

April 29, 2003

Dr. William D. Travers
Executive Director for Operations
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001

SUBJECT: NUREG-CR-6813, "ISSUES AND RECOMMENDATIONS FOR ADVANCEMENT OF PRA TECHNOLOGY IN RISK-INFORMED DECISION MAKING"

Dear Dr. Travers:

We have undertaken an effort to assess the agency's needs for improved Probabilistic Risk Assessment (PRA) technology to risk inform its regulations. As part of this effort, we commissioned Karl N. Fleming of Technology Insights to prepare the attached report on issues whose resolution would increase the use of risk information in regulatory decisions. This report is based on the author's extensive experience as a practitioner and a participant in the development of the American Society of Mechanical Engineers (ASME) Standard for Probabilistic Risk Assessment for Nuclear Power Plant Applications. In addition, Mr. Fleming conducted interviews with PRA practitioners and decision makers from NRC staff and selected industry representatives. This report has been published as NUREG/CR-6813, "Issues and Recommendations for Advancement of PRA Technology in Risk-Informed Decision Making."

Based on the information gathered during the interviews, his reviews of a number of risk-informed initiatives, and the experience in performing and reviewing PRAs, the author identified a set of recurrent issues that arise in the use of PRAs for risk-informed decision making. Obviously any such list only represents a "snapshot" at a particular time, since many of the issues are being addressed in ongoing activities such as the standards development, the industry peer-review process, the NRC coherence program, and the development of Draft Regulatory Guide DG-1122, "Determining the Technical Adequacy of Probabilistic Risk Assessment Results for Risk-Informed Activities."

The attached report groups the identified issues into the following general categories:

- Use of limited-scope PRAs in risk-informed applications submitted in accordance with Regulatory Guide (RG) 1.174, "An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis," to quantify full-scope metrics
- Lack of completeness within the specified scope
- Model-to-plant fidelity issues
- Lack of, or inadequate, treatment of uncertainties
- Quantification issues (e.g., error due to cut-set truncation)
- Multi-unit site modeling issues
- Lack of treatment of aging effects
- Issues with the use and interpretation of risk metrics
- Lack of coherence between probabilistic and deterministic safety approaches

The interviews conducted by the author identified completeness as the most important area. Issues associated with completeness include the following notable examples:

- Lack of criteria for and consistency in evaluating the impact of missing elements in scope on the application of RG-1.174
- Lack of acknowledgment or consideration of limitations in the PRAs used in submittals
- Inadequate justification and documentation for screening events from a PRA
- Lack of incorporation of operating experience in PRAs
- Inadequate treatment of common-cause failures
- Lack of detailed review by plant personnel to ensure fidelity with plant systems, operator actions, etc.

The author also makes the observation that while valid technical arguments can be made to justify the exclusion of portions of a full-scope PRA model for some applications, resources must be continually expended by both the NRC and its licensees to determine the validity of decisions that are based on an incomplete model. The author further notes that at some point it becomes reasonable to ask whether these burdens are comparable to the effort needed to develop a full-scope PRA.

We believe that this report will serve as a useful resource in the agency's ongoing effort to risk inform its regulations.

Sincerely,

/RA/

Mario V. Bonaca
Chairman

Attachment: NUREG/CR-6813, "Issues and Recommendations for Advancement of PRA Technology in Risk-Informed Decision Making," April 2003.



Issues and Recommendations for Advancement of PRA Technology in Risk-Informed Decision Making



Technology Insights



**Advisory Committee on Reactor Safeguards
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001**



Issues and Recommendations for Advancement of PRA Technology in Risk-Informed Decision Making

Manuscript Completed: January 2003

Date Published: April 2003

Prepared by
K.N. Fleming

Technology Insights
6540 Lusk Blvd., Suite C-102
San Diego, CA 92121

H.P. Nourbakhsh, NRC Project Manager

Prepared for
Advisory Committee on Reactor Safeguards
U.S. Nuclear Regulatory Commission
Washington, DC 2555-0001
NRC Job Code B1564



ABSTRACT

The purpose of this report is to assess the adequacy of PRA for use in regulatory decisions and provide recommendations for its advancements. The insights and recommendations documented in this report were developed by conducting interviews, examining case studies in risk-informed regulation, and by applying experience in developing and applying PRA technology and participating in PRA peer reviews. A number of insights were developed from the review of the recent Davis-Besse vessel head degradation and previous risk-informed and deterministic safety evaluations of the Alloy 600 nozzle cracking issue. Using the author's experience in performing and reviewing several of the existing industry PRAs, a number of technical issues were identified that help defines the current state of the art in PRA technology. The results and conclusions of this report include a number of recommendations intended to resolve some of the issues that were identified and to advance the use of PRAs in risk-informed decision making.

CONTENTS

	Page
ABSTRACT	iii
EXECUTIVE SUMMARY	vii
FOREWORD	ix
ACRONYMS	xi
1 INTRODUCTION	1
1.1 Purpose	1
1.2 Current U.S. Status of PRA Applications	1
1.2.1 Historical Perspective of U.S. PRA Development	1
1.2.2 Efforts to Achieve and Confirm PRA Quality	4
1.3 Approach to this Project	7
1.4 Organization of the Report	7
2. INSIGHTS FROM ACRS AND NRC STAFF INTERVIEWS	9
2.1 Interviewees	9
2.2 Key Results of Interviews	9
3. INSIGHTS FROM SELECTED RISK-INFORMED EVALUATIONS	13
3.1 Davis-Besse Vessel Head Degradation	13
3.2 Callaway Steam Generator Electro-Sleeving	17
3.3 Risk-informed Emergency Diesel Generator AOT Extensions	19
3.4 Risk-informed Inservice Inspection of Piping Systems (RI-ISI)	20
3.5 Less Successful Risk-Informed Applications	20
4. TECHNICAL ISSUES IN PRA FOR RISK-INFORMED DECISION MAKING	23
4.1 Use of Limited Scope PRAs in RG 1.174 Applications	23
4.2 Lack of Completeness within the Specified Scope	32
4.3 Model to Plant Fidelity Issues	34

CONTENTS

	Page
4.4 Treatment of Uncertainties	34
4.5 Quantification Issues	37
4.6 Multi-Unit Site Issues	40
4.7 Lack of Treatment of Aging Issues	40
4.8 Lack of Coherence Between Deterministic and Probabilistic Safety Approaches	41
4.9 Impact of Peer Review Follow-Up and the PRA Standards	42
5. CONCLUSIONS AND RECOMMENDATIONS	43
5.1 Conclusions	43
5.1.1 Success Stories in Risk-Informed Regulation	43
5.1.2 Difficulties Encountered in Selected Risk-Informed Evaluations	43
5.2 Recommendations	46
6. REFERENCES	49

FIGURES

4.1 Uncertainty in Change in CDF for BWR Weld Overlay Example	36
---	----

TABLES

4.1 Summary of Frequent Issues Identified in PRA and RIR Submittal Reviews	24
--	----

EXECUTIVE SUMMARY

The purpose of this report is to assess the adequacy of PRA for use in regulatory decisions and to provide recommendations for its advancements. The report was prepared for the U.S. Nuclear Regulatory Commission (NRC) Advisory Committee on Reactor Safeguards (ACRS) to support the development of a document on such advancements of PRA technology under development by the ACRS. The insights and recommendations documented in this report are those of the author and were developed by conducting interviews with NRC staff and selected industry representatives, examining case studies in risk-informed regulation, and by applying the author's experience in developing and applying PRA technology and participating in PRA peer reviews.

The interviews that were conducted identified both positive and negative aspects of the use of PRA to support risk-informed change requests that have been submitted and reviewed to date. There is a wide consensus that implementation of Regulatory Guide 1.174 and other risk-informed initiatives such as the risk-informed oversight process have been successful in terms of both safety enhancement and burden reduction. Some of the issues that were raised in these interviews include the lack of completeness in scope and level of detail of existing PRAs, inadequate treatment of uncertainties, and difficulties in reaching a consensus on the role of PRA peer reviews and formulation of standards to assure adequate PRA quality.

A number of insights were developed from the review of the recent Davis-Besse vessel head degradation and previous risk-informed and deterministic¹ safety evaluations of the Alloy 600 nozzle cracking issue. The most important of these insights are that epistemic uncertainty was inadequately considered in previous risk-informed and deterministic safety evaluations of the Alloy 600 cracking issue, and that a more complete risk-informed evaluation would need to address several issues that were ignored in the previous evaluations. These issues include consideration of a broader set of scenarios, use of alternative hypotheses about the progression of damage from Alloy 600 cracking, and a more critical evaluation of the capability of visual inspections to provide a backstop for unexpected damage. Insights were also developed from a review of the Callaway Steam Generator Electro-Sleeving submittal including the refutation of the notion that maintaining the status quo under the existing deterministic regulatory requirements is automatically justified.

Using the author's experience in performing and reviewing several of the existing industry PRAs, a number of technical issues were identified including those in the following categories, each of which include several specific issues:

- Use of limited scope PRAs in RG 1.174 applications to quantify full scope metrics
- Lack of completeness within the specified scope
- Model to plant fidelity issues
- Lack of or inadequate uncertainty treatment
- Quantification issues

¹ This term, deterministic, is used to mean "other than probabilistic". All non probabilistic analyses and evaluations that are performed to ensure that a plant's licensing basis are met are classified here as deterministic analyses. It is quiet ironic that the dictionary definition of deterministic, the capability to predict the outcome from knowledge of the antecedent causes, is nearly the opposite of uncertainty. The use of this label to describe traditional safety analyses that were established prior to the advent of PRA suggests a lack of appreciation of uncertainties inherent in any predictive analysis. This issue is developed further in this report.

- Multi-unit site modeling and quantification issues
- Lack of capability to treat aging effects on risk results
- Issues with use and interpretation of risk metrics
- Lack of coherence between probabilistic and deterministic safety approaches

The results and conclusions of this report include a number of recommendations intended to resolve some of the issues that were identified and to advance the use of PRAs in risk-informed decision making. These recommendations include:

- A Proposal for Industry and NRC collaboration to update the PRA Procedures Guide (NUREG/CR-2300). The updated procedure guide should also provide the guidance on how to maintain or to upgrade a PRA in order to meet the requirements of the PRA standards.
- Development of a Handbook on Treatment of Uncertainties in PRA and Risk-Informed Decision Making
- Adoption of consistent definition of risk and fundamental safety questions for deterministic as well as probabilistic safety evaluations
- Development of generic estimates of risk contributors to supplement incomplete PRA scope and to facilitate risk insights about unique plant features
- Development of a program to validate PRA results

FOREWORD

This report was prepared for the NRC Advisory Committee on Reactor Safeguards (ACRS). The information in this report will be considered by the ACRS members in the development of a document on advancements of PRA technology in risk- informed decision making. The views expressed in this report are solely those of the author and do not necessarily represent the views of the ACRS.

ACRONYMS

Acronym	Definition
ACRS	Advisory Committee on Reactor Safeguards
ANS	American Nuclear Society
AOT	Allowed Outage Time
ASEP	Accident Sequence Evaluation Program
ASME	American Society of Mechanical Engineers
ASP	Accident Sequence Precursor
BDD	Binary Decision Diagram
BWR	Boiling Water Reactor
BWROG	BWR Owners' Group
CCF	Common Cause Failure
CDF	Core Damage Frequency
CRDM	Control Rod Drive Mechanism
DBA	Design Basis Accident
ECCS	Emergency Core Cooling System
EDG	Emergency Diesel Generator
EPRI	Electric Power Research Institute
HRA	Human Reliability Analysis
IEEE	Institute of Electrical and Electronic Engineering
ISLOCA	Interfacing Systems LOCA
ICCDP	Incremental Conditional Core Damage Probability
IGSCC	Intergranular Stress Corrosion Cracking
IPE	Individual Plant Examination
IPEEE	Individual Plant Examination for External Events
IREP	Interim Reliability Evaluation Program
LERF	Large Early Release Frequency
LOCA	Loss-of-Coolant Accident
LWR	Light Water Reactor
MGL	Multiple Greek Letter
NEI	Nuclear Energy Institute
NRC	Nuclear Regulatory Commission
NREP	National Reliability Evaluation Program
PORV	Power Operated Relief Valve
PRA	Probabilistic Risk Assessment
PWR	Pressurized Water Reactor
PWSCC	Primary Water Stress Corrosion Cracking
RAI	Request for Additional Information
RAW	Risk Achievement Worth
RCP	Reactor Coolant Pump
RG	Regulatory Guide
RI-ISI	Risk-informed Inservice Inspection
RI-IST	Risk-informed Inservice Testing
RI-TS	Risk-informed Technical specification
RSS	Reactor Safety Study

Acronym**Definition**

ROP	Reactor Oversight Process
RVLIS	Reactor Vessel Level Indicator System
SALP	Systematic Assessment of Licensee Performance
SDP	Significance Determination Process
SG	Steam Generator
SSC	System, Structure, and Component
STPNOC	South Texas Project Nuclear Operating Company
TF	Thermal Fatigue

1. INTRODUCTION

1.1 Purpose

The purpose of this report is to assess the adequacy of PRA for use in regulatory decisions and provide recommendations for its advancements. The report was prepared for the NRC Advisory Committee on Reactor Safeguards (ACRS) in support of an effort to develop a document on advancements of PRA technology to support risk-informed decision making. The insights and recommendations documented in this report were developed by conducting interviews with NRC staff and selected industry representatives, examining case studies in risk-informed regulation, and by applying the author's experience in developing and applying PRA technology and participating in PRA peer reviews. This report discusses: the enhancements that need to be incorporated in PRAs to make them more complete; the current elements of PRA that need to be improved; how uncertainty (both aleatory and epistemic) should be developed in the PRA and used in the decision making process, and some insights on how deterministic safety evaluations could be improved using risk insights. The report also includes suggestions on criteria to be used by the staff in making decisions regarding the need for risk analysis to support a regulatory decision.

1.2 Current U.S. Status of PRA Applications

1.2.1 Historical Perspective of U.S. PRA Development

Reactor Safety Study (1975)

The genesis of PRA as it has been applied to Light Water Reactors (LWRs) is the Reactor Safety Study (RSS) that was completed in 1975 [1]. Many of the methods and risk insights that we still use today were introduced in that landmark study. The accident sequences that were found to make the most important contributions to the risk of a severe accident were found not to be correlated to the design basis accidents. Hence, small break loss of coolant accidents were found to be much more significant than those initiated by large pipe breaks. Moreover, the concept that conservative safety analyses of design basis accidents can establish an upper bound on the risk to public health and safety was refuted. In fact, the risks calculated in the RSS were found to be completely determined by severe core damage accidents that exceed the design basis envelope. Although it was obvious that the consequences of a severe core damage event would exceed those of a design basis event, a key insight here was that the frequency of severe core damage events was much higher than expected using traditional defense-in-depth thinking. Such thinking suggested that accidents require the postulation of failure of several independent systems and fission product barriers following an initiating event. Such thinking supported the qualitative judgments made in the deterministic framework of safety analysis that accidents more severe than the design basis accidents were so unlikely as to be negligible in the definition of the design basis and associated general design criteria. With the benefits of insights from application of PRA technology as well as lessons learned from various incidents and accidents such as the Browns Ferry Fire and the Three Mile Island accident, it is now clear to most PRA practitioners that severe core damage events are in fact much more likely than the design basis accidents. If one were to go back and assess the frequency of the DBA initiating events and realistic probabilities of the specific success and failure combinations that are assumed in the safety analysis reports, it is clear that frequencies much lower than currently calculated core damage frequencies would result.

Many forget the reason why the RSS was performed: to support congressional debate on the renewal of the Price Anderson Act which limits liability from reactor accidents. Two decades would pass until any serious effort to risk-inform the regulatory process was made. Both pro and anti nuclear advocates used information presented in WASH-1400 to support their respective arguments that nuclear power was safe and unsafe, respectively. This political debate fueled much controversy about the usefulness of PRA in the decision making process that the benefits of increased use of PRA could not be fully appreciated. The NRC asked Hal Lewis to lead a Committee to review and clarify the achievements and limitations of the study [2]. That led to an NRC policy statement on the use of PRA [3] that placed so much emphasis on the uncertainties in PRA results. This policy statement pretty much put PRA out of business as a tool to resolve safety and licensing issues in the regulatory process.

