NUREG-1050
Draft Report

Probabilistic Risk Assessment (PRA): Status Report and Guidance for Regulatory Application

Draft Report for Comment

U.S. Nuclear Regulatory Commission

Office of Nuclear Regulatory Research

DO NOT MICROFILM COVER

NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 1717 H Street, N.W.
   Washington, DC 20555
2. The NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission,
   Washington, DC 20555
3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC Office of Inspection and Enforcement bulletins, circulars, information notices, inspection and investigation notices; Licensee Event Reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the NRC/GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, and NRC booklets and brochures. Also available are Regulatory Guides, NRC regulations in the *Code of Federal Regulations,* and *Nuclear Regulatory Commission Issuances.*

Documents available from the National Technical Information Service include NUREG series reports and technical reports prepared by other federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal and periodical articles, and transactions. *Federal Register* notices, federal and state legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, and are available there for reference use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

NUREG-1050
Draft Report

Probabilistic Risk Assessment (PRA): Status Report and Guidance for Regulatory Application

Draft Report for Comment

U.S. Nuclear Regulatory Commission

Office of Nuclear Regulatory Research

DO NOT MICROFILM COVER

NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 1717 H Street, N.W.
   Washington, DC 20555
2. The NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission,
   Washington, DC 20555
3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC Office of Inspection and Enforcement bulletins, circulars, information notices, inspection and investigation notices; Licensee Event Reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the NRC/GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, and NRC booklets and brochures. Also available are Regulatory Guides, NRC regulations in the *Code of Federal Regulations,* and *Nuclear Regulatory Commission Issuances.*

Documents available from the National Technical Information Service include NUREG series reports and technical reports prepared by other federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal and periodical articles, and transactions. *Federal Register* notices, federal and state legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, and are available there for reference use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

NUREG--1050 (DRAFT)

DE84 900605

Probabilistic Risk Assessment (PRA): Status Report and Guidance for Regulatory Application

Draft Report for Comment

Manuscript Completed: January 1984
Date Published: February 1984

Division of Risk Analysis
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555

NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 1717 H Street, N.W.
   Washington, DC 20555
2. The NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission,
   Washington, DC 20555
3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC Office of Inspection and Enforcement bulletins, circulars, information notices, inspection and investigation notices; Licensee Event Reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the NRC/GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, and NRC booklets and brochures. Also available are Regulatory Guides, NRC regulations in the *Code of Federal Regulations,* and *Nuclear Regulatory Commission Issuances.*

Documents available from the National Technical Information Service include NUREG series reports and technical reports prepared by other federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal and periodical articles, and transactions. *Federal Register* notices, federal and state legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, and are available there for reference use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

GPO Printed copy price: _____

ABSTRACT

This document describes the current status of the methodologies used in probabilistic risk assessment (PRA) and provides guidance for the application of the results of PRAs to the nuclear reactor regulatory process. The PRA studies that have been completed or are underway are reviewed. The levels of maturity of the methodologies used in a PRA are discussed. Insights derived from PRAs are listed. The potential uses of PRA results for regulatory purposes are discussed.

CONTENTS

CONTENTS (Continued)

CONTENTS (Continued)

ILLUSTRATIONS

LIST OF TABLES

PREFACE

The Nuclear Regulatory Commission (NRC) is faced with many types of decisions in discharging its legal responsibilities for the regulation of nuclear power plants. These may be categorized as follows:

1. How safe should plants be?
2. How safe are they?
3. Does the safety of plants need to be improved?
4. How should the desired level of safety be ensured during the lifetime of the plant?
5. What issues require research to improve the state of knowledge and enhance effective regulation?

The first question involves sociopolitical considerations. In the past, safety levels have been qualitatively based on judgment. The safety levels would also be affected by any safety goals that might be implemented by the NRC in the future. This document does not address the first question.

The central aim of this document is to evaluate the level of development of probabilistic risk assessment (PRA) to determine how this analytical tool should be used in regulation as an aid to answering questions two through four, as well as to assess the likelihood that more research will improve the usefulness of PRA.

The probabilistic methods used in PRA cover a wide range of technical disciplines, from statistics to human-behavior sciences. Deciding how PRA should be used by the NRC in deciding regulatory issues requires an understanding of the existing information base and a knowledge of the methods used in performing a PRA. Therefore, this document provides

an overview of the level of maturity of PRA, the uncertainties in PRA that confront the regulatory decisionmaker, and the research under way to improve the methods, reduce the uncertainties, and allow more effective decisionmaking in the face of remaining uncertainties.

Historically, safety questions have been answered using conservative deterministic techniques, and safety systems have relied on defense in depth. Much of the conservatism arises from a healthy regulatory caution generated by the uncertainty associated with the current knowledge of phenomenology and of plant response to accidents and transients. PRAs generate many insights to aid the decisionmaker, which derive from a realistic integral view of plant design and operation. The PRAs suffer from the same substantial uncertainties as do deterministic analyses, but they attempt to address them more explicitly, add discipline to the evaluation of the operation of a plant, and result in a more complete understanding of risk-important systems and functions, interactions among systems, and the importance of human actions.

Uncertainties must be considered carefully before a decision is reached. The fact that PRAs provide a mechanism to display areas of uncertainty (more so than do conventional deterministic analyses) is actually a strength of PRA rather than a weakness. The weakness that must be guarded against is the tendency to take the PRA point estimates as a given. One of the principal advantages of PRA is the potential for providing an additional qualitative and quantitative perspective of the overall importance of uncertainties. Proper consideration of these uncertainties can enhance engineering judgment. This report attempts to provide general guidance for the use of PRA in regulatory decisionmaking.

PRA has become a widely used discipline in regulation practiced by both the NRC and the nuclear industry, in the regulatory arena, and touches a wide range of issues and decisions. The Reactor Safety Study (RSS) provided important insights, including the fact that small-break loss of coolant accidents (LOCA) and transients, rather than the large LOCA, are now estimated to be the principal contributors to risk.

The growing library of PRAs provides a rich information base of risk and reliability insights that are relevant to the NRC mission, but these insights have not been used in a comprehensive way. Therefore, this document distills this information and provides an overall perspective of the insights that PRAs have provided in the past.

This document is timely. It marks the end of a decade since the Reactor Safety Study (WASH-1400) was published and comes at the time when the pressures are great to increase the use of PRAs in regulation, e.g., their use in most unresolved safety issues, the assignment of priorities to generic safety issues, and the consideration of the broad severe core damage issue. Thus, now is the proper time to pause and delineate carefully the role that the assessment of risk and reliability should play in the evaluation of reactor safety and in regulatory decisions.

ACKNOWLEDGMENTS

Preparing this reference document on probabilistic risk assessment (PRA) has been a challenge. The state of the art of PRA, insights gained from the results of PRAs to date, and the appropriate uses of PRA in the regulation of nuclear power are complex and controversial subjects. It is to the credit of all who participated in writing this document that substantive concensus was achieved.

Princpal credit goes to Joseph A. Murphy, whose technical guidance and long hours of dedicated labor made this document possible. However, essentially equal credit must go to the other contributing writers, as follows:

David C. Aldrich, Sandia National Laboratories
Allan S. Benjamin, Sandia National Laboratories
Robert J. Budnitz, Future Resources Association
Jack W. Hickman, Sandia National Laboratories
Vojin Joksimovich, NUS Corporation
William E. Vesely, Battelle-Columbus Laboratory
Jonathan Young, Energy, Inc.

Please note that the word "authors" is not used. The effort was truly cooperative, with substantial criticisms and rewrites of each individual's contributions resulting in an amalgamation of knowledge and viewpoints and a higher quality document than would have otherwise been possible.

Credit also goes to Richard Denning, Batelle-Columbus Laboratory and Dan Alpert, Sandia National Laboratories for their valuable input in the area of accident phenomena and consequence analysis.

Last, but certainly not least, credit must be given to Pamela S. Bloomfield, Pamela L. Foust, and Sheri Hichar for their untiring effort in preparing drafts on a very short time schedule; to Ausra M. Richards for her competent performance as technical editor; and to the many people in the NRC who read and commented upon the initial drafts of this document. Without their work, the effort would have gone for naught.

This document will now undergo peer review through the process of public comment and an NRC workshop to be held on February 22-23, 1984. Chairmen of the two panels of the workshop will be Dr. Herbert J. C. Kouts, Brookhaven National Laboratory, and Dr. Norman C. Rasmussen, Massuchesetts Institute of Technology. Panelists will be knowledgeable persons in government, the nuclear industry, and academia, as well as consultants and contractors expert in the performance of PRAs. It is anticipated that this peer review process will enhance the quality and utility of the final document to be published in May. The final step in the peer review process will be an independent review of the final document by the National Science Foundation, to start in the late spring of 1984.

Malcolm L. Ernst, Acting Director
Division of Risk Analysis
Office of Nuclear Regulatory Research

REFERENCES

The background references for the material contained in this draft are incomplete. We have indicated by an asterisk where we intend to make additional references. Complete reference information will be incorporated in the final report.

PROBABILISTIC RISK ASSESSMENT:
STATUS REPORT AND GUIDANCE FOR REGULATORY APPLICATION

1. INTRODUCTION

In the plan to evaluate the NRC's Safety Goal Policy Statement (issued for comment NUREG-0880, Revision 1, dated May 1983), the Office of Nuclear Regulatory Research was directed "to collect available information on PRA studies and prepare a reference document that describes the current status of knowledge concerning the risks of plants licensed in the U.S. It is essential that a reference document be prepared and receive peer review so that the staff, licensees, and public have a common base of information on the dominant contributors to the probability of core melt and to the public risk due to radiation from serious nuclear accidents, the strengths and weaknesses of current plant designs and operations, and the usefulness of PRA and the safety goals in assessing such strengths and weaknesses." This report, presenting the current state of the art of PRA and guidance for its potential uses in the regulatory process, has been prepared in response to that directive.

This document discusses the purpose and content of a PRA and identifies the PRAs, and many other probabilistic studies, performed to date (Chapter 3). It then discusses the level of maturity of, and the uncertainties associated with, the various elements of PRA methodology (Chapter 4). Chapter 5 discusses the results obtained from PRAs to date; the generic insights that can be derived from the studies of dominant accident sequences and the systems, functions, and human actions found to be important from the perspective of core damage or risk; and insights relative to areas amenable to improvement and to means for preventing a degradation of plant safety with time. The final portion of the document

(Chapter 6) discusses potential uses of PRA in regulation, whether or not used in conjunction with safety goals, and presents important considerations in using the results of PRAs in decisionmaking. Chapter 2 provides a summary of the document, including all of the important findings.

The three appendixes provide extended coverage of the material contained in Chapters 3 through 5. These appendixes are written in jargon familiar to the PRA practitioner and are designed to provide technical credibility to the document. To improve the readability of Chapters 3, 4, and 5, most of the detailed references were omitted, but they are provided in the appendixes.

The state of knowledge necessary for performing certain steps of a PRA is rapidly evolving. Excellent progress has been made in the understanding of (1) the phenomenology associated with severe core-damage accidents, including accident progression; (2) behavior in the fuel, the reactor-coolant system, and the containment; and (3) the performance of the containment under the varied temperature and pressure conditions that can occur in severe core damage accidents. Research in these areas is being conducted by both government and industry, and it is reasonable to expect better understanding of the phenomena in the future. As some of the uncertainties are narrowed and estimates are improved, the insights and recommendations provided herein may also change. Thus, updates of the information presented here may be desirable from time to time as the state of the art progresses.

2. SUMMARY

2.1 Introduction

This chapter provides a summary aimed principally at those decisionmakers who do not have the time to study Chapters 3 through 6 in detail. Some descriptive narrative is provided, but much of the summary is in the form of listings of the more important findings. The reader is referred to the individual chapters for more detailed findings and supporting rationale and to the appendixes for a fuller understanding of the technical bases.

2.2 What Is PRA and How Can It Be Used?

PRA is an analysis that identifies and delineates the combinations of events that, if they occur, will lead to a severe accident (i.e., core melt), estimates the frequency of occurrence for each combination, and then estimates the consequences. As practiced in the field of nuclear power, PRAs focus on core-damage accidents, since they pose the greatest potential risk to the public.

The PRA integrates into a uniform methodology the relevant information about plant design, operating practices, operating history, component reliability, human reliability, the physical progression of core-melt accidents, and potential environmental and health effects in a realistic manner. It uses both logic models and physical models. The logic models depict the combinations of events that could result in a core-damage accident and can be used to determine the frequencies associated with each combination. The physical models depict the progression of the resulting accidents and the damage. For example, the combinations of events that can lead to LOCA, and the probabilities that these combinations will occur, are identified by a logic model, while the analysis of containment response to the accident is based on a physical model. The risk associated with any type of accident is the combination

(the product) of the frequency of occurrence and the resulting damage. The information extracted from a PRA in the form of predicted frequency of occurrence, resulting damage, and risk provides quantitative and qualitative insights into the aspects of plant design and operation that are the most significant contributors to risk.

The public health effects and economic losses resulting from a core-damage accident, which may also involve the release of radionuclides into the environment, can be assessed by means of environmental transport, protective action response, and consequence models. The environmental transport models use site-specific data to predict the spread and fallout of the released radionuclides. The consequence models use local demographic data to predict the health effects expected to occur in the surrounding population. Throughout the analysis, realistic assumptions and criteria are used. When information is lacking or controversy exists, the individual analysts may introduce conservatisms, increase uncertainties, or evaluate bounds, but the goal of the PRA is to produce an analysis that is as realistic as possible. An integral part of the risk-assessment process should be an uncertainty analysis, which includes not only uncertainties in the data but also uncertainties arising from modeling assumptions.

A number of studies of varying scope have been completed already . Almost a dozen studies have assessed core-damage sequences, and some have evaluated the containment response as well. Another dozen studies have gone further and assessed the public risk . Many other PRA studies of much narrower scope have been performed. For example, several years ago, the NRC studied the reliability characteristics of all auxiliary-feedwater systems, using a simplified, prescriptive analytical approach. Also, PRA techniques have been used to study specific accident sequences, such as anticipated transients without scram (ATWS).

Athough the purposes of these assessments varied considerably, each study had one or more of the following objectives or end uses in mind:

1. Identification and assessment of dominant contributors to core damage or risk.
2. Assessment of the plant-specific importance of TMI-related requirements and issues.
3. Assessment of risks at sites with high population densities.
4. Assessment of specific generic safety issues.
5. Training of plant personnel.
6. Development and integration of PRA methodology.
7. Training in the performance of PRAs.
8. Assignment of priorities in the use of resources.
9. Assessment of operating experience and events.
10. Improvement of operating, testing, and maintenance procedures.
11. Development of technical information to support recommendations on siting criteria.
12. Evaluation of emergency-response procedures.

2.3 Maturity of PRA

The level of maturity in the various elements of a PRA study depends, to a large degree, on the desired end uses of the analyses. It also depends on the nature and degree of uncertainties in the results and the degree of realism in the models, because both must be reasonably appropriate for the desired application. Uncertainties arise because the available data are insufficient to allow some parameter to be characterized with the desired precision, there is no consensus in the technical community on the issue, or the facts are simply unknown. One must also recognize that the realism of a model may be decreased by the introduction of conservative estimates as a substitute for unknown information or merely for the purpose of simplifying the model.

2.3.1 Plant Modeling and Model Evaluation

The methods currently in use (event and fault trees) are basically the same as those used in the Reactor Safety Study (RSS), although refinements have been made to improve the scope and depth of modeling. This aspect of PRA is generally considered to be mature, except that significant improvements are still expected in the qualitative and quantitative treatment of common cause failures. The major limitations are the following:

1. Completeness. Some types of events (e.g., sabotage) are explicitly excluded from present-day PRAs, principally because of difficulty in quantifying the initiating event. Also, certain events that were not identified in the model might occur. However, considering the variety of existing PRAs, the fact that a substantial base of operating experience is available for analysis and the fact that the understanding of applicable physical processes and system characteristics is believed to be quite complete, completeness does not appear to be the principal limitation of a PRA.

2. Representativeness. The degree to which plant models represent plant behavior is a problem that cannot be definitively determined at this time. However, it is believed that current models contain a conservative bias that is intentionally inserted by the analysts when phenomena are poorly understood. For example, success-failure criteria are often taken from information in the Final Safety Analysis Report, which has a strong conservative bias.

3. Validity. Many elements of a PRA can be validated through the use of operating or experimental data or through reliable analysis. However, the validation of the frequency

of rare events depicted in system modeling is not subject to experimental validation.

2.3.2 Data

Since the RSS, data on initiating events has improved, but in general the generic data base has not changed much and only limited causal data is available. Some PRAs use extensive studies of plant-specific data to augment the generic data base. However, there is no standard guidance for the use of generic versus plant-specific data. In general, the data base and treatment of data can be considered to be reasonably mature. To improve the data base significantly would require a substantial effort that would have to be supported by the industry. The uncertainty in the data base is random and varies with the type of data.

In general, for most PRA estimates of core-melt frequency or risk from internal initiators, the upper and lower bounds (approximately the 95th and 5th percentile values) vary by factors of 6 to 10 around the central estimate.

2.3.3 Human Errors

Past experience has shown that human actions can be important in the initiation of accident sequences, can cause failures of systems or functions given a random initiating event, or conversely can rectify or mitigate an accident sequence once initiated (recovery). The current methodology is reasonably mature, except for the treatment of cognitive errors. The methodology is a refined and formalized version of that used in RSS, and its results are more replicable. However, an empirical data base is not now available; this could be rectified by additional research. One of the biggest areas of uncertainty is the proper diagnosis of accidents, and substantial research is under way to improve diagnosis capability

using observed events. In general, the uncertainties associated with procedural human errors are somewhat larger than those associated with data, but both are the same order of magnitude.

2.3.4 Accident Processes and Source Terms

The uncertainties surrounding the estimation of source terms is not random in nature but arise basically from a lack of knowledge. The following phenomena are all sources of significant uncertainty: (1) core damage, (2) in-vessel and ex-vessel core melt, (3) in-vessel and in-containment fission product transport, and (4) temperature and pressure threats to containment integrity. All these phenomena are being studied extensively at the present time. The best indications are that the methods used in RSS for estimating source terms are conservative and that the uncertainties are larger than those associated with modeling, data, and human error. The source term increases dramatically if the containment fails, especially if the failure occurs early. If containment integrity is maintained for several hours after core melt, then natural and engineered mechanisms (e.g., deposition, condensation, filtration) can significantly reduce the quantity and radioactivity of the aerosols released to the atmosphere. Assessment of source terms will be more mature within the next year or two.

2.3.5 Offsite Consequence Analysis

The estimation of offsite radiological consequences is relatively mature. Since RSS improvements have been made in modeling capabilities, model evaluation studies have been performed, and models have been applied to provide guidance in areas such as emergency planning and reactor siting. Uncertainties in consequence estimates remain large, however, and stem principally from uncertainties in (1) the magnitude of the source term which influences all consequences, (2) the form and effectiveness of emergency response which can make a

large difference in predicted early health effects, (3) the dose-response relationships for somatic and genetic effects, and (4) the modeling of deposition process, including the possibility of condensation and rainout of moisture in the released plume.

2.3.6 Accidents from External Initiators

External initiators include seismic events, fires and floods inside the plant, external floods, high winds, aircraft, barge, and ship collisions, noxious or explosive gases offsite, etc. These are in contrast to "internal accident initiators" which are caused by active or passive plant equipment failures, operator errors, and/or loss of offsite power.

The ability to consider external initiators has undergone major advances since publication of RSS. Much developmental work is presently in progress, and much still needs to be done. In general, the uncertainties associated with the calculated risks of external initiators are much larger than those associated with internal initiators. The principal uncertainties lie with development of the hazards curve (i.e., the frequency of occurrence of an event exceeding a given magnitude, e.g., the likelihood of a seismic event exceeding an acceleration of 0.5 g). The methodology for assessment of seismic events, internal fires and floods, and high winds has reasonably matured for qualitative assessments but not for quantitative application. Therefore, little confidence should be placed in any estimates of the risk from external initiators compared to those from internal initiators.

The risks from other external initiators are generally considered to be low, because of either the very long recurrence time associated with the event or the NRC's deterministic treatment of these areas. However, additional research is

needed to develop screening criteria for selecting from these potential accident initiators those that might need to be considered in risk assessments.

2.4 Summary of Major Insights Regarding PRA Methodology

- In general, the methodology can be considered relatively mature for essentially all qualitative and quantitative applications of the results of systems analyses to the understanding of accidents resulting from internal initiators and, given a source term, to the analyses of consequences.
- PRA is not a mature technology if the desired application is the strict comparison of quantitative estimates with regulatory numerical criteria for the purpose of determining compliance.
- Quantitative estimates of core-damage frequency from internal initiators attributable to accidents have an uncertainty of perhaps an order of magnitude on either side of a point estimate. The uncertainties in risk estimates are larger than those drawn from the systems (core damage) analyses, because of source term uncertainties. Thus, conclusions drawn from core damage analyses are generally more robust than those drawn from quantitative risk analyses.
- Quantitative estimates of core-damage frequency or risk due to externally initiated accidents have significantly larger uncertainties than those due to internally initiated accidents and the two should not be compared with any confidence.
- Generally, the uncertainties are reduced when PRA results are used in a relative sense rather than in an absolute sense. Thus, one may expect a greater degree of confidence in identifying on a relative basis the dominant accident sequences due to internal initiators.

2.5 Core Damage and Risk Insights

Several important insights gained from the aggregate results of all PRAs performed to date are covered in the following subsections.

2.5.1 Broad Insights Regarding Core Damage and Risk

1. The estimated likelihood of accidents leading to core damage are generally higher than had been thought prior to the publication of the RSS.

2. The range of core damage frequency estimates in the current library of PRAs covers about two orders of magnitude (about 10^{-5} per year to 10^{-3} per year). Variability of results has been examined yielding the conclusion that it is possible to uncover general reasons for the variability that are attributable to plant design, operation, site characteristics, scope of the studies, PRA methods employed, and analytical assumptions postulated. However, plant-specific design and operational differences would make it difficult to predict with confidence the estimated core-melt frequency of a plant without performing a plant-specific PRA.

3. Most core-melt accidents do not lead to very large offsite consequences. A very wide range of potential consequences for core melt accidents seems to exist, depending on many factors. The fraction of core-melt accidents that might lead to large offsite consequences generally involves the early failure of containment in relation to the time of core melt (e.g., either before or just after core melt) or containment bypass.

4. Plants meeting all applicable NRC regulatory requirements have been found to vary significantly in terms of quantitative measures of risk and in terms of the specifics of the key accident sequences that dominate risk.

5. The results of the PRAs indicate that public health accidents beyond the design basis are the principal contributor to public health risk. This indicates that the designers, operators, and regulators have been generally effective in reducing the risks from anticipated operational occurrences and design basis accidents.

6. In the case of a low-likelihood accident involving a major offsite radiological release, PRA has provided important insights about the nature of offsite consequences by showing that:

   - Latent cancer risk is an important element of the offsite risk. Earlier thinking had been that prompt radiation-induced fatalities and offsite property contamination were the principal offsite concerns.
   - Estimated onsite economic damage generally is much larger than estimated offsite property damage, given a core melt.
   - Differences from site to site for estimated prompt fatality risks and offsite property risks are very great, but for latent cancer risks and onsite property damage risks they are not very great.

7. PRA has enabled the qualitative assessment of the importance of the various contributors to core damage frequency and risk. Among the important findings are:

   - Operational considerations are important to overall risk and may be comparable to the importance of design considerations. Human errors play an important role in overall reactor safety.
   - Containment performance is a key element in estimating overall risk to the public.

- Small LOCAs and transients are dominant accident and risk contributors in most PRAs; large LOCAs are usually not important contributors to overall risk.
- Earthquakes and internal fires seem to play an important role in risk, although this conclusion is very plant-specific and uncertain.
- At this time, the uncertainties in estimating the risk from external initiators are sufficiently large that comparisons with the risk from internal accident initiators are tenuous at best.
- Airborne radionuclide pathways are by far more important contributors to offsite risk than liquid pathways.

8. A study has been made of all licensee event reports (LERs) from 1969-1981 (the precursor study). The LERs that appeared to have some risk significance, given occurrence of the event, were evaluated using probabilistic techniques to estimate the likelihood of core melt. Insights from this study to date are not much different from those gleaned from existing PRAs, if the failures contributing to the TMI accident, the Browns Ferry fire, and the Crystal River transient are assumed to have been reasonably remedied and auxiliary feedwater improvements are assumed to have been made as required after the TMI accident.

9. While much attention is normally placed on dominant accident sequences and ways to reduce risk further, one of the most important insights gained from PRAs is the need to maintain the reliability of risk-important systems and components at or near the levels assumed in the PRA. Degradation of such systems or components can sharply increase risk or the chance of core melt. These systems

and components must be identified in order to focus attention on proper maintenance and surveillance testing procedures.

2.5.2 Insights Regarding Accident Sequences

1. Systems that are important for assuring reliable operation and preventing core damage accidents are not necessarily the same as those which are important in reducing offsite risk.

2. PRAs have revealed that a few groups of accident sequence types tend to dominate the risks in all plants studied. However, the dominant accident sequences can be expected to be different for different plants; the reasons for the dominance are plant specific and relate to design and operational differences. Different assumptions made by the PRA analysts also can explain some of the differences observed in PRA results.

3. Some of the key accident sequences are generic, while others constitute safety issues that are quite plant-specific. Many times the PRAs have been useful in suggesting cost-effective remedies.

4. Despite the plant-specific differences and resultant uncertainties in estimating accident sequence frequency, generic studies to support regulatory decisionmaking generally can be accomplished effectively by grouping the plants into classes which have similar accident sequences that dominate core-melt frequency or risk.

5. The failure of long-term heat removal is a large functional contributor to core melt frequency for both PWRs and BWRs. It is mainly associated with LOCAs in PWRs and with transients in BWRs.

6. The accident sequences that appear to be dominant release contributors are either those that enable radioactivity to bypass the containment or those that result in containment failure before or shortly after core melt. This early failure may result from major common-cause initiators, such as unrecovered loss of offsite power, fires, or earthquakes, or it may result from combinations of system failures.

