed Applications", Report NUREG-1602, draft report issued for comment (June 1997)

42. NRC/DG-1061, 1997: U.S. Nuclear Regulatory Commission, "An Approach for Using PRA in Risk-Informed Decisions on Plant-Specific Changes to the Current Licensing Basis", Draft Regulatory Guide issued for comment (June 1997)

43. Orvis and Moieni, 1990: D.D. Orvis and P. Moieni, "Seismic Margin Review of Plant Hatch Unit 1: Systems Analysis", Report NUREG/CR-5632, Accident Prevention Group and U.S. Nuclear Regulatory Commission (1990)

44. PG&E, 1988: Pacific Gas and Electric Company, "Long Term Seismic Program Final Report --- Chapter 6, Probabilistic Risk Analysis", work also accomplished in part by Pickard Lowe & Garrick (1988)

45. Prassinos/Margins, 1986: P.G. Prassinos, M.K. Ravindra, and J.B. Savy, "Recommendations to the Nuclear Regulatory Commission on Trial Guidelines for Seismic Margin Reviews of Nuclear Power Plants", Report NUREG/CR-4482, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1986)

46. Ravindra, 1984: M.K. Ravindra, H. Banon, R.H. Sues, and R.D. Thrasher, "Sensitivity Studies of Seismic Risk Models", Report EPRI-NP-3562, Electric Power Research Institute (1984)

47. Ravindra, 1985: M.K. Ravindra, R.P. Kennedy, and R.H. Sues, "Dominant Contributors to Seismic Risk: An Appraisal", Report EPRI-NP-4168, Electric Power Research Institute (1985)

48. Reed, 1985: J.W. Reed, M.W. McCann, J. Ihara, and H. Hadidi-Tamjed, "Analytical Techniques for Performing Probabilistic Seismic Risk Assessment of Nuclear Power Plants", Proceedings of 4th International Conference on Structural Safety and Reliability, Kobe, Japan, Vol. III, p. 253 (1985)

49. Reed and Kennedy, 1994: J.W. Reed and R.P. Kennedy, "Methodology for Developing Seismic Fragilities", Electric Power Research Institute, Report TR-103959 (1994)

50. RMIEP, 1987: J.E. Wells et al., "Seismic Analysis of the LaSalle Unit 2 Nuclear Power Plant -- Risk Methodology Integration Evaluation Program (RMIEP)", Report NUREG/CR-4832, Vol. 8, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (197)

51. Shieh, 1985: L.C. Shieh, J.J. Johnson, J.E. Wells, J.C. Chen, and P.D. Smith, "Simplified Seismic PRA: Procedures and Limitations", Report NUREG/CR-4331, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1985)

52. Spurgin and Moieni, 1991: A.J. Spurgin and P. Moieni, "An Evaluation of Current Human Reliability Assessment Methods", in "Proceedings of International Conference on Probability Safety Assessment and Management, Beverly Hills, 1991", p. 821, Elsevier Publishers, New York (1991)

53. SQUG-GIP, 1992: Winston and Strawn, EQE Inc., MPR Associates Inc., Stevenson & Associates, and URS/Blume, "Generic Implementation Procedure (GIP) for Seismic Verification of Nuclear Plant Equipment, published by Seismic Qualification Utility Group (SQUG), Revision 2 (1992)

54. SSMRP, 1981: P. D. Smith et al., "Seismic Safety Margins Research Program, Phase I Final Report", Report NUREG/CR-2015, in 10 volumes, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1981)

55. SSMRP, 1986: G.E. Cummings, "Summary Report on the Seismic Safety Margins Research Program", Report NUREG/CR-4431, Lawrence Livermore National Laboratory and U.S. Nuclear Regulatory Commission (1986)

56. SSRAP, 1992: Senior Seismic Review and Advisory Panel (R.P Kennedy, W.A. von Riesemann, L.A. Wyllie Jr., A.J. Schiff, and P. Ibanez), "Use of Seismic Experience and Test Data to Show Ruggedness of Equipment in Nuclear Power Plants", Report SAND92-0140, Sandia National Laboratories (1992)