Pre-IPE Era (1975 to 1992)

The current set of industry PRAs was developed over three phases: In the pre-IPE phase, which began shortly after the NRC completed its landmark Reactor Safety Study in the mid 1970's, there were a limited number of plant specific PRAs to address special issues such as challenges to NRC decisions to license several plants near large population centers (Zion and Indian Point, Limerick), emergency planning issues (Seabrook, Shoreham), a seismic design and siting controversy (Diablo Canyon), and the issue of whether to restart TMI-1 after the accident on Unit 2. During this period several plants unilaterally decided to supplement the required deterministic safety evaluations by performing plant specific PRAs to develop insights in support of the design and safety improvements of the facility (e.g., South Texas, Oyster Creek, Susquehanna, Midland, Millstone).

Northeast Utilities was the first known licensee to use its plant PRAs for the Millstone Units and Connecticut Yankee during this period to implement a policy on risk management. This policy included the identification and resolution of plant specific vulnerabilities, the use of risk insights to allocate resources for plant improvements, and the use of risk information to prioritize and support safety-related activities at all units. This policy also included self-imposed limits on continued plant operation based on an evaluation of core damage frequency that is not unlike the current NRC Significance Determination Process (SDP). The PRA manager was a key player on the management team for this licensee's nuclear facilities.

It was during this phase of PRA development that the treatment of external events such as fires and seismic events were introduced into PRA and somewhat later accident sequences initiated during low power and shutdown modes and internal floods were incorporated into the PRA technology palette. Improved methods for the treatment of common cause failure and human reliability analysis were also developed during this period.

A comprehensive description of the state-of-the-art PRA technology as it was developed in the middle of this phase of the PRA history is found in the PRA Procedures Guide [4]. A unique quality of this work was a broad participation of PRA experts from industry and the NRC and extensive international peer review. Somewhat later in this period, the NRC performed a major update of their version of full scope PRAs on five plants covering a representative set of reactor and containment types from the entire industry [5]. These PRAs included a full treatment of internal and external events and no doubt went further to incorporate both epistemic and aleatory uncertainties than PRAs had done previously. This uncertainty analysis was supported by an extensive research program and associated substantial budget. Subsequently, two of the plant PRAs covered in this work were expanded to include accident sequences from low power and shutdown modes. To this day, the NUREG-1150 PRAs represents the NRC's most outstanding contribution to the state-of-the-art of PRA.

Post IPE Era (1992 to 1995)

Efforts to conduct PRAs on most of the remaining plants did not occur until the NRC requirement to perform the Individual Plant Examinations (IPEs) and Individual Plant Examinations for External Events (IPEEEs) which occurred in the late 1980's and were completed by the mid 1990's. Prior to the issuance of Generic Letter 88-20, which requested a systematic examination for severe accident vulnerabilities, the NRC tried unsuccessfully to initiate requirements for having each of the plants perform a full scope plant specific PRA that exceeded the IPE requirements, but were unable to define a legal basis because PRAs were not needed to meet existing licensing requirements. While essentially all the IPEs and many of the IPEEEs were submitted using the form of a PRA, most were greatly simplified in relation to what was then regarded as a state-of-the-art PRA. Hence, the standard for acceptance was limited to a perceived capability to identify severe accident vulnerabilities. The scope, level of detail, and level of completeness of the submittals varied widely.

After the IPEs and IPEEEs were completed and approved, there were two camps in the industry including a smaller proactive camp that continued to support the maintenance, upgrades, and updates of more complete PRA models and a much larger camp that had developed more simplified PRAs for the IPEs. This larger camp was somewhat reluctant to make investments into upgrades until a return on that investment could be more easily visualized. Uncertainties about the success of efforts to risk-inform the regulatory process tended to maintain the relative size of these camps. The larger camp may have been intimidated by the magnitude of the investments made by the proactive camp in building and maintaining a PRA capability, or may have overestimated how large these investments would be. In assessing the cost of these enhancements, the larger group may not have appreciated the efficiencies that were realized in reducing the costs of performing PRAs as a result of software and methods enhancements and increased competition among PRA consultants. The burden of PRA methods development was greatly reduced during this era so that PRA project completion costs were more predictable in comparison with prior eras.

Risk-informed Regulation Era (1995 to present)

The current era began when the NRC issued a revised policy statement in 1995 on the use of PRA in decision making that included a more positive view on the role that risk information should have in supporting regulatory decisions [6]. The key events that had a significant impact on utility decisions to make significant investments in their PRAs during this period were the issuance of RG 1.174 [7] and associated application specific regulatory guides and standard review plans. For the first time in these regulatory guides, the NRC provided clear criteria for the review of risk-informed changes to the licensing basis including quantitative risk acceptance guidelines for judging whether a calculated change in core damage frequency (CDF) or large early release frequency (LERF) would be considered large enough to impact the decision. Prior to these guides, industry attempts to present risk based or risk-informed arguments to get relief from a regulation were very difficult to develop and for the NRC staff to review, due to lack of criteria for judging "how safe is safe enough," and lack of standards for PRA quality. These factors yielded long and costly Requests for Additional Information (RAIs), RAI responses, and associated staff reviews and a high level of inconsistency in acceptance criteria.

1.2.2 Efforts to Achieve and Confirm PRA Quality

During each of the previous eras, the technology of PRA developed and matured. There were several attempts to write down PRA procedures by different bodies and eventually the technology developed to the point where standards for PRA were developed. Some of the most important of these efforts are summarized here.

PRA Procedures Guide (NUREG/CR-2300)

Following the completion of the Reactor Safety Study, the NRC was responsible for developing a number of useful guidance documents for performing PRAs, such as the IREP and NREP PRA Procedures Guides issued by NRC contractors. The first significant effort to integrate both industry and NRC contributions to PRA technology development was the effort to complete the PRA Procedures Guide published in NUREG/CR-2300 in 1982. There was quite broad participation by industry, NRC and NRC contractors in preparation of this guide and substantial resources were invested by the NRC, ANS, and IEEE to sponsor the authors who wrote this guide. The participants were initially requested to develop a PRA standard. However, the PRA practitioners who were tasked to write this guide developed a consensus that it was premature to attempt standardizing the PRA methodology. Instead a compendium of methods was described that provided an excellent snapshot of the state-of-the-art of PRA in that time frame. This guide invented the concept of dividing the PRA into the Level 1/Level 2/Level 3 framework, included methods for the treatment of internal and external events such as seismic events, fires and floods, but predated the time frame when PRAs were expanded to consider accidents initiated at shutdown. This guidebook was very useful to support PRAs that were performed in the 1980 and 1990's and in fact much of it is still relevant today.

South Texas PRA and Risk-informed Initiatives

One of the major success stories in risk-informed regulation was the recent NRC approval of the request by South Texas Project Nuclear Operating Company (STPNOC) to gain exemption from special treatment requirements in 10 CFR Part 50 for safety related components that were shown to be of low risk significance. A foundation for this success was the high level of confidence obtained among the NRC staff in the quality of the STPNOC PRA which provided the risk insights for this exemption request as well as many previous risk-informed initiatives. STPNOC was the first plant to provide a risk-informed basis for most of the Allowed Outage Times (AOTs) and Surveillance Test Intervals in the Technical Specifications. In its July 23, 2001 letter to Chairman Meserve, regarding the exemption for Special Treatment, the ACRS stated:

“ STPNOC has developed a state-of-the-art PRA in which the licensee, the regulators, and the public can have confidence. The staff engaged an independent contractor to perform a review of the STPNOC PRA and their report indicates that the STPNOC PRA is of good quality.”

STPNOC had unilaterally decided to perform a PRA as early as 1983 for the initial purpose of getting a second opinion on the adequacy of the design to support the plant licensing prior to start of commercial operation. The plant owners took a proactive approach to risk-informed regulation by notifying the NRC staff of their intent to use the PRA to justify the development of risk-informed technical specifications appropriate for a unique level of safety system redundancy among U.S. plants. The plant owners paid for an independent technical review by Sandia National Laboratories as well as for significant upgrades to the original internal fire analysis to address Sandia review comments. This review was much more detailed

than the staff and contractor reviews that were performed for the IPEs and IPEEEs due to the limited purpose of these examinations. While the PRA that was reviewed by Sandia was limited to Level 1 treatment of internal and external events, including seismic events, the PRA was subsequently expanded to include a full Level 2 treatment of all events in full power, low power, and shutdown modes. Hence, it is one of the few, if not the only plant PRA that can address RG 1.174 risk acceptance guidelines in a fully quantitative manner. The STPNOC PRA includes a state-of-the-art treatment of both aleatory and selected epistemic uncertainties. There is no question that the full scope nature of the assessment and the exposure of the PRA to a detailed peer review supported the level of confidence credited in the ACRS statement.

Industry PRA Peer Reviews and PRA Standards

As noted in Section 1, there were a number of other plants which at one time or another performed very detailed state-of-the-art PRAs for various mostly special reasons. But most plants did not perform a PRA of any kind until the IPE and IPEEE requirements were issued. In addition to that many of these plants did not maintain the PRAs once the IPE/IPEEE requirements were met. As the future of risk-informed regulation was still being formulated, there was a perception shared by many that the investments made in the case of South Texas were prohibitively expensive and that such investments would have to be made before any other plant could achieve a level of benefit that would offset the investment. A related concern was that benefits of risk-informed regulation at STPNOC were thought by some to be uniquely large due to the three trains design and the decisions made in the licensing of that plant to use of a conservative classification of safety-related components relative to most other plants. Finally, the movement toward a deregulated utility market made it much more difficult to justify investments in operating plants without near-term payback in comparison with plants under construction. These concerns were probably exaggerated as the investments made by STPNOC were not as great as perceived by many in the industry. The STPNOC PRA was phased over a number of years to ensure maximum utility participation, technology transfer, and utility ownership. Nonetheless, other licensees were reluctant to follow the success path to achieving high confidence in their PRAs that was blazed by South Texas.

Concerns about the cost-benefit justification for upgrading PRAs were part of the motivation for several industry initiatives that suggested a more incremental approach to investing in PRAs to support risk-informed regulation. These initiatives included the NEI PRA peer review program that was initiated by the Boiling Water Reactor Owners' Group (BWROG) and subsequently adopted by the other three owners groups [8], as well as the EPRI PSA Applications Guide [9]. A key philosophy of these initiatives was to define the uses and applications that existing PRAs could support, identify the strengths and weaknesses of the existing PRAs, and then to define an optimum path to making selective upgrades and applications so that the benefits from the investments could be realized incrementally. While there was a sense of respect for what STPNOC had achieved in its PRA program, the rest of the industry pretty much decided to get to the same point along a path of incremental investments and hopefully benefits. It is doubtful whether there would have been nearly as many submittals in response to RG 1.174 if the utilities had been required to develop South Texas Level PRA programs prior to the risk-informed submittals. On the other hand, no plant that adopted the incremental approach is likely to achieve the benefits that South Texas has already realized by following the more proactive paths to success.

The industry PRA peer review process has had a major impact in upgrading the quality and consistency of the existing industry PRAs relative to the IPE/IPEEE submittals. This stems, in part, from a standardized process and checklist for reviewing the PRAs, participation of both PRA consultants and PRA practitioners from plants in the same peer group, and a documented set of strengths and weaknesses for each plant using a peer group consensus process. Of less importance to its success is the use of a numerical grading system for assessing the capability of each PRA element and its sub-elements despite the fact that this

aspect has been the focus of much of the discussion about the merits of this program. Some of the areas where improvements to PRA technology have been made as a result of these peer reviews include:

- More consistency in the treatment of generic issues such as reactor coolant pump (RCP) seal LOCA impact on the time to uncover the core in PWRs
- More consistent and defensible treatment of electric power recovery
- More consistent treatment of initiating event frequencies for transients and loss of coolant accidents (LOCAs) and other generic data parameters
- Exchange of information among plants on PRA best practices and a forum to transcend close ties between specific groups of utilities and specific consultants with differing views on PRA methodology

While the NRC staff is reviewing this process and addressing its role in future risk-informed regulation applications, as evidenced in draft Regulatory Guide DG-1122, in the opinion of the author, the staff cannot fully appreciate the PRA improvements resulting from these peer reviews for several reasons. These include: minimal NRC participation on these reviews, limited publication of peer review results in NRC submittals and other public domain forums, and the fact that the NRC staff generally does not have access to entire PRAs to assist in their review of specific RG 1.174 submittals which only present the results of certain calculations that are performed to support the submittals. The sheer volume of submittals to review and finite resources available to support the reviews would preclude a full PRA review by the staff in each instance even if they had access to the completed PRAs. If such reviews were carried out, the staff would have a greater appreciation of the upgrades that have been made since the IPEs and IPEEEs were submitted.

Despite the PRA improvements attributed to the peer review process, there are some very significant issues and deficiencies being identified in these peer reviews and some of these are discussed in Section 4. The effort to address these issues is still a work in process and varies from plant to plant depending on the risk-informed applications being pursued and the RAIs received from the NRC staff during reviews. Most of these issues were identified through the author's participation in the performance and reviews of industry PRAs and others were identified during this project as discussed in Section 1.3 below.

The other significant initiative to address PRA quality is the development of industry standards for PRA. It remains to be seen what impact these standards will have as they have not yet been utilized in a significant way. It will be necessary to have some case studies in risk-informed regulation where these standards and the associated utility self assessments are used to establish the appropriate role of the standards. There are indeed gaps between the criteria used in the peer reviews and the requirements listed in the current versions of the standards. Hence, PRAs that are upgraded to address these gaps should result in further enhancements to the scope and quality of the existing PRAs.

Unfortunately, the consensus process that the industry has developed to write standards takes a long time from initiation to approval by the American National Standards Institute (ANSI). The American Society of Mechanical Engineers (ASME) PRA standard was initiated in 1997 and was only approved by ANSI in 2002. The American Nuclear Society (ANS) standard on external events should be ready in 2003 but the ANS standards for low power and shutdown and fire PRAs are lagging behind in the schedule. The ASME PRA Standard was limited in scope to internal events including floods but excluding fires, full power initial conditions, and limited treatment of Level 2 that is sufficient for estimating LERF. The level of coordination

among these standards writing efforts is lacking and it is unclear how PRA practitioners will be able to perform a full scope PRA with different standards written by different organizations. This lack of coordination is responsible for inconsistencies and gaps. The inconsistencies include different approaches to inclusion of tutorial guidance (less in ASME relative to ANS) and lack of balance in level of detail. For example, there is one small section of a few pages with requirements for internal flooding in the ASME standard and entire separate standards for fires and external events. The gaps include a lack of Level 2 treatment in the ANS standards for low power and shutdown and external events, external events requirements that do not address shutdown modes, and vice versa. There is not currently any effort to develop standards for Level 2 PRAs beyond determination of LERF, nor any standard for extending the PRA to Level 3. In the future, there needs to be an umbrella standard that provides an integrated perspective for performing a full scope PRA and fills in the missing gaps as well as better guidance for use of partial scope PRAs to evaluate risk metrics intended all risk contributors. In addition to the standards, the industry is developing a self assessment process to address the gaps between the ASME PRA standard and the NEI peer review process. There may be a long way to go before the utilities and the NRC will have a single set of well-integrated documents to address PRA quality issues. The level of resources that will be needed to achieve the consensus on PRA quality may in fact exceed that of having each plant follow the South Texas road map.

1.3 Approach to this Project

The approach that was followed to develop technical insights for this report was to conduct interviews with NRC staff and selected industry representatives, review selected case studies in risk-informed regulation including some that created some of the technical issues discussed in this paper, and incorporate insights from PRAs that were performed and reviewed by the author. From these sources a set of technical issues was developed that provide the basis for recommendations to the ACRS to consider in the formulation of their vision for the future direction of PRA technology in risk-informed decision making.