7. PRAs have indicated situations where system success criteria based on licensing considerations may be overly conservative for realistic severe accident estimations. Resolution of these questions may result in an overall estimation of lower risk from some accidents.

8. The dependency of multiple systems on a common service (e.g., pump cooling or room cooling) is a major contributor to accident sequences. However, these sequences generally include a long delay before radioactive material is released, providing the plant operator with the opportunity to recover from initial support system failures.

2.5.3 Additional Insights on External Initiators

1. The results of the analysis of external initiators seem highly plant specific. For seismic events, the specifics of one plant's PRA results do not seem to be transferable to another plant even though it may be a similar type. Although the specifics are different for fires, their general character is similar in that for important fires there is major involvement with cables or control area affecting multiple redundant safety systems. For flooding events, the results do not seem to be transferable from plant to plant.

2. For seismic events:

   - The significant seismic contributors to core damage or risk as identified to date are generally larger than the safe shutdown earthquake (SSE).
   - Most of the major contributors are the result of plant structures failing under seismic loads and then disrupting the operation of safety systems.
   - Local ground-subsoil geological conditions are important in all seismic analyses accomplished in PRAs to date.

3. Most of the fires that have been found to be important to risk are those whose likelihood and/or severity are substantially reduced by the new NRC regulatory approach now being implemented.

4. For high winds, metal-sided structures are more fragile than concrete structures or equipment and are more likely to fail and compromise overall plant safety.

5. Those external initiators involving the plant site, such as seismic, external flooding, and high winds, are generally accompanied by loss of offsite power, which contributes to associated system unavailabilities.

2.6 Uses of PRA in Regulation

2.6.1 Introduction

Clearly, PRA techniques generate useful information and insights regarding the design and operation of a nuclear power plant, which can be useful to the regulator in the decision process by providing an improved understanding of the full range of accident sequences and their relative importance. Equally as clearly, many limitations in our knowledge exist, which lead to uncertainty in deterministic analyses as well

as in the quantification of the risk associated with nuclear power plants; and these uncertainties cannot be resolved by a PRA.

The regulatory decisionmaker must evaluate each analysis, whether deterministic or probabilistic, and judge whether the assumptions and boundary conditions employed are sufficiently valid and the results sufficiently robust to justify its use in making regulatory decisions. No technical analyses, deterministic or probabilistic, are ever formally complete or completely certain. In most instances the uncertainties identified by PRAs are also inherent, but not identified, in the more deterministic analyses. Therefore, it is important that the decisionmaker understand all significant strengths and limitations so as to make more effective use of all available analyses, including the information contained in PRAs. There are many types of regulatory decisions, and the weight given to the quantitative PRA results should vary depending on the degree of precision necessary.

2.6.2 Allocation of Resources

Because of its integrated nature and greater reliance on realistic information, a PRA presents the best available information concerning the specific ways in which the critical safety functions at nuclear power plants can fail to be performed. This is true even though the models are incomplete and uncertainties are associated with quantification of the models. The PRA information should be used, where appropriate, to guide and focus a wide spectrum of activities designed to improve the state of knowledge regarding the safety of individual nuclear power plants, and the nuclear industry as a whole. The resources of the NRC, as well as those of the industry, are limited; application of PRA techniques or insights from previous studies provides one more useful tool to permit the decisionmaker to allocate these resources to

areas most likely to reduce risk or to define or limit the uncertainties.

The nature of the decisions necessary to allocate regulatory resources does not require great precision in PRA results. It is sufficient to place the research and the efforts directed toward resolving generic safety issues into broad categories of risk impact (e.g., high, medium, and low). The reasoning is that a potential safety issue would not be dismissed unless it were clearly of low risk. Thus, one or more completed PRA studies can be used for this categorization even though they do not fully represent the characteristics of some plants, provided the nature of these differences is reasonably understood and can be qualitatively evaluated. One of the most important benefits of the use of PRA to set priorities for effort is the documentation of a comprehensive and disciplined analysis of the safety issue, which enhances debate on the merits of specific aspects of the issue and reduces the reliance on more subjective judgments.

Information from PRAs should also be used appropriately to guide the overall direction of inspection and enforcement efforts. A catalog of information derived from PRAs indicates that certain surveillance tests and maintenance activities are significant contributors to plant risk and frequency of plant damage. If a generic risk profile is available, it can be used to identify critical surveillance testing and maintenance activities that have the potential, if not done properly, of significantly altering the predicted plant risk or severe core-damage frequency. Generation of such information for each class of operating plants should assist a reactor inspector in focusing inspection effort on the critical activities at each facility. In a similar manner, these generic

insights (available by reactor class) provide valuable information to both the licensee and the regulator in understanding, and allocating resources to correct, potentially significant operational occurrences at a plant, even if a plant-specific PRA is not available.

2.6.3 Generic Regulatory Applications

The insights derived from PRAs provide additional information to aid in rule-making and the development of regulatory guides and branch technical positions. Such activities would be aimed at either reducing risk or relaxing regulatory requirements that do not have a significant impact on risk. The catalog of plant system, component, and operational practices that have been found in the completed PRAs to have a significant impact on risk or core melt frequency can lead to the development of risk-significant insights for the generic classes of plants. It is possible, however, that the number of plant classes (or surrogates) derived in such a study may be large, since many of the risk-significant features of a plant occur in the balance-of-plant where there is less standardization of design.

The degree of detail necessary in establishing the generic classes depends on the nature of the decision being made. In general, the decisionmaker will not rely on small differences in numerical results and will temper the insights gained from PRAs with sound engineering judgment. Thus, sorting the reactor population into a large number of generic classes will not often be necessary. Many times the qualitative insights drawn from PRAs could be more important than the quantitative insights.

If the estimates from a plant-specific PRA of core damage frequency and risk are very low, regulatory requirements could be varied, however, clearcut ground rules would be needed. This would provide for relief from regulatory requirements

when variations in plant design clearly reduce the risk. Much information can be gained from limited studies of specific issues using simplified systems reliability analyses. While these limited studies are insufficient to predict accurately the absolute level of risk, they can identify problems relatively, as was done in the plant-specific studies of auxiliary feedwater systems (AFWS).

2.6.4 Plant-Specific Regulatory Applications

Verification that a given level of safety (or risk) is likely to be achieved is a reasonable use of PRA results. Thus, the "bottom-line" numerical results of a PRA are useful, provided they are generated with care, their uncertainties and biases are clearly understood, and they are used in conjunction with other conventional regulatory tools. The information presented in a PRA can be a useful tool for the direction of regulatory attention and resources. However, the quantitative results of a PRA cannot be used in a compliance versus noncompliance sense. The stated uncertainties in a plant-specific PRA, compounded by the inability to quantify modeling uncertainties in any but a subjective manner, make it very difficult to determine formally with any degree of confidence that a specific safety limit (in terms of public risk or frequency of core melting) is met.

A plant-specific PRA, performed early in the design process, can yield a tremendous amount of information about expected plant performance that would be useful to designers as they perform their detailed design. Such information would also be of use to the regulator. Because of the lack of specific design details in some areas, as well as the lack of plant-specific data, the results of such analyses cannot be considered a true prediction of plant risk or of the frequency of

core damage. Rather, such analyses generate useful information on potential weaknesses in the design and allow evaluations of the efficacy of corrective design modifications. A plant-specific risk study can be used to evaluate the importance of operating events and to assess the safety of the plant when equipment is not operable. Also, a catalog of accident sequences and their associated relative probabilities can be used to train emergency response personnel. This could lead, for example, to improving the criteria for the declaration of site or general emergencies and in developing guides for the diagnosis and prognosis of accidents. The models generated also provide the tools with which to reduce allowable outage times and surveillance intervals and can be used in evaluating the advisability of plant shutdown when equipment is out of service beyond the times allowed in current technical specifications.

After a plant-specific PRA has been performed, steps should be taken to monitor the performance of the plant to assure that the level of safety estimated in the study is maintained. The PRA should be a living document that is used and appropriately updated, rather than a completed project which sits on a shelf. The PRA should be used in the context of a safety or reliability assurance program to evaluate operational occurrences and to check the significance of operational data as they are acquired.

2.7 PRA and Regulatory Decisionmaking

The previous section provided recommendations for the uses of PRA in regulation as well as some general cautions. However, the role to be played by PRA in regulatory decisions is not clear. There can be no magic formula for decisionmaking, and the weight to be given to any type of information including PRAs will vary from case to case depending on the nature of

the issue, the results of the PRA, the nature of other information, and other considerations which could affect the overall judgment.

One of the more difficult problems facing the PRA analyst is the display of the results of the PRA, and the uncertainty and sensitivity analyses, in such a way as to communicate effectively to the decisionmaker. It would be impracticable to portray all possible permutations. Therefore, a selection must be made of those uncertainties most important to the decision, and the best style of presentation must be chosen to communicate this information. A more standardized format needs to be developed.

Some important factors that should be considered in determining the weight to be given to PRA results in any decision are:

1. Do the scope and depth of the PRA study reasonably match the needs of the decision?

2. Do results of peer reviews conducted on the study add or subtract from the strength of the results?

3. The qualitative insights from the study. For example, do the qualitative insights as to the nature of the dominant accident sequences appear reasonable from an operational or engineering sense? This includes an assessment of the degree of realism associated with the study.

4. The impact of alternative regulatory actions on the estimated risk, together with the ease and costs of implementation.

5. The magnitude of the quantitative estimates, as well as the results of sensitivity analyses and the bounds and likely biases of the major uncertainties surrounding the

point estimates. Where the reasonable upper bound of the PRA estimate indicates that the issue does not warrant regulatory attention, then substantial weight may be given to the quantitative PRA results. Similarly, the quantitative results may be given substantial weight in a decision to take regulatory action, if the lower bound estimate indicates a safety concern. Between these extremes the quantitative results cannot be the principal basis for making a decision, but the qualitative and quantitative results can provide unique perspectives and information to the decisionmaker on the integral performance of the plant.

One major concern exists with regard to the use of PRA results in decision making; too often the decisionmaker has the tendency either to go too quickly to the bottom line (which is the weakest part of a PRA), or to dismiss the PRA entirely as being too uncertain. Neither path is appropriate. Safety goals, or other types of numerical criteria, tend to drive the user to the bottom line in spite of all the cautions to the contrary. Therefore, such numerical criteria need to be constructed and implemented as to minimize this tendency. The decisionmaker must devote attention to the design and operational insights derived from the analyses.

The most important product obtained from PRA is the framework of engineering logic generated in constructing the models, not the precise numbers resulting from the mathematical manipulations of these models. The patterns, ranges, and relative behaviors which are obtained can be used to develop insights into the design and operation of a plant. These insights can only be gained from an integrated, consistent approach such as PRA. Therefore, the performance of the PRAs and the display of results and uncertainties must be constructed so that they provide convenient and scrutable stopping places, encourage

viewing these insights and understanding the underlying assumptions and uncertainties, and discourage undue fixation on the bottom line. Only then will regulation be able to draw fully upon the potential benefits of PRA as an information source and regulatory tool.

3. CURRENT STATUS OF PROBABILISTIC RISK ANALYSIS USAGE

3.1 Probabilistic Risk Assessment

In the United States the application of PRA techniques in the study of nuclear power safety started essentially with the NRC-sponsored Reactor Safety Study (NUREG/CR-75/014) and continued with the Reactor Safety Study Methodology Applications Program (RSSMAP) and a few limited applications in the mid-1970s. PRA was not widely used until after the accident at Three Mile Island (TMI). The TMI accident resulted in a realization that accidents other than design-basis accidents needed to be addressed more thoroughly in the regulatory process. Since transients, small LOCAs and human errors identified in the RSS as major contributors to risk and had contributed to the TMI accident, attention focused was on PRA. Furthermore, the President's Commission (Kemeny) and the NRC-sponsored Rogovin investigation (NUREG/CR-1250) strongly encouraged the use of PRA techniques in the regulation of nuclear power. Numerous plant-specific PRAs and generic studies using PRA techniques were undertaken as a result.

3.2 Description of PRA

The objective of PRA is to identify and delineate the combinations of events that, if they occur, could lead to undesirable public consequences and to estimate the magnitude of those consequences. The PRAs performed to evaluate nuclear reactor safety focus on core-damage and core-melt accidents because these accidents are expected to pose the greatest potential risk to public health and safety. Relevant information about plant design, operating practices, operating history, component reliability, human reliability, the physical progression of the accident, and potential environmental and health

effects is processed through various analytical models to obtain an estimate of plant safety. Both logic models depicting combinations of events that could result in core damage or core melt and physical models depicting the progression of accidents resulting from these combinations of events are used. The models are evaluated probabilistically to provide both qualitative and quantitative insights as to the level of risk and to identify the design, site, or operational characteristics that are the most important to risk.

3.2.1 PRA Study Process

A PRA is a multidisciplinary study involving a team of individuals with differing expertise. The major steps in the analysis are shown in Figure 3-1. The analysis involves developing a set of possible accident sequences and estimating their outcomes. To this end, several sets of models are developed and analyzed, depending on the scope and the objectives of the study. Among them are models related to plant systems, to the response of the containment, and to offsite consequences.

Plant-system models generally consist of event trees, which depict initiating events and combinations of system successes and failures, and fault trees, which depict ways in which the system failures represented in the event trees can occur. These models are analyzed to estimate the frequency of each accident sequence.

The containment models represent the events that occur during the accident but before the release of radioactive material from the containment. They cover the physical processes induced by each accident sequence in the core, in the reactor-coolant system, and in the containment as well as the transport and deposition of radionuclides inside the containment. The analysis examines the response of the containment to

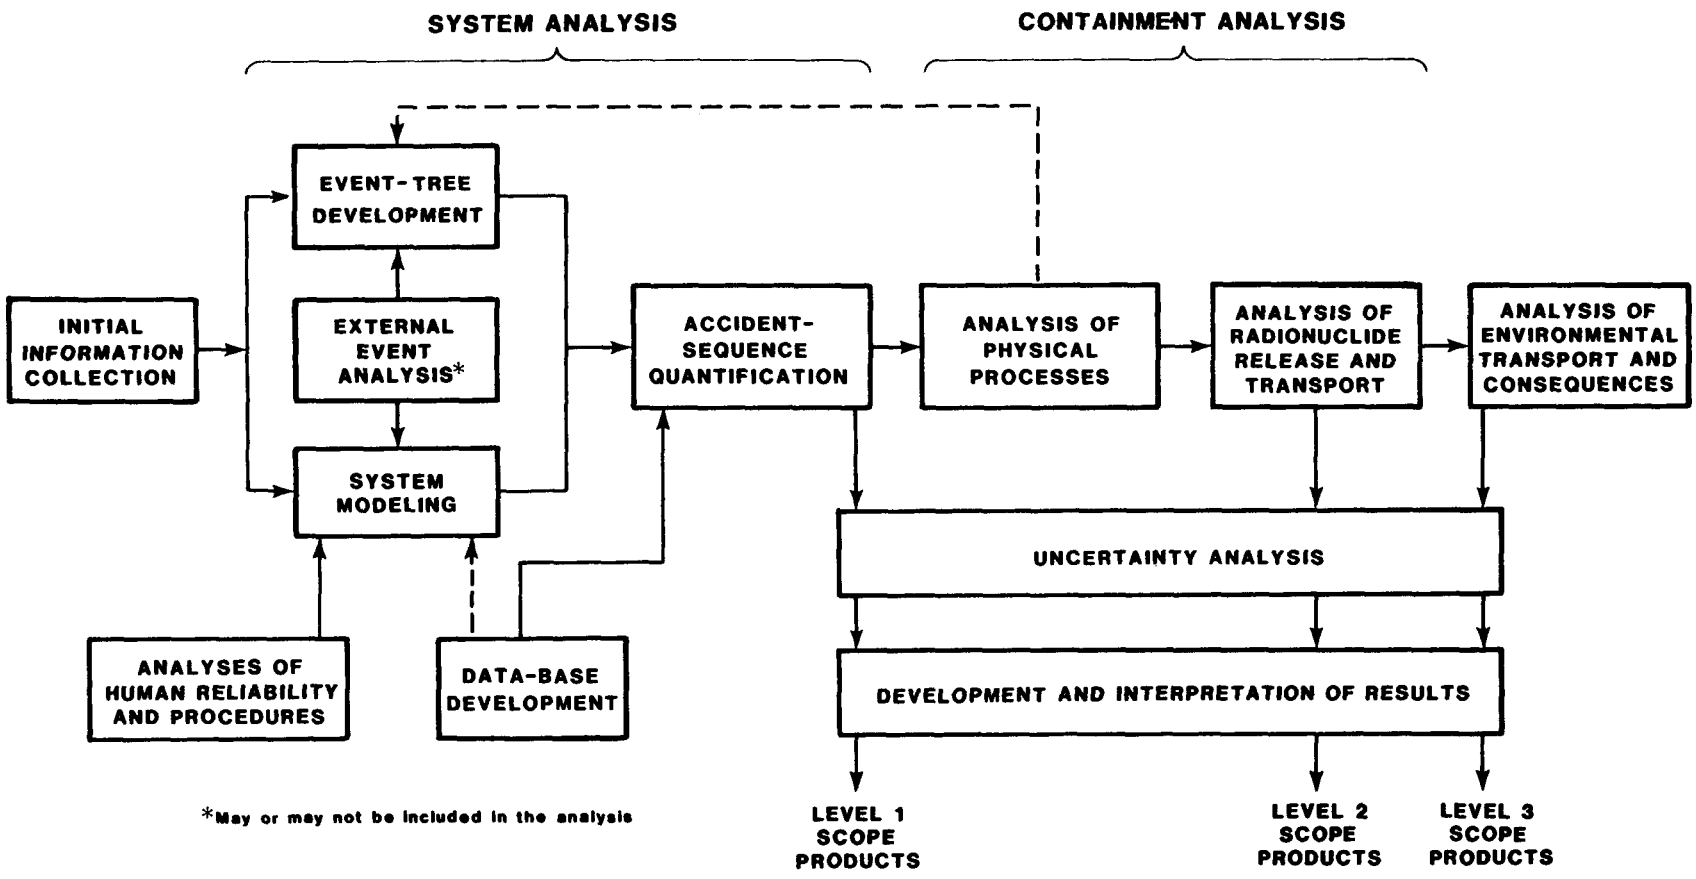

Figure 3-1. Risk Assessment Procedure

these processes, including possible failure modes, and evaluates the releases of radionuclides to the environment.

The offsite consequences of the accident in terms of public-health effects and economic losses is estimated by means of environmental transport and consequence models. These models use meteorological data (and sometimes topographic data as well) to assess the transport of radionuclides from the site. Local demographic data and health-effects models are then used to calculate the consequences to the surrounding population.

An integral part of the risk-assessment process is an uncertainty analysis. Uncertainties in the data and uncertainties arising from modeling assumptions are propagated through the analysis to estimate the uncertainties in the PRA results.

The results of the risk assessment are analyzed and interpreted to identify the plant features and operational practices that are the most significant contributors to the frequency of core melt and to risk. They can also be used to generate a variety of qualitative information on the events and failures associated with various consequences. Throughout the analysis, realistic assumptions and criteria should be used. When information is lacking or controversy exists, it may be necessary to introduce conservatisms or to evaluate bounds, but the goal of a PRA study is to perform an analysis that is as realistic as possible.

3.2.2 Levels of Scope in PRA Studies

The scope of PRA studies varies considerably, depending on the objective. The most common objective is the estimation of core-melt frequency. The PRA Procedures Guide (NUREG/CR-2300) termed this a Level-1 PRA, which consists of an assessment of plant design and operation, emphasizing sequences that could lead to a core melt. External events, such as floods or earthquakes, may or may not be included. The result

is a list of the most-probable core-melt sequences, their frequencies, and insights into their causes. Such a scope provides an assessment of plant safety and of the adequacy of plant design and operating procedures from the perspective of preventing core melt, but it does not permit an assessment of containment response or the public risk associated with the plant.

In addition to the analyses performed in a Level-1 PRA, a Level-2 PRA analyzes the physical accident phenomena, the response of the containment, and the release of radionuclides into the environment. This type of a study does not provide a full assessment of public risk, because offsite consequences are not assessed. It does, however, provide insights into risk by generating relative frequencies and source terms of various magnitudes (release categories).

A Level-3 PRA analyzes the transport of radionuclides in the environment and assesses the public-health and economic consequences of accidents in addition to performing the analyses of Level 1 and 2. A Level-3 study provides a full assessment of public risks.

The level of effort varies with the scope and the depth of the analysis. A summary of past experience is presented in Figure 3-2. Note that the largest variability in effort lies in a Level-2 analysis. It seems reasonable to expect that the efforts expended in this area will diminish significantly in the future as ongoing research on accident phenomena and source terms is completed.

3.3 PRA Studies Performed to Date

The major PRA studies completed for specific U.S. nuclear power plants are listed in Table 3-1 (Level-3 studies) and Table 3-2 (Level-1 and -2 studies). A number of studies are

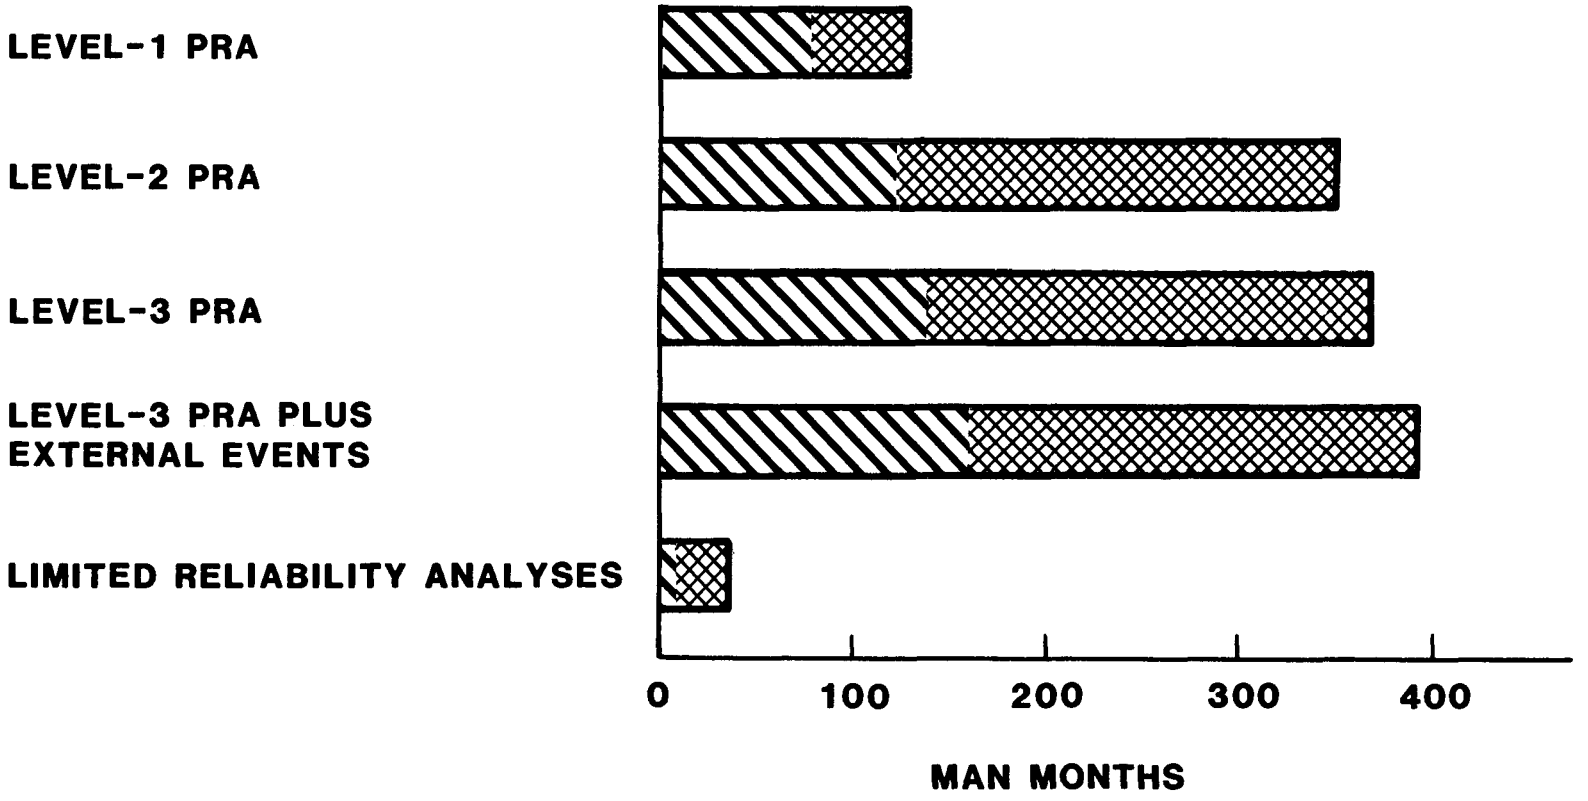

NOTE: DATA TAKEN FROM PRA PROCEDURES GUIDE (NUREG/CR-2300): LEVELS 1, 2 AND 3 ARE DEFINED IN THE PRA PROCEDURES GUIDE. CROSS-HATCHING INDICATES RANGE OF ESTIMATE.

Figure 3-2. Level of Effort Required

Table 3-1

Completed Full-Scope (Level-3) PRAs

Plant	Issuance	Operating License	Rating (MWe)	NSSS/AE[1]	Containment	Sponsor	Report
Surry 1	1975	1972	788	W/S&W	Dry-Cylinder	NRC	NUREG-75/014 (WASH-1400)
Peach Bottom 2	1975	1973	1065	GE/Bechtel	Mark I	NRC	NUREG-75/014 (WASH-1400)
Big Rock Point	1981	1962	71	GE/Bechtel	Dry-Sphere	Utility	USNRC Docket 55-155
Zion 1 & 2	1981	1973	1040	W/S&L	Dry-Cylinder	Utility	USNRC Docket 50-295
Indian Pt. 2 & 3	1982	1973	873	W/UE&C	Dry-Cylinder	Utility	USNRC Dockets 50-247 and 50-286
Yankee Rowe	1982	1960	175	W/S&W	Dry-Sphere	Utility	USNRC Docket 50-29
Limerick 1 & 2	1983	(1985)	1055	GE/Bechtel	Mark II	Utility	USNRC Docket 50-352
Shoreham	1983	(1984)	819	GE/S&W	Mark II	Utility	USNRC Dockets 50-322 and 50-353
Millstone 3*	1983	(1986)	1150	W/S&W	Dry-Cylinder	Utility	Controlled document
Susquehanna 1*	1983	1983	1050	GE/Bechtel	Mark II	Utility	Draft
Oconee 3*	1983	1973	860	B&W/Duke	Dry-Cylinder	EPRI/NSAC	Draft

*Completed but not yet publicly available.