1.4 Organization of the Report

The results of the interviews that were performed with NRC staff and selected industry representatives are presented in Section 2. In Section 3, a number of case studies in risk-informed regulations are reviewed, including several that were identified in the interviews to develop further insights. A brief review of ACRS reports from 1997 through 2001 that addressed PRA issues was performed to supplement the insights from the interviews. A number of technical issues that were identified from the interviews and the author's PRA and review experience are identified and discussed in Section 4. The recommendations on the future direction of PRA in risk-informed decision making to address these issues are presented in Section 5.

2. INSIGHTS FROM ACRS AND NRC STAFF INTERVIEWS

2.1 Interviewees

A total of about 20 interviews were conducted with key members of the NRC staff and selected industry representatives including the Nuclear Energy Institute staff. The NRC staff members who participated in the interviews included most of senior management and PRA staff from both the Office of Nuclear Regulatory Research (RES) and the Office of Nuclear Reactor Regulation (NRR). ACRS input was obtained during two ACRS meetings and separate discussions with some ACRS members. The interviews ranged in duration from about 30 minutes to 90 minutes. Some were conducted in person and others via phone. During the interviews, the following questions were addressed in each of the interviews, although the time spent on each question varied according to the interests and background of the participants:

- What are the success stories in risk-informed regulation why were they successful, and what were the roles and uses of the PRAs in contributing to the success?
- What are examples of risk-informed regulation that were not successful and why?
- What examples can be cited where there were inadequate or inappropriate uses of PRA in risk-informed submittals?
- What examples can be cited where risk-informed evaluations yielded insights about the limitations of traditional deterministic approaches to safety assessments?
- What are the most important strengths of PRA technology and what does it bring to the decision making process that was previously missing?
- What are the limitations of PRA technology that have the most impact on staff reviews?
- What is each participant's vision for the advancement of the use of risk insights in decision making?

2.2 Key Results of Interviews

In response to questions about where risk-informed regulation is working, success stories and the expectations for future successes, each interview contributed some positive input on the success of risk-informed regulation, including the following typical responses.

- There was general support for the movement to risk inform the licensing process and high expectations for increasing use of risk-informed approaches in the licensing process by NRC senior management.
- There have been several hundred successful submittals and approvals for risk-informed changes to licensing basis per RG 1.174 and strong consensus that the development and use of this regulatory guide was a major success story. The majority of these submittals were to request changes to technical specifications and to ASME Section XI in-service inspection requirements for piping systems.

- At least some if not most risk-informed applications (e.g. RI-ISI and configuration and operational risk management) are resulting in real safety improvements.
- NRC staff acceptance (albeit highly qualified) of NEI PRA peer review process and ASME PRA standard (via DG-1122)
- The new risk-informed Regulatory Oversight Process (ROP) is viewed as major improvement over the traditional Systematic Assessment of Licensee Performance (SALP) approach.

In response to the questions that were geared to extract problem areas with the use of PRAs in decision making, the following issues were identified:

- Perceived reluctance on the part of the industry to make investments to improve the scope and quality of industry PRAs that will be needed to continue the advancement of risk-informed regulation
- Lack of the consensus between the industry and the NRC staff that the quality, scope, and level of detail of different parts of a PRA model or for PRAs used in different applications can vary from application to application
- Insight that staff's segregation of IPEs and IPEEEs and reluctance to integrate different components of PRA in previous NRC-Sponsored PRA projects may have contributed to industry reluctance to perform full scope PRAs (all modes, all events, Level 3)
- Lack of or inadequate treatment of aleatory and epistemic uncertainties in PRAs; lack of appreciation for limitations of PRA in submittals
- Lack of treatment of aging effects and passive component reliability issues in PRAs (e.g., Davis-Besse)
- Uncertainties due to events not considered or screened out of PRA models (e.g., LOCAs via component failures other than pipes)
- Continued reluctance by some NRC staff to accept PRA results in decision making
- Licensees not doing enough to defend the quality of existing PRAs; staff lack of appreciation for PRA upgrades that have been made since the IPE and IPEEE submittals
- Need to address risk aspects of non-risk-informed license change requests when based solely on meeting current licensing requirements (Callaway electro-sleeving of SG tubes)
- Lack of incorporation of insights from the Accident Sequence Precursor (ASP) program (One interviewee noted that as much as 20% of events evaluated in ASP program points to initiating events and accident sequences not modeled in existing PRAs)
- Lack of clear vision for future of risk-informed regulation from senior NRC staff
- Lack of consistency in applying RG 1.174 due to lack of prescriptive guidance
- Need for more consistency and better guidance and training for staff uses of PRA

- Need for better approaches for incorporating uncertainty into cost-benefit and multi-attribute decision making
- Problems associated with options for voluntary risk-informed initiatives via RG 1.174 vs. continued compliance with deterministically derived requirements, and an involuntary risk-based oversight process. This has led to confusion about whether mere compliance with regulations is adequate without consideration of risk aspects of a plant change (e.g., Callaway electro-sleeving).

In addition to the above key points, the interviews cited a number of specific examples where risk-informed regulation was particularly successful as well as others that highlighted issues which need to be resolved to advance the cause of risk-informed regulatory reform. Risk-informed in-service inspection of piping systems was often cited as a major success story as the majority of licensees now have an approved risk-informed inspection program to replace their ASME Section XI programs for non destructive examinations of welds in Class 1 and 2 piping systems. There has been a large number of risk-informed changes to plant technical specifications that successfully applied the applicable regulatory guides. The two examples most often cited as pointing to problems and limitations of risk-informed regulation were the request to delay the requested inspection for Alloy 600 cracking in control rod drive penetrations at Davis-Besse and the request to incorporate a new method of repairing damaged steam generator tubes at Callaway.

From the interviews and the above listed insights from these interviews, a list of risk-informed submittals and staff evaluation reports that highlighted both successful and unsuccessful examples of risk-informed regulation was developed. The review of these documents provided another input to the delineation of issues and recommendations for this project. A summary of the results of these document reviews is provided in Section 3.

3. INSIGHTS FROM SELECTED RISK-INFORMED EVALUATIONS

In the NRC staff interviews, a number of specific examples were cited that exposed technical issues with the risk-informed process as well as some of the more significant success stories in application of RG 1.174. Several of these examples are discussed in the sections below.

3.1 Davis-Besse Vessel Head Degradation [10]

A number of interviewees cited various aspects of the Davis-Besse vessel head degradation event as challenging the wisdom and exposing the limitations of risk-informed regulation. Others interviewed saw little if any relevance of this event to risk-informed regulation. One aspect of the event that caused concern was the licensee's risk-informed request to delay the timing of the inspection to investigate possible Alloy 600 cracking in the Control Rod Drive Mechanism (CRDM) nozzles for several months and the staff's partial approval of that request (the staff permitted a delay roughly half of that requested).

NRC had issued IE Bulletin 2001-1 requesting that certain plants conduct inspections of control rod drive nozzles after finding cracks in these nozzles at Oconee earlier that year. Davis-Besse requested that these inspections be delayed approximately three months until the next scheduled refueling outage in March 2002. The supplemental information that was provided to support this request [11] included a risk-informed evaluation to estimate the potential risk impacts of delaying the exams as well as a number of engineering calculations that supported the request. The PRA evaluation included an estimate of the probability of a medium size LOCA that was within the design capabilities of the emergency core cooling system which was represented as a scenario that would set an upper bound on the consequences of potential accidents that could occur during the period in which the exam was delayed. The likelihood of boric acid induced wastage of the reactor vessel head carbon steel from leakage of reactor coolant through possible stress corrosion cracking was qualitatively dismissed in this evaluation just as this mechanism had been dismissed in previous NRC and industry examinations of Alloy 600 cracking in French and U.S PWRs dating back to the 1980's. The arguments used to dismiss this wastage scenario are viewed by the author as classic modeling assumptions of the type often used in PRA and various deterministic safety evaluations. Hence, this is an example of an epistemic uncertainty although it was not treated as such in the supporting PRA. In the PRA, it was simply an unverified assumption supporting the calculated risk with a level of certainty of 100%. As it has been extensively reported since then [10], the subsequent inspection at Davis-Besse surprisingly discovered that significant wastage of the vessel head material had in fact already occurred and had been in progress for at least several years prior to the risk-informed transaction that resulted in a small delay of the inspection.

A second risk-informed aspect of this incident is NRC's risk characterization of the event as part of the risk-informed Significance Determination Process (SDP) which had not been completed at the time of this writing. The risk characterization of the issue for the inspection delay request was defined in terms of a potentially small increase in the likelihood of a small to medium LOCA whose size would be limited by axial cracking and would be bounded by the dimensions of a single CRDM nozzle which would be well within the mitigation capability of the ECCS. Therefore, the increased risk of a core damage event over such a short period was calculated to be small. This risk characterization was based in part on several modeling assumptions that were not introduced in the PRA but rather were developed as part of industry and NRC deterministic evaluations of a stress corrosion cracking damage mechanism. This mechanism was responsible for a number of cracks and leaks in a number of PWR pressure boundary components dating back to the 1980's. These components were Alloy 600 nozzles associated with pressurizer instrument lines, hot leg piping connections to the reactor vessel, and CRDM nozzles in the vessel head. The risk

characterization that was attempted did not identify the potential for a larger opening in the pressure boundary due to extensive external corrosion of the vessel head material as was surprisingly identified during the Davis-Besse inspection. Indeed, all involved in evaluating Alloy 600 cracking issues seemed to be very surprised to learn about this corrosion.

The risk characterization of the CRDM nozzle cracking issue as being confined to a small increase in the likelihood of an axial crack in the CRDM tube is not unlike the characterization used by both the industry and staff in addressing the Alloy 600 cracking issue when it first surfaced in service experience data in the mid 1980's. The first observed crack in a CRDM nozzle occurred at Bugey in France back in 1991. Some five years before that, leaks in Alloy 600 nozzles for instrument lines on pressurizers were found to be caused by the same, primary water stress corrosion cracking (PWSCC), damage mechanism. In fact, during the industry and NRC staff management of this PWSCC issue, the potential for extensive external vessel head wastage due to corrosion from boric acid leakage from the reactor coolant system had been considered but was dismissed based on several unverified assumptions, that appear to have been made based on "deterministic" evaluations of this damage mechanism and extrapolations of limited service experience. As noted in NRC's Lessons Learned report on this issue, the French safety authorities were sufficiently concerned about the potential for significant external corrosion resulting from CRDM nozzle cracking that they required more extensive head inspections at every refueling outage. Interestingly, the French lacked confidence in the same type of modeling assumptions that were used to dismiss the wastage scenario back here in the U.S. Eventually the time and costs of these inspections in France led to a decision to replace the vessel heads in their whole fleet of reactors.

The U.S. "deterministic" evaluations of CRDM nozzle cracking had applied what a PRA analyst would describe as "modeling assumptions" about the behavior of boric acid leaking from the system and the capability of visual examinations to detect leakage before safety margins would be compromised. These modeling assumptions included the use of crack propagation models used to determine the growth rate of axial cracking, for ruling out the possibility for circumferential cracking, and the time available for inspections. The deterministic evaluation included an additional assumption that any leaking boric acid would immediately flash to steam leaving dry boric acid crystals on the vessel head surface. In comparison with boric acid in the liquid form, such crystals were known to have very slow corrosion rates supporting the further assumption that visual inspections would be a sufficient backstop to prevent significant corrosion. After having only considered the slow corrosion rates supported by the assumed axial cracking mode and the acid crystallization assumption, the time to significant degradation was estimated to be several years which was judged sufficient to implement effective inspections and an adequate justification to delay volumetric examination of the suspected cracks. This information was used to support an argument that visual external surface exams would be sufficient to identify significant degradation long before safety margins were compromised. An additional "modeling" assumption that was made was that any leaks could be effectively identified via visual exams despite the facts that significant quantities of boric acid crystals might mask the condition of the vessel head surface and that such inspection could only be performed with the reactor vessel in a cold and depressurized state.

As noted in NRC's Lessons Learned Report, the French were not willing to adopt several of these modeling assumptions in their more conservative treatment of this issue. In particular, the French were very concerned about uncertainties associated with inputs to the crack propagation models used to support the U.S. evaluations as several of the inputs to these models could not be estimated with sufficient accuracy, including details of the stress fields and inside surface temperatures of the vessel head. Hence, a large part of the explanation for the different approaches to treating the Alloy 600 cracking issue in France and the U.S. can be explained in the different approaches employed in the application of modeling assumptions, or what has been termed as "epistemic" uncertainties.

The author's view of the several implications of this event on risk-informed regulation are as follows:

- The uses of modeling assumptions and expert opinion in the deterministic safety evaluations of the Alloy 600 cracking issue prior to the Davis-Besse event are analogous to PRA modeling assumptions that are the source of epistemic uncertainties. With the benefit of hindsight several of these modeling assumptions have resulted in an optimistic evaluation of the safety significance of Alloy 600 corrosion cracking. The risk-informed evaluation was remiss in not including at least a sensitivity study to examine the impact of alternative modeling assumptions on the behavior of underlying damage mechanism. The structural reliability engineers only offered one set of models to predict the behavior of cracks, boric acid crystals, and the ability of visual exams to pick up damage before it progressed significantly. The PRA evaluation was completely dependent on the validity of these models. This is a striking example of how epistemic uncertainties are not only available to challenge the results of a PRA, but also the validity of the conclusions derived from so-called deterministic safety evaluations.
- A PRA is only as good as the deterministic knowledge that is available to support the assumptions in the model. This deterministic knowledge includes the results of relevant deterministic safety evaluations that provide the technical foundation for the PRA. Regulators who are concerned about the uncertainties inherent in PRA results need to be just as concerned about uncertainties in all the deterministic safety evaluations that are performed to demonstrate that deterministic requirements are met and provide key inputs to the PRA. Unfortunately, deterministic safety evaluations do not seem to be held to the same level of accountability to address uncertainties as is the case with PRA.
- The deterministic safety evaluations of the Alloy 600 cracking issue prior to the discovery of damage at Davis-Besse had to make assessments similar to the questions addressed in a PRA. These questions include the need to identify the relevant accident sequences, and to make at least qualitative judgments about the frequency of the sequences (credible vs. incredible) and the sequence consequences. These common aspects of deterministic and probabilistic safety evaluations need to be better appreciated before they can become better integrated in the risk-informed regulatory process. Coherence requires that both approaches adopt a common definition of risk.
- The risk-informed argument to delay the vessel head inspection by a few months at Davis-Besse incorporated the same naive modeling assumptions that were employed in the previous U.S. deterministic evaluations of this issue. If the risk evaluation were updated to incorporate what we know today about the potential for external head corrosion, the risk impact of delaying the inspection would have increased a lot over what had been predicted, but the conclusion would likely have been similar, namely, that the risk impact of waiting a few months or weeks is small, albeit highly uncertain. It appears that external corrosion of the vessel head was going on for at least four years and an additional few months or so may not have been significant, especially since the vessel head still appeared to be capable of maintaining the pressure boundary, albeit with reduced margins, at the time of the discovery. It seems as though this event has much more serious implications on the deterministic approach to safety assessment as risk-informed decision making has played only a minor role in this instance. The validity of all previous evaluations that depended on the capability to predict the consequences of stress corrosion cracking of Alloy 600 nozzles is now open to question.
- The lessons learned from the Davis Besse head degradation should be incorporated into risk-informed ISI evaluations because such evaluations may not adequately consider the possibility that damage mechanisms creating cracks in the reactor coolant pressure boundary can lead to external wastage of ferritic components on the pressure boundary.

- In making a risk characterization of such an event, there are a number of specific factors that need to be taken into account. First it should be understood that the probability that the Davis-Besse event progresses to core damage is zero, because the head degradation was in fact discovered prior to any accident and the risk exposure from this specific incident is now terminated as the vessel head is being replaced with a new one. In making a risk characterization, it is assumed that some other plant experiences a similar damage mechanism or that this damage mechanism happens in the future at Davis-Besse and one is trying to predict the probability that the degradation leads to an accident prior to it being discovered and repaired. Some of the questions that such a risk characterization should consider are as follows:
 - ▶ What are the possible initiating events and event sequences that should be considered in the evaluation? Although the scenario involving cracks leading to wastage was initially considered back in the 1980's both in France and the U.S., this scenario was dismissed which is the PRA equivalent of assuming that the frequency of the scenario is zero – an unattainable state in the PRA world. The idea that a medium LOCA with no degradation of mitigation via the ECCS would set a conservative bound on the consequences of cracking was an assumption that we now know is just wrong. The capability to predict the progression of this Alloy 600 damage mechanism with existing fracture mechanics technology was grossly overstated.
 - ▶ How long does the degradation proceed until it is discovered? There were several events that keyed the timing of discovery at Davis-Besse. One is the timing of the discovery of CRDM cracking at Oconee in 2001 that led to the order to inspect Davis-Besse and the other plants in its peer group. For example, there is some probability that Oconee would have been discovered earlier as well as later than it was, given that some four years transpired between the issuance of Generic Letter 97-01 and the discovery of extensive circumferential cracking at Oconee. Another is the time of the next convenient outage at Davis-Besse relative to the time of the order. This would have framed the options for a risk-informed request to delay the exam which could have been significantly longer than the requested three months in this event.
 - ▶ What is the probability that the damage is detected in the next attempted inspection? It is now obvious that this was optimistically treated up to and including the request for the inspection delay again due to overconfidence in the ability to detect cracking from visual exams and in the modeling of boric acid crystal formation. There is some chance that the efforts to repair the cracks at Davis-Besse would not have identified the external corrosion damage, in which case there could have been at least another fuel cycle of plant operation available for additional degradation of the vessel head and exposure for an accident initiating event. This would have led to additional time for the wastage of the vessel head material to progress. The initiating event for a pressure boundary failure scenario could be failure at the site of the corrosion, or failure as a result of pressure transient from another initiating event. It is conceivable that the time window for this “gedanken” experiment could be much larger than 4 years.
 - ▶ What are the possible initiating events that could have occurred if the detection had been delayed? Small, medium or large LOCA? Excessive (Beyond ECCS capability) LOCA? Ejection of one or more CRDM nozzles? Transient with increasing primary system pressure challenging liner membrane? The author does not have sufficient information to determine

all the possibilities, but the existing documentation appears to have very little discussion of the possible initiating events beyond the obvious LOCA candidates. Any good PRA should include more than a single sequence in the quantification of risks.

- ▶ What are some of the conditioning events to be considered in the event trees for each possible initiating event? Does a component on the pressure boundary become a missile challenging containment? Does the vessel head fail in response to a pressure transient? Does the loss of coolant flow from the vessel head over the long term following RCS depressurization and emptying of the Borated Water Storage Tank permit ECCS recirculation? Is there ECCS sump blockage due to damaged vessel head insulation? Is sump strainer blockage more or less likely than the case with a pipe break? Does the reactor vessel level indicator system (RVLIS) provide correct indications of reactor water level following ECCS operation? These are questions that need to be answered in constructing a PRA model of this degradation incident. It is not immediately clear that the risk characterization as a potential small or medium LOCA is sufficiently complete to capture all the risk implications.
- ▶ The application of any modeling assumptions about the degradation mechanism and the rate at which various damage states may occur on the pressure boundary need to be treated very carefully and not just with the “best estimate” assumptions used in the Alloy 600 cracking deterministic safety evaluations. Alternate models should be considered (epistemic uncertainties) and at a minimum, uncertainties in the model inputs such as stresses and temperatures need to be quantified or at least examined via sensitivity analyses.
- ▶ The characterization of the inspection process in the Davis-Besse evaluation was too simplistic and was approached as an exercise in human reliability analysis. The probability that a given level of damage does not show sufficient evidence to identify the damage was not and should be considered. This characterization is of course tied to the characterization of the behavior of boric acid crystals.

3.2 Callaway Steam Generator Electro-Sleeving [12]

The Callaway Plant requested approval of repairs to damaged steam generator tubes using a proprietary electro-sleeving process developed by Framatome Technologies Inc. This is an example of a licensing amendment request that was not a risk-informed submittal, but rather a change in the method for repairing steam generators which was requested for acceptance on the basis that the new method was in conformance to existing, deterministic licensing requirements. The NRC staff agreed that the proposed change meets the existing licensing requirements, but also considered the possible impact of the change on the capability to prevent a containment bypass during a beyond the design basis severe core damage accident. Specifically the capability of the steam generator tubes repaired by this method to prevent a thermal creep rupture failure during a high pressure core damage event with dry steam generators was investigated. In the staff evaluation, an estimate of the possible increase in LERF was estimated, and these estimates fell into the range of interest for increased management attention according to Regulatory Guide 1.174 risk acceptance guidelines. As a relatively high increase in LERF was estimated the staff was not united behind the decision to grant the license request. Interestingly, since core damage accidents are not within the design basis the use of RG 1.174 to investigate the impact of the change on LERF, the sleeving method was being evaluated for a range of conditions that exceed the design basis.

This license request and the risk-informed evaluation performed by the staff expose a fundamental limitation in the deterministic licensing requirements and the notion that simply meeting the existing deterministic requirements is automatically justified. Neither steam generators nor containments were specifically designed to withstand the conditions of severe core damage events with or without a large early release but rather were designed to perform specific safety functions during design basis accidents. At one time, there was an unresolved safety issue with regard to the capabilities of plants to withstand beyond design basis severe core damage events but this issue was resolved in part by the IPE/IPEEE program. Now that this issue is resolved, any plant change that is deemed to meet the current licensing requirements is regarded as acceptable whether or not the change represents an increase in a severe accident risk metric such as CDF or LERF. Since the current licensing requirements do not explicitly address the capability of safety-related components to perform functions during severe core damage events, this change was found to meet these requirements. The root cause of the staff anxiety and lack of a consensus behind the granting of this request is easy to understand. There must be other cases involving non-risk-informed change requests that could have been treated in the same way but were not. In the opinion of the author, the root of the problem is a weakness in any deterministic evaluation that chooses to ignore changes in severe accident risk metrics. While the question of whether plants designed against the deterministic safety requirements exhibit an acceptable level of risk of severe accidents was addressed for the entire industry at a given point in time, the inevitable changes in plant design and operation have led and will continue to lead to changes in the risk profiles.

The decision by staff in this case to introduce a risk-informed evaluation was very astute as the decision to approve the request was able to benefit from the associated risk insights. However, any plant change has the potential to result in changes the CDF and LERF, whether it is presented as a risk-informed submittal or not. So a key issue for the staff is to decide under what criteria a risk impact evaluation needs to be performed for requested changes that are not submitted as risk-informed.

Several of the people interviewed in the preparation of this paper pointed to this review as exposing a difficult area of risk-informed regulation. Indeed, the review points more directly to weaknesses in attempts to perform deterministic evaluations of plant changes without considering risk impacts. This example and its results challenge the mindset that:

- There is no need to justify the default position that meeting the current licensing requirements is always justified.
- Risk-informed regulation should be voluntary.
- The burden of proof in risk-informed submittals is to justify the change and not to justify the status quo.
- Any change meeting the existing requirements or decision to maintain the status quo is automatically justified.

Part of the anxiety caused by the review of this submittal is due to another weakness of deterministic safety evaluations that are typically performed to determine whether the existing regulatory requirements are being met, and that is the fixed or static nature of the design basis accidents used in these evaluations. Since the design basis accidents are fixed, the question of whether a design change may alter the logic for the selection of the design basis accidents is never asked. It would seem prudent in any safety evaluation to consider whether the change would impact the frequency or mitigation capability for any class of reactor accidents whether in the original design basis or not. There does not appear to be a traceable way to review

whether a change actually invalidates the selection of the design basis accidents, as the basis for this selection of accidents was never documented in way that the logic could be retraced or checked.

In summary, this submittal and its review point to a strength of the risk-informed approach and a weakness to any deterministic safety evaluation that simply checks whether the existing regulations are being met.

3.3 Risk-informed Emergency Diesel Generator AOT Extensions

Many of the successful applications of RG 1.174 decision making involved risk-informed justifications to increase the allowed outage times of safety-related components such as emergency diesel generators (EDGs) from a typical level of three days to 10 to 14 days. There are a number of reasons why this application of risk-informed regulation has been quite successful including the following:

- Increasing the EDG AOT from a typical value of 3 days to a requested level 10 to 14 days is very beneficial to the licensees as it provides the capability to perform EDG maintenance on-line and reduces the workload for maintenance during refueling outages. This was a key factor in the industry efforts to reduce refueling outage unavailability from the historical 60+day level to the sub 20 day level that is becoming more common. Hence, the staff received many requests to change the EDG AOTs to 10 to 14 days.
- Even though nearly all of the licensees used a full power internal event PRA to perform the risk impact evaluations, qualitative reasoning was effective in showing that risk changes during outages would be a decrease in risk as the unavailability of the EDGs during the outages would actually decrease. Hence the change in risk calculated with the at power PRA model would be conservative as these beneficial changes during the outages were being ignored. In the case of external events, which are normally omitted in these risk calculations, qualitative arguments were made and accepted by the staff that changes in the AOT would not have a significant impact on the risk of external events, as basic events associated with EDG maintenance unavailability were not risk-significant.
- The cause and effect relationship between changes in AOTs and changes in risk metrics such as CDF and LERF are fairly well known and localized to changes in maintenance unavailability. While the staff introduced some rather arcane risk metrics (incremental conditional core damage probability or ICCDP) to evaluate changes in AOT, the task of computing these risk metrics is rather straightforward.
- Licensees were able to show that the risk changes of increased allowed outage times could be offset or minimized by compensatory measures such as no switch-yard maintenance during EDG maintenance and tracked using configuration risk management tools.
- Essentially all the submittals were able to show that the calculated risk increases were very small or small according to RG 1.174 risk acceptance guidelines.

3.4 Risk-informed Inservice Inspection of Piping Systems (RI-ISI)

There was a somewhat different set of reasons why this example of risk-informed regulation was successful, including the following:

- As with the EDG AOT example, this application was favored by the industry because of a large benefit associated with burden reduction as these applications justified a large net reduction in Section XI imposed volumetric examinations in Class 1 and 2 piping systems. In addition to eliminating the cost of these exams, there was a significant reduction in personnel radiation exposures from elimination of exams in high radiation areas.
- Unlike the EDG AOT example, it was much more difficult to show a cause and effect relationship between changing ISI exam locations and changes in risk metrics. This was due to the fact that the PRA models did not include a detailed representation of the exam locations on passive components, and the difficulty in predicting the impact of changes in pipe rupture frequency due to changes in the ISI program. This added difficulty was overcome with special risk-informed methodologies supported by large research programs at ASME and EPRI. Many person years of effort was invested by the industry and the NRC staff to approve these methodologies, whereas there was comparatively little new methodology needed to address the EDG AOT issues.
- The changes in CDF and LERF calculated using the Westinghouse Owners' Group (WOG) and EPRI RI-ISI methodologies were very small in relation to RG 1.174 risk acceptance guidelines. Hence the conclusions that the risk impacts are acceptable were perceived to be less susceptible to variations in PRA assumptions and inputs.
- Insights from service experience and results of many years in performing ISI programs led the staff and the industry to conclude that the relationship between Section XI ISI programs and piping reliability performance was very weak. It was generally recognized that Section XI was producing very little evidence of damage in high stress locations while additional augmented ISI programs had to be added to address damage mechanisms that showed up in service experiences that were not understood when Section XI ISI was introduced. Hence, on a qualitative basis it was easy to show that the changes in risk from RI-ISI were small and very likely to result in a net reduction to risk metrics.

As noted in the previous sections, the implications of the Davis-Besse vessel head degradation on previous and future RI-ISI evaluations need to be determined.

3.5 Less Successful Risk-informed Applications

There were other risk-informed applications that for one reason or another were not viewed as successful as the EDG AOT and RI-ISI examples. Some of the reasons are listed as follows:

- Risk-informed Inservice Testing of pumps and valves was not very popular with the licensees because of poor experience by licensees that performed the pilot studies, namely that the benefits were perceived to be small in relation to the investments needed to support the application.

- Some licensees did not follow the guidance in RG 1.174 and the application specific guides and, as a result, were not accepted by the staff. (Sequoyah RI-IST).

4. TECHNICAL ISSUES IN PRA FOR RISK-INFORMED DECISION MAKING

From the staff interviews, reviews of selected risk-informed initiatives discussed in the previous sections, and the author's experience in performing and reviewing PRAs, a set of technical issues in PRA for risk-informed decision making was developed. These issues are summarized in Table 4-1. This is not regarded as a complete list but rather a representative set of issues that at least some of the existing industry PRAs exhibit to varying degrees. It should be regarded a set of issues developed at a "snapshot" in time as many issues are being addressed in current activities to upgrade the PRAs. The list is presented to help define the current state-of-the-art of PRA and point to areas of improvement that would enhance the risk-informed decision making process. The issues fit into the following general categories that are discussed in the following sections.

- Use Of Limited Scope PRAs In RG 1.174 Applications
- Lack Of Completeness Within The Specified Scope
- Model To Plant Fidelity Issues
- Lack Of Uncertainty Treatment
- Quantification Issues
- Multi-Unit Site Issues
- Lack Of Capability To Treat Aging Effects On Risk Results
- Risk Metric Issues
- Lack of Coherence Between Probabilistic and Deterministic Safety Approaches

4.1 Use of Limited Scope PRAs in RG 1.174 Applications

A typical industry PRA has several of the following limitations in PRA scope

- No quantitative PRA for external events such as seismic events which is limited to a seismic margins evaluation
- A simplified screening level analysis for internal fires that cannot be directly compared to the results of the internal events analysis
- A simplified screening level analysis for internal flooding that is based on out of date pipe failure rates, does not include significant consideration of human error induced flooding, and likely has screened out fire protection system piping based on inadequately documented assumptions
- No PRA to estimate the annual average CDF or LERF from low power and shutdown events
- A Level 2 PRA treatment that typically includes a simplified and conservative estimation of LERF
- Point estimate quantification of CDF and LERF with little or no quantitative treatment of uncertainties

**Table 4-1 Summary of Frequent Issues Identified in PRA
and RIR Submittal Reviews**

Areas of Difficulty	Specific Issues	Comments
Use of limited scope PRAs in RG 1.174 applications	<p>Lack of criteria and consistency for use of qualitative risk impact evaluations for contributors missing in PRA scope</p> <p>Lack of acknowledgment or consideration of PRA limitations used in submittals</p>	<p>Range of treatment includes no consideration of out of scope contribution to risk metrics, brief “arm waving” statements, to thoughtful and logical discussion that provides significant support for the risk impact conclusions</p> <p>Unlikely that this will ever change unless licensees are asked to provide this. Ironically, addition of this perspective would do more to build trust than to undermine the PRA</p>
Lack of completeness within the specified scope	<p>Inadequate treatment of support system initiating events</p> <p>Inadequate justification / documentation for events screened out of the PRA model</p> <p>Inadequate resolution of accident sequences and dependencies in event sequence modeling</p> <p>Inadequate treatment of dependencies in event sequence quantification</p> <p>Inadequate common cause failure treatment</p>	<p>Variability in treatment of support system initiators is much larger than can be justified by the variability in plant designs. Few plants have systematically examined dual bus initiators but those that did so have identified significant risk contributions. Initiators caused by combinations of faults and unavailability states in different support systems are normally overlooked.</p> <p>PRA documentation of the early stages normally not carried forward. Lack of review by those with intimate knowledge of the plant.</p> <p>Lack of consistency in level of detail in the event tree modeling to pick up dependencies such as dual unit interactions, transient induced LOCAs via PORV lifting, and use of inappropriate criteria for terminating accident sequences</p> <p>Lack of defensible method for treating dependencies between two or more human actions in the same time frame is the biggest issue here.</p> <p>There is still a wide variability in PRAs in the treatment and coverage in CCF components and failure modes. Most are using the NRC generic estimates of MGL parameters and very little are applying existing guidelines to make CCF parameter estimates plant specific. Current NRC-Sponsored methodology does not address plant to plant variability and this is important for at least EDGs. Many models of support system initiating event frequencies do not include CCF treatment.</p>