[1]NSSS--Nuclear Steam System Supplier; AE--Architect-Engineer.

Table 3-2

Completed Level-1 and -2 PRAs

Plant	Issuance	Operating License	Rating (MWe)	NSSS/AE[1]	Containment	Sponsor (program)	Report
Oconee 3	1981	1973	860	B&W/Duke	Dry-Cylinder	NRC (RSSMAP)	NUREG/CR-1659
Sequoyah 1	1981	1981	1148	W/TVA	Ice Condenser	NRC (RSSMAP)	NUREG/CR-1659
Grand Gulf 1	1981	1982	1250	GE/Bechtel	Mark III	NRC (RSSMAP)	NUREG/CR-1659
Calvert Cliffs 1	1981	1974	845	CE/Bechtel	Dry-Cylinder	NRC (RSSMAP)	NUREG/CR-1659
Crystal River 3	1982	1976	797	B&W/Gilbert	Dry-Cylinder	NRC (IREP)	NUREG/CR-2515
Browns Ferry 1	1982	1973	1065	GE/TVA	Mark I	NRC (IREP)	NUREG/CR-2802
Arkansas 1	1982	1974	836	B&W/Bechtel	Dry-Cylinder	NRC (IREP)	NUREG/CR-2787
Millstone 1	1983	1970	652	GE/EBASCO	Mark I	NRC (IREP)	NUREG/CR-3085
Calvert Cliffs 2	1983	1974	845	CE/Bechtel	Dry-Cylinder	NRC (IREP)	Draft

[1]NSSS--Nuclear Steam System Supplier; AE--Architect-Engineer.

currently in progress and a number of reviews have been completed (see Sections 3.4 and 3.5).

3.3.1 Completed Level-3 Studies

As of late-1983, 12 full-scope (Level-3) PRA studies had been completed for U.S. light-water reactors (LWRs). These are listed in Table 3-1. Two of the plants listed in Table 3-1, Surry, Unit 1 and Peach Bottom, Unit 2, were analyzed in the RSS, using hypothetical composite sites for the estimate of offsite consequences. The objective of this NRC-sponsored study was to make a realistic quantitative estimate of the risks from commercial U.S. nuclear power plants and to compare them with other, nonnuclear, societal risks. As is now well known, the RSS broke much new ground in the development of quantitative risk-assessment techniques. All PRAs performed since the RSS have used same basic methodology, although they have incorporated significant improvements in some of the specific techniques.

After the TMI accident, the NRC staff examined the risk posed by a number of nuclear power plants, concentrating on those that are located close to major population centers. One of these plants was the Zion Station. The NRC staff assessment, which was of very limited scope, was performed by postulating that the PWR plant analyzed in the RSS was located at the Zion site. The NRC staff concluded that, if the RSS PWR plant were located at the site of the Zion Station, it might represent a large fraction of the total societal risk from all U.S. nuclear power plants then in operation because of nearby population densities. To obtain a comprehensive assessment of the safety of the Zion Station, Commonwealth Edison commissioned a much more thorough study (Zion).

Practically identical reasons motivated Consolidated Edison and the Power Authority of the State of New York (PASNY) to

commission studies for Indian Point Units 2 and 3, respectively (Indian Point 2, 3).

Like the Zion and Indian Point units, the Limerick plant and Millstone Unit 3 are located close to population centers. The PRA studies (Limerick 1981 and 1983, NUREG/CR-3085) for these plants were therefore requested by the NRC, which was concerned that the operation of those plants might also represent a disproportionately high level of societal risk. The Limerick and Millstone 3 plants were still under construction which allow greater flexibility to accommodate design changes, if necessary.

The motivation for the Big Rock Point PRA was somewhat different. Consumers Power Company was confronted with the potential need to implement a wide array of plant modifications arising from the Systematic Evaluation Program (SEP) and post-TMI requirements. Internal estimates of the costs associated with the implementation of these requirements ran as high as $125 million. Since Big Rock Point is a small plant (240 MWt) that is remotely sited, Consumers Power did not believe that the risk of operating the plant would be sufficiently large to warrant expending that amount to keep the plant in operation. A PRA study (Big Rock) was chosen as the tool by which current and future regulatory requirements could best be assessed. It was judged that a PRA would focus attention on those plant features that contribute most significantly to public risk and would identify cost-effective modifications that might be implemented voluntarily.

The remaining Level-3 PRAs were initiated by the nuclear industry. The studies sponsored by the Yankee Atomic Electric Company (Yankee Rowe), the Long Island Lighting Company (Shoreham), the Pennsylvania Power Light Company (Susquehanna) and the Nuclear Safety Analysis Center had similar objectives. In general, the sponsors wanted to estimate the risk of the

plant and to identify the plant characteristics that are the most important to risk.

3.3.2 Completed Level-1 and -2 PRAs

Four of the completed studies listed in Table 3-2 (Sequoyah, Oconee, Calvert Cliffs, and Grand Gulf) were performed under the auspices of the NRC's Reactor Safety Study Methodology Applications Program (RSSMAP). The objectives of this early program were to apply RSS methods to reactor and containment designs different from those studied at Surry and Peach Bottom, in order to determine the sensitivity of dominant accident sequences to plant design features. These limited-scope studies did not include external events or risk estimates and involved only a few man-years of effort. They can be viewed as limited-budget Level-2 studies.

The NRC's Interim Reliability Evaluation Program (IREP) has provided five PRA studies (Crystal River, Arkansas Nuclear One, Browns Ferry, Millstone, and Calvert Cliffs). IREP evaluations were limited in scope in that they were carried forward only to the point of estimating core-melt frequencies. They did not consider external events, their containment analyses were limited, and risk estimates were not included. The IREP studies can be viewed as expanded Level-1 studies.

Of the PRA studies performed overseas, the Sizewell B study, sponsored by England's Central Electricity Generating Board, can be viewed as a Level-2 study, while the Ringhals 2 study, sponsored by the Swedish State Power Board, was a Level-1 study.

3.4 PRA Studies Under Way

A list of PRA studies that are under way is provided in Table 3-3. These studies are all sponsored by utility companies or the Electric Power Research Institute and vary in scope, motivation, and level of completion. A number of PRA studies are

Table 3-3

PRA Studies Under Way

Plant	Start Date	Operating License	Rating (MWe)	NSSS/AE	Containment	Sponsor
Oyster Creek	1977	1969	620	GE/B&R	Mark I	Utility
Sequoyah 1	1979	1981	1148	W/TVA	Ice Condenser	EPRI/TVA
Browns Ferry 1	1980	1973	1065	GE/TVA	Mark I	Utility
Midland	1981	(1984)	852	B&W/Bechtel	Dry, cylinder	Utility
Seabrook	1982	(1984)	1150	W/UE&C	Dry, cylinder	Utility
McGuire 1	1982	1981	1180	W/Duke	Ice condenser	Utility
Bellefonte	1982	(1985)	1213	B&W/TVA	Dry, cylinder	Utility
TMI 1	1983	1974	792	B&W/Gilbert	Dry, cylinder	Utility

also being performed in other countries (e.g., Italy, Taiwan, Japan, Switzerland).

3.5 PRA Study Reviews

Virtually all the PRA studies completed to date have recognized the need for assurance of technical quality. The studies received both internal and external reviews. The internal reviews typically consisted of intradisciplinary and interdisciplinary reviews as well as those provided by a separate peer review group. Once a utility-sponsored study is submitted to the NRC, the NRC conducts an extensive review, which may include the participation of national laboratories and various highly specialized consultants. These reviews are necessarily restricted to a single PRA and are not generally intended to provide generic insights into nuclear risks and PRA methodology.

Several reviews have, however, been conducted with the objective of providing broader insights. They consist of (1) an NRC-sponsored review of the RSS by the Lewis Committee; (2) a review by the Government Accounting Office (GAO) of PRA techniques, with emphasis on safety assessments performed in the Indian Point studies; (3) an NRC-sponsored review of the Indian Point study; (4) an EPRI-sponsored review of five PRA studies; (5) the NRC-sponsored Accident Sequence Evaluation Program; and (6) the industry-sponsored Industry Degraded Core (IDCOR) program. These reviews are briefly described below; a more detailed statement of objectives and important conclusions are presented in Appendix A.

3.5.1 Reviews of Specific Studies

3.5.1.1 Reviews of RSS -- The NRC established an independent advisory group under Prof. Lewis of the University of California at Santa Barbara to assess the key aspects of the RSS (NUREG/CR-0400). The report concluded that the RSS was

"a substantial advance over previous attempts to estimate the risks of the nuclear option." The Lewis point criticized some of the analytical techniques used in the RSS and concluded that the uncertainty ranges on results were larger than stated in the report. However, it strongly recommended that PRA techniques be used to reexamine existing NRC regulations and practices to make them more rational and left little doubt that the PRA approach was extremely useful.

3.5.1.2 GAO Review of the Indian Point Study -- The GAO review was requested by the House of Representatives Committee on Energy and Commerce. The review (GAO) was divided into two phases, with phase I concentrating on PRA techniques as they apply to the Indian Point study and phase II to assessing the state of the art in PRA as well as NRC's use of risk assessment. The following paragraph is representative of phase I findings:

> "The Indian Point PRA is a comprehensive risk assessment which assesses plant systems performance, the ability of the plant to contain radioactivity, and the consequences of potential accidents. While many analysts consider the Indian Point PRA to be the state of the art in risk assessment, it suffers from the same fundamental problems as all PRAs: uncertainty and incomparability of results. Also, although the study identified the dominant contributors to risk, it did not identify the precise level of risk from operating the Indian Point nuclear power plants."

3.5.1.3 NRC-Sponsored Review of Indian Point Study (IPPSS) -- This review (NUREG/CR-2934) was commissioned by the NRC in preparation for the Atomic Safety and Licensing Board (ASLB) hearings on Indian Point. The findings of the review led to

revisions in more than half the IPPSS damage states. However, the revised damage-state frequency estimates were within a factor of 2 of the IPPSS estimate for all but two damage states.

3.5.1.4 NRC-Sponsored Review of Zion PRA -- The accident-sequence analysis portion of the PRA was reviewed and the results were published.* Among the review results were the identification of new sequences and analytical inadequacies which were corrected by the plant owners. The review pointed out the importance and uncertainty associated with support system success criteria, models for the recovery of offsite AC power, and the risk due to external events.

3.5.1.5 NRC-Sponsored Review of Limerick PRA -- In this published review of the Limerick PRA,* revised dominant accident sequences were identified and are providing a basis for reasonable safety improvements through voluntary modifications made by the plant owners. In addition, review results are being used in accident evaluation for draft environmental statements and as a source of information for generic considerations in the severe-accident generic decisionmaking.

3.5.1.6 NRC-Sponsored Review of Big Rock Point PRA -- The results of this review* were used to address the need for implementing several post-TMI requirements (NUREG-0737) and to evaluate alternative cost-effective modifi- cations for reducing public risk.

3.5.2 Reviews of Multiple Studies

3.5.2.1 EPRI Review of Five PRA Studies -- Although a knowledge of the results, insights, applications, and efficacy of various PRA methods could be of significant value to nuclear

utilities and regulators, PRA studies usually produce multi-volume reports that are difficult to comprehend and assess without extensive and dedicated scrutiny. Realizing that there is considerable interest in, and controversy about, results and their validity, the Electric Power Research Institute initiated in 1982 a review and comparison of five PRA studies: Big Rock Point, Zion, Limerick, Grand Gulf, and Arkansas Nuclear One. The overall objective was to provide a summary of the five studies together with an interpretation of the results that would be of interest to both technical specialists and persons involved in management.*

3.5.2.2 Accident Sequence Evaluation Program (ASEP) -- This program, sponsored by the NRC as part of the ongoing Severe Accident Research Plan (SARP), is intended to identify the accident sequences that have the greatest potential for dominating core-melt frequency or risk in LWR power plants, to identify the range of frequencies for these accident sequences, to determine the plant characteristics or uncertainties that most affect these frequencies, and to establish the specific plant characteristics, uncertainties, and frequency ranges that apply to specific classes of plants.* The ASEP uses not only existing NRC and industry-sponsored PRAs but also a number of generic studies, most of which have been sponsored by the NRC. The results have been used to update the accident-frequency predictions in some of the earlier PRAs. Some of these studies are described in Section 3.6.

3.5.2.3 Industry Degraded Core Program (IDCOR) -- Fourteen PRA studies were reviewed in the ongoing industry-sponsored program undertaken to assist the utility industry in developing a technical position on issues related to severe-accident rulemaking.* The purpose of the PRA review is to provide information regarding the perception of risks associated with severe accidents, the basis for initial

investigation of accident processes and phenomena, and the potential impact on risk of various proposed changes to plant design or operation. This information will help in deciding what regulations might reduce risk.

3.6 Studies of Special Issues

Since most regulatory actions involve decisionmaking that affects all plants or large classes of plants, a number of so-called generic studies, usually addressing special issues, have been performed. These studies have drawn from the large information base (models, accident sequences, risk profiles, and insights about the plant characteristics that are the most important contributors to risk) created by the NRC- and industry-sponsored PRAs and have used PRA techniques as their principal analytical tool. Thus, these studies have played an equally large role in the use of PRA in the regulatory process. A few of them are described below.

3.6.1 Risk-Based Categorization of NRC Technical and Generic Issues

Perhaps the first well-known use of PRA insights in the regulatory process occurred in 1978 when the Probabilistic Analysis Staff performed a study (SECY/78-7616) to categorize the existing technical and generic issues facing the NRC. The primary objective was to assist in identifying the task-action plan issues that have the greatest safety significance on a relative risk basis. One hundred thirty-three task action items were reviewed and assigned to four broad categories ranging from those having high risk significance to those not directly relevant to risk. Of the 133 items, 16 fell in the high-risk categories. The ranking aided the selection of the generic issues that would be designated "unresolved safety issues." This effort was recently redone by the NRC's Office of Nuclear Reactor Regulation (NRR) to include all TMI action

plan issues and issues identified since the TMI accident. The most recent effort, similar to the earlier effort, developed and quantified the accident sequences associated with each issue.

3.6.2 Value-Impact Assessment of Alternative Containment Concepts

Another regulatory use of PRA techniques also occurred in 1978. The NRC was then considering the underground siting of nuclear power plants, an issue that had been raised by environmental organizations. Under the sponsorship of the Office of Nuclear Regulatory Research, a study (NUREG/CR-0165) was undertaken to compare the relative value and cost of alternative containment concepts "between the present regulations and underground siting that could add to plant safety." Using insights from the RSS, the study considered nine alternative designs found to be the "logical alternatives." Filtered atmospheric venting was found to be the design alternative most promising on a value-impact scale. This study contributed to the subsequent focusing of containment research on filtered vents and the diminution of interest in the underground siting of nuclear power plants.

3.6.3 NRC Auxiliary Feedwater Studies

After the TMI accident, the NRC's Office of Nuclear Reactor Regulation sponsored a series of studies to review the design of auxiliary feedwater systems in U.S. PWRs. These studies used PRA techniques to identify potential failures that could dominate the unreliability of auxiliary feedwater systems during transients caused by a loss of main feedwater, including the station blackout sequence.* (The ability to cope with this particular sequence had not been a licensing requirement for the earlier licensed plants.) This study, which demonstrated the value of applying PRA techniques when

at the system level, led to changes in the safety review process. A quantitative requirement on auxiliary-feedwater availability was added to in the standard review plan, and studies of auxiliary-feedwater reliability have become a routine requirement for licensing.

3.6.4 Analysis of DC Power Supply Requirements

This study was undertaken as part of the NRC's generic safety task A-30, "Adequacy of Safety Related DC Power Supplies."* The issue stemmed from the dependence of decay-heat-removal systems on DC power supply systems, which nominally meet the single-failure criterion. The failure of DC power supplies affects the ability to cool the reactor core. It was found that DC power-related accident sequences could represent a significant contribution to the total core-damage frequency. It was also found that this contribution could be substantially reduced by the implementation of design and procedural requirements, including the prohibition of certain design features and operational practices, augmentation of test and maintenance activities, and staggering test and maintenance activities to reduce human errors.

3.6.5 Station Blackout

Two studies addressed the unresolved safety issue A-44, "Station Blackout." Together they provide the technical bases for resolving the A-44 issue. The first study, "The Reliability of Emergency AC Power Systems in Nuclear Power Plants,"* when combined with the relevant loss-of-offsite-power frequency, provides estimates of station-blackout frequencies for 18 nuclear power plants and 10 generic designs. The study also identified the design and operational features that are the most important to the reliability of AC power systems.

The second study, "Station Blackout Accident Analysis," (NUREG/CR-3226) focused on the relative importance to risk of

station-blackout events and the plant design and operational features that would reduce this risk.

The technical bases supplied by these PRA-type special issue studies are currently being used to formulate the NRC strategy for resolving of the station-blackout issue.

3.6.6 Precursors to Potential Severe Core-Damage Accidents

This study (NUREG/CR-2497) is applying PRA techniques to operating experience to identify the high-risk features of plant design and operation. The operating-experience base is derived from the licensee event reports (LERs) of operational events that have occurred in U.S. nuclear power plants. The events of interest are multiple events that, when coupled with postulated events, result in plant conditions that could eventually result in severe core damage.

The precursor study is a long-range study that is still under way. In the first 2.5 years, 169 significant precursors were identified for the 432 reactor-years of operating experience represented by LERs submitted from 1969 to 1979; preliminary findings show, 56 precursors for 126 reactor-years of operating experience for 1980-1981. The results were used to analyze accident sequences and estimate core-melt frequencies for operating plants. One objective of the precursor study is to compare these results with the estimates made in existing PRAs.

3.6.7 Anticipated Transients without Scram (ATWS)

The NRC staff evaluation of anticipated transients without scram in NUREG-0460* was one of the first applications of PRA techniques to an unresolved safety issue. The evaluation highlighted the relative frequency of severe ATWS events for various reactor types and estimated the expected reduction in frequency for various postulated plant modifications. The

study also proposed quantitative goals for resolving this issue.

Other notable examples of PRA application to the ATWS issue are the NRC-sponsored survey and critique of the reactor protection system (RPS) and the quantitative evaluation of proposed ATWS-related modifications sponsored by a consortium of U.S. utilities.* The RPS survey reviewed some 16 reliability studies, mostly in published PRAs, to compare the predicted failure probability per unit demand, the anticipated-transient frequency, and primary influences on RPS unavailability. There was a surprising degree of agreement among the 16 studies. The second study quantified the relative improvement to be gained by implementing a set of recommendations being proposed by the utility consortium in an ATWS petition to the NRC.

3.6.8 Pressurized Thermal Shock

In addressing pressurized thermal shock, probabilistic assessments were used to derive screening criteria to identify operating plants needing modification. The owners groups associated with the different PWR designs submitted estimates of frequencies of severe overcooling events. Analytical efforts using PRA techniques are continuing to evaluate the risk significance of this issue.

3.6.9 Addition of Pilot-Operated Relief Valves to Combustion Engineering Plants

The purpose of this study was to determine the change in risk from the addition of pilot-operated relief valves (PORVs) to those Combustion Engineering plants that do not have PORVs.* The study indicated that for certain plants an appreciable fraction (40 to 50 percent) of the risk reduction came from the additional pressure relief for ATWS sequences and the remainder from the addition of feed-and-bleed capability,

reducing the frequency of core-melt sequences involving the loss of decay-heat-removal capability.

3.6.10 BWR Water Level--Inadequate Core Cooling

PRA techniques were used in the analysis of TMI Action Item II.F.2, BWR water level--inadequate core cooling. The results indicate that there is no need for additional instrumentation for detecting inadequate core cooling in the BWRs.* The study showed that improvements in existing systems for water-level measurement and improvements in operator performance (Shoreham and Limerick), could make the predicted core-damage frequency due to failure in water-level measurements in the plants analyzed much smaller than the total core-damage frequency predicted in recent PRAs for BWRs.

3.6.11 Scram Discharge Volume

An analysis of pipe breaks in the BWR scram system indicated that the postulated sequence of events is not a dominant contributor to core-melt frequency.* It was based on the assumptions that the failure frequency of the scram discharge volume (SDV) pipe is about 10^{-4} per plant-year and that the operability of required mitigation equipment is not degraded by the resultant adverse environment.

3.6.12 Other TMI Action Plan Items

PRA provided tools for the analysis of TMI action items II.K.3.2 and II.K.3.17, (i.e., the frequencies of LOCAs caused by stuck-open pressurizer PORVs and outages of ECCS respectively).

The results of II.K.3.2 indicated that the frequency of small LOCAs from stuck-open PORVs, with the PORVs operated as they are at present, was in the range of the small-LOCA frequency in the RSS and that no additional measures to reduce the PORV-

LOCA frequency are required.* The purpose of the data collection under item II.K.3.17 was to determine whether there is a need for cumulative outage requirements in the technical specifications and which plants had a significantly greater than average cumulative ECCS outage time.

3.6.13 Evaluation of Exemptions from Limiting Conditions for Operation, Technical Specification Changes, and Surveillance Requirements

Probabilistic models have been used by the NRC staff to perform sensitivity studies for providing insights into the bases for limiting conditions for operation (LCOs), LCO extensions, and testing and maintenance requirements. Some specific examples include allowed outage times for auxiliary feedwater systems and diesel-generator LCO extensions.

3.6.14 Waterhammer

Unresolved safety issue A-1 deals with the potential impact of waterhammer events in operating reactors. A fairly large number of reported waterhammer events in recent years have caused concern regarding the ability of plant systems and safety features to respond adequately. Several existing plant-specific risk assessments were reevaluated to determine the risk importance of this issue. The study* showed that the inclusion of waterhammer data caused virtually no change in the quantification of dominant accident sequences. These results were used as part of a value-impact analysis in support of the resolution of issue A-1, which will be documented shortly.

3.6.15 Toughness of Supports for Steam Generators and Reactor-Coolant Pumps

The low fracture toughness of the supports for steam generators and reactor-coolant pumps is unresolved safety issue A-12. PRA techniques were used to simulate support-structure

failures during an earthquake. The results showed that backfits to operating plants were unwarranted and that for new plants the regulatory requirements are cost effective.*

3.6.16 Seismic Design Criteria

PRA techniques were used to estimate the incremental risk due to changes in seismic criteria. The results* showed that proposed changes would not affect the plant risk significantly.

3.6.17 Containment Sump Performance

As part of its effort to resolve unresolved safety issue A-43, the NRC staff performed a limited risk assessment to gain insights into the potential for risk reduction. Issue A-43 deals with the possibility that, after a LOCA in a PWR, the recirculation sump will be blocked by debris from damaged pipe insulation. A parametric study* was performed for various frequencies of sump blockage, coupled with an engineering evaluation of debris generation in a high-energy pipe break, to arrive at realistic estimates of the core melt contribution due to this issue. Preliminary results indicate that the risk-reduction potential is very dependent on plant-specific design features like the type and location of insulation. The resolution of this issue will be documented shortly.

3.6.18 Draft Environmental Statements

Accident evaluations in draft environmental statements include classes of accidents that are beyond the design basis. Plant-specific PRAs, when available, and existing PRAs of similar plants are now being used to estimate the risk-significant accident sequences and plant characteristics.

3.6.19 Selected Topics in the Systematic Evaluation Program

To support the integrated assessment phase of the Systematic Evaluation Program (SEP), analyses were performed to determine

the risk significance of selected SEP topics. Proposed modifications that would upgrade the plant to current licensing criteria were evaluated to determine their effect on core-melt frequency and risk. The results were considered in arriving at backfit decisions. Many issues, such as loose-parts monitoring and RCS leak detection, were found to have low risk importance for virtually all the plants reviewed. Other issues (e.g., DC power availability, fire protection, and recirculation switchover in PWRs) were often found to have high risk importance. These studies have provided useful insights and allowed resources to be applied to the areas where the greatest reduction in risk could be achieved.

3.6.20 Emergency Planning and Response

Several studies have been performed to provide guidance emergency planning and response. Their results formed the basis for the implementation of emergency planning zones for the plume-exposure pathway and for NRC staff recommendations regarding the use of thyroid-blocking agents.

3.6.21 Reactor Siting

A study was performed* to develop bases for formulating new regulations for siting nuclear power plants. Generic and site-specific calculations were performed to evaluate the sensitivity of predicted consequences to variations in source terms, population distribution, weather conditions, and emergency response. The study concluded that estimates of offsite consequences are strongly dependent on population distribution, but relatively insensitive to variations in weather conditions. Predicted early health effects are very sensitive to source-term magnitudes and the timing of emergency response.

3.6.22 Economic Risks

Several studies have examined economic consequences and risks.* Their results indicate that economic risks are dominated by relatively high-frequency forced outages and that the economic losses predicted for the owners of the plant generally exceed offsite economic consequences.

3.6.23 Filtered-Vented Containments

This study* was the first to use PRA methods and results for discriminating between alternative design options and operational strategies. In general, it was determined that a filtered-vented containment may be cost effective for BWRs, but not for PWRs.

3.6.24 Reduction of Severe-Accident Risk

The Severe Accident Risk Reduction Program is being conducted to provide a basis for severe-accident decisionmaking.* Through interactions with others and independent analysis, it is incorporating insights from various research programs to investigate the use of generic plant categories to evaluate the benefits and costs of proposed safety features.

3.7 Concluding Remarks

As can be seen from the preceding discussions PRA is now widely used by both the NRC and the nuclear industry to analyze a wide range of issues and decisions. Besides providing a useful regulatory tool, PRA studies have provided a rich base of information and insights that are relevant to the NRC mission. As has been seen from the special issue studies, this information base is already starting to affect the regulatory process. Deciding how PRA should be used in the future to help the decisionmaker address regulatory issues requires

not only an understanding of the existing information base, but also a perspective on the methods used; the latter cover a wide range of technical disciplines, from statistics to human-behavior sciences.