Table 4-1 Summary of Frequent Issues Identified in PRA and RIR Submittal Reviews (cont'd)

Areas of Difficulty	Specific Issues	Comments
Model to Plant Fidelity Issues	Lack of review by system engineers, operators and plant personnel	In order to ensure model to plant fidelity, it is necessary that personnel with intimate knowledge of the plant and the procedures review certain aspects of the PRA such as system notebooks, operator action treatment, etc. Not only does this support PRA quality but it also facilitates PRA technology transfer to plant personnel and supports effective risk management. This has been done to varying degrees and even when done, is not always periodically updated.
Lack of uncertainty treatment	<p>Lack of quantification of parametric (aleatory) uncertainty</p> <p>Uncertainty due to SSCs not modeled or screened out of PRA model</p> <p>Lack of quantification and sensitivity analysis of modeling (epistemic) uncertainties</p>	<p>Point estimates of CDF contributors may not represent means if uncertainties on PRA input data are not assessed to ensure input point estimates are means. If steps are taken to ensure inputs are means, point estimates of CDF are reasonable estimates of the means if “typical” cut-sets dominate, e.g., combinations of independent failure events and unavailabilities. Ability to safely propagate mean point estimates is suspect in case of uncertainties in some of the time dependent models such as RCP seal LOCA time to core damage vs. time to recover offsite power models. Many “best estimate” assumptions regarding engineering calculations probably closer to medians or modes rather than means. Point estimates of ISLOCA frequencies often underestimated due to state of knowledge dependence among multiple check valve failure rates. Errors of a factor 5 to 10 commonly result from this mistake.</p> <p>Possibilities for failures in reactor vessels such as Davis-Besse head corrosion not considered in current PRA models. Bases for screening out SSCs and events from PRA are not very well documented.</p> <p>Robustness of the quantitative results may be suspect unless key modeling issues identified and examined via sensitivity analysis. When there exist alternative and plausible hypotheses about specific modeling assumptions and a reasonable treatment of expert opinion, selected epistemic uncertainties should be treated quantitatively: Examples include RCP seal LOCA models and curve fits to industry data for time to restore offsite power. Care needs to be taken to prevent this from becoming an open ended and counterproductive exercise.</p>

Table 4-1 Summary of Frequent Issues Identified in PRA and RIR Submittal Reviews (cont'd)

Areas of Difficulty	Specific Issues	Comments
Lack of uncertainty treatment (cont'd)	<p>Treatment of time dependent failure rates in Bayesian Updating</p> <p>Use of uncertainties in decision-making process</p> <p>Generic treatment of uncertainties</p>	<p>There are many cases in which Bayes updating with plant specific data is being performed over many years of plant operation within which there have been significant changes in plant management, maintenance practices, etc. and any time dependent trends in SSC performance are being masked. The industry lacks tools to perform time-trend analysis with Bayes' updating.</p> <p>While many complain that not enough is done to quantify uncertainties in PRA, it is not clear how such uncertainty information will be used in decision making. Safety goals, quantitative health objectives and probabilistic criteria in regulatory guides ask for mean values and have presumably accounted for uncertainties in setting the criteria.</p> <p>In lieu of a full quantification of uncertainties in existing PRAs, are there approaches to quantify epistemic and aleatory uncertainties on a generic basis?</p>
Quantification and Risk Metric Issues	<p>Lack of validation of PRA results for CDF and LERF</p> <p>Uncertainty due to cut-set truncation tools</p>	<p>In the IPE/IPEEE era, it was generally accepted by many experienced PRA practitioners that variability in plant PRA results for CDF, LERF, and other common risk metrics were more due to analysts driven factors, such as differences in assumptions, modeling treatment, methodology etc., than actual physical plant differences. Although this less true today, PRA results tend to be benchmarked against other PRA results than to any objective data. Much more could be done along the lines of NRC ASP program to benchmark PRA models against industry data so that bottom line PRA results could be taken more seriously.</p> <p>In linked fault tree codes, since truncation is performed prior to Boolean reduction, the accident frequency associated with truncation is unknown. Robust conclusions that risk changes are less than RG 1.174 acceptance guidelines are difficult to meet since the frequency and nature of the truncated model are unknown. In event tree linking, the problem is less severe since the truncated accident frequency is actually quantified, though needs to be managed to control magnitude. New quantification tools such as binary decision diagram (BDD) method are capable of solving this problem but few if any commercially available tools are available to solve this. Calculation of RAW values for risk classification of SSC type of applications (Option 2 of SECY-98-300) can be significantly impacted by truncation, e.g., some SSC calculated as RAWs being less than 2.0 are actually greater than 2.0.</p>

Table 4-1 Summary of Frequent Issues Identified in PRA and RIR Submittal Reviews (cont'd)

Areas of Difficulty	Specific Issues	Comments
Quantification and Risk Metric Issues (cont'd)	Lack of capability and effort applied to eliminate logic errors from complex logic models	Contemporary event tree-fault tree logic models are very large and complex. In addition the models are developed, modified, expanded, and applied over long periods of time by different people and contractors, etc. It is very difficult to perform basic logic error debugging. Logic errors may exist latent in the models for periods of several years and not identified until the model is exercised in certain ways and in applications. In linked fault tree models there are too many cut-sets to review for that technique to be successful in identifying a large fraction of the errors. Some of the logic errors result from incrementally logical patches that create unexpected illogical interactions with other model elements modified at different times. There is a need for better tools and better modeling building guidelines to be able to perform more effective logic error reviews of the models.
	Preoccupation with bottom line numbers and lack of effort to develop risk insights	Despite the number of speeches, papers, and strong consensus among PRA experts that risk insights are more important than the bottom line numbers, many of the current industry PRAs have devoted only limited effort to derive risk insights from the results. Evidence of this condition are very limited PRA results summaries, results tables pasted directly from computer outputs with no discussion or documented review, limited analysis of risk contributions, inadequate accounting for unusual results and no comparative perspectives.
	Use of fault trees to model initiating event frequencies	The correct methodology for use of fault trees to model initiating event frequencies is applied in very few of the existing industry PRAs. Problem areas are incorrect application of tools set up to model unavailability, lack of enumeration of failure modes, misapplication of the 24-hour mission time (vs. an 8760-hour mission time), and inability to link the dependencies with fault trees in the same systems present in the model for mitigation functions.
	Adequacy of CDF and LERF risk metrics	Today we have roughly 100 reactor units each with a calculated CDF typically in the range of 1E-4 to 1E-5 per reactor year. When we are down to the last year of operation of the last reactor, our industry risk, which will be two orders of magnitude less than today in the absence of significant ageing effects, yet these individual plant risk metrics would be unchanged. Is there merit in the development and use of industry wide risk metrics such as the probability of a core damage event over the remaining reactor year population to guide NRC decision making on issues that involve the whole fleet of plants?

Table 4-1 Summary of Frequent Issues Identified in PRA and RIR Submittal Reviews (cont'd)

Areas of Difficulty	Specific Issues	Comments
<p>Lack of capability to treat aging effects on risk results</p>	<p>Lack of explicit representation of passive components and safety features in base PRA models</p> <p>Lack of questioning of constant failure rate assumption and treatment of time dependent failure rates</p> <p>Large uncertainties about specific degradation mechanisms</p>	<p>Many passive component failures screened out of initiating events and event sequence models, others represented implicitly; strong tendency not to revisit these early PRA model decisions after years of updates</p> <p>Few outside the NRC Operating Experience Risk Analysis Branch perform any trending analysis and uncritically apply the constant failure rate assumption masking possible temporal trends in equipment behavior. Current models for LOCA frequencies (NUREG/CR-5750) assume LOCAs would be dominated by several pipe failure damage mechanisms (TF for PWRs and IGSCC for BWRs) however these and many other observed pipe damage mechanisms are inherently aging effects. Hence LOCA frequencies should be increasing with time. Also since most plants have performed RI-ISI on Class 1 and 2 piping, perhaps we should have plant specific LOCA frequencies?</p> <p>Poor understanding of damage mechanisms tends to lead to an underestimation of the time available for inspection and to incorrect conclusions about the effectiveness of specific inspection techniques, e.g., Davis-Besse Head Degradation.</p>
<p>Lack of Coherence Between Deterministic and Probabilistic Safety Approaches</p>	<p>Need for consistent definition of risk</p> <p>Need for consistent set of accident scenarios</p> <p>Need for criteria to evaluate risk impacts of plant changes that meet existing requirements</p>	<p>The PRA definition of risk [17] should be adopted for use in both deterministic and probabilistic evaluations of a safety issue.</p> <p>Design basis accidents are essentially generic and the basis for their selection is obscured. DBAs are artificially constrained by the single failure criterion while PRA results are dominated by risks that involve more than a single failure. PRA provides a more complete representation of accident sequences but is skewed away from the design basis. PRA results strongly suggest that risk significance of accident sequences are to be highly plant specific and hence, design basis accidents should also be considered plant specific.</p> <p>The NRC and the industry need to have a predicable process for deciding when to invoke risk considerations in non-risk-informed change requests.</p>

Table 4-1 Summary of Frequent Issues Identified in PRA and RIR Submittal Reviews (cont'd)

Areas of Difficulty	Specific Issues	Comments
<p>Lack of Coherence Between Deterministic and Probabilistic Safety Approaches (cont'd)</p>	<p>Need for consistent treatment of uncertainties including need to address epistemic uncertainties in deterministic safety evaluations</p> <p>Lack of clear criteria for evaluating impact of changes on defense in depth and safety margins.</p>	<p>Uncertainties need to be given greater emphasis in traditional deterministic safety evaluations especially modeling assumptions for damage mechanisms for passive components and structures.</p> <p>All current RG 1.174 submittals must provide arguments that defense in depth and safety margins are maintained. However, defense in depth has not been defined with sufficient clarity to be able to determine when a sufficient level of defense in depth has been provided. In addition, safety margins are not well enough defined to evaluate when they are indeed adequate. As a result, these aspects of risk-informed decision making, though conceptually sound, are nothing more than arm waving.</p>

These typical scope PRAs are being used in conjunction with qualitative arguments of varying degrees of rigor to address the risk acceptance guidelines of RG 1.174 which beg for a full scope PRA evaluation of the change. Although such valid arguments have been constructed and accepted by the NRC staff, there is currently inadequate guidance to consider the impact of PRA scope limitations on the validity of the conclusions that are being drawn. The following points need further investigation.

- It is reasonable to exclude external events from the scope of a PRA if there is some convincing evidence to support the hypothesis that the within scope contributors dominate the baseline CDF and LERF results as well as the changes to CDF and LERF in the application for focused RG 1.174 (Option 1) applications. It is extremely doubtful if such arguments can be made for the case of fires for many of the older plants with weak physical separation of redundant systems, structures, and components (SSCs). It would be interesting to identify how many risk-informed applications were performed of older plants based on PRAs that did not include the fire contribution. This situation is compounded when an older plant is claiming very low CDF values from internal events, e.g., less than 1×10^{-5} per year.
- For Option 2 of SECY-98-300 or Option 3 applications exclusion of external events from the PRA models should only be justifiable when it can be shown that such excluded contributions make small contributions to risk. If the combination of a very low calculated CDF and generic insights that missing contributors may dominate exists, such a hybrid approach is not nearly as trustworthy as developing a more complete PRA.
- NEI has proposed an approach to classify SSCs making use of the best information available including PRAs, seismic margins, simplified fire PRAs, etc. in which case conservative treatment is applied to treat the PRA scope limitations [13]. While this approach appears reasonable, there are pitfalls to such a hybrid approach. The obvious shortcoming is that the conservative treatment to compensate for less than full scope PRA treatment runs contrary to the goal of focusing resources on the areas that are most important to safety. Another concern is that some risk contributions may fall through the cracks. In full scope and fully integrated PRAs of fires for example, there are some risk significant sequences involving combinations of failures from fires and other independent events. For example, a fire may disable one train of safety significant SSCs and another redundant train could fail independently. The risk significance of the SSCs involved in the independent events may be significantly increased in relation to the internal events and it is unlikely that a simplified treatment will pick this up.
- For plants that did not perform a seismic PRA and have calculated a relatively low CDF or LERF from internal events, it is not clear whether seismic events dominate or not or whether they even make risk significant contributions. In these applications, it is reasonable to exclude the seismic PRA if it can be shown that the proposed change does not impact the capability to mitigate the consequences of a seismic event in focused Option 1 applications. Since SSC failures caused by a seismic event are expected to be dominated by common cause failures, such arguments should not be too difficult to construct. For Option 2 of SECY-98-300 and 3, the same comment made above for fires seems to apply. For Option 2 of SECY-98-300 or 3 applications safety related SSCs that support safety functions during any mode should not be placed in a low safety significant category unless there is an external event PRA or an equivalent external event exclusion argument to support it in all applicable operating and shutdown modes.
- For Option 1 applications that can be shown not to impact the capability to mitigate accidents initiated during shutdown, it is reasonable to treat the risk contributions from shutdown in a qualitative manner. For Option 2 of SECY-98-300 or 3 applications safety SSCs that support shutdown safety functions should not be placed in a low safety significant category unless there is a shutdown PRA or equivalent shutdown PRA exclusion argument to support it.
- While there are valid technical arguments that can be made to justify the exclusion of some portions of a full scope PRA model for risk-informed regulation, there are resources that must be continually applied by the licensee and the NRC to check the validity of the risk-informed decisions in light of the use of an

incomplete PRA model. At some point, it is reasonable to ask whether these additional resources are small or large in relation to the use of full scope PRA to start with.

- Consideration should be given to the development of generic estimates of risk contributions to various elements of a full scope PRA for use in lieu of a full scope PRA. The idea is that if a partial scope PRA is used only part of the risk increase 'budget' should be used in decision making. This might require placement of plants into several categories to be able to account for general characteristics such as seismic siting, age of plant in relation to fire protection requirements etc.

4.2 Lack of Completeness Within the Specified Scope

There are a number of issues identified in Table 4-1 assigned to this category including lack of adequate treatment of support system failures as initiating events, improper treatment of dependencies and common cause failures, improperly terminated event sequences that do not fully meet the success criteria, and several other issues that all lead to understatement of the core damage frequency. Some elaboration of these deficiencies is provided below.

Support System Initiators

There is a disturbingly large variability in how thoroughly support system initiators are treated among the current industry PRAs. Some initiating event lists are limited to loss of offsite power, loss of service water and loss of component cooling water whereas others include very detailed treatment of support system failure modes including loss of single trains of many mechanical and electrical support systems and in some cases common cause failures of multiple buses are included with interesting dependent failure interactions identified. The fact that support system to front line system interfaces are highly plant specific makes it difficult to achieve standardization. However there is a good deal of emphasis of this in the ASME PRA standard and that should help improve this issue.

Almost all of the plant PRAs are struggling with the issue of how to develop and quantify fault trees to model the system failure modes that represent initiating event frequencies and most of the existing PRA software is not designed for this purpose. It is extremely difficult to use the existing fault tree linking tools to link fault tree models of system failure modes as initiating events to those for system failure to perform mitigation functions. The issue of how to treat common cause failures of normally operating systems, in these system initiating event models, is very evident in these models. Existing PRA guidance is lacking in this area. Of particular concern is that plant PRAs that do a good job in modeling support system initiating events tend to find that these events dominate the internal CDF profiles.

SSCs Screened Out of PRA Models

In many RG 1.174 such as RI-IST and Option 2 of SECY-98-300, much of the effort to perform the application stems from back-fitting explanation for why most of the plant SSCs are not explicitly represented in the PRA models. While in most cases, there exist good justifications for why SSCs have been screened out, it is important to know what the justifications were to understand how to resolve the status of the SSC in a particular application. If the justification is that failure of the SSC would not lead to an initiating event or contribution to loss or degradation of a function needed to mitigate an initiating event, that is information that needs to be retained in the application. If the justification is that the probability of failure is expected to be very small, it may be necessary to perform a consequence analysis of failure in order to establish the effective Risk Achievement Worth of the SSC for the application. A superior approach is to include in the PRA documentation the documentation that justifies the exclusion of every SSC that is not modeled in the PRA. This approach would minimize the resources needed to perform and review subsequent risk-informed applications.

Resolution of Sequences and Dependencies

Given resolution of the previously mentioned completeness issues in the selection of initiating events, there is a separate issue of lacking completeness in the definition of accident sequences in the PRA models. The two most frequent causes of this shortcoming are inappropriate success criteria for terminating the development of accident sequences and inadequate treatment of dependencies and interactions that bear on the consideration of safety functions that need to be mitigated. There are frequent examples in which event sequences are terminated as “successful” end states at the end of a 24-hour period without achievement of stable plant conditions. Examples of the unstable conditions include failure to isolate leaking Steam Generators with continuing inventory loss bypassing the containment and decreasing coolant inventory that is still above the active fuel after 24 hours. Examples of missing dependencies include failure to consider the probability that a transient could develop into a LOCA via pressure increase lifting the pressurizer PORVs, and failure to consider various multi-unit interactions such as dual unit vs. single unit loss of offsite power events.

Human Action Dependencies

One of the most common difficulties shared by most if not all existing PRAs is the lack of adequate treatment of dependencies between two or more human actions in the same event sequence or cut-set. While some PRAs have performed sensitivity studies to identify where multiple human actions have been applied, there is inadequate guidance on how to quantify the probability of human errors given knowledge that other human errors are postulated in the same sequence and time frame. This issue compounds the long term problem that there is a lack of a consensus on the appropriate human reliability technique to model and to quantify a single human action and is likely more important than the apparently insoluble “errors of commission” and “organization factors” problems.

Common Cause Failure Treatment

Unlike the situation with human reliability, there has been a basic agreement on the overall methodology for treatment of common cause failures for quite some time. The NRC work over the past decade to collect a comprehensive common cause database has been an excellent advancement to this area, however the use of this database to gain insights on improvements to the methodology has been lacking. Unfortunately there are still many plant PRAs that have not taken full advantage of the methods and data that are available for this important contributor. Some of the major deficiencies in a typical current PRA include:

- Inadequate coverage of components and failure modes in the existing common cause failure models
- The tendency to lift generic estimates of beta factors and other parameters from the existing NUREGs vs. use of the available methods to develop plant specific parameter estimates
- The need to treat asymmetric component configurations as opposed to the symmetry assumptions in CCF models
- Lack of treatment of plant to plant variability in developing CCF parameter uncertainty distributions (especially an issue with emergency diesel-generators)
- Tendency for lack of inclusion of CCF contributions to evaluation of SSC risk importance measures

As the issues in this lack of completeness category are resolved, there should be an upward trend of the CDF and LERF results at the affected plants as all of these lead to an understatement of risk levels. However, there will be an understandable reluctance to announce changes to previous CDF results that have been cited in risk-informed submittals. There is an attendant issue that frequent changes to the published CDF results may undermine credibility and create a bookkeeping nightmare to have to go back and check each previous risk-informed decision to see if it has been impacted by the PRA update. A major

challenge for the NRC and the industry is to avoid any disincentives to incorporate new and improved knowledge into the PRAs while ensuring that risk levels are being properly managed.

In one of the NRC staff interviews it was pointed out that roughly 20% of the events being classified as “accident precursors” as part of the ASP can be associated with initiating events, accident sequences, or plant conditions that are not normally modeled PRAs. This points to a need to consider new efforts to benchmark PRA procedures against operating experience.

4.3 Model to Plant Fidelity Issues

A major challenge to the success of risk-informed regulation is the establishment of appropriate links between the PRA group and the plant management organization to ensure proper PRA model configuration control. Plants have addressed this issue with varying degrees of success. The most successful plants in this regard have periodic reviews of selected portions of the PRA performed by system engineers and plant operators to ensure that the PRA model reflects the as-built and as-operated plant. PRA groups often struggle with limited resources to keep various elements of their PRA models up to date. Those who win this struggle enjoy a secondary benefit of PRA technology transfer out of the PRA group which is necessary to achieve the full benefits of PRA as a plant risk management tool.

4.4 Treatment of Uncertainties

As noted earlier, most plant PRAs have not routinely included a thorough treatment of uncertainties in the PRA. A thorough treatment within the state-of-the-art would include a thorough quantification of parametric uncertainties, use of mean values for each uncertain parameter for all point estimate quantification steps, quantification of selected epistemic uncertainties where sufficient information is available, and sensitivity studies to address key modeling assumptions and epistemic uncertainties that are not readily amenable to quantification. Examples of epistemic or modeling uncertainties that are well within the state-of-the-art for quantification include uncertainty in PWR reactor coolant pump seal LOCA performance under loss of seal injection and loss of heat removal conditions, uncertainty in fitting curves to time to restore offsite power data, alternative hypotheses about success criteria, and many of the models for severe accident challenges to containment integrity. The use of alternative seismic attenuation models in the development of uncertainties in the seismic hazard curve is another well-known example where epistemic uncertainties are quantified on a “routine” basis. One of the issues here is lack of criteria for deciding which epistemic uncertainties to quantify and which to be relegated to sensitivity study treatment. This problem is exacerbated by the disturbing lack of effort being applied in most industry PRAs in developing the summary reports that are supposed to develop risk insights from the results.

In the NRC staff interviews, there was a general consensus that this issue is much less important than the limited PRA scope issue, but this is still an important issue to be addressed. The NRC safety goals and risk-informed decision criteria have been developed under the assumption that mean estimates of CDF and LERF would be provided and that such estimates would consider a thorough treatment of uncertainties. As noted earlier, a relatively small fraction of the industry PRAs have put an emphasis on uncertainty quantification in their PRAs. There are several explanations for this including:

- Some of the PRA consultants from which the licensees acquired training and technology for performing PRA have instilled a mindset among many licensees that uncertainty quantification is not necessary and not useful for decision making. It is also unclear to many how quantified uncertainties would be used by the NRC in risk-informed decision making.

- Some of the PRA software that is being used makes it difficult to perform uncertainty analysis, other software is only capable of limited treatment, and in a few cases, uncertainty analysis is an integral part of the implementation of the software. There are commercially available tools that can be used to compensate for these software deficiencies.
- PRA training and associated software tools to perform Bayes' updating of generic distributions with plant specific data is surprisingly lacking. There is only one PSA software tool that has a built in capability to perform Bayes' updating of distributions.
- The best case study readily available where both epistemic and aleatory uncertainties were quantified was in NUREG-1150. There was a very large research project that funded the cost of this work, and while the information is available to support industry PRAs, the costs of repeating the expert elicitation exercise in NUREG-1150 in industry PRAs are prohibitively high.

Despite the above reasons, there are some very strong motivations to improve the treatment of uncertainties in PRA for future risk-informed decision making. Some of these motivations are listed as follows:

Means vs. Point Estimates

This issue was addressed in a 1997 ACRS Letter on treatment of uncertainties vs. point values in the PRA related decision making process [14]. A minimum level of PRA uncertainty analysis is needed to justify the assumption that the point estimates of a PRA provide reasonable estimates of mean CDF and LERF. In a calculation sense, a PRA comprises sums of product terms where each product term consists of an initiating event frequency and one or more basic event probabilities. In the case where each term of the product is computed from a set of mutually independent parameters, the mean sequence frequency is equal to the product of the event frequencies and probabilities obtained by the mean parameter values. We speak about means in this context as parameters of an underlying uncertainty distribution. Whenever the sequence cut-set is not the product of events computed with independent parameters, the mean of the sequence frequency may differ from the mean point estimate. There are several situations where this independent assumption cannot be supported creating the potential for significant differences between the mean sequence frequency and the mean point estimate:

- The state of knowledge dependence where the sequence frequency is the product or higher power exponents of a basic event probability of frequency. The most important example of this situation in terms of its impact on calculating mean risk metrics such as LERF is the interfacing systems LOCA sequence of comprising two or more valve failures where each failure is computed using the same state of knowledge based failure rate. Examples have been found in industry peer reviews where failure to account for this led to an underestimation of LERF by factors ranging from three to ten. There are also important CDF sequences with this problem such as the traditional station blackout sequence with two or more independent diesel generator failures.
- Time coupled dependencies in basic events whose failure probability is a function of time and there is uncertainty in both the time available and time necessary to perform the action. In some cases such as a PWR station blackout, there may be two or more events in the same sequence that involve such a time dependency.
- PRA success criteria that are subject to uncertainty are often based on so-called "best estimate" thermal hydraulic analysis. It is well known in the field of expert elicitation that "best estimate" is often better correlated to the median than to the mean if such an estimate were to be replaced by a full uncertainty treatment.

PRAs that present point estimates as representing mean values of CDF and LERF should at a minimum perform reviews to show that state of knowledge uncertainties and uncertainties in time dependent models have been taken into account in the mean point estimates.

Treatment of Temporal Variations in PRA Parameters

Despite the lack of acceptance of Bayes' treatment of uncertainties by the NRC in earlier phases of PRA development, there has been more widespread acceptance of Bayes' updating as a means for developing uncertainty distributions for component failure rates and initiating event frequencies when both generic and plant specific evidence must be taken into account. One technical issue that the industry PRAs are struggling with is how to perform Bayes' updates without masking temporal variations in failure rates that may occur over several decades of a plant and industry lifetimes. Often plant specific evidence over periods of 10 to 20 years is being collected yielding very narrow updated uncertainty distributions in the failure rate estimates. This treatment often ignores the possibility that changes in equipment performance or maintenance practice may suggest different failure rate in different time periods yielding an artificially high confidence in the central tendency of the failure rate distribution. Tools to perform the data analyses that identify trends as well as determine the appropriate intervals over which to average the data are needed. Otherwise Bayes' updates will mask an important variability that will be missing in the uncertainty analysis.

Enhanced Information for Decision Making

The important argument to be made for a more complete treatment of uncertainties in the PRA is to enhance the robustness of the PRA information for the decision making process. The more sources of uncertainty that are reflected in the results of a PRA, the less margin needs to be applied by the decision maker to provide confidence that risk levels will be maintained to an acceptable level.

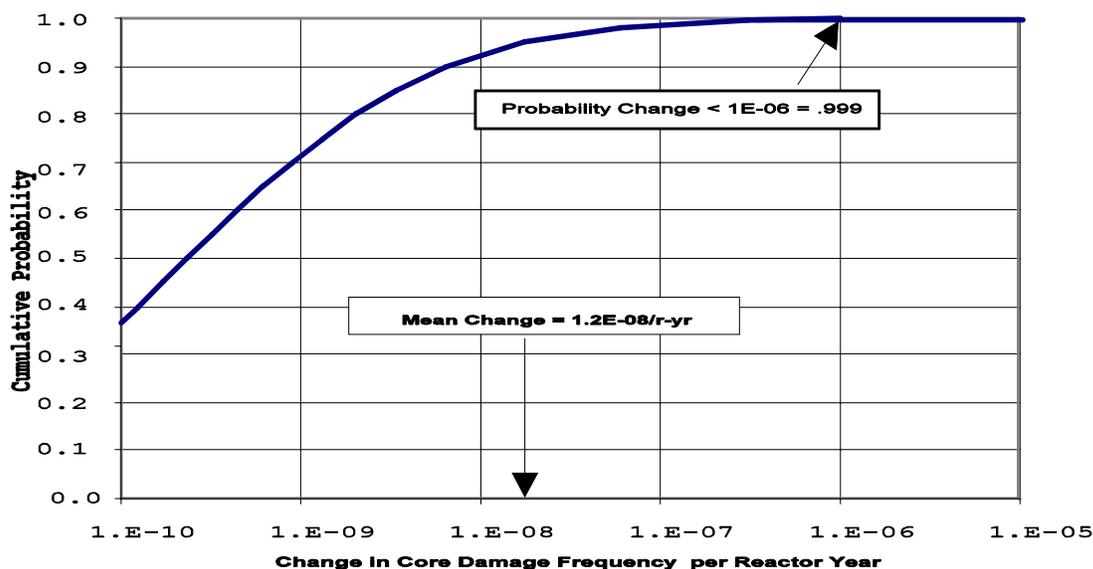


Figure 4-1 Uncertainty in Change in CDF for BWR Weld Overlay Example [15]

In the case of developing the baseline risk levels for a plant, it is not clear how useful knowledge of the quantitative impact of uncertainties is to decision makers. The primary benefits of a good uncertainty analysis are to provide a convincing case that the results of the PRA for the mean CDF and LERF are valid and to develop a sense of confidence that uncertainties have been carefully considered in the development of the PRA models and their quantification. In the case of a change in risk analysis such as those performed

for Regulatory Guide 1.174 applications, an uncertainty analysis in the change in CDF and LERF can be very helpful in demonstrating that risk changes are indeed small and that such conclusions are robust in light of the underlying uncertainties. An example of this type of uncertainty analysis was performed to address an issue associated with inspection requirements for welds in BWR piping that had been repaired with weld overlays to address inter-granular Stress Corrosion cracking problems [15]. The question addressed in this examination was the change in risk associated with a proposed relaxation of the inspection requirements in these repaired welds. There are large uncertainties associated with this change in risk analysis due to large uncertainties in estimating weld failure rates, uncertainties in how much these failure rates change as a function of ISI program, and uncertainties in the consequences of weld failure as measured by the conditional probability of core damage given pipe failure as well as the plant-to-plant variability in this risk metric. The uncertainty in the change in CDF resulting from the proposed inspection program change is presented in Figure 4-1. As seen in this figure, the uncertainties are very large, several orders of magnitude, in the context of trying to estimate the point value of the change. On the other hand, there was still a high degree of confidence (>.999) that the change was less than the risk acceptance guideline from RG 1.174 of 1×10^{-6} per year. One important characteristic in a change in risk uncertainty analysis is that only those sources of uncertainty that participate in the event sequences impacted by the change are reflected in the analysis as the remained PRA uncertainties cancel.

4.5 Quantification Issues

The existing PRAs exhibit a variety of quantification issues that need to be understood in the associated risk-informed applications. Several of these are discussed in the following:

Lack of Validation of PRA Results

In one of the NRC staff interviews, it was pointed out that roughly 20% of the events classified as accident precursors in the NRC accident precursor program refer to initiating events and accident sequences that are not included in PRAs. There are additional discrepancies between PRA results and service experience realities that are presented in a series of reports published by the NRC on system reliability and availability experiences. It is not clear that this discrepancy is being addressed in any of the industry efforts to achieve PRA quality. This points to the question of how PRA results are validated.

Another indication of a validation issue is the interesting trend in the updated PRA results since the IPE and IPEEE submittals. At the time of the IPE and IPEEE submittals, the results reported for CDF and LERF spanned several orders of magnitude with PWR CDF results from greater than 1×10^{-4} to less than 1×10^{-6} per year with BWR results spanning a range of somewhat lower values. In the current updated PRA results the upper bounds of these distributions have been reduced as a result of plant changes to reduce high risk contributions and changes to PRA modeling assumptions etc. The interesting consequence of these updates is that the variability in CDF and LERF results has been greatly reduced.