4. LEVEL OF MATURITY IN ELEMENTS OF PRA STUDIES

4.1 Characteristics of Maturity

A PRA study is multidisciplinary. Depending on its scope, it requires analyses of plant systems, human behavior, the physical progression of core-melt accidents and the relevant phenomena, health effects, and seismic hazards, to name a few. However, not all of the methods used by the various disciplines have reached the same level of maturity. For example, the methods of reliability analysis have been used in some form since World War II, whereas, the method used for a analysis of core-melt progression are new and unique to reactor technology.

PRA studies need not be full risk assessments. Thus, the question as to how PRA should be used in the regulatory process must consider what parts of the PRA exhibit the greatest strengths and what parts are weaker; in other words, what level of maturity has been reached by the different methods used in PRA. The level of maturity depends on several characteristics: the stability of the method, the degree of realism, the degree of uncertainty, whether major progress is desired to improve the method, and whether it is feasible to achieve that progress, especially in the short term. Using these indicators of maturity, the overall level of maturity of each PRA element can be gauged. Decisions based on mature methodologies, in general, are expected to elicit more confidence than decisions based on less mature approaches.

A "stable" method is one that has not changed for a considerable period of time. An "unstable" method is one subject to fast-moving developments. Applying, or using the results of, an unstable method requires greater caution. This does not imply that a stable method is always highly accurate and free of error.

The realism of a method refers to the extent that approximations or conservatisms may have been knowingly or unknowingly introduced because of unknowns, or merely to simplify the models, or perhaps because of error. Whether the ultimate "result" is accurate within its stated uncertainties or is conservative or nonconservative.

It is important to recognize that uncertainties are not unique to PRA; they reflect a lack of data or experience or a lack of knowledge about system response, human behavior, or accident phenomena. These uncertainties are present in estimates made by means of PRA techniques, deterministic modeling, or so-called engineering judgment. They reflect current experience and knowledge, and the state of the overall technology. PRA displays its uncertainties explicitly and, in so doing, focuses attention on them. Thus, PRA analyses display uncertainties more explicitly than do other analytical approaches, even though the extent of the uncertainty is the same in all cases.

Displaying the uncertainties, as PRA does, provides important information to the decisionmaker. A proper uncertainty analysis can provide an estimate of how this lack of experience and/or knowledge affects engineering insights drawn from PRA. This is done by propagating uncertainties through the analysis or by performing sensitivity analyses. Thus the treatment of uncertainties should logically be considered a strength of PRA rather than a limitation.

The remaining sections of this chapter address the level of maturity of the various PRA methods and practices (Sections 4.2 through 4.5) and of uncertainty analysis itself (Section 4.6).

4.2 Plant-System Modeling

Plant-system models delineate the behavior of plant systems in response to potential initiating events, the outcome being either a successful termination of unfavorable sequences or progression to core damage or core melt. The models discussed in this section include not only the representation of hardware behavior but also the contributions of human interactions with the system and the evaluation of the data used as inputs to quantification. Quantification is discussed in Section 4.2.3. Systems modeling of the response to external initiators is considered in Section 4.5.

4.2.1 System Modeling in Event Trees and Fault Trees

The logic models (event trees and fault trees) used for plant systems are basically the same as those used in the RSS. Their application has become more efficient and effective, and the modeling techniques have been improved in both scope and depth. The major refinements lie in the ability to model interactions between initiating events and mitigating systems. These improvements derive from a better understanding of the root causes of initiating events, of plant responses to support-system failures, and of the interactions and dependences among mainline systems and support systems. Model evaluation techniques have improved so that the analyst can increase the scope and depth of analysis, if necessary.

The major limitations remaining are those associated with completeness and with the accuracy of the representation of plant behavior, including human interactions with the system; the latter are discussed in Section 4.2.2.

There is always the possibility that the PRA models are incomplete. The analyst may not have identified or adequately defined certain events. Some events are specifically excluded from the models because they are known or thought to be highly improbable. The estimate of improbability is based on

a combination of experience and theory. The existence of important omissions remains a concern. Although completeness cannot be demonstrated, except within the very rough bounds of operating experience, the consensus of the PRA community is that most of the major insights obtained from PRA are valid and will remain valid even if new accident sequences and sequence dependences are identified and added to the model. (If the major concern is with quantitative results such as core-melt frequency or risk, then the impact of discovering an important new sequence would be great.)

Analyses of operating history are being performed to determine the interactions and dependent failures that have occurred. This knowledge is fed back to the analyst so that modeling techniques can be improved. Thus, with time, we anticipate that even the impact of possible incompleteness on the overall numerical values will decrease somewhat.

The accuracy with which a model represents true plant behavior is difficult to assess. Modeling, by its nature, implies an abstraction and an approximation of physical reality. The validity of the representation inevitably becomes an issue under these circumstances. The extent of the problem cannot be quantified at this time because there are few references against which accuracy can be gauged. There will be some improvement in the ability to evaluate this aspect as more operational data are acquired over the years.

The present status is that a somewhat conservative bias appears to have been intentionally inserted into the modeling process at those points where phenomena are poorly understood. This bias is most evident in two aspects of system analysis: (1) in the sometimes conservative choice of success criteria for various functions (e.g., whether one or two parallel pumps of three are required to perform a specific function) and (2) in the often conservative approximations used for

modeling the occurrence, time sequence, and effect of thermal-hydraulics and related phenomena. Specifically, since either thermal-hydraulic analyses are not available for some of the specific transient and LOCA sequences studied or the validity of certain analyses is questionable, there is a need to introduce simplifications (typically conservative ones) to allow the analyses to proceed. The conservatisms (and possibly some nonconservatisms) arising from these approximations have not been quantified. However, as mentioned in discussing incompleteness, the resulting limitations do not negate the usefulness of model results if the limitations are properly recognized when interpreting the results.

Another issue is the extent to which dependent failures are properly modeled or quantified. While this issue is complex, the consensus is that there has been significant progress in this area in the years since the RSS. If a decisionmaker in any particular application is aware of the possible uncertainties arising from this aspect, it is unlikely that the insights from these analyses would be invalidated.

Because PRA modeling involves very rare core-melt events, complete validation in an experimental or experiential sense is not achievable. However, validation either through operating and experimental data or through reliable analysis is feasible for many subelements of the analysis and has been partly accomplished.

Historical data analysis, including the analysis of accident precursors, is a means of validating sequence and dependence modeling to some extent. Some important aspects of PRA system modeling cannot be experimentally validated. Fortunately, this does not seem to represent a major limitation to the usefulness of most PRA insights.

4.2.2 Human Interactions

Human interactions are important in the operation, control, maintenance, and testing of equipment in practically any industrial activity. While beneficial, human interactions (e.g., repair and recovery operations) often enable systems to achieve an extremely high availability, they also contribute to the accident frequency. In the cases of the airline and chemical industries, it has been estimated that human interaction has contributed as much as 90% of the accidents.

Similar findings have emerged from some PRA studies for nuclear power plants. Past PRA studies have found that both beneficial and detrimental contributions of human interactions impact the ordering of dominant sequences and the risk of the plant. For example, the studies invariably have included human actions that can result in the unavailability of plant systems before an initiating event or that can cause an initiating event to occur. Beneficial human interactions include the diagnosis of the nature of an accident and recovery from an accident sequence. Clearly, PRA techniques provide a framework for assessing the importance of human interactions in a spectrum of accident sequences.

The definition of specific accident sequences in PRA studies offers the analysts the opportunity to investigate how human interactions affect the estimates of core-melt frequency or risk. For example, the uncertainties in the quantitative impact can be assessed, the ways that operators affect the course of an accident can be described, and the importance of human interactions in a particular sequence can be quantified.

The basic techniques for analyzing most human interactions was first developed in the RSS. These techniques have since been refined and formalized to improve the understanding of the human effect on plant safety. Furthermore, the identified

limitations in the techniques have stimulated numerous suggestions from people outside the PRA field for further improvements. Hence, the analyses of human interactions in PRA studies are undergoing rapid improvement.

The usefulness of human-interaction analysis in PRA studies in defining corrective solutions to reduce risk has been limited. The following issues contribute to this limitation:

1. Human behavior has been recognized as a complex subject for centuries and does not lend itself to simple models like those for component reliability. This makes the analysis of human interactions more dependent on the judgment of the analysts.

2. Human impacts have been described as just success and failure states to match the logic used for equipment failures. This method does not account for the full range of human interactions.

3. Generic human-error probabilities have been applied on a judgmental basis, because a simplified model of the various parameters that affect human performance has not yet been fully developed.

4. The analysis of dependences has concentrated on dependences between operators rather than dependences between operators and the plant.

5. Techniques for considering the sensitivities and uncertainties address primarily uncertainties in the quantification of data as opposed to alternative logic models for incorporating human interactions.

While it is difficult to state the impact of these current limitations on the risk profiles, many analysts feel that

they are within the stated uncertainty bounds. These limitations are more likely to affect the estimates of accident frequency than the estimates of consequences. The uncertainties associated with the data base affect the confidence that can be placed in calculated human-error probabilities. However, PRA allows a relativistic interpretation of the results so that an order-of-magnitude level of accuracy may be sufficient to identify potentially significant human errors and to determine their importance in accident sequences. In any event, experience and data collected from other sources seem to indicate that most human-reliability data are adequate. The higher error rates (approximately 1×10^{-3} per act) associated with procedural errors are believed to have error factors of roughly 3. (The error factor is the ratio of the upper bound to the median. See Section 4.2.3.) The error factors increase with decreasing error rates. For single or multiple human-error probabilities associated with procedural errors at 1×10^{-5}, the error factor is in the vicinity of 10.

The analysis of human interactions in PRA studies is clearly a developing art. Improvements in future PRA studies are likely to include the following:

1. Development of interim methods for considering the importance of operator decisionmaking under accident conditions.

2. Development of certain representations of the time-dependent effect of human interactions on the success or failure of a system or safety function.

3. Use of a more structured technique for developing data from expert opinions.

4. Development of more systematic approaches for incorporating human interactions into the PRA framework.

5. Better integration of the system and human-reliability analyses.

6. Collection of data from training simulators to verify some of the judgmental data and support the development of simple models of human behavior.

These improvements should substantially increase the understanding of possible human behavior under accident conditions.

Limitations to the detailed description of human interactions will still exist, and they should be recognized. Both the qualitative description of the human interaction logic and the quantitative assessment of those actions rely on the virtually untested judgment of experts. Additional work is needed to develop simple mathematical models of human performance and to identify the parameters that affect human performance in accident situations.

The future outlook for the modeling of human interactions in PRA studies can be viewed as excellent. However, the depth of the techniques must be expanded so that the effect of changes in design, procedures, operations, training, and the like can be measured in terms of a change in a risk parameter, such as core-melt frequency. Then trade-offs or options for changing the risk profile can be identified. To do this, the methods for identifying the key human interactions, for developing logic structures to integrate human interactions with the system failure logic, and for collecting data suitable for their quantification must be strengthened. These improvements will have to be achieved before the associated uncertainties can be narrowed substantially .

4.2.3 Data Impacts to System Modeling and Quantification

Once the models of system and plant performance are completed, data of various types are needed to quantify the models and estimate accident frequencies or system failure probabilities. In general, the PRAs have relied heavily on generic data, supplemented with plant-specific data. The amount of analyzed data on transient initiating events has significantly increased since the publication of the RSS. The amount of data for LOCA initiating events is sparse and the accuracy and precision of the estimates of frequencies for these events have not improved. Generic component failure data, applicable to general PRA evaluations, have not significantly improved in accuracy or precision. They have benefited from studies of licensee event reports in some cases; however, few causal data are available, and the overall understanding of root failure causes has not improved significantly. The availability of dependent-failure data has improved marginally, and quantitative estimates in this area remain largely subjective. The improvements that have occurred have not had a major effect on either the numerical results of PRAs or on the insights.

Uncertainties in data are generally expressed as error factors, where the error factor is the ratio of the upper-bound estimate to the median value or best-estimate value. The range, which is the ratio of the upper-bound estimate to the lower-bound estimate, is the square of the error factor. This assumes that the median value is the geometric mean between the upper bound and the lower bound. Generally, the upper-bound estimate is approximately the 95th percentile; the lower bound is approximately the 5th percentile. Therefore, there is a 90 percent probability (or confidence) that the actual value of a parameter will fall within the range.

Two conceptually different sources of uncertainty appear in the data base for component and system failures. The first

is a natural variability of failure rates in the existing population for each piece of hardware; the second arises from imperfect knowledge of the actual behavior. The latter source can be reduced over time as more data are accumulated; the former will remain a source of uncertainty, though not a source of errors.

For present data on transient initiating events, the error factors are considered to be small, about a factor of 2 to 3. For data on LOCA initiating events, the error factors are thought to be 3 to 10. Component-failure rates have error factors of approximately 3 for active components (pumps, valves, etc.) and 10 to 20 for passive components (pipes, wires, etc.). The failure rates for passive components are generally substantially lower than those for active components, even considering the uncertainties.

Data on test and maintenance intervals and durations generally have error factors of about 2 or less. Corrective-maintenance intervals and maintenance durations have larger error factors.

Error factors for common-cause probabilities involving two or more coupled failures generally increase as the probability decreases. For common-cause probabilities in the vicinity of 1×10^{-5}, the error factor is approximately 3 to 10. These error factors apply to the probability of multiple failures (i.e., unavailabilities).

The above error factor values represent gross averages of error factors extracted from several completed PRAs and applied to internal initiating events only. PRAs that use plant-specific data generally have smaller uncertainties (error factors) than those that use only generic data. Uncertainties arising from data uncertainties are generally the

only ones that are explicitly quantified in a PRA. Uncertainties in dependent-failure and human-error probabilities often dominate the data uncertainties associated with calculated system unavailabilities and accident frequencies. The specific effect of data uncertainties, however, depends on the application. When a single contributor dominates the risk, then the uncertainties in data for the contributor will have significant effects on the results. When many unrelated contributors contribute equally to the risk, then the data uncertainties from any one contributor will not have large effects.

The quantification of systems models involves two steps: the reduction of the trees to what have usually been called "minimum cut sets" (i.e., the minimum groups of failures needed to describe the particular branch of the tree) and the actual quantification itself. The first step typically involves Boolean algebra and is not usually controversial; nor does it introduce important additional uncertainties. The second step is more problematic: each contributing factor has a distribution, then the resulting quantified result must also be expressed as a distribution. Several methods have been used, ranging from simple approximations (e.g., using log-normal input distributions) to carrying out Monte Carlo calculations to generate the output distributions. The consensus is that no substantial problem exists with the methodology in this aspect, although there is strong discussion about the best way to display and interpret the output "results" (single-value estimates, full distributions, ranges, etc.) in a manner that adequately conveys their meaning to a decision-maker.

4.2.4 Concluding Remarks on System Modeling

System modeling in PRA studies is usually considered to be very mature, significantly more mature than the study of accident progression phenomena, and comparable in maturity to the

analysis of offsite consequences. The techniques of fault trees and event trees have advanced considerably since their initial application in the RSS, and a variety of approaches to their use are available. The conclusions drawn from system modeling are generally quite solid, even though issues about the completeness of the analysis persist.

The treatment of the underlying assumptions in systems analysis (e.g., success criteria, thermal-hydraulics phenomena) is an unresolved issue. Also, some improvements are still desirable in the failure-rate data base, since the ranges (and error factors) are quite broad for some important areas. Progress has been made recently in the collection and analyses of certain data, but more is needed. The improvements in data have not changed the insights very much; it is believed that the conservatisms and the incompleteness of modeling have not affected these insights much either.

The modeling of human interactions introduces substantial uncertainty. However, great progress has been made recently in this difficult area, and much more work is now under way. Within a very few years this aspect of system modeling is expected to become more systematic and the results more reproducible. Overall, the results of work under way can be expected to yield important insights.

In summary, the whole area of PRA system modeling has advanced significantly since the RSS. The conclusions and insights it affords are usually reasonably sound, if appropriate consideration is given to the uncertainties and if great numerical accuracy is not required for the particular application. Most important, system modeling has provided insights about the relationships among systems, failures, and phenomena that could not have been obtained in any other way.

4.3 Accident Progression, Containment Response, and Source Terms

In-plant accident processes are analyzed to determine the thermal-hydraulic response of the plant to accident sequences, the progression of severe accidents, containment performance under severe accident loadings, and the characteristics of radionuclide releases to the environment (source terms) for accident sequences or groups of sequences. The analyses include a wide range of phenomena, some of which are not well understood.

The characteristics of radionuclide releases to the environment are described in terms of various timing and location parameters, the thermal energy release rate, and the quantities of radionuclides released. The quantities of radionuclides available for release from the plant depend on the processes by which radionuclides are released from the fuel and transported through the reactor-coolant system, the containment, and possibly buildings external to the containment before reaching the environment. Analyses have shown that both natural and engineered retention mechanisms can significantly reduce the inventory of radionuclides in the containment atmosphere if enough time is available for those mechanisms to act. Therefore, source terms are strongly affected by whether or not the containment fails and, if it fails, by the time and mode of failure.

The capabilities of in-plant consequence analysis have improved substantially since the RSS and are currently rapidly changing. Since the TMI accident, severe-accident research has expanded broadly, the aim being not so much to improve PRA but to acquire information about severe-accident behavior for possible use in plant regulation. Large experimental and mechanistic code-development efforts have been initiated or redirected to explore important severe-accident phenomena.

Advances have also been made in the codes and methods used for developing and quantifying containment event trees.

Shortly after the TMI accident, questions were raised about the realism of the methods used to analyze source terms in the RSS and subsequent PRAs. In 1981, the NRC published an evaluation of "The Technical Bases for Estimating Fission Product Behavior During LWR Accidents" (NUREG-0772). As a result of deficiencies identified in that and other reviews, a number of research programs have been undertaken to improve the ability to model radionuclide release and transport in severe accidents. The NRC's Source Term Reassessment Study, IDCOR, and other studies are evaluating the effect of improved analysis capabilities on predicted source terms.

Many uncertainties are associated with the predictions of severe-accident progression, containment response, and source terms. The NRC and other organizations are making considerable effort to better define and reduce those uncertainties. Presently, few sensitivity studies exist, the validation of models and codes for the broad range of severe-accident phenomena is extremely limited, and quantitative uncertainty estimates are not available. As a minimum, current research can be expected to provide a better characterization of source-term uncertainties and in some important areas reduce the conservatisms in PRA analyses.

Since the analysis of in-plant consequences is rapidly changing, the method is unstable. Indeed, developments are occurring so rapidly that, for a PRA being undertaken today, it is difficult to recommend a set of computer codes. A key issue is the depth of analysis required for radionuclide behavior. The decision will depend on the extent to which uncertainties are reduced through the use of complex models and on the degree of potential biases associated with simpler models. Major advances are currently being made in the understanding

of processes controlling radionuclide release and transport. However, processes that are closely coupled to the progress of extensive fuel damage, such as the release of the less volatile radionuclides from fuel or the generation of hydrogen during core slumping, will likely always have large uncertainties because of the difficulties associated with experimental validation.

4.4 Offsite-Consequence Analysis

Offsite-consequence analyses attempt to predict the frequency distribution of possible consequences for core-melt accidents. Models have been developed which describe the transport, dispersion, and deposition of radioactive materials and then predict their resulting interactions with and influence on the environment and man. Consequences can include early fatalities and injuries, latent-cancer fatalities, genetic effects, land contamination, and economic costs.

The first comprehensive assessment of consequences was performed in the RSS. Since that study, modeling capabilities have been improved, model-evaluation studies have been performed, and existing models have been applied to provide guidance for planning and decisionmaking in areas such as emergency planning and reactor siting. In addition, studies have been performed to examine the importance of potential consequences resulting from releases of radioactive materials to liquid pathways.

Uncertainties in offsite-consequence predictions have not yet been assessed comprehensively. What currently exists is a large body of parametric (or sensitivity) analyses in which consequences are calculated for a range of plausible values of a key parameter or model. The PRA Procedures Guide (NUREG/CR-2300) made a tentative listing of the relative contribution to total uncertainty of the major parameters and models in an

offsite-consequence analysis. Important contributors to uncertainty were:

1. The magnitude of the source term, which strongly influences all consequences.

2. The form and effectiveness of emergency response, which can make a large difference in predicted early health effects.

3. The rate of dry deposition (fallout during rainless periods) of particulate matter from the plume, which affects the distances to which land-use restrictions or crop impoundments may be required and the intensity and dispersion of early health effects.

4. The modeling of wet deposition (washout by rainfall), which affects the low-probability, high-consequence end (tails) of the distributions of all consequences.

5. The dose-response relationships for somatic and genetic effects.

It also appears that the condensation of moisture in the released plume could have a significant impact on resulting consequences.

Even though the uncertainties in offsite-consequence analyses have not been thoroughly examined, their general magnitude can be inferred from the results of many existing sensitivity studies. Major contributors to uncertainty, as well as the magnitude of uncertainties, depend strongly on assumptions about source terms and site characteristics. For estimates of the consequences resulting from very large source terms at a highly populated site, and given that the source term is

known (i.e., source term uncertainty is not included), the following crude estimates of uncertainties can be made:

1. Mean early fatalities could range from approximately a factor of 10 above present "best" estimates to nearly zero. This broad range is in large part due to uncertainty in the effectiveness of short-term emergency response near the plant.

2. The uncertainty in the mean predicted population dose (man-rem) is estimated to be a factor of 3 or 4, while the uncertainty in the predicted mean number of latent cancer deaths (which depends on the population dose) is approximately a factor of 10.

3. In general, the uncertainties are larger in the magnitude of the extremely low-probability, high-consequence portion ("tails") of predicted consequence-frequency curves.

Ongoing research in the United States and several foreign countries will improve consequence-analysis capabilities for use in PRA and other applications. Efforts are focused on quantifying and, where possible, reducing uncertainties. Although uncertainties are likely to remain quite large, a thorough examination of their origin and magnitude will provide both a firmer basis for the application of offsite-consequence analysis and a better understanding of its limitations.

4.5 Accidents from External Initiating Events

PRA analysts have conveniently assigned accidents resulting from various "external initiating events" into a separate category, principally because the method for treating the

them is different from the method for treating so-called internal events.** The external events are

1. Earthquakes
2. Internally initiated fires
3. Floods
4. High winds--tornadoes, hurricanes
5. Other: aircraft, barge, and ship collisions,
   truck, train, pipeline accidents
   external fires
   volcanoes
   turbine missiles
   lightning

The analysis of external events has seen major advances since the RSS. The basic approach taken in the probabilistic analysis of initiating events consists of quantifying the expected frequency of the initiating event, determining its effects on various pieces of equipment, and determining the resulting effect of any degradation or failures on plant performance.

Much active developmental work is in progress, and abilities in this area should continue to improve. However, the uncertainties associated with such analyses are still significantly larger than those associated with internal initiating events, principally because of uncertainties associated with the development of the hazards curves (i.e., the frequency of occurrence of an event exceeding a given magnitude). At the present time, even with the increased maturity, the methodologies have not progressed to the point where great credence can be given to quantitative risk estimates, particularly in

---

**Both "internal initiating event" and "external initiating events" are misnomers, since the former category is usually taken to include accidents starting with the loss of offsite electric power, while the latter usually includes internally initiated fires and floods.

comparison with internal events. The principal benefit of external event analyses lies in the qualitative assessment of their effect on components, systems, and structures and the relative importance to safety of such functions. For each of the major external initiating events the sections that follow discuss the level of maturity of the analysis.

4.5.1 Earthquakes

The intrinsic problem with seismic hazard analysis is that the dominant contributors to reactor risk come from earthquakes significantly larger than the earthquakes used as the design basis, and the frequency of recurrence of these large earthquakes are difficult to estimate because there are no historical records. This is especially true in the eastern United States but are also true in California. Thus, the estimates of frequency of occurrence have major uncertainties and these uncertainties in input data create large uncertainties in the final results. Extrapolations to larger earthquakes, with very long recurrence intervals, present another problem. It is believed that for any given site or seismic province there is a maximum earthquake motion that can be sustained by the specific geologic features; thus the extrapolations are usually cut off in one fashion or another. It turns out that the numerical results of some recent PRAs may be more sensitive to the nature of the cutoff procedure than to any other of the several assumptions.

Still another uncertainty comes from efforts to characterize the motion associated with these large earthquakes in terms of physical parameters, such as frequency-dependent acceleration, velocity, and displacement; the shape of the motion in time; and the energy dispersion. Sometimes the historical record can be of value in characterizing the expected motion of future earthquakes; but unfortunately, for many of the important historical earthquakes, records are limited to known

structural damage and area over which the motion was felt, and even these can be less than reliable.

The PRA analysis requires a probabilistic description of the transmission of earthquake motion through the rockbed and soil to the building substructure and the interaction of the motion with that substructure, with the goal of characterizing the motion felt by safety-related structures and equipment. Regarding substructure interaction, advances have been made under the auspices of the NRC's research program and, though much remains to be accomplished, substructure interactions now contribute less than do other elements of the analysis to the overall uncertainties.

The fragility of structures and equipment is the other major element of seismic PRA. Here significant progress has been achieved and a large number of reactor studies are completed. The fragility analysis relies partly on a data base that is neither strong nor extensive, and partly on analytical methods designed to overcome many of the weaknesses in the data base. One roadblock to precise fragility analysis is the inadequate characterization of the input motion. Various surrogate accelerations have been proposed and used parametrically, each with limitations that offset some of the advantages.