There is a large fraction of the current PWRs with updated mean CDF results in the narrow range of 1×10^{-5} per year to 2×10^{-5} per year including the results for STP however most of the plants in this group do not have some of the plant features that would suggest a lower than average CDF such as is the case at STP. These features include three and four train redundancies of electrical and mechanical systems, high physical separation of redundant trains, capability of preventing seal LOCA conditions during a station blackout, and very low level of seismic hazards. A large fraction of the PWRs who report essentially the same CDF values have two train redundancy, offer less physical separation, have no way to protect the seals during a station blackout, and are sited in areas with higher seismicity than central Texas. Of course, most of these plants have not even performed a seismic PRA nor a fire PRA that is as detailed and realistic as the internal events analysis.

One of the weaknesses of the industry peer review process and the current PRA standards is the lack of a structured process to ensure consistency between results for CDF and LERF and plant features

responsible for the deviations from typical CDF and LERF results across the industry. So, while PRA consistency has been improved, we are still not at the desired point where the variability of PRA results or lack thereof is driven by physical variability between the plants. In addition, there may be an unconscious bias that is introduced by industry efforts to reduce the calculated risk of high risk contributors with no counterbalancing effort to question whether some lesser risk contributors might have been optimistically treated. The case for advancing the use of PRAs in risk-informed decision making could be strengthened if more was done to validate the results of PRAs so that results variability was not driven by differences in the treatment of generic issues.

Cut-set Truncation Uncertainty

The existing PRA quantification software employ some form of truncation in the algorithms for quantifying the accident sequence frequencies. The CDF and LERF estimates presented in the PRA reports are based on the sequences and cut-sets left over after a truncation process in which sequence and cut-sets with frequencies below a user-defined cutoff have been deleted from the model. Those software that employ the fault tree linking technique (as opposed to event tree linking) quantify individual cut-sets prior to the completing the last steps in Boolean reduction, in which case the error introduced by truncation is not determined and the upper bounds that can be estimated are often too large to dismiss. In some cases, it may be difficult to show that the error due to truncation is small in relation to RG 1.174 risk acceptance guidelines. In addition, the PRA software tools are programmed to compute risk achievement worths (RAWs) from the truncated model in which case RAW values for some SSC may be underestimated. Guidance could be improved for how to better manage truncation uncertainty in processing the results. There is a relatively new technique referred to as the Binary Decision Diagram (BDD) that is capable of solving fault trees without truncation or approximation which may be the longer term solution to this issue. In the mean time, it is necessary to take some extra steps to ensure that conclusions regarding risk significance of SSCs are robust in light of truncation uncertainty.

Logic Error Debugging

A very disturbing problem is that the event tree/ fault tree logic for a state-of-the-art PRA is so complex that it is very difficult to review the PRA model to identify simple errors in the logic. This is a limitation that could be minimized with better tools and better guidance on a structured process to build the logic model. In the current PRA models, there is too much reliance on the review of individual cut-sets as a means to ensure proper logic. Typically only 100 or so cut-sets for CDF and LERF are even presented in a PRA summary report and in some cases this many only represent a few percent of the total CDF and LERF estimates. Logic errors that incorrectly suppress the risk contribution of sequence cut-sets are not reliably identified using this process. The linked fault tree models are simply too large to rely on manual review of the trees as only a small portion of the logic can be seen and assimilated at one time. In the industry PRA peer reviews, there were many examples of illogical cut-sets identified in the presented results. Not enough effort is being applied to perform logic error reviews and the tools are not very good at assisting with this task.

Bottom Line Numbers vs. Risk Insights

We have had conferences, seminars, and workshops on PRA since the late seventies and at all of those many speakers have pointed out that the most important outputs of a PRA are the risk insights and these insights are much more important than the bottom line numerical results. If you took a poll among PRA practitioners and other stakeholders who use PRA results, more than 90% would likely agree with such a statement. Unfortunately, this insight is not being put to practice very much in each successive update in a typical industry PRA. The results summaries that are prepared to explain the most recent PRA updates tend to be very brief, include only point estimates of CDF and LERF results, the percent contribution to these results from accident sequences with common initiating events, and a bar chart or two with some results of the risk importance analysis. Absent from most of these summaries are more extensive analyses of risk contributions, insights about the key plant features responsible for the specifics in the numerical results, explanation for the changes since the previous PRA updates, and information about uncertainties and

sensitivities. It is very typical that a results update review include only the top 100 cut-sets, even though this only normally accounts for a small fraction of the CDF or LERF. This issue ties in with the previous comments on logic error debugging in the sense that the summaries are indicative of a lack of effort to interpret the results. It is common that when a significant effort is expended to develop insights that logic errors and other problems with the PRA model are identified.

Use of Fault Trees to Model Initiating Event Frequencies

As noted earlier, the existing PRA guidance and software for developing and quantifying fault trees are focused on the fault trees for calculating the failure probability of system functions following an initiating event. They are also being used for the purpose of modeling system failure modes that represent initiating event frequencies. From a systems reliability viewpoint, this can be explained as the difference between calculating the system failure intensity or hazard rate vs. the system unavailability. The problems that are being encountered in this application include:

- Incomplete enumeration of system failure modes
- Use of the wrong “mission time” (24 hours vs. 8760 hours)
- Inadequate consideration of common cause failures in normally operating vs. standby systems
- Lack of capability to link these fault trees with fault trees for mitigation functions for the same systems
- Omission of SSC failure contributions to initiating event frequencies from SSC importance evaluations.

Adequacy of CDF and LERF as Risk Metrics

The current set of risk metrics, CDF and LERF, were designed to address a range of applications for a reactor unit in which the capability to prevent core damage accidents and large early releases could be evaluated using existing plant PRAs. One motivation for using LERF in lieu of a full Level 2 or Level 3 PRA is to minimize the extent to which a risk-informed evaluation would depend on the modeling of severe accident phenomena and resulting source terms that still remain controversial. Once one extends into a Level 2 PRA there is lack of agreement between the industry and the NRC on which severe accident phenomena to consider and how to model them. It is clear however that the range of PRA applications could be expanded if the PRAs were extended to Level 2 or Level 3.

There are additional limitations of the CDF and LERF risk metrics including the lack of treatment of event sequences with no core damage accidents which if added could be useful to provide more visibility of the design basis accidents in the PRA and to expand the range of PRA applications to include investment risk. As discussed below, these metrics also fail to capture the unique risk aspects of multi-unit sites.

A final limitation of these individual reactor risk metrics is that they fail to capture the total risk of having an accident across the reactor population and over the remaining lifetimes of the reactors. One could introduce some industry wide metrics such as the probability of an accident over the remaining reactor years in the existing reactor lifetimes, for example.

Use and Interpretation of Risk Importance Measures

With the NRC staff approval of the STP special treatment exemption and the industry interest in Option 2 of SECY-98-300 to risk informing 10 CFR Part 50, there is expected to be increased use of risk importance measures to help define categories of risk significance for plant SSCs. There is a host of technical issues with risk importance measures and a definite lack of a consensus among PRA practitioners on how to treat them. Some of these issues are delineated in Reference [20] and include:

- Failure to include initiating event frequency contributions in risk importance metric and lack of definition of risk achievement worth (RAW) for initiating event contributors
- Difficulties in mapping basic event risk importance to SSC importance

- Impact of cut-set truncation on importance measures, especially RAW
- Lack of independence between Fussell-Vesely and RAW for the same basic event yet presentation of results in orthogonal graphics
- Lack of visibility of functional dependencies which help determine risk importance
- Lack of visibility of equipment importance which is “buried” in operator action models
- Confusion on how to incorporate common cause basic event importance in computing SSC importance
- Inability of risk importance metrics to capture risk impacts that influence multiple basic events
- Lack of agreement of whether PRA supplying risk importance measures need to be complete and full scope or whether external events and low power and shutdown states need to be included

4.6 Multi-Unit Site Issues

As mentioned earlier, it is difficult to identify and adequately treat dependencies that exist between systems at multi-unit sites, particularly those with highly convoluted support system dependencies (systems and subsystems shared by different units). There are initiating events that may or may not impact two or more units at the same site, human action dependencies in deciding how to deploy equipment and personnel to support all plants on the site, and the possibility of accidents involving two or more reactors. The risk metrics that are being employed such as CDF and LERF are being developed either for one representative reactor unit, or for each reactor independently. Multi-unit reactor accident consequences are currently being ignored and there is no consideration that the frequency of core damage per site year will be increased due to independent contributions from each reactor at the site.

4.7 Lack of Treatment of Aging Issues

Although there has been research into the question of aging effects of SSC components in PRA, current PRAs continue to assume that initiating event frequencies and component failure rates are constant in time, without necessarily making any tests of this hypothesis, and such rates are assumed to occupy the flat region of the well known “bathtub curve” of failure rates as a function of the lifetime. Current PRA models for passive components are essentially limited to the treatment of pipe breaks and various damage and fragility models used for external events analysis. The most recent work in developing estimates for piping system failure rates [16] is based on service data with piping failures due to various degradation mechanisms. However the assumption that pipe breaks are dominated by degradation mechanisms is inconsistent with the assumption of constant failure rates. This suggests the need to investigate whether PRAs should be employing plant age dependent failure rates for pipe failure rates used for both LOCA initiating event frequencies and internal flood frequencies.

As evident in the earlier discussion on the Davis-Besse vessel head degradation, there are uncertainties in the existing models for predicting the degradation rates and failure modes of degradation mechanisms of passive components. This issue is of particular importance in RI-ISI applications that attempt to introduce such models for characterizing the failure rates of piping system components. Fortunately, the existing approach for risk-informed regulation includes a requirement to monitor the performance of SSCs whose requirements have been relaxed in the application. Obviously, there is an important feedback loop between this monitoring process and future enhancements to these degradation models that can reduce this modeling uncertainty.

4.8 Lack of Coherence Between Deterministic and Probabilistic Safety Approaches

In a number of the NRC staff interviews, it was pointed out that they were very actively engaged in an activity to address the lack of coherence between the probabilistic approach to safety analyses and the so-called deterministic approach to safety assessment. It was not evident from these particular interviews where this was heading or what conclusions might be reached from this effort, however the very fact that this lack of coherence was being addressed is viewed as a positive development. Insights gained from the review of specific risk-informed evaluations as described in Section 3 can be used to develop some observations about the current lack of coherence between these approaches and some suggestions as to what could be done to achieve greater coherence. The current lack of coherence between these approaches is evidenced by the following:

- Lack of consistency between the accident sequences considered: predefined design basis accidents limited to single failures in active safety systems for deterministic evaluations vs. a systematic enumeration of accident sequences with all logical combinations of failures and successes of safety and non-safety systems in PRAs.
- Different approaches to treatment of uncertainties: subjective application of safety margins, conservative assumptions, and invocation of “defense-in-depth” arguments vs. attempts to quantify uncertainties in assignment of accident frequencies and consequences.
- Questionable effectiveness of deterministic evaluations of defense-in-depth and safety margins in Regulatory Guide 1.174 applications.
- Different uses and perhaps different definitions of the concept of “risk”: Vague references to ensuring that there is “no undue risk to public safety” in justifying decisions made in the deterministic arena vs. use of an accepted quantitative definition of risk such as that of Reference [17].

Lack of Common Accident Scenarios

This contribution to lack of coherence was evident in the Callaway steam generator example and came up in a different way in the Davis-Besse head degradation issue. Callaway had requested approval of a new way to repair its steam generators on the basis that the current regulatory requirements (tied to the design basis accidents) were met. The staff decided to evaluate the capabilities of this repair strategy to cope with conditions that could only occur during beyond design basis core damage accidents. The different conclusions that were reached in these evaluation can be largely attributed to the use of different scenarios to base the evaluations. In the Davis-Besse instance, both the deterministic and probabilistic evaluations were based on an incomplete representation of scenarios and hence the possibility for a scenario more severe than a medium size LOCA was erroneously dismissed. The inability of the design basis accident to fully capture the safety significance of an issue is central to this lack of coherence issue. To achieve coherence on this point will require that we adopt the more complete representation of scenarios that are provided in a state-of-the-art and full scope PRA. A particularly archaic relic that should be re-examined is the single failure criterion that helps define the current design basis accident envelope.

Lack of Common Approach to Treating Uncertainties

Both the deterministic and probabilistic evaluations of the Alloy 600 cracking issue made by both the industry and the NRC staff exhibited an inability to address or fully appreciate uncertainties in our ability to predict the consequences of primary water stress corrosion cracking phenomena in CRDM nozzles and hence could not fully connect the dots between a crack propagation phenomena and significant wastage of reactor pressure vessel material. Another source of uncertainty that was inadequately addressed in these evaluations was the capability of existing visual exams to determine the state of the reactor coolant pressure boundary under piles of boric acid crystals.

The use of the term “deterministic evaluation” to describe the method of performing safety evaluations that preceded the application of PRA reflects a kind of arrogance about the capability of the regulatory process to address uncertainties. It is ironic that a concern about uncertainties is often cited by critics of the movement to utilize PRA results to guide decisions, yet uncertainties are only addressed implicitly in the traditional approach to safety assessment. The conclusions of any so-called deterministic safety evaluation are subject to the same sources of uncertainty that are available to support or challenge the results of a PRA as was clearly revealed in the Davis-Besse incident. It seems that a greater awareness of epistemic uncertainties associated with deterministic evaluations would contribute to a greater degree of coherence.

Lack of Clear Definition of Defense-in-depth and Safety Margin

This issue arises from the fact that any successful RG 1.174 application is required to demonstrate that the requested change meets not only probabilistic risk criteria for judging the acceptable risk impacts but also the continued adherence to the principles of defense in depth and maintenance of safety margins. No doubt all successful submittals under RG 1.74 made claims that defense in depth and safety margins were maintained but it is questionable whether these aspects of the evaluation are particularly meaningful. The problem is that there are no clear criteria available to judge the sufficiency or adequacy of a given safety margin or application of defense in depth. Some of the issues with current definitions of defense-in-depth are discussed in Reference [18]. While one can cite examples where safety margins have been applied in the development of existing regulatory requirements, it is not very clear how to predict whether any change to an existing safety margin is to be considered acceptable. The author is unaware of any risk-informed initiatives that were not approved or seriously questioned because of inadequate treatment of defense in depth or safety margins. The author agrees that these are important principles but they need to be more clearly defined so that any two analysts are likely to get the same result in evaluating their adequacy in a risk-informed or deterministic evaluation.

Lack of a Consistent Definition of Risk

This last issue of coherence points to the need for a consistent definition of risk that can be applied in both deterministic and probabilistic safety evaluations of regulatory issues. An impediment to the current level of incoherence is the lack of a clear definition of risk in the current regulatory requirements. Greater coherence will require that we adopt a common definition of risk for all aspects of safety assessment. The definition of risk proposed by Kaplan and Garrick [17] is widely accepted in the PRA community and is the basis for the risk definition in the ASME PRA standard [19]. The author can identify no reason why this same definition or similar cannot be adopted for deterministic evaluation purposes. Use of a consistent definition of risk would also facilitate the use of a common set of scenarios to use in safety evaluations.

4.9 Impact of Peer Review Follow-up and the PRA Standards

The technical issues discussed in this section are based on results of previous risk-informed evaluations and industry PRA peer reviews. As each of these issues has appeared in one or more safety evaluation reports and PRA peer review Fact and Observations, there is a high expectation that many of the issues will be resolved to varying degrees as plants incorporate changes to address these issues in ongoing and future PRA updates. In addition, the ASME and ANS standards will be available to support the near term PRA updates. It is unknown to what extent these issues will be resolved in future PRA updates and upgrades for risk-informed applications. In the next section, recommendations are presented for steps that can be taken by the industry and the NRC to address these issues to the extent needed to advance risk-informed decision making.