The final element of seismic PRA analysis is linking the failures of structures and equipment into a system analysis of the plant. Here the event trees developed in the other parts of a Level-3 PRA study usually provide the basis for the analysis. Unfortunately, this part of the problem is not as easy to accomplish in practice as an analogous internal-events analysis, because a large earthquake can cause numerous failures that compromise redundant systems. The analyses completed to date, however, have found that the seismic risks are often dominated by accident sequences involving only a

few important seismic failure modes, often failures of structures. This simplifies the analysis somewhat, but only to the extent that one can make (and defend) the simplifying assumption that failure of the single or few components or structures actually brings the plant to core melt. This assumption has often been made with the full knowledge that it is conservative, sometimes highly so if quality assurance deficiencies in manufacturing, construction, and design are unimportant. Also, the human factors have not been treated with the desired thoroughness. For example, recovery by the operators has not yet been well treated in the studies. For all of these reasons, the analyses tend to be conservative in character rather than realistic in a numerical sense.

Overall, a consensus prevails that the uncertainties in the final results (core-melt frequency, offsite risks) remain quite large for seismic PRA analyses. Error factors of 10 to 30 (implying ranges of about 100 to 1000 for the 5 to 95 percent confidence interval) might be reasonable at present. Of course, these large numerical ranges for quantitative results do not compromise the highly significant engineering insights obtained by these PRA methods and many of these insights are new or unavailable with traditional methods. In particular, the system vulnerabilities and common-cause dependences that are revealed by the analyses can point the decisionmaker to the need for careful consideration using other approaches and other decision criteria.

4.5.2 Internal Fires

Only recently has the probabilistic analysis of internal fires become an accepted part of full-scale PRA studies. Only a very few such probabilistic studies of fires have been made, and the methodology used is by no means mature. Although the applications have been few, the literature covering both methodological improvements and applications is growing. The

state of the art has improved enough to allow the inclusion of internal fires as part of a plant-specific PRA. Furthermore, important new research is under way.

The initial phase of a probabilistic fire analysis is the identification of critical areas. The criterion for a critical area is whether a fire in that area could compromise important safety equipment. In practice this criterion is narrowed to emphasize areas where multiple equipment could be compromised--in particular, several trains of redundant equipment to perform the same safety function. Identification often begins with the fire zones delineated in the more classical regulatory analysis, supplemented or modified by a walk-through to identify areas with a potential for cross-zone fire spread and where transiently present fuels might supplement fuels always present in a zone. While this part of the analysis remains more an art than a science, the consensus is that it is reasonably mature, in the sense that uncertainties introduced in the PRA results are thought to be smaller than uncertainties arising from other aspects. Barrier identification feeds the fault-tree analysis by identifying candidate physical areas of concern for fire or fire coupled with a relatively high probability of random failures.

The next phase of the fire analysis is the estimation of the frequency of fire initiation for each critical zone. The frequency may depend on location within the critical zone if fuel-loading conditions, cross-zone spreading potential, or other circumstances require that level of detail. While a historical data base is available for fires initiated in various areas, the data base is not fully adequate as an empirical basis for working out the desired initiation frequency. The analyst must bring to bear location-specific information from his walk-through and from his background knowledge, and hence this part of the analysis contributes important subjective uncertainties. Despite these uncertainties, the ability of a

skilled analyst to rank the important potential fire-initiating locations is generally quite good, and it is now accepted that a good analyst will usually not overlook important zones.

The more difficult aspects of the probabilistic analysis occur in the next phase of study: the estimation of the likelihood that equipment will be disabled by a fire. This problem is compounded by uncertainties in the modeling of detection and suppression systems, actual fuel availability (amount and character of transient fuel, etc.), the stochastic nature of fire growth over time, the size of the affected secondary zone where hot gases can cause equipment failure or induce secondary fires, and access for firefighting. A number of important models have been developed over the years to assist the analyst in calculating the likely progression of the fire, but in even the best cases the quantitative uncertainties remain large.

In even the most advanced cases, the available models are only approximate in character and are not capable of accurately modeling fire spread in, for instance, a compartment full of crowded objects in a unique configuration. Also, the analysis of failure modes for components exposed to the whole spectrum of combustion products needs more methodological development and more test data. Finally, additional methods need to be developed for treating the intercompartmental spread of fire and combustion products. Even more important, no firm estimate is yet available for the uncertainty introduced by these incomplete capabilities. Research on this question, along with studies of the expected success of the analyst in evaluating fire detection and suppression, may help in determining the achievable accuracy of probabilistic fire analyses.

It is still too early to judge the achievable accuracy of the results of fire PRA with regard to core melt frequency or risk. The uncertainties are quite large, at least as large as those for the analysis of internal events. Of course, the engineering insights obtained have already been very useful from the few PRA studies performed to date and are in no way invalidated by the large uncertainties in the quantitative results. These large uncertainties will probably be narrowed somewhat by the results of current research, but it is too early to predict the effects of this research.

4.5.3 High Winds

For tornadoes and hurricanes, the state of maturity of PRA analysis remains modest. However, some useful insights have been gained, despite the rather sparse experience to date.

An important difference between hurricanes and tornadoes is the character of their winds. Hurricanes tend to produce mainly straight winds, with velocities rarely exceeding 130 mph. Tornadoes can produce wind speeds much higher than 200 mph, characterized by the more familiar "twister" wind forms. A further distinction is that tornadoes can frequently pick up objects and throw them great distances, making it imperative to study tornado missiles.

Several methods are available to estimate the likelihood of a tornado or hurricane producing a wind speed in excess of a particular value at a reactor site. These methods rely mainly on historical records that exist nearly everywhere in the United States with enough data to provide a useful starting point. Uncertainties in the analysis are still significant contributors to the overall uncertainty of the PRA results. For example, at a given site, the estimated annual frequencies of exceeding a wind speed of 120 mph might differ by as much as an order of magnitude for different analytical methods.

The problem of determining the likelihood of tornado missiles of various sizes is also quite difficult. The analyst must determine how many objects of various sizes are available in the vicinity for the wind to pick up (e.g., telephone poles, automobiles, trees, and even heavier objects).

The fragility analysis takes several forms and involves assumptions that have uncertainties. First and foremost is the assumption, now commonly accepted as a certainty, that offsite power and any onsite power openly exposed to high winds will be lost. Second, the structures must be analyzed. In this analysis, the metal-sided buildings are typically found to be much more vulnerable than concrete-sided ones, with failure modes including buckling, pressure collapse, and the tearing away of corners. The third and most difficult analysis is the fragility of equipment within a building. It is usually assumed that the wind-induced failure of any building implies the failure of all enclosed equipment. This assumption is probably conservative, though not necessarily always so, and with only the roughest ability to quantify the uncertainties introduced.

While engineering insights are available concerning vulnerabilities, the estimates of core-melt frequency or risk from high winds are highly uncertain.

4.5.4 Flooding

Flooding as a cause of severe accidents has not been analyzed in most full-scope PRAs, even though flooding clearly poses a potentially serious challenge to the overall safety of nuclear power plants. The methods used in the few completed studies have been limited to internal flooding and are quite similar to the approach used in studying internal fires. Thus the analyst must first identify critical areas, then estimate the probability that a flood might occur, and determine how long the flood might continue before it is stopped or its flood-

waters are drained. Finally, he must estimate the effect on critical safety functions and integrate the predicted harm to safety functions into an overall system study that uses the event-tree approach.

Flooding analysis is complicated by several factors. The fragility of safety equipment exposed to a spray-type flood from a pipe break is very difficult to analyze quantitatively. Flood-induced corrosion can compromise the ability of safety equipment to remain operable over the very long recovery period after a particular flood has been nominally "controlled." Another flaw in the analysis is the limited ability to quantify partial blockages of drains or sumps that are relied on to carry away floodwaters. Finally, flooding (especially from an external source) can randomly deposit solid matter like sludge, silt, or even sizable objects onto reactor plant equipment causing problems difficult to analyze.

Compounding the analytical problem is the issue of spray flooding from a pipe or tank leak, which could cause electrical failures at nearby locations. The data base and analytical methods for coping with this issue are not well developed. The possibility also exists that unusual dependences among equipment (e.g., spatial colocation of electrical or support equipment) will cause additional vulnerabilities. Difficulties in modeling human intervention can also complicate the analysis.

4.5.5 Other External Events

Several other categories of external events may pose a hazard to a nuclear power reactor: aircraft impacts; barge and ship collisions; truck, train, and pipeline accidents; external fires; volcanoes; turbine missiles; and lightning. Typically, these external events are analyzed probabilistically by performing a bounding analysis on their frequency of occurrence. An estimate is made of whether the initiating event is serious

enough to merit "concern." This estimate is usually semiquantitative. For example, the analyst might consider how large an external fire must be to compromise important safety equipment. The analyst next estimates quantitatively the likelihood that such a large fire might occur. If the frequency is low or can be bounded well enough, the analysis ends with the statement that the effect under study "does not contribute significantly to the overall uncertainty."

This approach is fully adequate if the analyst performs his work well, but several pitfalls could lead an analyst astray:

1. The analyst might choose the wrong initiating event. For example, a low consequence, but highly probable, initiating event may contribute a large amount of risk, while an event with more serious consequences, but a very low frequency may result in insignificant risk.

2. The analyst might overlook some coupled failure modes.

3. The estimated frequency of occurrence might be badly in error because there is no historical record and the extrapolation procedure is erroneous or because the historical data base is actually erroneous or inapplicable.

The main insights gained from the analyses performed on these initiating events are that, generally, they have minor risk significance and few have required further study. This insight is quite important, because it indicates the effectiveness of the deterministic design and operational requirements in ensuring plant adequacy in these areas. The design and regulatory approaches seem to be adequately conservative.

4.5.6 Sabotage

Sabotage as an initiating event has not been traditionally included in PRAs, but the threat of sabotage has long been

recognized and treated outside the PRA arena. PRA techniques have occasionally been used to do various vital-area and penetration analyses related to sabotage, but the risk of sabotage itself has never been estimated, principally because of difficulties in quantifying the threat frequency.

4.6 Uncertainty Analysis

The preceding sections have discussed the sources of uncertainty in PRA results (parameter variation, modeling, completeness) and, where possible, have indicated the general uncertainty bounds associated with the various elements of a PRA. Uncertainty analysis provides a framework for properly combining and describing the uncertainties associated with various elements of the analysis to determine the overall uncertainties associated with the results (e.g., risk) or intermediate quantities (e.g., sequence frequency).

Risk analysts are only at the threshold of performing comprehensive uncertainty analyses. A variety of techniques have been used or proposed. However, many are still being developed and, in general, the methods have not been applied in all their combinations for all parts of the PRA. The quantification of completeness uncertainties is particularly difficult and cannot be rigorous because it examines the limits of knowledge only from the side where something is known. Where it has been addressed, it was based on subjective judgment. Justification of these subjective judgments by evidence or even by consensus of experts is difficult. However, completeness uncertainties can be minimized by thorough analysis and detailed peer review.

Uncertainties associated with modeling can be addressed through sensitivity studies that show the variation of results with different models. However, the incorporation of such

results into an overall uncertainty analysis is still in its infancy, though various approaches have been taken in the past.

Means for analyzing the uncertainties associated with data variability are reasonably mature. Several methods are available for estimating uncertainties in basic data and propagating them through the analysis. While they differ in philosophical approaches, they may produce similar results, particularly when the data base is large.

Probabilistic analyses have often considered only data uncertainties. When completeness and modeling uncertainties have been incorporated, they relied heavily on subjective judgments, which are difficult to validate. The decisionmaker must therefore be aware of the elements considered in the uncertainty analysis and, more important, the elements not considered. The decisionmaker must also understand the framework in which the uncertainties are displayed.

Significant efforts are currently under way to improve the ability to perform meaningful uncertainty analyses and methodological improvements can be expected in the next few years. However, while improved methods will be available, many of the present difficulties associated with quantifying uncertainties about completeness will remain, because of the lack of knowledge.

Since subjective judgment is required in assessing the uncertainties for many of the contributing elements, the analyst should present the basis for his subjective judgments. He should also perform sensitivity studies if reasonable variations in the subjective judgments of peer reviewers could lead to significantly different insights or results that could affect the decision process.

4.7 Conclusions

This chapter has discussed the level of maturity of the various methodologies that comprise the discipline of PRA. It is important to recognize that the level of experience or knowledge (i.e., the data base) differs with different parts of the PRA. Thus, the reliance the decisionmaker places on the PRA insights should likewise differ with different parts of the analysis. As the state of the art exists today, several conclusions can be reached.

The methodology and information base for the system analysis of internal events are reasonably mature and stable. A relatively high level of confidence can be placed in insights about the relative importance of plant characteristics, dominant accident sequences, and core-melt frequencies. Since the weakest part is the human-reliability analysis, most confidence can be placed in conclusions that do not involve human-error uncertainties.

The weakness in the analysis of the frequency of external events, the lack of data on component response to these initiating events, and the lack of uniformity in analytical methods for external events results in less confidence in insights derived from this part of the PRA. Little confidence should be placed in insights that are derived from relative comparisons of internal and external event analyses or across different types of external events. More confidence can be placed in insights that are based on comparisons within the analysis for a specific initiating event. That is, one can have more confidence in the identification of the accident sequence that dominates core-melt frequency from among all the accident sequences that result from seismic initiating events than one could have in identifying the dominant sequence from a mixture of seismic, fire, and small-LOCA initiating events.

The analysis of accident progression, containment response, and source terms currently represents the most uncertain and most unstable part of a PRA study, so much so that it is not even clear what is the best approach in many cases. This situation can be expected to continue until the ongoing source-term work has stabilized predictions somewhat. Very limited confidence should be placed at this time on insights derived from this part of the analysis.

On the other hand, offsite-consequence analysis is relatively mature. Although it does not have the long experience of system analysis, a relatively high level of confidence can be placed on insights derived. However, offsite consequence analysis cannot overcome the uncertainties associated with the source term. Also, the stochastic uncertainties associated with weather conditions are quite large, so the actual consequences associated with an offsite release of radionuclides are quite uncertain.

In a field as complex as nuclear reactor safety and PRA, generalizations are difficult to support. Nevertheless, the motivation to do so is strong. Therefore, in the interest of crisp communications and at the risk of oversimplification, Table 4-1 summarizes confidence levels for various PRA activities.

Table 4-1

PRA Activities and Related Confidence Levels

PRA Activity	Level of Confidence in Insights
System analysis for internal events	Relatively high
System analysis for external events	Modest for insights applicable to specific external events; very limited for insights spanning different external events or internal events
Accident progression, containment response, and source terms	Very limited
Offsite consequences	Relatively high, but applicable only to assumed source term

5. INSIGHTS GAINED FROM PROBABILISTIC RISK ANALYSIS

5.1 Scope

The results of a PRA generally provide information about plant design, operation, and safety that previously may not have been identified explicitly . This information enhances the understanding of such diverse items as design weaknesses and strengths, the importance of assumptions about accident phenomena, operational deficiencies and strengths, sources of uncertainty and their implications, levels of risk, and the relative importance of contributors to plant risk. Other types of analysis may provide some of this information, but it has been the PRA (e.g., the analysis process itself) that has directed the examination of potential accidents and plant risk in a disciplined, quantitative way and has presented this information in an integrated context. PRAs develop insights and findings at many levels of resolution. This chapter presents some of those insights at four levels:

1. Global findings
2. Insights into plant risk
3. Insights into dominant accident sequences
4. Important findings about systems, functions, and human reliability.

5.2 Global Findings from PRA

In addition to plant-specific and generic insights, the PRAs performed to date have yielded certain global insights that apply not only to those plants analyzed but also to all current nuclear power plants, based on our knowledge of their general design and operating characteristics. These global insights are summarized below.

1. The estimated frequency of core-melt is generally higher than had been thought before the RSS.

2. Most large core-melt accidents are not likely to lead to very large offsite consequences. A very wide range of potential consequences seems to exist for core-melt accidents, depending on many factors. The small fraction of core-melt accidents that might lead to large offsite consequences generally involve an early failure of containment in relation to the time of core melt or containment bypass (e.g., either before or just after core melt).

3. The specifics of dominant accident sequences and the estimates of risk vary significantly from plant to plant even though each plant meets all applicable NRC regulatory requirements.

4. PRA has identified the following insights about offsite consequences:

   a. Latent-cancer risk is an important element of public risk: before the RSS, concern had centered on early fatalities and offsite property contamination.
   b. For core-melt accidents, the estimated offsite economic losses are generally much smaller than the estimated plant losses.
   c. In contrast to public health risks, economic risks from LWR operation are dominated by relatively high frequency forced outages.
   d. Estimated risks of early fatalities and injuries are very sensitive to source-term magnitudes and the timing of emergency response of civil agencies.
   e. Estimates of early fatalities and property losses differ greatly from one site to another, but site-to-site differences are not very great for latent cancers and onsite property damage.

5. PRA has enabled the qualitative assessment of the ways in which various elements of reactor safety contribute to risk. Among the important findings are the following:

   a. Operational considerations are important to overall risk and may be comparable in importance to design features.
   b. Human errors play an important role in overall reactor safety.
   c. Containment response plays a key role in the overall public risk.
   d. Small LOCAs and transients are dominant contributors to risk in most PRAs; LOCAs are usually not important contributors to overall risk.
   e. Earthquakes and internal fires seem to play an important role in plant risk, although this conclusion is very plant-specific. The uncertainties in estimating the risk from external events and internal fires are so large that comparisons with the risk from internal events are tenuous at best.
   f. As contributors to risk, airborne pathways are much more important than liquid pathways.

6. Accidents beyond the design basis are the principal contributors to public risk. This indicates that the designers and regulators have been generally effective in reducing the risks from expected operational occurrences and design-basis accidents.

7. A few groups of accident sequence types tend to dominate the risk in all plants studied. However, the reasons for the dominance are plant-specific and relate not only to design and operational differences but also to assumptions about the effectiveness of procedures and operator actions.

8. Some of the key accident sequences are highly plant-specific, while others are generic. PRA results have been useful not only in identifying these dominant contributors but in suggesting cost-effective approaches to remedying them, if necessary.

9. The plant-to-plant variability in PRA results is expected, partly because of the considerable variability in the specifics of plant design, operation, and siting. The results developed by PRAs reflect these plant-to-plant differences.

10. Systems that are important to reliable operation and the prevention of accidents are not necessarily the same as those that are important in risk mitigation.

5.3 Insights into Plant Risk

The analysts performing PRA studies gain valuable engineering and safety insights. Conceptual insights are the most important benefits of PRAs, and the most general of these is the entirely new way of thinking about reactor safety in a logic structure that transcends normal design practices and regulatory processes. PRA methods introduce much-needed realism into safety evaluations, in contrast with deterministic analysis which uses a conservative, qualitative approach that can mask important matters. The results of several studies, including the RSS, indicate important distinctions between contributors to different types of outcomes of potential accidents. The risk cannot be measured in terms of any single indicator, and changes in plant configuration that significantly affect one indicator may or may not affect the others. For example, a modification that reduces the frequency of core melt may not affect public risk and vice versa.

The results of PRA studies are expressed in terms of core-melt frequencies, frequencies of radionuclide releases of various

magnitudes, or curves presenting the frequencies of occurrence of different consequences (e.g., early and latent fatalities), depending on the level of the PRA. These are further discussed below.

5.3.1 Core-Melt Frequencies

The estimates of core-melt frequency range from 10^{-5} to about 10^{-3} per year. Variability of results is attributed to plant design, operation, site characteristics, PRA methods, analytical assumptions, and the scope of the studies. Sensitivity studies in each of these areas would be useful. At this time, caution must be exercised in comparing the quantitative results of various PRAs.

One of the results of a PRA study is the identification of the accident sequences that are the dominant contributors to core-melt frequency. An analysis of several published PRAs has shown the relative contribution to core-melt frequency of several salient features of the dominant accident sequences. This is shown in Figure 5-1 for both PWRs and BWRs.

Figure 5-1 provides a number of illuminating insights:

1. The split between LOCA and transient contributors to core-melt frequency is about equal for PWRs and about 10:90 for BWRs. However, some recently completed studies for newer PWRs indicate ratios similar to those for BWRs.

2. The failure of long-term decay heat removal is a major functional contributor to core-melt frequency for both PWRs and BWRs. It is associated with LOCAs in PWRs and with transients in BWRs.

3. Anticipated transients without scram are small contributors to core-melt frequency in PWRs but significant contributors in BWRs.

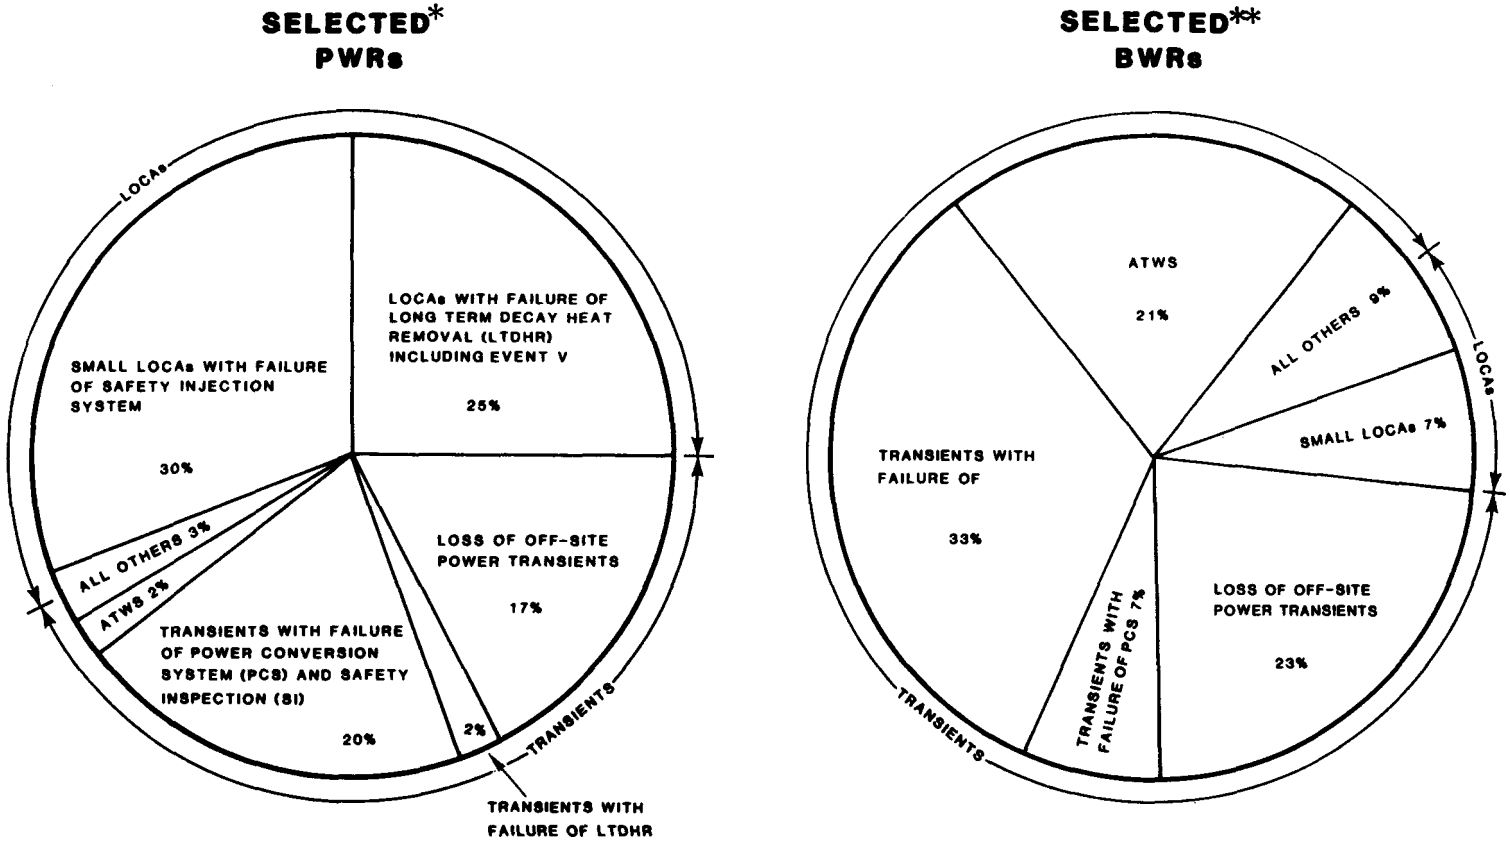

* Arkansas-1, Oconee-3, Sequoyah-1, Surry, Zion

** Grand Gulf, Limerick, Peach Bottom-2

"All Others" contain roughly equal mixes of LOCA and transients

NOMENCLATURE:

LOCA	Loss of Coolant Accident	PCS	Power Conversion System
SI	Safety Injection	LTDHR	Long Term Decay Heat Removal
ATWS	Anticipated Transients Without Scram	Event V	Interfacing Systems LOCA

Figure 5-1. Relative Core-Melt Frequency Contributions

4. As a group, small LOCAs with failure of long-term decay heat removal are large contributors to core-melt frequency for PWRs.

5.3.2 Radionuclide Releases

The results of many studies indicate that the dominant core-melt and dominant radionuclide release sequences largely coincide. This coincidence results from the conclusion that each core-melt sequence leads to a containment failure with a fairly high likelihood of a large radionuclide release to the atmosphere. Hence, the core-melt sequences with the higher frequencies generally yield higher frequencies of significant releases. A departure from this trend is seen in the Zion study, which did not find that all core-melt sequences automatically lead to containment failure.

The studies surveyed generally show that public risk is less sensitive to plant-system unavailabilities than to core-melt frequency. This is because risk is controlled more by the capability of the containment to withstand challenges to its integrity than by the unavailabilities of the safety systems that protect the core integrity. The frequency of significant radionuclide release, tends to decrease as the containments become stronger.

The accident sequences that appear to emerge as dominant contributors to release are those in which radioactive material bypasses the containment or the containment fails concurrently with (or shortly after) core melt. This early containment failure may be caused by major common-cause initiating events, such as unrecovered losses of offsite power, fires, or earthquakes. Such sequences are not necessarily the dominant contributors to core-melt frequency (e.g., interfacing-system LOCA). The ranking of core-melt and significant release sequences is shown in Table 5-1.

Table 5-1

Comparison of Core-Melt and Release Sequences

Sequence	Core-Melt Ranking	Significant-Release Ranking
	ZION STUDY	
Small LOCA: LTDHR failure	1	4
Seismic AC power loss	2	1
AC power loss and AFWS failure	13	2
Interfacing-system LOCA	16	3
	INDIAN POINT-2	
Seismic loss of control or power	1	3
Fires in electrical tunnel and switchgear room	2	4
Seismic (direct) containment failure	21	1
Interfacing-system LOCA	24	2
	INDIAN POINT-3	
Small LOCA failure of high-pressure recirculation	1	4
Fires in switchgear room and cable-spreading room	2	3
Interfacing-system LOCA	15	1
Seismic containment failure	37	2

Figure 5-2 provides frequency ranges for categorized radionuclide-release fractions from selected PRA studies. Three illustrative cases are displayed: (1) severe containment-failure modes (i.e., early overpressurization or containment bypass); (2) late containment failure; and (3) containment remains intact despite core melt. Only the nuclides most important from the standpoint of health effects are included--the noble gases (Xe, Kr), iodine (I), cesium (Cs), and tellurium (Te).

Figure 5-2 graphically displays some of the insights about source terms gained since publication of the RSS. The following points emerge:

1. Core melts may not always result in containment failure. For those that do not, the retention properties of the containment are substantial.

2. If the containment fails a long time after core melt, only moderate release fractions result. The range between the predictions of various studies is extremely wide for these cases and further resolution from current analytical and/or experimental programs should be illuminating.

3. Only containment bypass, early overpressurization sequences, or sequences involving common-cause containment and core cooling failures lead to large releases. Because of the existence of dose thresholds, the occurrence of early health effects is generally limited to these containment failure modes.

5.3.3 Offsite Consequences

PRA studies have provided a number of significant insights into severe offsite consequences. Several of these were listed in Section 5.2, and others are discussed in Appendices A and B. Clearly, all consequences are sensitive to the

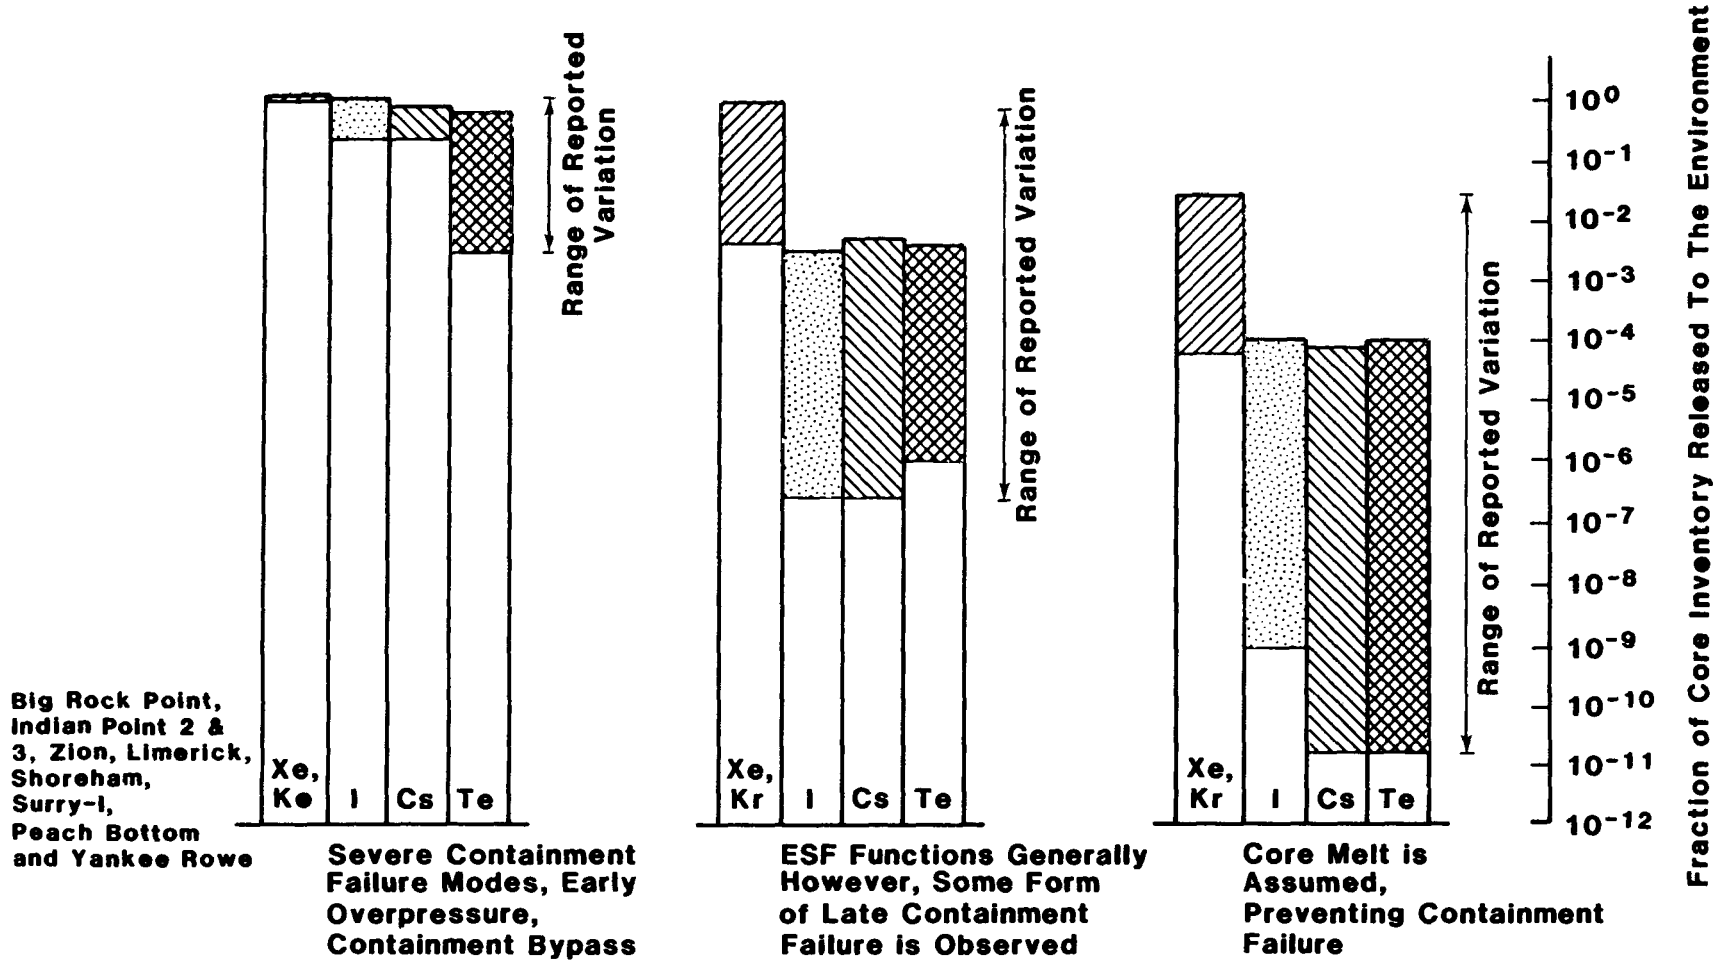

Figure 5-2. Range of Radionuclide Release Fractions from Listed PRA Studies

amount of radioactive material that could be released during an accident (the source term). However, early fatalities and injuries are particularly sensitive because of the existence of dose thresholds for these effects. If potential source terms are substantially smaller (by at least one order of magnitude), then the risk of early health effects generally would no longer be a principal concern. Nonetheless, the consequences of such accidents could still be large; the nature of the risk would shift to latent health effects and on the more localized problem of land contamination. Land contamination is roughly proportional to the quantity of long-lived radionuclides (mainly cesium) released. Tradeoffs between decontaminating an area, barring its use (interdiction), and a possible increased risk of cancer would need to be considered. In the limit, release of only the noble gases (krypton and xenon) could still result in significant offsite radiation exposure and population dose.

In addition to the source-term magnitude, the estimated number of early health effects is very sensitive to assumptions about the nature and effectiveness of potential emergency measures. For large releases of radioactive material, prompt evacuation and sheltering are potentially effective means of reducing the numbers of early health effects. Latent-cancer fatalities are not as sensitive to emergency response assumptions because larger areas and longer times are involved.

The weather (wind speed, rain, or dry weather) at the time of the accident can have a very large effect on offsite consequences. The variation in weather from site to site does not appear to affect the total risk appreciably because the probabilities of weather types that contribute the most to variation in consequences are not significantly different in different climates. However, total risk depends strongly on site characteristics (e.g., population density, land use); these considerations are important for reactor siting.

5.3.4 Insights into External Events

PRA studies have provided a new understanding of the importance of external events to public risk. In addition, specific insights into system response and methodology application have been derived. Some of the most significant insights are summarized below.

The results of the analysis of external events seem highly unpredictable. For seismic events the specifics of one plant's PRA results do not seem to be transferable to another plant, even though the plants may be similar. Although the specifics are different, the general character of fires is similar; major cable or control areas are involved and multiple redundant safety systems are affected. For flooding events, the results do not seem transferable from plant to plant.

Detailed analyses of external events have identified some accident sequences initiated by these events as important contributors to either core-melt frequency or to risk. Thus, the conclusion in the RSS that external events contributed only about 25 percent to plant risk has perhaps not been borne out; but this is difficult to confirm because of the uncertainties and potential uncertainties in analysis of external events.

For seismic events the following conclusions as indicated:

1. Earthquakes significantly larger than the safe-shutdown earthquake (SSE) are the significant seismic contributors to plant risk.

2. Most of the major accident sequences initiated by earthquakes involve seismic failures of plant structures and subsequent effects on the operability of safety systems.

3. Local ground and subsoil conditions have been an important issue in all PRAs investigating seismic events.

4. Earthquakes usually result in the loss of offsite power, which affects, the availability of systems important to safety.

Most of the fires found to be important to risk are those whose likelihood and/or severity are substantially reduced by the new NRC regulatory approach now being implemented (Appendix R and associated guides and standards).

For high winds, metal-sided structures are more fragile than other structures and most equipment, and are more likely to fail and compromise overall plant safety. Like earthquakes, high winds generally cause losses of offsite power, affecting system availabilities.

5.4 Dominant Accident Sequences

The RSS showed that the risk posed by the two plants that had been studied stemmed primarily from a few accident sequences. The relevance of these dominant accident sequences was immediately recognized. Uncertainties in the frequency or consequences estimated for these sequences would have the greatest effect on risk estimates. To achieve a significant reduction in risk, potential backfits or improvements in future designs would have to reduce either the frequency or the consequences of the dominant accident sequences. Thus, the understanding of risk, and the ability to effectively reduce risk, hinges on an understanding of the accident sequences that dominate risk.

The importance of the generic nature of the dominant accident sequences identified in the RSS was also recognized. If, for example, the dominant accident sequences were common for all

PWRs, then regulatory decisions or design alternatives reducing the risk from the dominant accident sequences would be effective for all PWRs. On the other hand, if PWRs in general had different dominant accident sequences, then a regulatory decision or design alternative could have different effects, and perhaps substantially different effects, for different PWRs.

Shortly after the RSS was published, the NRC instituted a program to address the similarities of dominant accident sequences for different PWRs and BWRs. The results of the program indicated that the dominant accident sequences are not consistent across very broad classes of plants (e.g., all PWRs or all BWRs) because each plant is unique and may exhibit accident sequences that are peculiar to its individual design and operation. The dominant sequences may be categorized according to the sequence of plant functions that failed (as opposed to the sequence of specific events that occurred). Two accidents may have different sequences of specific events yet have the same sequence of functional failures. Many dominant accident sequences can then be placed in the same functional sequence category.

This task is proceeding under sponsorship of the Office of Nuclear Regulatory Research in the Accident Sequence Evaluation Program (ASEP), which is concentrating on the accident sequences that dominate core-melt frequency. This program has reviewed a dozen existing PRAs, including both NRC and industry-sponsored studies. The dominant accident sequences identified have been assigned to functional sequence categories. As a group, these accident sequences summarize the functional sequences that have been found to dominate core-melt frequency in past PRAs. The industry-sponsored IDCOR program has also been assessing what functional accident sequences are the most important. These programs reflect the

current knowledge of which accident sequences have the greatest potential for dominating core-melt frequency in the current LWR population.

The plant functions used to prevent core melt or mitigate consequences differ with the initiating event, which is usually a LOCA or a transient. LOCAs are component or piping failures that result in a loss of cooling water from the reactor-coolant system. For LOCAs, the common set of functions performed by the mitigating systems is as follows:

1. Render reactor subcritical
2. Remove decay heat (core cooling)
3. Protect containment from overpressure caused by steam evolution
4. Scrub radioactive material from the containment atmosphere.

Transient events, as the term is used in PRA, are events which cause one or more physical parameters of the plant to exceed the normal operating range and for which prompt achievement of reactor subcriticality (scram) is desired. For transients, the common set of functions performed by the mitigating systems is as follows:

1. Render reactor subcritical
2. Remove core decay heat (core cooling)
3. Protect reactor-coolant system from overpressure failure
4. Protect containment from overpressure caused by steam evolution
5. Scrub radioactive material from the containment atmosphere.

Functional accident sequences can be defined in terms of the initiating event (transient or LOCA) and then by the subsequent functional failures. This approach was used in listing

the PRA functional accident sequences in Tables 5-2 and 5-3. The tables also show the range of accident-sequence frequencies that have been reported as central estimates in past PRAs. Appendix C describes some of the major design differences and uncertainties that contribute to the wide variations in frequencies among PRAs.

These functional sequence categories are those shown to dominate core-melt frequency and/or risk for a large number of plants. The category that actually dominates at a given plant depends on the features of that particular plant. Some of the more important features that tend to affect the dominance of accident sequences are listed as "major uncertainties" in Tables 5-2 and 5-3. It should be recognized that the specific component failure modes and human errors involved in these functional accident sequences can be expected to vary considerably from plant to plant. It must also be recognized that new functional accident sequences may be found in plants that have not yet been assessed by PRA methods. The above categories reflect current knowledge.

These sequences differ slightly from those described in Figure 5-1. First of all, they are described functionally rather than in terms of system behavior; second, Tables 5-2 and 5-3 show the range of frequencies that have been estimated, whereas Figure 5-1 shows the average. The functional sequences described in Tables 5-2 and 5-3 are the internally initiated sequences that are believed to have the greatest potential for dominating core-melt frequency or risk in LWR power plants today.

Very few of the PRAs performed to date included external events. Those that did indicate that the dominant accident sequences are quite plant-specific. Because of insufficient information, it is not possible at present to develop a meaningful list of externally initiated accident sequences that

Table 5-2

Functional Accident Sequence Categories (PWR)

Sequence Category	Freq. Range x 10^{-6}	Major Uncertainties	Comment
1) Transient Loss of Reactor Subcriticality	60 1	RPS Reliability RCS Ability to Withstand Pressure Spike	ATWS Rule Pending
2) Transient Loss of Integrity Loss of Core Cooling	30 <1	PORV Demand Rate HPIS Availability Necessity to Switch-Over to Recirc.	TMI Fixes (Raising PORV Set Point and Anticipatory AFWS Start Signal) Should Reduce Sequence Freq.
3) Transient Loss of Core Cooling	1000 0.1	Feed and Bleed Capability AFWS Availability	TMI Fixes Have Called for Many Improvements in AFWS Availability
4) Transient Loss of Core cooling Loss of Containment Heat Removal	100 0.2	Redundancy of AC Power Sources Battery, CST Depletion Times Possibility of Induced RCS Pump Seal Leak Long Term Ventilation Loss Effects AFWS Availability	NRC Position Statement-Forthcoming
5) LOCA Loss of Core Cooling	200 ≤0.4	LOCA Frequency ECCS Success Criteria ECCS Redundancy	Small LOCA May Be Higher Than Thought Due to RCP Seal Leaks TMI fixes Stressed Better Procedures for Small LOCA
6) LOCA Loss of Core Cooling Loss of Containment Heat Removal	6 <1	LOCA Frequency ECCS Success Criteria ECCS Redundancy	Small LOCA May Be Higher Thought Due to RCP Seal Leaks TMI Fixes Stressed Better Procedures for Small LOCA

Table 5-3

Functional Accident Sequences Categories (BWR)

Sequence Category	Freq. Range x 10^{-6}	Major Uncertainties	Comment
1) Transient Loss of Reactor Subcriticality	50 0.1	RPS Reliability Adequacy of ECCS Unknown Phenomenology in RCS Ability of Open to Control Water Level	ATWS Rule Pending
2) Transient Loss of RCS Integrity Loss of Core Cooling	70 <0.2	ECCS Availability Operator Procedures for ADS SRV Demand Rate	
3) Transient Loss of RCS Integrity Loss of Containment Cooling	1000 0.1	RHR Availability SHV Demand Rate	Estimated Time to Core Melt Appear Longer Than Previously Expected, Thus Longer Times for Recovery
4) Transient Loss of Core Cooling	700 0.2	ECCS Availability Operator Procedure for ADS	Station Blackout Rules Pending
5) Transients Loss of Containment Cooling	100 <0.4	RHR Availability ECCS Success Criteria ECCS Redundancy	Estimated Time to Core Melt Appear Longer Than Previously Expected, Thus Longer Times for Recovery
6) LOCA Loss of Containment Cooling	5 <0.1	RHR Availability Time Available for Recovery	Estimated Time to Core Melt Appear Longer Than Previously Expected, Thus Longer Times for Recovery

RPS - Reactor Protection System
SRV - Safety Relief Valve
RHR - Reactor Heat Removal

have the greatest potential for dominating core-melt frequency or risk in LWR plants.

In performing generic studies, analysts would like to identify which plants can be expected to have similar dominant accident sequences. Different PRAs sometimes identify different dominant accident sequences. (The difference may reflect increased knowledge, particularly when the results of recent PRAs are compared with those of earlier studies.) To remove this difference, the NRC is sponsoring research to update some of the older PRAs. The remaining differences reflect actual design or operational differences among plants and must be recognized and considered in generic regulatory decisionmaking. This can be done by identifying the plant features that most affect the frequency and consequences of these potential dominant accident sequences, identifying the plants that share these common features, and assigning the plants to groups with common characteristics for the features that have the greatest effect on the dominant accident sequences. The analysis necessary for such an approach is being sponsored by the NRC. Its result will be the grouping of plants into categories for which generic regulatory decisions would be expected to exert similar effects in terms of risk reduction.

In summary, the following points can be made:

1. Dominant accident sequences can be expected to differ for different plants and have relatively wide frequency ranges as a result of differences in plant design, and operation, and siting.

2. Despite plant differences that affect sequence frequencies, generic studies that support regulatory decisionmaking can be performed effectively by assigning plants to categories with similar dominant accident sequences.

5.5 Important System Dependences, Functions, Systems, and Human Interactions

In addition to the overall quantitative results for core-melt frequency and risk, PRA studies provide information on the individual contributors to these results. They indicate which aspects of plant design, operation, and siting are important to risk and how they are related. Some of these aspects have been shown to have a general applicability as patterns in the PRA results have emerged; this is especially true on the functional level. Other findings of PRAs relate to systems, human interactions, and dependences. These findings, described in the sections that follow, can assist in decisions regarding potential risk-reduction modifications, the assignment of reliability-assurance priorities, and the evaluation of the results of new assessments.

5.5.1 Important Specific Findings

Many of the contributors to core-melt frequency or risk result from interactions among systems, events, and phenomena. Interactions or dependences pertinent to external events are described in Section 5.3 and Appendix C. Additional insights into dependences and their treatment in PRAs are as follows:

1. The dependence of multiple systems on a common service such as pump cooling or room cooling is a major contributor to accident sequences. However, in these sequences a long time is generally available before the release of radioactive material, which gives the operator the opportunity to recover from initial support-system failures.

2. The importance of recirculation failures to core-melt frequency in PWRs depends on the ability to use high-pressure and low-pressure systems independently and the mode of switchover to recirculation (manual or automatic).

The propensity for human error in switchover from injection to recirculation under LOCA conditions is an important consideration.

3. Some systems that are cross-connected between units at multiple-unit sites improve the availability of support systems because of improved flexibility and thus diminish the effect of support-system failures on accident-sequence frequencies.

4. For BWRs, the loss of long-term containment heat removal was considered important in past PRAs because it eventually resulted in the failure of coolant makeup systems. However, because of the long times involved, the operators have considerable time for recovery actions, which can reduce the importance of these accidents to core melt.

5. Operator error is often a significant contributor to coolant-injection failures in the case of transients and small LOCAs in BWRs. This happens because the operator may fail to initiate depressurization if the high-pressure systems are unavailable, and automatic depressurization may occur too late to protect the core.

6. High-pressure events contribute more to overall risk than do low pressure events, particularly in PWRs.

7. In BWRs, the progression of low-pressure events is much slower than it is in PWRs.

8. The risk of low pressure events in BWRs can be reduced significantly by recovery actions.

In addition to these findings on dependences, some PRA results imply that plant analyses and understanding may be adjusted

in the future to reflect current perceptions more accurately. Some of these perceptions are summarized below.

First, recovery actions can have a dramatic effect on accident progression. Thus, current assessments of core-melt frequency and release fractions should consider potential operator recovery actions.

Second, at PWRs, long-term accidents like small LOCAs may not require switchover to recirculation. These accidents may be mitigated by switching to closed-cycle cooling (e.g., the residual-heat-removal mode). This additional possibility may reduce the core-melt frequency estimated for these accidents.

Third, PRAs have indicated that system-success criteria based on licensing considerations may be too conservative for severe-accident prevention. For instance, if, in PWRs, the decay-heat-removal heat exchangers are rejecting heat from the recirculating emergency coolant stream, containment systems (fan coolers and/or sprays) may not be necessary for successful accident mitigation. Other accident sequences have revealed similar findings about system-success criteria. The resolution of these questions may lower the overall risk that is estimated for some accidents.

5.5.2 Relative Importance of Systems

In investigating design alternatives and in establishing surveillance programs, analysts must understand which systems are most important. It is not easy to establish the relative importance of systems because (1) many ways exist for defining importance; (2) the importance of systems depends on the dominant accident sequences, which differ for different plants; and (3) the systems are interdependent, particularly the support systems that provide electric power, cooling, control, and other functions that support the main systems.

Nevertheless, attempts have been made to clarify the subject. Under NRC sponsorship, Battelle Columbus Laboratories (1982) issued a draft report that addresses two risk-importance measures to evaluate a feature's importance in further reducing the risk and its importance in maintaining the risk level. One of the importance measures, called the feature's "risk-reduction worth," was developed for use in assigning priorities to future improvements. The second type of importance measure, called the feature's "risk-achievement worth," was developed for assigning priorities to features that are most important in reliability assurance and risk maintenance.

The Battelle Columbus study applied the risk-worth measures to the four plants studied in RSSMAP: Oconee, Grand Gulf, Calvert Cliffs, and Sequoyah. The four plants employ light-water reactors of the two major types (BWR and PWR), the four types of nuclear steam supply systems (General Electric, Westinghouse, Babcock & Wilcox, and Combustion Engineering), and three containment types (large dry, Mark III BWR, and ice condenser). The four studies provide a good opportunity for comparing the importance of systems because the PRA methods for the studies were generally similar.

Reproduced as Figure 5-3 is a graph showing the risk-achievement ratios and the risk-reduction ratios with core-melt frequency as the risk measure. The risk-achievement ratios are graphed above the dividing line and indicate the factor by which core-melt frequency would increase if the system had a failure probability of unity (that is, it was never operable). The risk-reduction worths are graphed below the dividing line and indicate the factor by which core-melt frequency could be reduced at the plant by improving system reliability. Also shown is human action identified by RSSMAP as having the largest risk-achievement worth.