5. CONCLUSIONS AND RECOMMENDATIONS

5.1 Conclusions

Some of the conclusions derived in the preparation of this report are listed as follows:

5.1.1 Success Stories in Risk-informed Regulation

There are number of noteworthy success stories in risk-informed regulation thus far especially the development and application of Regulatory Guide 1.174 to guiding risk-informed changes to technical specifications and to in-service inspection requirements for piping systems. Several hundred licensing basis change requests under Regulatory Guide 1.174 have been submitted and successfully implemented. Many of these have led to improvements in safety while bringing about a better allocation of resources with respect to the contributions to severe accident risk. Safety improvements have resulted from a greater awareness of the risk impacts of design features and operational issues, and use of compensating measures to offset any small risk increases associated with the requested changes. In particular there seems to be a broad consensus on the following points:

- Implementation of Regulatory Guide 1.174 and associated application specific guides and standards has been a success for all stakeholders.
- The risk-informed oversight process has been a major improvement over the previous SALP program and has resulted in a much greater awareness of the risk significance of deviations in licensee and plant performance.
- Efforts to Implement Paragraph a(4) of the Maintenance Rule and associated industry initiatives to develop and apply configuration risk management tools have led to operational safety enhancements by evaluating risk impacts of plant configuration changes that are permitted by the technical specifications.
- There is a general consensus on the need to provide and assure the technical adequacy and quality of industry PRAs to support risk-informed regulation, however differences remain between the industry and the NRC about the adequacy and sufficiency of the NEI PRA Peer Review process and the ASME PRA standard to establish the minimum acceptance standards for PRAs.
- Although there are different perspectives on the extent of improvements that have been made in industry PRAs since the IPEs and IPEEs were submitted, there is broad agreement that the quality of existing industry PRAs has increased since those submittals.

5.1.2 Difficulties Encountered in Selected Risk-informed Evaluations

Insights from Interviews

The interviews with the NRC staff and nuclear industry representatives and the review of selected risk-informed submittals identified several areas of difficulty and a variety of technical issues that resulted from risk-informed initiatives to date. The major areas of difficulty included:

- Concerns about the adequacy of the scope, level of detail, and quality of existing PRAs. These concerns are discussed in greater detail below:
- Inadequate treatment of uncertainties in existing PRA and in reviews of risk-informed submittals

- Lack of a consensus between the industry and the NRC staff on the minimum capabilities of existing PRAs that are needed to support risk-informed decision making. This is in part reflected by the difficulties in getting agreement on the adequacy of the NEI PRA Peer Review program and the ASME PRA Standard.
- Lack of consistency in submittals and reviews under Regulatory Guide 1.174, particularly how the missing elements in the scope of PRAs are handled. The lack of well-defined acceptance criteria, to determine whether defense-in-depth and safety margins are adequately evaluated in these submittals, raises questions about the meaningfulness of these evaluations in the submittals and the reviews.
- The NRC received and provided partial approval to a risk-informed request to delay the vessel head inspection at Davis-Besse to address Alloy 600 cracking in the CRDM nozzles. With the benefit of hindsight following the discovery of significant vessel head corrosion during the subsequent inspection, this risk-informed evaluation as well as earlier deterministic evaluations of Alloy 600 cracking dating back to the 1980's depended on unverified modeling assumptions that turned out to be wrong. An important lesson from this event, discussed in greater detail below, is the fact that epistemic uncertainties are available to challenge both risk-informed evaluations and so-called deterministic safety evaluations.
- A very interesting development occurred during the NRC staff review of a request by Callaway to approve a new approach to repairing damage in Steam Generators on the basis that the approach meets all the existing deterministic requirements. The NRC staff decision to apply the principles of Regulatory Guide 1,174 during this review yielded risk insights that were important in the evaluation. This example challenges the notion that simply meeting the existing requirements is automatically justified.

Treatment of Uncertainties in Risk-informed and Deterministic Evaluations

All the decisions made in the regulatory process whether risk-informed or not are made in the face of uncertainties and within the boundaries of the state of knowledge of nuclear power plants and how they behave under both normal and accident conditions. A good quality PRA will utilize all of the relevant state of knowledge that supports the existing deterministic safety evaluations of a licensed facility. The validity of the PRA is obviously dependent on the validity of the supporting information. Both deterministic and probabilistic safety evaluations must deal with the same sources of uncertainties which are available to challenge and support the results and conclusions of the evaluations and the decisions that are made to license and regulate the licensed facility. The reliance on modeling assumptions to support the deterministic and probabilistic evaluations and regulatory decisions associated with the Alloy 600 nozzle cracking issue is a striking example of an epistemic uncertainty. This example demonstrates that so called deterministic safety evaluations are vulnerable to the same types of uncertainties that we seek to address quantitatively in a state-of-the-art PRA. Unfortunately, decisions that have been made as part of the so-called deterministic approach to safety have not been held to the same degree of accountability on the question of uncertainties as has been the case with PRA. The primary reason why the topic of uncertainties naturally comes up in PRA is that the fundamental safety questions, discussed in more detail below, naturally expose these uncertainties in attempting to answer the questions.

Fundamental Safety Questions

Both deterministic and probabilistic safety evaluations must deal with the same fundamental safety questions that are addressed in a PRA and provide the framework for the currently accepted definition of risk [17] but deal with them in different ways. These questions are: what can go wrong? (i.e., what are the relevant scenarios), what is the likelihood? (or simply, is it credible?), and what are the consequences? In the deterministic approach to safety, these questions were addressed historically by defining the design basis accidents and developing a set of ground rules for performing safety analyses for these accidents with

only qualitative evaluations to determine the likelihood. When new issues arise such as the Alloy 600 cracking issue, or plant changes that are requested within the existing regulatory requirements, there is a potential for introducing new accident sequences and for changing the frequencies of the previous set of accidents. Hence, even if the original safety analysis was successful in assuring adequate safety, which has not ever been demonstrated, this adequacy should be open to question whenever changes to the boundary conditions are introduced. Hence any change to our state of knowledge whether discovered in a PRA has the potential to challenge the basis of the existing regulations. It would seem that any deterministic or probabilistic evaluation of a change to a plant change or change to our state of knowledge about a plant would benefit by addressing this common set of fundamental safety questions. When changes are introduced, we need to understand if there are any new sequences introduced, what their frequencies and consequences are, and what are the changes in existing sequence frequencies and consequences. This should be fundamental to any deterministic and probabilistic evaluation. In addition, adoption of these fundamental questions would help achieve coherence between risk-informed and deterministic evaluations.

Significant Resources Invested to Achieve PRA Quality

The industry and the NRC have invested a large amount of resources in updating and upgrading the quality and scope of PRAs since the time of the IPE and IPEEE submittals. Most of the improvement to date can be attributed to efforts to apply the guidelines in RG 1.174 and as a result of the industry PRA peer reviews. There are expectations for continued improvements from the more recent efforts by ASME and ANS to develop PRA standards, however, this is still to be demonstrated. The industry peer review process, the PRA standards, and specific plants whose PRAs have been subjected to rigorous peer reviews such as STPNOC have provided a good description of some of the important attributes of a PRA that has sufficient quality to support risk-informed decisions such as RG 1.174 and Option 2 of SECY-98-300 applications.

Technical Issues to Resolve for Future Risk-informed Decisions

A number of technical issues have been identified in PRA peer reviews and safety evaluations of technical issues and risk-informed submittals whose resolution would greatly enhance the capabilities of the current industry PRAs to support risk-informed decision making. The most important of these issues include :

- Use Of Limited Scope PRAs In RG 1.174 Applications
- Lack of Completeness Within the Specified Scope
- Model to Plant Fidelity Issues
- Lack of Uncertainty Treatment
- Quantification Issues
- Multi-Unit Site Issues
- Lack of Capability to Treat Aging Effects on Risk Results
- Risk Metric Issues
- Lack of Coherence Between Probabilistic and Deterministic Safety Approaches

The current lack of coherence between these approaches is evidenced by the following:

- Lack of consistency between the accident sequences considered: predefined design basis accidents limited to single failures in active safety systems for deterministic evaluations vs. a systematic enumeration of accident sequences with all logical combinations of failures and successes of safety and non-safety systems in PRAs.
- Different approaches to treatment of uncertainties: subjective application of safety margins, conservative assumptions, and invocation of “defense-in-depth”: arguments vs. attempts to quantify uncertainties in assignment of accident frequencies and consequences.
- Questionable effectiveness of deterministic evaluations of defense-in-depth and safety margins in Regulatory Guide 1.174 applications.

- Different uses and perhaps different definitions of the concept of “risk”: Vague references to ensuring that there is “no undue risk to public safety” in justifying decisions made in the deterministic arena vs. use of an accepted quantitative definition of risk such as that of Reference [17].

5.2 Recommendations

In order to advance the capability of PRAs to support risk-informed decision making and to develop a clean interface between deterministic and probabilistic safety evaluations, the following recommendations are made.

Updated PRA Procedures Guide

The technical issues in existing PRAs discussed in this report suggest that it is appropriate to consider an update of the PRA Procedures Guide in NUREG/CR-2300. That particular reference is cited because it was the last time the industry and the NRC collaborated to develop a comprehensive set of PRA procedures and guidance. The need for this is supported by the following points:

- There have been many developments in PRA technology since publication of the previous guide.
- The current resources including the standards and peer review process are by design lacking in guidance on how to perform the PRA tasks, yet specific examples have been described in this report where guidance is lacking, e.g. treatment of uncertainties, dependencies, HRA, quantification issues, etc.
- Existing guidance for specific issues such as HRA, CCF modeling, etc has not been correlated to the requirements in the standards and the issues found in the peer reviews.
- It is an appropriate time to attempt a collaborative effort on the development of PRA guidance rather than continue to trend toward parallel and uncoordinated efforts by the industry and the NRC. The updated procedure guide should also provide the guidance on how to maintain or to upgrade a PRA in order to meet the requirements of the PRA standards.

Uncertainty Analysis and Treatment in Decision Making Handbook

This could be a separate item or rolled into the scope of the recommended PRA Procedures Guide update. The idea would be to put together a handbook that includes procedures for performing uncertainty and sensitivity analysis using a practical set of examples and an example PRA model to work with. Both aleatory and epistemic uncertainties would be addressed in the quantification of distributions for risk metrics such as CDF, LERF, and CCDF curves. The handbook should address not only how uncertainties treated in the PRA but how they impact decision making with examples to show the pitfalls if uncertainties are inadequately addressed.

Guidelines for Deterministic Safety Evaluations

Regulatory Guide 1.174 was originally designed for voluntary risk-informed license amendments and relief requests and was recently augmented to give the staff the option to use it under so-called special circumstances. This regulatory guide addresses both the risk aspects and deterministic safety evaluation aspects of a risk-informed decision. A further development of the concepts in this regulatory guide could be devised that would provide guidance for a safety evaluation, be it labeled as risk-informed or not. Development of such a guide could include a more logical basis for addressing deterministic safety evaluation principles such as defense-in-depth, safety margins, and the reactor safety cornerstones as well as criteria for assessing whether these principles are adequately addressed in an evaluation. At the same time, the guide could provide a more integrated discussion of how epistemic and aleatory uncertainties and

the three fundamental safety questions are addressed in both risk-informed and deterministic safety evaluations. Such a guide could eventually replace the current problematic assumption that by simply meeting the existing regulations is necessarily adequate in non-risk-informed decisions.

Generic Estimates of Risk Contributors from Missing PRA Scope

A good application of the expert elicitation process that was performed to support NUREG-1150 would be to develop a generic set of CDF and LERF risk estimates for various elements of a PRA that would be organized in such a manner to provide surrogate risk estimates for parts of a PRA work scope that were missing in a given application. The first step would be to develop a template for a full scope PRA that would delineate the initial power levels and plant states, internal events, external events such as seismic events, internal hazards such as fires and floods, etc. The next step would be to develop plant categories that would capture the general level of protection against seismic events, fires, etc. as well as siting considerations that are judged to be responsible for variations in these risk contributors. The expert elicitation process would also provide a generic set of results and inputs for the internal events PRA which could be used as a reference for identifying risk insights about unique plant features.

PRA Validation Program

An important goal of further advancements in PRA is the achievement of a sufficient level of standardization so that variations in PRA results from plant to plant are dominated by physical differences in the plants and operational differences such that analysts driven variations are minimized. In the opinion of the author, we are far from meeting that goal as evidenced by the technical issues delineated in Table 4-1. More efforts are needed to validate PRAs by using the insights from programs such as the NRC Accident Precursor Program and other structured reviews of plant service experience.

Use of a Consistent Definition of Risk

An important issue of coherence points to the need for a consistent definition of risk that can be applied in both deterministic and probabilistic safety evaluations of regulatory issues. An impediment to the current level of incoherence is the lack of a clear definition of risk in the current regulatory requirements. Greater coherence will require that we adopt a common definition of risk for all aspects of safety assessment. The definition of risk proposed by Kaplan and Garrick [17] is widely accepted in the PRA community and is the basis for the risk definition in the ASME PRA standard [19]. The author can identify no reason why this same definition or similar cannot be adopted for deterministic evaluation purposes. Use of a consistent definition of risk would also facilitate the use of a common set of fundamental safety questions that are used to frame all deterministic and probabilistic evaluations.

6. References

1. U.S. Nuclear Regulatory Commission (USNRC), "Reactor Safety Study, An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants," WASH-1400 (NUREG-75/014), October 1975.
2. Lewis, H. W. et al., "Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission," NUREG/CR-0400, September 1978.
3. USNRC, "NRC Statement on Risk Assessment and the Reactor Safety Study Report (WASH-1400) in Light of the Risk Assessment Review Group Report," January 18, 1979.
4. USNRC, "PRA Procedures Guide, A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants," Final Report, Vol. 1-2, NUREG/CR-2300, January 1983.
5. USNRC, "Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants," NUREG-1150, December 1990.
6. USNRC, "Use of Probabilistic Risk Assessment Methods in Nuclear Regulatory Activities; Final Policy Statement," Federal Register, Vol. 60, No. 158, August 16, 1995.
7. USNRC, "An Approach for Using Probabilistic Risk Assessment in Risk-informed Decisions on Plant Specific Changes to the Current Licensing Basis," Regulatory Guide 1.174, June 1998.
8. Nuclear Energy Institute, "Probabilistic Risk Assessment (PRA) Peer Review Process Guidance", Nuclear Energy Institute Report, NEI-0002, 2000.
9. True D.E. et al., EPRI PRA Applications Guide, EPRI TR 105396 1995.
10. USNRC, "Degradation of the Davis-Besse Nuclear Power Station Reactor Pressure Vessel Head Lessons Learned Task Force Report", September 30, 2002.
- [11] Letter from First Energy to U.S. Nuclear Regulatory Commission, "Supplemental Information in Response to NRC Bulletin 2001-01, Circumferential Cracking of Reactor Pressure Vessel Head Penetration Nozzles", Docket 50-346, October 17, 2001.
12. SECY-99-199, "Electrosleeve Amendment Issued to Union Electric Company for Callaway Plant, Unit 1"
13. Nuclear Energy Institute, "10 CFR 50.69 SSC Categorization Guideline", NEI-00-04, Draft Revision C, June 2002.
14. ACRS Letter to Chairman Jackson on Treatment of Uncertainties vs. point values in the PRA related decision making process, December 16, 1997.
15. Fleming K. N. and J. Mitman, "A Quantitative Assessment of a Risk-Informed Inspection Strategy for BWR Weld Overlays", Proceedings of the 8th International Conference on Nuclear Engineering, Baltimore MD, April 2-6, 2000.
16. Poloski, J.P., et al., "Rates of Initiating Events at U.S. Nuclear Power Plants: 1987-1995," NUREG/CR-5750, INEL/EXT-98-00401, February 1999.

17. Kaplan S., and B.J. Garrick, " On the Quantitative Definition of Risk", Risk Analysis, Vol. 1, No. 1, 1981.
18. Fleming, K.N., and F. A. Silady , "A Risk-informed Framework for Defense in Depth for Advanced and Existing Reactors", Reliability Engineering and System Safety 78 pp. 205–225, 2002.
19. American Society of Mechanical Engineers, "Standard for Probabilistic Risk Assessment for Nuclear Power Plant Applications", ASME-RAS-2002, April 5, 2002.
20. Fleming, K.N., Developing Useful Insights and Avoiding Misleading Conclusions from Risk Importance Measures in PSA Applications", Proceedings of PSA '96, Park City Utah, September 29, 1996.