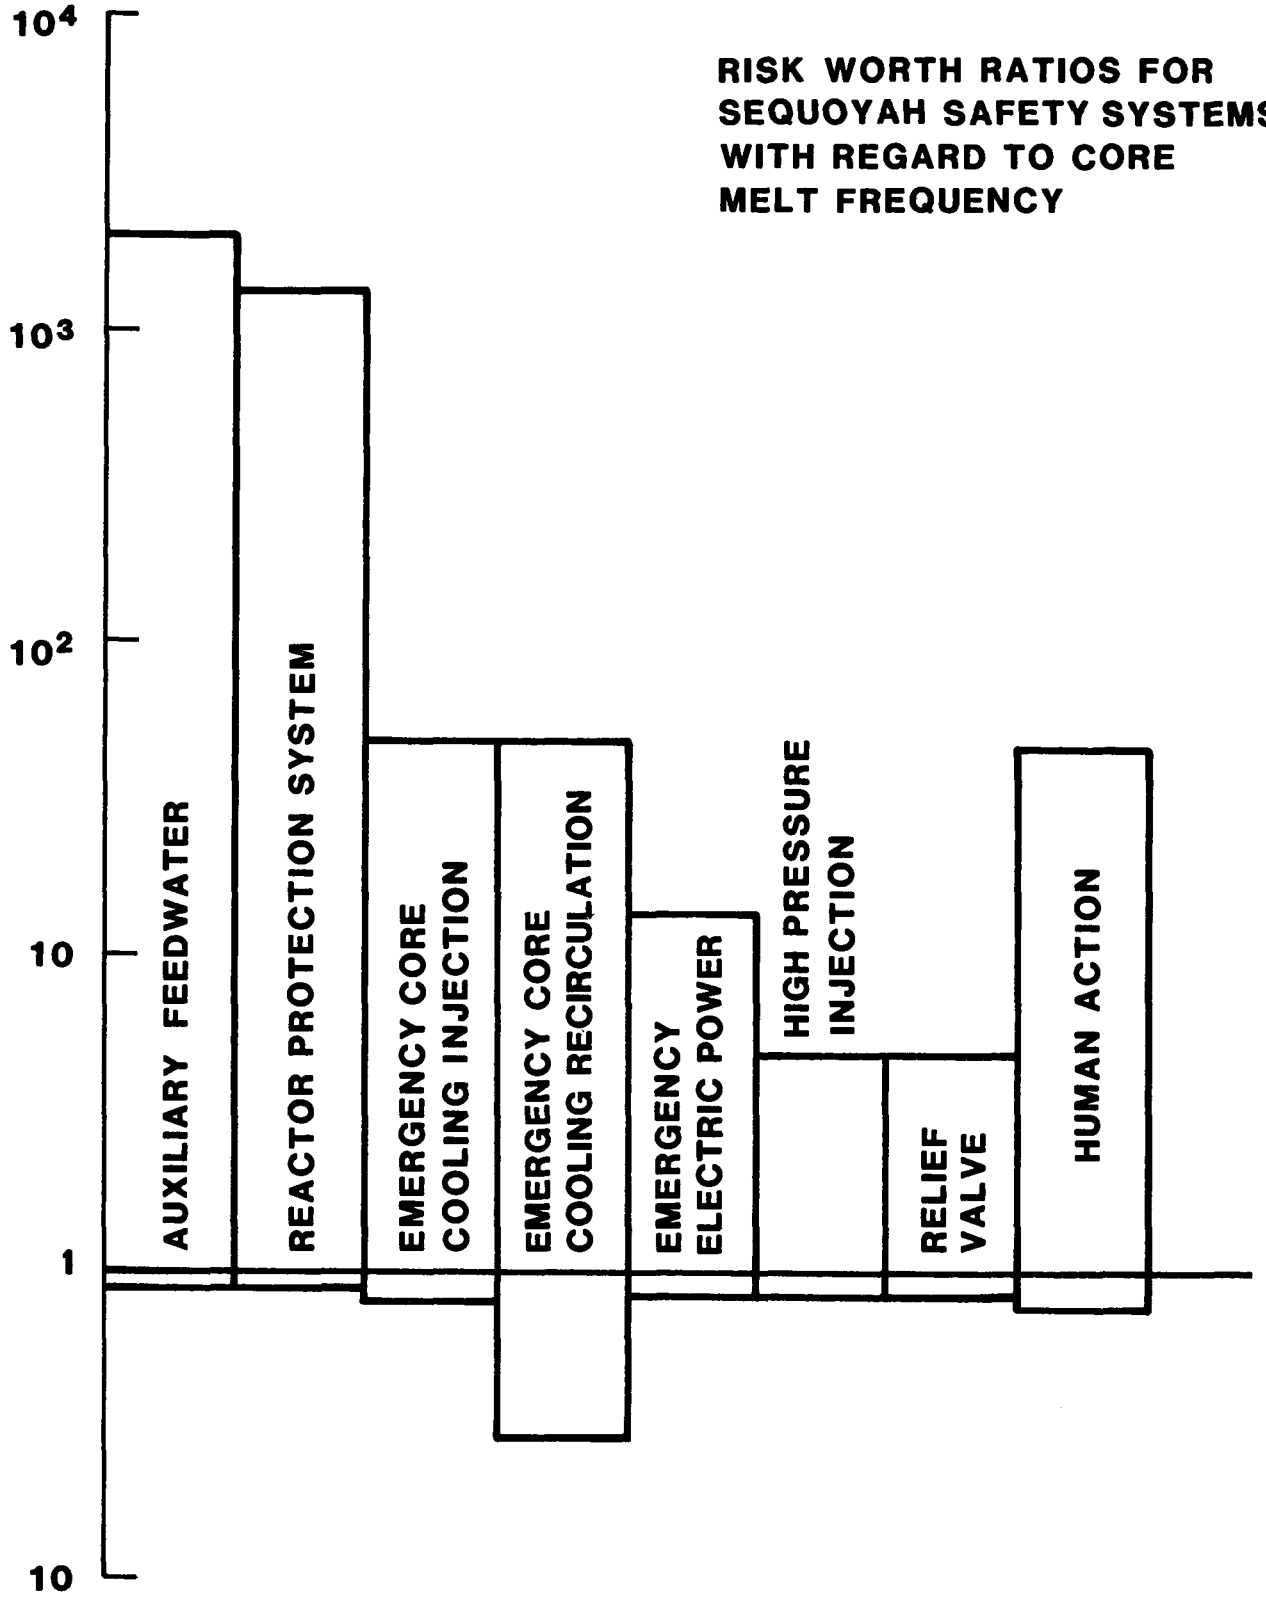

Figure 5-3. Risk Worth Ratios for Sequoyah Safety Systems with Regard to Core-Melt Frequency

Figure 5-3 shows that a very significant increase in core-melt frequency could occur if the reliability of important plant systems were allowed to deteriorate below that predicted by PRAs. The figure thus emphasizes the need for a sound reliability assurance program to ensure that this deterioration does not occur.

Another important study performed under the sponsorship of the Office of Nuclear Reactor Regulation reviewed 15 published PRAs and estimated the relative importance of systems from their contribution to the dominant accident sequences in Figure 5-3. Both BWRs and PWRs were considered. The results are shown in Figures 5-4 and 5-5. The arrows indicate that each system was not involved in the dominant accident sequences in at least one PRA.

The IDCOR program has also arrived at some generic conclusions about the relative importance of systems. For PWRs, the following systems are fairly consistently most important for all plants:

1. Auxiliary feedwater system
2. High-pressure injection system
3. Low-pressure recirculation system

For BWRs, less consistency was found but, in general, the following systems often appeared important:

1. Power-conversion system
2. High-pressure injection system
3. Reactor-core isolation cooling system
4. Reactor-protection system
5. Residual-heat-removal system

The modest differences in the above three studies reflect differences in measures of importance, differences in system

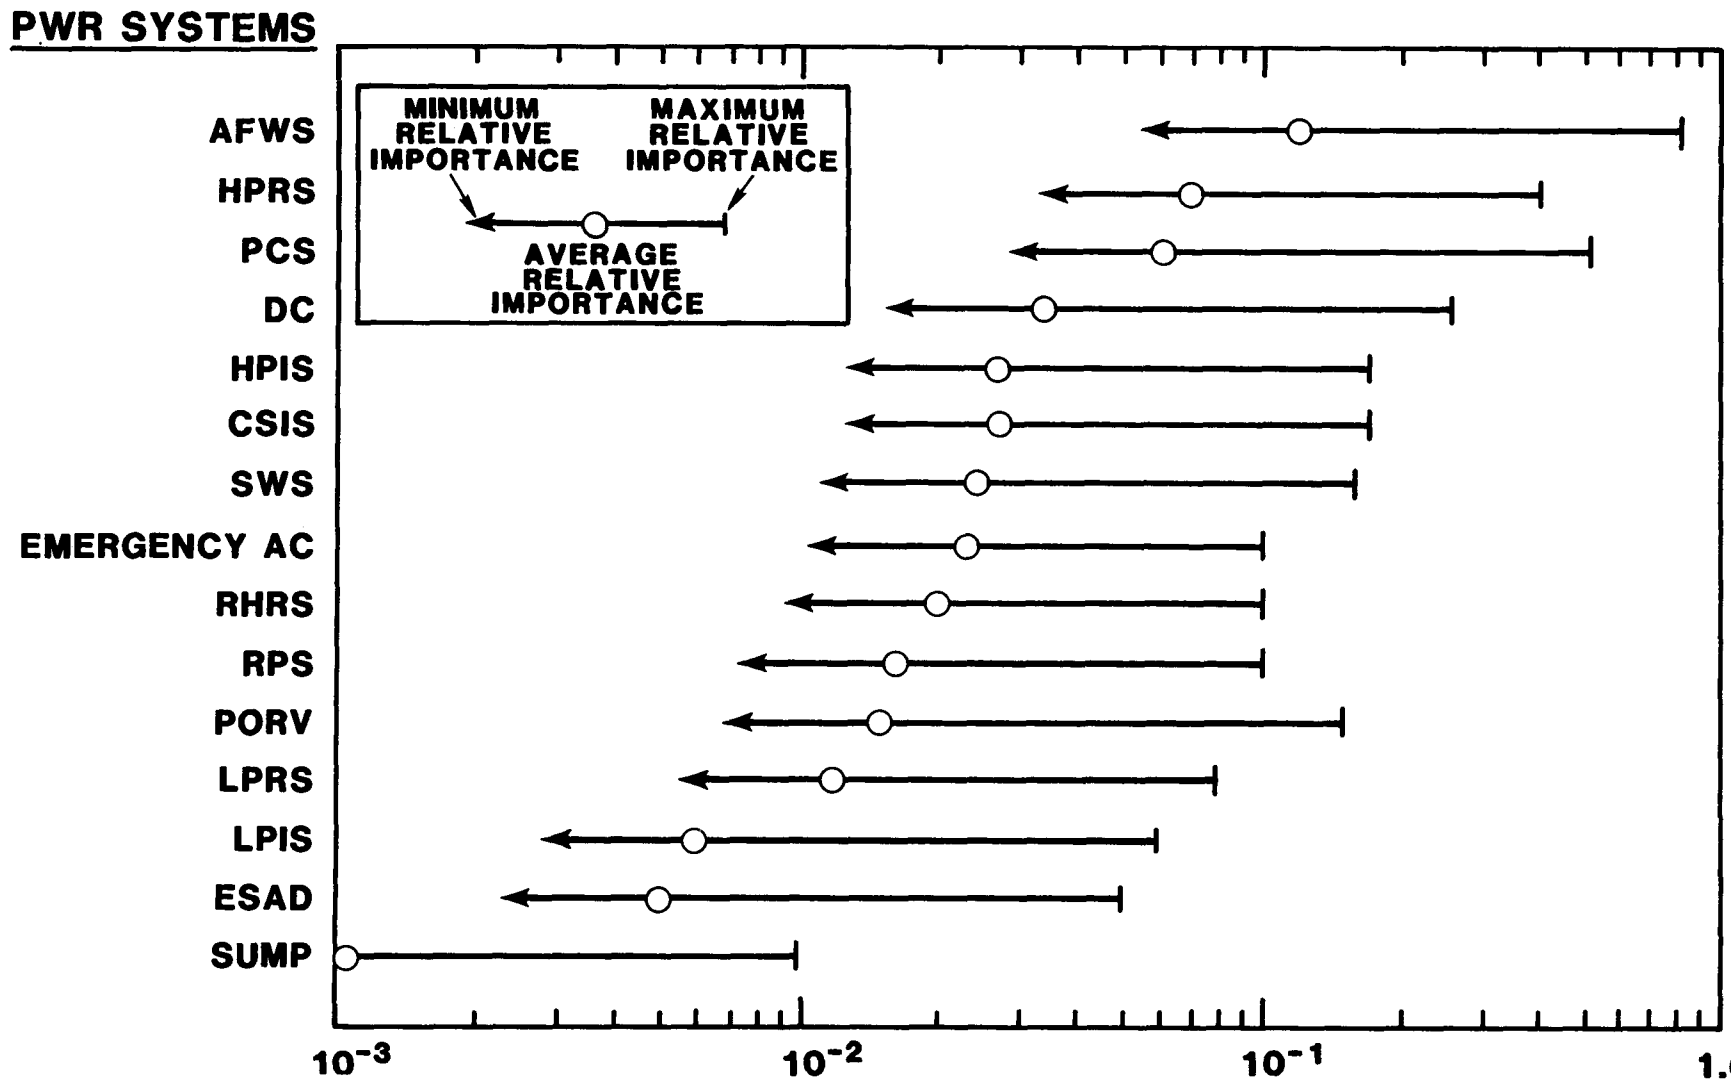

Figure 5-4. Relative Importance of PWR Systems Considering Dominant Accident Sequences from 15 PRAs

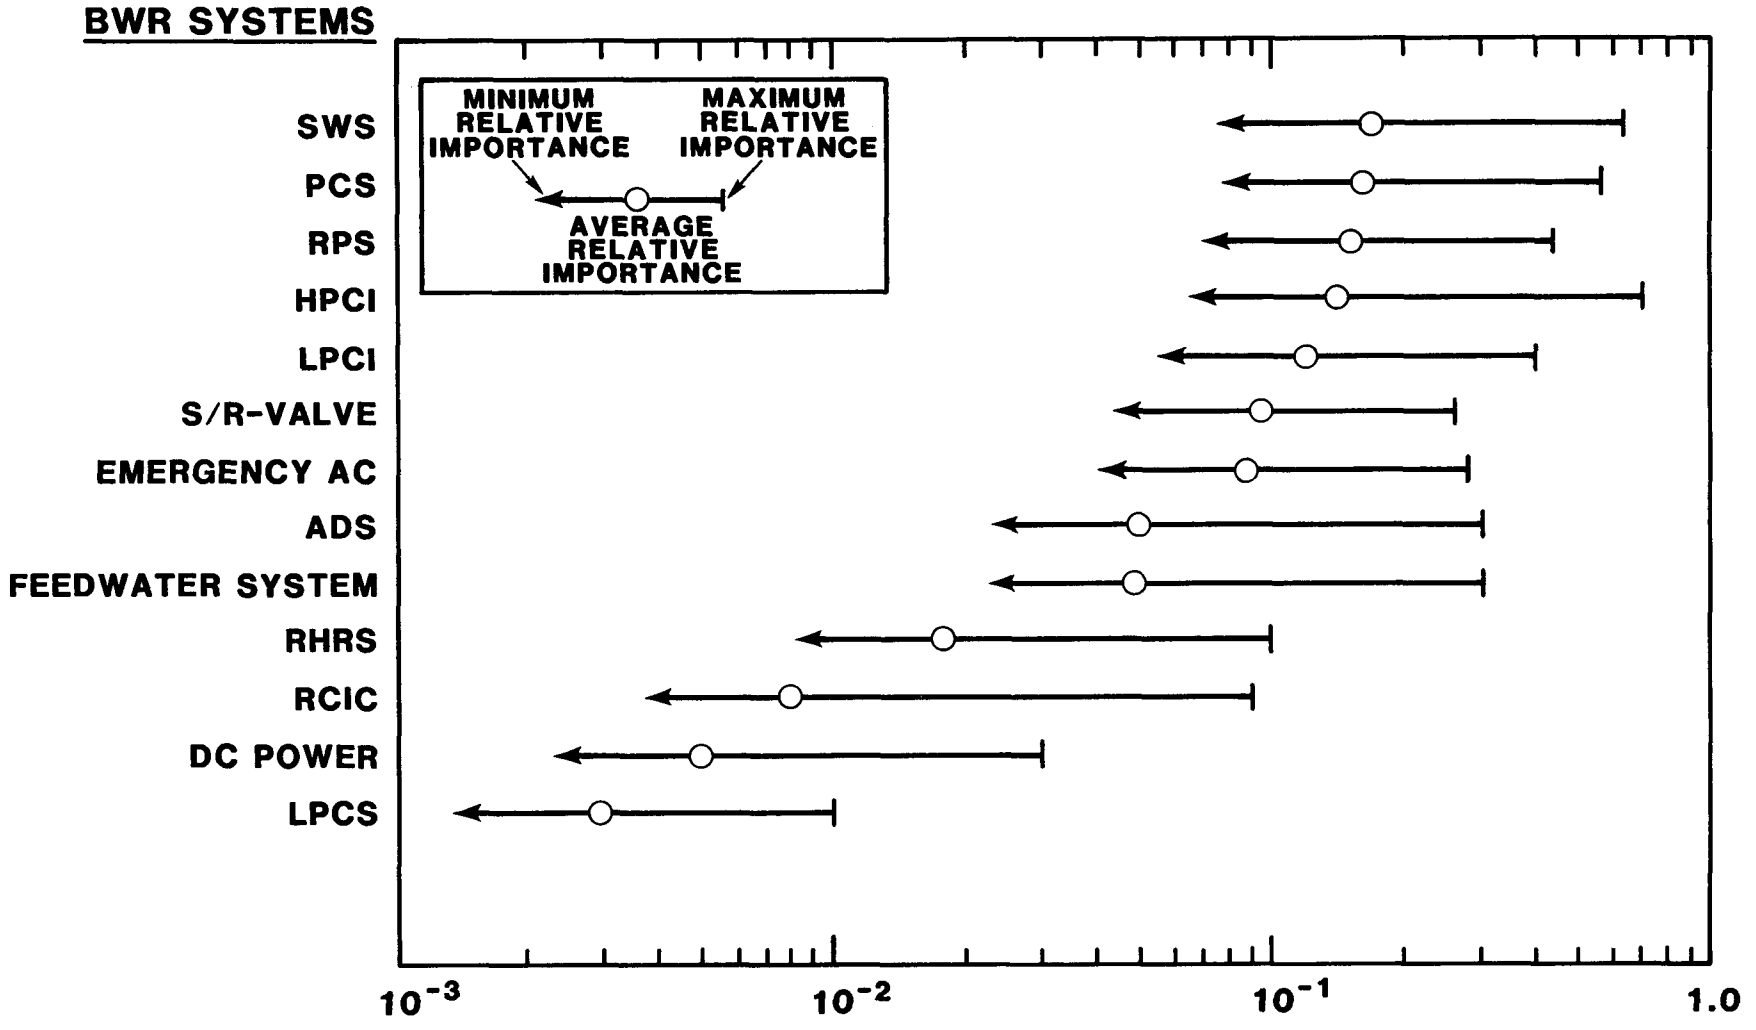

Figure 5-5. Relative Importance of BWR Systems Considering Dominant Accident Sequences from 15 PRAs

boundary definitions (i.e., whether support systems were included in the front-line-system definition), and differences in the range of PRAs considered and how much updating was done on older PRAs.

Two conclusions can be reached. First, the systems important from a risk-improvement standpoint may be different from those considered important in a reliability assurance program. Second, the relative system importance is very specific to the plant.

5.5.3 Relative Importance of Human Error

All of the importance studies described above found human errors to have relatively high importance in specific situations. As in the case of system importance, the specific human errors varied from plant to plant. Generic conclusions about the relative importance of these errors are difficult. The two human errors that appear the most consistently to be important in PRAs are failure to switch over properly to recirculation during PWR LOCA sequences and failure to initiate the automatic pressure relief system manually after the failure of high-pressure injection in small LOCAs in BWRs.

5.6 Insights from Precursor Studies

An ongoing study is examining operating experience data and assessing plant safety as it is reflected by the operating experience. A report based on analyses of operating data reported from 1969 to 1979 was published in 1982 and subjected to intensive peer review. Analyses of later operational events are continuing.

The work performed to date, viewed in light of the comments submitted during the peer review, supports the following insights:

1. Accident precursors can generally be assigned to one or another of the generic-accident sequence classes previously identified in PRAs. However, the precursors may include unique or unusual failures or interactions. This suggests that the limit of resolution of the PRA methodology may be at the system or component failure level, with a more limited capability to evaluate specific component-failure modes.

2. Many of the initiating-event frequencies and function unavailabilities developed from operating experience agree reasonably well with PRA results.

3. No evidence exists that the rate of occurrence of significant precursors varies with plant age.

4. The number of potential precursors does not vary significantly among reactor vendors or architect-engineers.

5. Human errors are involved in a significant percentage of major precursors. Operator errors of commission are not modeled well in PRAs.

6. Losses of offsite power and losses of feedwater contribute significantly to core-melt frequency, as predicted by PRAs. However, LOCAs do not seem to be as important as predicted by PRAs.

5.7 Insights Regarding Reliability Assurance

PRA presents a "snapshot" of the risk profile at a given plant at a given time. As time progresses, modifications to plant equipment or procedures (i.e., operating or maintenance practices) can change the risk profile. Furthermore, as operating data accumulate, the improved information base may suggest that the generic failure rates used for some components should

be modified or that the potential for dependent failures differs from the potential previously assessed. Thus, there is a need to update the analyses and to make the PRA essentially a "living" document that reflects the impact of plant modifications and acquired data.

As discussed in Chapter 4 and Appendix B, there are techniques that permit an analyst to measure the incremental effect of a degradation in a given safety function, system, or component. Such analyses permit the plant operator or owner and the NRC to focus inspection and quality-assurance activities on the plant features that could significantly increase the core-melt frequency or risk estimates of the PRA. The features identified by such an analysis may not necessarily be those that are major contributors to risk. Rather, they are the features that could become dominant, if their failure characteristics are degraded significantly in relation to those used in the analysis. Other importance measures are useful in identifying the important contributors to the assessed risk. These are useful in deciding where to improve the plant if a reduction in risk is desired.

The availability of an updated PRA would also make possible a means for interpreting the significance to risk (or core-melt frequency) of variations in component-failure rates as determined by acquired plant-specific data. Similarly, plant models could be compared with actual occurrences to ensure that they reflect the best information on plant performance and interactions between systems and components.

The use of PRA techniques alone will not constitute an adequate reliability assurance program. PRA techniques at present have limited application to such potential problems as improper design, faulty installation, or improper specifications of performance requirements. Thus, PRA techniques must be integrated with appropriate quality-assurance and quality-

control approaches for a comprehensive reliability or safety assurance program, with the PRA techniques providing key information about the risk impacts of reliability-assurance alternatives.

5.8 General Insights Regarding Improvement of Plant Safety

Next to an explicit quantification of public risk, the identification of specific safety concerns and the evaluation of possible solutions are probably the best recognized and most widely used applications of PRA. The performance of a PRA naturally leads to significant improvements in the understanding of the design and operation of the various systems, the response of the containment, and the role of plant operators under accident conditions. This understanding, in turn, often reveals design or procedural modifications that can improve safety.

Table 5-4 lists some of the plant-specific changes either that have resulted from PRA results or are planned. Each of these changes was made because a PRA identified a significant deficiency in the existing design of a plant system or in a maintenance procedure. The insights that led to these modifications originated from PRAs of all levels (as defined in Section 3.2.2). The use of PRA to reach generic insights regarding safety improvements is being investigated by several ongoing or recently completed programs sponsored by both the NRC and by the industry (see Sections 3.5 and 3.6). These insights span improvements in several areas, including (1) existing plant system design and maintenance procedures, (2) operating procedures for severe accidents, and (3) consequence mitigation add-ons. Each is explored in turn.

Generic insights regarding plant system design and maintenance procedures generally evolve when a plant-specific conclusion is replicated over a number of plants. For example, the

Table 5-4

Examples of Plant Modifications Made or Committed to Based on PRA Insights

Plant	Plant Modification	PRA	Level
Arkansas Nuclear One	Station battery test scheduling changed to reduce probability of common-mode failures	IREP[a]	1
Arkansas Nuclear One	AC and DC switchgear room cooler actuation circuitry test procedure established	IREP	1
Millstone	Logic changes made to emergency AC power load sequencer to eliminate single failure	IREP	1
Sequoyah	Procedures changed to ensure that upper compartment drain plugs are removed after refueling	RSSMAP[b]	2
Oconee	Procedure and hardware changes made to reduce the frequency of interfacing system LOCA	RSSMAP	2
Indian Point	Upgrading of charging-pump alternative shutdown power supply to reduce the probability of RCP seal failure	IPPSS[c]	3
Indian Point	Replacement of manual valves with motor-operated valves in fan-cooler service-water lines	IPPSS	3
Big Rock Point	Hardware modification to restrict flow in reject line between condenser hotwell and condensate storage tank	BRP[d]	3

[a]Interim Reliability Evaluation Program.
[b]Reactor Safety Study Methodology Applications Program.
[c]Indian Point Probabilistic Safety Study.
[d]Big Rock Point Probabilistic Risk Assessment.

importance of the interfacing-system LOCA was originally identified in the RSS and was replicated for the plants analyzed in RSSMAP. This inferred the need for increased attention to maintenance procedures for the interfacing system check valves. In general, the results obtained in the Severe Accident Risk Reduction Program (SARRP) and the IDCOR program indicate that a modest overall reduction in core-melt frequency may be possible from specific hardware or maintenance improvements to existing systems. Such improvements include modifications to the auxiliary feedwater systems, improvements in emergency AC power systems, modifications to the reactor-protection system, improved maintenance for ice-condenser floor drains, etc.

Insights about operating procedures for severe accidents generally come from PRA findings regarding the progression of dominant accident sequences and the role of the operator during these sequences. For example, the following types of operator error have been found to be important in many PRAs:

1. Failure to realign the emergency core-cooling system manually from the injection mode to the recirculation mode when the water inventory in the refueling water storage tank falls below a set level (PWRs).

2. Failure to initiate the bleed-and-feed mode in PWRs or to actuate the automatic depressurization system in BWRs when the reaction-coolant system is at high pressure during accidents initiated by transients.

3. Failure to initiate the liquid-poison injection system or to insert control rods manually during accidents involving a failure of the reactor-protection system in BWRs.

Recognition of the importance of specific operator actions can be a vital first step toward defining both appropriate

procedures for the management of severe accidents and appropriate approaches for operator training. To date, PRA insights into man-machine interfaces have not been used as effectively as they could be.

Insights about consequence mitigation systems draw from PRA findings about the types of loading that pose the most serious threats to containment integrity. Currently, the uncertainties regarding containment loading and response are large; various task forces and projects at the NRC and within the industry are addressing the problem. Nevertheless, the following insights on containment-failure modes and applicable mitigation approaches have tentatively emerged:

1. For the strongest PWR containments, most of the offsite risk (health effects and property damage) is associated with the containment overpressurization from steam after core melt. Often this results from the loss of all AC power. Applicable mitigation systems might include a low-volume filtered vent or an AC-independent containment cooling or spray system.

2. For the less strong PWR containments, including subatmospheric containments and ice condensers, NRC-sponsored studies have identified early failures from hydrogen burning and out-of-vessel "steam spikes" as being potentially important. Thus, hydrogen control and containment-water management may be effective mitigation approaches for these containments. The IDCOR program has expressed a dissenting viewpoint, however.

3. For BWR containments, overpressurization from steam before core melt and overpressurization from steam and noncondensable gases after core melt are important contributors to risk. A significant fraction of the risk is attributed

to anticipated transients without scram (ATWS). Mitigation approaches could include a high-volume unfiltered vent (for the ATWS) together with those mentioned above for the strong PWR containments.

6. REGULATORY USES OF PROBABILISTIC RISK ANALYSES

6.1 Introduction

The evolution of PRA methods for the study of reactor safety has been rapid in the past few years. Today, the number of full-scope PRAs completed or under way is large, the number of skilled practitioners has grown rapidly to many dozens, and the applications of PRA have spread to include many (perhaps most) of the important areas of concern in reactor safety. This is a remarkably rapid growth, especially considering the history of PRA. This section describes the development of regulatory usage of PRA over the last decade.

Chapters 4 and 5 clearly show that PRA techniques generate useful information about, and insights into, the design and operation of a nuclear power plant by providing an improved understanding of the full range of accident sequences and techniques for assessing their importance. The regulator can use this information to supplement the decision process. Also, clearly, many limitations in our knowledge lead to uncertainty in the quantification of the risk which must be reflected in the use of the results of PRAs. Proper use of the results of PRA in the regulatory process should emphasize the applications that rely heavily on well-established methods and minimize the uses that rely heavily on the methods that are weak.

The regulatory decisionmaker must evaluate each analysis, whether deterministic or probabilistic, and judge whether the assumptions and boundary conditions employed by the analyst are sufficiently valid and the results sufficiently robust to justify using of the analysis in making regulatory decisions. No technical analyses, whether deterministic or probabilistic, are ever formally complete or completely certain. In most instances, the uncertainties identified in PRAs are also

inherent in deterministic analyses. Therefore, the decision-maker must understand all significant uncertainties so as to optimize the use of all available analyses, including the information contained in PRAs.

Because PRA results are often associated with large uncertainties, the uncertainties should be considered in the decision process. The uncertainty range for a given result, such as for the core-damage frequency, defines the range within which the true value is likely to fall with some associated confidence or probability. Any point estimate, such as a median or a mean, is simply one value chosen from the range of possible values. Because of the large uncertainties and large ranges that generally exist for PRA results, the use of point estimates alone can lead to less informed decisions. While some theoretical approaches are available for explicitly incorporating uncertainties in decisionmaking, they have not been fully applied to PRAs. The PRAs that have been performed and have been reviewed in the preceding sections indicate that the following guidance should be followed for uncertainties:

First, because of the arbitrariness in the details of the statistical techniques, the details associated with a calculated uncertainty range are not very meaningful. Neither the uncertainty range nor the probability distribution of values through the range are known as precisely as the PRA results might indicate. An uncertainty range can thus be viewed as a range within which the true value can be expected to lie with some high but fuzzy confidence, such as roughly 90 percent, with some unknown distribution.

Second, the uncertainty ranges that were generally estimated for core-melt frequencies and risks in past PRAs were applicable to data uncertainties (i.e., imprecisions in statistical estimation, uncertainties in data extrapolation, and unit-to-

unit variations). Uncertainties attributable to modeling and assumptions were not generally included in available PRA uncertainty analyses, and their impacts need to be considered separately in sensitivity analyses.

Third, when making comparisons with another risk-based analysis or with a criterion or goal, the calculated uncertainty bounds must be carefully examined. If uncertainty bounds (modified to incorporate the effects of uncertainties attributable to modeling and assumptions) do not overlap, the decisionmaker can assign high confidence to the results, provided they have been subjected to an adequate peer review.

Obviously, the decisionmaker does not require perfect information, and it would be inappropriate to dismiss PRA information simply because overlap occurs. In any decision process, all available information should be considered, and credibility should be based not only on the estimated statistical accuracy but also on the judgments of technical experts, and the degree of conservatism appropriate for the decision. The current state of knowledge incidates that generally a difference in point estimates (single values, medians, or means) of an order of magnitude or more is likely to be very significant, whereas differences of less than a factor of 3 will generally overlap (i.e., the upper bound of the smaller estimate will be greater than the lower bound of the larger estimate) so that the difference in the two point estimates is less significant. Of course, where uncertainties are very large, the above general guidance would have to be used with caution. Even if a significant degree of overlap appears in the uncertainty bounds, the decisionmaker still may have useful information, but he or she must recognize the possibility that comparisons of point estimates may be in error, with this potential for error depending on the degree of overlap. The likelihood of improperly comparing items is difficult to quantify accurately in such a case, because the probability distribution of possible

values about the point estimate is unknown. Nevertheless, qualitative evaluation can lead to useful information.

Finally, where significant overlap occurs, a qualitative engineering evaluation of the dominant failure characteristics and the integral knowledge gained from a disciplined attempt to model plant performance may provide information useful to the decisionmaker. For example, if quantitative results are for practical purposes indistinguishable, but a plausible and previously unidentified failure path has been identified, the decisionmaker has acquired useful information. Clearly, in the discussions that follow PRA results and insights supplement the information that would be available to the decisionmaker from deterministic evaluations alone. However, the PRA results will not make the decision for the regulator, nor should they. Many types of regulatory decisions exist, and the weight given to the quantitative PRA results depends on the degree of precision necessary and other factors affecting the decision. Even the qualitative insights gained from PRAs are, in most cases, based to some degree on quantitative results.

This chapter is therefore organized in broad categories of potential usage, grouped in terms of the degree to which they require precision in quantitative results. These categories, which are discussed after a summary of past PRA applications, are assignments of priorities for the use of resources, generic regulatory applications, and plant-specific uses. Section 6.8 deals with decisionmaking, using PRA as one ingredient.

6.2 Past and Present Practices

The first comprehensive application of PRA techniques was the RSS, which is widely accepted as a revolutionary piece of work that broke new ground in many areas. The RSS was the first broad-scale application of event-tree and fault-tree methods

to a system as complex as a nuclear power plant. Its principal objective was to reach some meaningful conclusions about the risks of commercial nuclear power plants.

For various reasons, the RSS became one of the most controversial documents in the history of reactor safety. The report was attacked on several grounds. Its conclusions were criticized as being used by reactor proponents to "prove, once and for all," that reactors were "safe" and that the report's hidden agenda had been to show how safe reactors were regardless of the truth. These allegations were completely refuted (Lewis et al.). The uncertainties in the final results were also criticized as being understated.

The discourse about the RSS, important as it was to the public acceptance of reactors and to the credibility of the regulatory authorities, created a reluctance to use PRA methods, in spite of the fact that the RSS was an important source of information about reactor safety. It uncovered or illuminated potential safety issues that were not effectively used in the late 1970s. The impact of the controversy was demonstrated by the NRC's reaction to the Lewis report (Lewis et al.). The Commissioners asked the NRC staff to document where, if anywhere, they had relied on RSS results or insights in the years since its publication in 1975. The staff responded by producing a rather voluminous report outlining essentially every regulatory action in which the RSS had been cited, including letters to licensee representatives, hearing testimony, and more formal safety reports and decisions. The staff document, produced in early 1979 just before the TMI accident, exemplifies what the Lewis Committee called the "siege mentality." The staff concluded that, with only one or two exceptions, no RSS insights or results had been used as a substantive part of any staff decisions or actions. RSS results or methods were applied on a few occasions shortly after its publication, and these applications were important.

One was the technical basis for the revised evacuation planning guidelines of the "Emergency Planning Task Force" report where RSS results provided the basis for the 10- and 50-mile emergency planning zones for plume and ingestion exposures. Another was the assignment of risk-based priorities to the "unresolved generic safety issues." A third important issue was the analysis of the ATWS issue. Then came Three Mile Island (TMI). The accident revealed that perhaps reactors were not "safe enough"; that the regulatory system had some significant problems, cited in both the Kemeny and Rogovin investigations; that the probability of serious accidents was not vanishingly small; and that new approaches were needed. Suddenly, the potential value of PRA as a regulatory tool, and the insights of the RSS itself, became apparent to the reactor-safety community.

People observed that the RSS had found transients, small LOCAs, and human factors to be dominant contributors to the overall risk and that the TMI accident sequence contained all three of these. It became apparent that PRA methods could be used to allocate the limited resources available for the improvement of safety (the Lewis Committee had recommended this only a year earlier). Most important, the reactor community understood that the concept of accident-sequence analysis, as an intellectual discipline separate from other (equally valid) approaches to reactor safety analysis, provided insights that could not be obtained in any other way. The initial applications of PRA methods in the aftermath of TMI were specifically directed at issues of high immediate concern. For example, PRA methods were used to study the reliability of auxiliary feedwater systems in PWRs. The studies revealed that the availability-on-demand of systems that fully met regulatory requirements ranged from best to worst by more than two orders of magnitude. One result was that some auxiliary feedwater systems, in which at least one train was thought to be fully independent of AC power, were

discovered to lack that feature. As another example, PRA methods were used in the Rogovin Special Inquiry to study the phenomena involved in the TMI partial core degradation and the a-priori likelihood of the TMI accident.

Soon thereafter, the NRC staff initiated the Interim Reliability Evaluation Program (IREP), a series of plant-reliability studies more limited in scope than the full-scale RSS-type PRAs and intended to cover five operating reactor designs. These IREP studies were followed by full-scale utility-sponsored PRAs for three plants judged by the NRC to pose potentially unusual risks because of the high population densities near their sites: Limerick, Indian Point, and Zion. These privately sponsored studies represented an important breakthrough, since these were the first important studies sponsored by utilities and performed by analysts from the commerical sector.

Since the initiation of these studies in 1979-1980, utilities have undertaken several other studies. Sometimes the motivation was to prepare for possible new regulatory requirements, but sometimes the utility managements wanted to obtain PRA insights on their own merits.

Within the regulatory staff, probabilistic methods are being adopted for engineering decisionmaking, typically to provide insights that bolster or supplement the traditional regulatory methods. In several issues that have arisen recently (e.g., pressurized thermal shock, steam generator tube ruptures, loss of offsite power and station blackout, loss of shutdown-heat removal, and human reliability), the NRC used PRA insights to assist in decisionmaking. Of course, the ATWS issue had been treated probabilistically as well as deterministically since the mid- to late 1970s.

Current applications of PRA results in regulatory decision-making are becoming increasingly widespread, with PRA techniques being used as an analytical tool to provide additional perspectives to safety analysis. However, concern is increasing about uncertainties and the credibility of quantitative results. To accommodate the uncertainties, the results are most commonly used in the "high-medium-low" sense for assigning priorities to both generic and plant-specific safety issues and in considering regulatory revisions. One recent application of great importance is the heavy reliance by all parties in the special Indian Point ASLB hearing on the methods and results of PRA. Another is the use of PRA insights in the Systematic Evaluation Program review of the 10 oldest operating plants, to help in decisionmaking on backfits or procedural changes. A third application is the continuing use of probabilistic perspectives in assigning priorities to generic safety issues. A plant-specific application of some note was the Big Rock Point PRA: this utility-sponsored PRA was performed to demonstrate to the NRC that many suggested safety-related retrofits would not be cost-effective because of the specific design and size of the Big Rock Point station, and the NRC considered the results of the PRA in its deliberations. Another plant-specific application was at Indian Point, where PRA insights identified a few modifications and procedural changes that offered substantial safety benefits at modest cost.

Some areas where PRA might eventually contribute importantly are still evolving. These include accidents initiated by fires, where the first PRA applications on a broad systems level have shown the techniques to be useful but in need of further development, and accidents initiated by earthquakes, where substantial development has already occurred under NRC sponsorship and private sponsorship. Another example is the study of core-melt progression and radionuclide transport, where improvements in the modeling of physical phenomena and

containment responses and the incorporation of these analyses into probabilistic models is now in a very active stage of development. This is the "severe accident" arena, where regulatory alternatives are being actively developed by the NRC and where probabilistic methods and insights from numerous PRA studies are expected to make important contributions to decisionmaking.

To summarize the present situation, the NRC is using PRA methods and results in varying degrees within NRC in many generic regulatory applications and some plant-specific ones. The applications affect almost the whole technical spectrum of regulation. This situation is a remarkable, considering that the first application of PRA techniques occurred less than a decade ago and the controversy that attended the initial RSS application.

6.3 Use of PRA In The Regulatory Decisionmaking Process

The traditional regulatory process is based on the concept of defense in depth. Plant-design requirements have been derived primarily through the analysis of design-basis accidents, supplemented by requirements intended to ensure safety-system reliability (e.g., the single-failure criterion). Specific requirements are codified in regulations, technical specifications, and license conditions. Additional guidance regarding acceptable plant features is given in regulatory guides. The design-basis accidents are a set of accidents chosen to envelop credible accident conditions. The design and operation of the plant ensure that these accidents will not substantially degrade the core, and conservative estimates of the radiological impacts of such accidents must be limited to prescribed values using engineered safety features and appropriate siting. This approach has been successful in that it is widely recognized that accidents outside of the design-basis envelope dominate the estimated low levels of risk associated with nuclear power plants. The more probable causes

of such accidents are believed to originate from multiple failures or human errors that are outside the domain of either the single-failure criterion or the common-cause failure mechanisms currently addressed in the regulations (seismic qualification, safeguards, fire protection, etc.). While this fact does not negate the effectiveness of the NRC's regulatory practices, it does raise a question as to whether additional protection for accidents beyond the design basis should be provided. Thus, the current widespread interest in degraded core accidents and PRA. The NRC's traditional analytic process has been principally deterministic in nature. That is, it has relied primarily on conservative engineering analysis of LWR safety and performance. The process is intuitive to the extent that it relies on the engineering judgment of technical experts. One shortcoming of this process is that an effective means for conducting an integrated and systematic analysis of the plant is not included. PRA provides a means for conducting such an analysis. It gives the regulator a powerful additional tool for organizing information into a logical framework and providing insights into the complex interrelationships among systems in a nuclear power plant. PRA provides comprehensive models for identifying dominant contributors to reactor risk by performing a systematic analysis of the design and operation of a nuclear power plant from a risk perspective. The analysis is not constrained to design-basis events, but instead provides an integrated assessment of primary safety systems, support systems, and plant operations with respect to core damage, containment failure, and radiological consequences. This tool permits the analyst to investigate the nature of the residual risk and to understand the character, variety, and importance of the constituent elements of risk. It also provides the decisionmaker with a means for evaluating the reduction of net risk derived from potential alterations in design or operation of a plant.

PRA results do not and should not dictate decisions. PRAs do provide an additional source of information, and the weight given to that information depends on its credibility for the particular regulatory action under consideration. Other factors must be considered in reaching an estimation of the worthiness of a potential regulatory action.

Several elements consitutute the regulatory decision process. The first is to determine the analytical methods appropriate for the decision. This could include qualitative and quantitative analyses, deterministic and probabilistic analyses, assessments of operating experience, and value-impact assessments. After the appropriate methods have been identified, analyses are performed and information is gathered and assessed as to technical credibility, employing technical peer review as appropriate. The third step in the decision process is the synthesis of all the applicable information to gain insights into the safety significance of the issue, conceptualize alternative resolutions of the issue (including the "no action" alternative), and evaluate the impacts of the various alternatives.

The final step is to develop recommendations for regulatory application. This step must consider the information base with its inherent uncertainties in judging the credibility and merits of the various insights and alternatives inherent in the proposed recommendations. This step would also include further peer and public review with appropriate feedback loops for additional analysis and synthesis, as appropriate.

Decisions on regulatory requirements are made within the framework described in "Regulatory Analysis Guidelines of the

U.S. Nuclear Regulatory Commission," NUREG/CR-0058, January 1983. Analyses performed using these guidelines should establish the logical framework for selecting and comparing candidate regulatory alternatives. The centerpiece of the regulatory analysis is a thorough inquiry into the values and impacts of alternative regulatory resolutions of the issues under consideration.

The analysis must display all of the important values and impacts (and their uncertainties) associated with a proposed regulatory change in an organized and clearly understandable form for the decisionmaker and other interested parties. Information should be displayed so that the decisionmaker can clearly determine the sensitivity of any conclusion to variations in the important inputs affecting that conclusion. All assumptions underlying each conclusion and the information from which it is drawn must be explicit.

Some less obvious potential values and impacts exist that need to be considered in any decision based in part on the results of a PRA, or that would require the use of a PRA to implement the decision. These considerations include:

1. Altering the regulations: Substantial costs are associated with altering the character and content of the regulations in both the regulated industry and the NRC. Large changes in the skills required for compliance and inspection could introduce severe personnel and training burdens.

2. Regulatory flexibility and stability: Flexibility to accommodate new information on severe accident risk or on implementation costs, or margins to accommodate unpleasant surprises, can have a significant value. A regulation

that makes the body of requirements more stable and predictable, and reduces uncertainties about future requirements, saves real costs to the staff and licensees, and thus to the taxpayers and ratepayers.

3. Time for implementation: Reductions in risk at operating plants have more value if they are implemented promptly. For those plants in design or construction, costs tend to be lower, and thus the potential for cost-effective improvements in safety assurance is better if the decisions can be implemented as early as practicable.

4. The ability to verify compliance: Controversy, delays in implementation, and substantial costs, including costs associated with controversies or delays can arise from ambiguities in the meaning of compliance.

5. The impact of the safety feature on defense-in-depth: Requirements that strengthen defense-in-depth, or otherwise strengthen the diversity with which safety is assured, are preferred over those which concentrate protection in fewer safety functions.

6. The overall safety impact of the safety feature: Safety features that are effective for broad classes of accident sequences or root causes of accidents are preferred over those that are narrowly targeted on specific vulnerabilities.

Regulatory decisions have been and will continue to be made despite uncertainty. The decisionmaker must recognize clearly the nature and source of these uncertainties. Uncertainty reduction must be given serious consideration when assessing a regulatory action. A safety feature that has little effect upon the median estimates of reactor risk, but substantially

reduces the upper limit of the estimate, would have a positive value because of the enhanced confidence in public health and safety.

Because regulatory decisions invariably include uncertainty, insights gained from both deterministic and probabilistic risk analyses must be evaluated in this light. Some of the uncertainties arise from the PRA methodology itself and from the lack of a comprehensive data base. However, the most substantial contributors to the overall uncertainty in PRAs originate from the more or less irreducible stochastic elements of the data base, an imprecise knowledge of the plant response to a given stimulus, and the progression of an accident with time (i.e., core-melt progression, fission-product transport, and containment performance). These latter uncertainties are not restricted to the PRAs, but must also be considered by decisionmakers in evaluating recommendations arising from more deterministic analyses. For example, it is clear that operating experience is a major consideration in deterministic analyses as well as in PRA. PRAs can aid in showing the risk impacts of this lack of knowledge in relation to other factors that contribute to risk and uncertainty.

The most important product of a PRA is the framework of engineering logic generated in constructing the models, not the numerical estimates resulting from the mathematical manipulations of these models. The patterns, ranges, and relative behaviors, which can be gained only from an integrated consistent approach like PRA, can be used to develop insights into the design and operation of a plant. In reaching decisions, the regulator must compare the credibility of the information gained from PRA to the credibility of other sources of information. Thus, the decisionmaker must consider the magnitude of the estimated risk, the degree of uncertainty associated with it, the results of sensitivity analyses, and

the net estimated effect of proposed alterations on the overall safety of the plant. Considered in this way, the PRA insights will provide the decisionmaker with a more comprehensive contextual definition of the problem and a better understanding of the impact on public risk, which is quantitative, albeit somewhat uncertain. Thus, the PRA permits a direct debate on important questions of risk rather than relying largely on subjective judgments.

One of the more difficult problems facing the PRA analyst is to display the results of the PRA and the results of uncertainty and sensitivity analyses in a way that communicates effectively to the decisionmaker. It would be impossible to portray the impact of all permutations of uncertainty in a scrutable manner. Therefore, the analyst must select the uncertainties that are most important to the decision and choose a style of presentation that will communicate this information best. A more standardized approach to this problem needs to be developed.

The final question facing the decisionmaker, assuming he or she has all of the information from PRAs in a scrutable form, is what weight to give to the qualitative and quantitative PRA insights versus all of the other available and pertinent information. There is no cookbook answer to this question, because it will depend heavily on the nature of the issue, the results of the PRA, the nature of other information, and other factors that could affect the overall judgment. However, some characteristics of the PRA results and study process that would be considered are:

1. The scope and depth of the PRA (i.e., does the nature of the PRA study reasonably match the needs of the decision).

2. The results of peer reviews, which could add to, or subtract from the credibility of the PRA results.

3. The qualitative insights obtained from the study. For example, do the qualitative insights into the dominant accident sequences appear reasonable from an operational or engineering sense? This includes an assessment of the degree of realism associated with the study.

4. The impact of alternative regulatory actions on the estimated risk, together with the ease and costs of their implementation.

The magnitude of the quantitative estimates must be considered, as well as the results of sensitivity analyses and the bounds and likely biases of the major uncertainties surrounding the point estimates. Where the reasonable upper bound of the PRA estimate indicates that the issue does not warrant regulatory attention, substantial weight may be given to the quantitative PRA results. Similarly, the quantitative results may be given substantial weight in a decision to take regulatory action, if the lower-bound estimate indicates a safety concern. Between these extremes, the quantitative results cannot be the principal basis for a decision, but the qualitative and quantitative results can provide unique perspectives and information to the decisionmaker on the integral performance of the plant.

One major concern exists over the use of PRA results in decisionmaking: all too often the decisionmaker tends either to go too quickly to the bottom line (which is the weakest part of a PRA) or to dismiss the PRA results entirely as being too uncertain. Neither path is appropriate. Safety goals or other types of numerical criteria tend to emphasize the bottom line in spite of all the cautions to the contrary. Therefore, such numerical criteria need to be constructed and implemented so as to minimize this tendency. The decisionmaker must pause along the way and pay particular attention to the design and operation insights derived from the analyses. Therefore, the

performance of the PRAs and the display of results and uncertainties must provide convenient and scrutable stopping places for the decisionmaker, to encourage viewing these insights and understanding the underlying assumptions and uncertainties, and to discourage undue fixation on the bottom line. Only then will the regulator be able to draw fully on the potential benefits of PRA as a source of information and a regulatory tool.

6.4 Assignment of Priorities

Even considering the incompleteness of PRA models and the uncertainties associated with the quantification of the models, a PRA, because of its integrated nature and greater reliance on realistic information, presents the best available information concerning the specific ways in which the critical safety functions at nuclear power plants can fail. This information can be used to guide and focus a wide spectrum of activities designed to improve the state of knowledge regarding the safety of individual nuclear power plants as well as that of the nuclear industry as a whole. The resources of both the NRC and the industry are limited, and the application of PRA techniques or insights provides one more useful tool to permit the decisionmaker to allocate these resources to areas most likely to reduce risk or to limit the uncertainties or to define uncertainties more clearly.

Chapter 5 discusses those items that have importance with respect to either plant risk or the frequency of core melt in published PRA results. While the completeness of such a listing cannot be assured for plants that have not been analyzed, these items have been found to affect significantly either the predicted frequency of core melt or the risk associated with a given plant. Such items should be examined to see whether they are generic and are likely to affect other plants of similar or even dissimilar design.

The nature of the decisions necessary to allocate regulatory resources does not require great precision in PRA results. It is sufficient to assign research and the efforts used to resolve generic safety issues to broad categories of risk impact (e.g., high, medium, and low). The reasoning is that a potential safety issue would not be dismissed unless it were clearly of low risk. Thus, one or more completed PRA studies can be selected as surrogates for the purpose of assigning priorities, even though, clearly, they do not fully represent the characteristics of some plants, provided the nature of these differences are reasonably understood and can be qualitatively evaluated. A given issue can then be evaluated in terms of the number of plants affected, the risk impacts on each plant, the effect of modifications in reducing the risk, and the effect of additional knowledge on improving the prediction of plant risk or core-melt frequency or in reducing or defining more clearly the associated uncertainties. These generic measures of significance, combined appropriately with other information (e.g., cost of resolving the issue) can be used to evaluate the issue under consideration. Obviously, a principal source of uncertainty may lie in the use of a representative plant model (a "surrogate") to represent a broad class of reactors.

The uncertainties involved in the measure used for assigning priorities are such that only large (order of magnitude) variations should be considered important. Thus, if core-melt frequency were the measure, it would be improper, based principally on the estimated frequency, to conclude that an issue associated with an estimated core-melt frequency of 3×10^{-5} per reactor-year is significantly more important than one associated with a core-melt frequency of 1×10^{-5} per reactor-year. However, it would normally be appropriate to assign priorities on the basis that a risk measure of 10^{-4} is more important than one of 10^{-6}.

As with any priority-assignment method, the final results must be tempered with an engineering evaluation of the reasonableness of the assignment, and the PRA-based analysis can serve as only one ingredient of the overall effort. One of the most important benefits of using PRA to assign priorities is the documentation of a comprehensive and disciplined analysis of the issue, which enhances debate on the merits of specific aspects of the issue and reduces reliance on more subjective judgments. Clearly, some issues would be very difficult to quantify with reasonable accuracy and the assignment of priorities to these issues would have to be based largely on subjective judgment.

Information from PRAs can also be used to guide the overall allocation of resources in inspection and enforcement programs. A catalog of information derived from PRAs indicates that certain surveillance tests and maintenance activities are significant contributors to the estimated frequency of plant damage or to risk. If a class-generic risk profile is available, it could be used to determine importance measures regarding critical surveillance testing and maintenance activities that can, if not done properly, significantly alter the predicted core-melt frequency or risk. These importance measures could be used in assigning priorities for inspection auditing, the training of operators and maintenance personnel, and reliability assurance program requirements. The generation of such information for each class of operating plant provides a rough ordering of important operating activities that should assist a reactor inspector in efficiently directing the inspection effort at a given facility. Similarly, generic insights (available by reactor class) assist both the licensee and the regulator in identifying and preventing potentially significant operational occurrences at a plant, even if a plant-specific PRA is not available.

6.5 Generic Regulatory Applications

Perhaps the greatest utility of PRA techniques to the regulators lies in providing insights that aid in rulemaking and the development of safety guides and branch technical positions. Such activities could be aimed at either reducing risk or relaxing regulatory requirements when they do not have a significant impact on risk. The catalog of plant systems, components, and operational practices that have had a significant impact on core-melt frequency or risk in various PRA studies can lead to generic insights for each of a variety of classes of plants. The resulting number of plant classes (or surrogates) may be large, however, because many of the risk-significant features of the plant occur in the balance of the plant, where the design is less standardized.

The use of surrogates to represent classes of plants for generic regulatory activities entails modeling uncertainty, because subtle system and human interactions may have a pronounced effect on the actual risk of a specific plant. Therefore, the possible existence of risk outliers precludes the confident use of the surrogate approach to estimate "bottom-line" risk or core-melt frequency for plants that have not been subjected to a detailed PRA. The presence of a plant-specific risk outlier does not mean that the absolute risk of the regulatory issue under review would be affected, although the relative importance of the issue probably would be affected. However, even the absolute importance of the issue could be affected. This means that regulatory decisions that are based to some extent on generic PRA results should consider whether deviations based, in part, on the results of a plant-specific PRA should be permitted. As an extension of surrogate or plant-class type of analyses, insights can be obtained for a given type of accident sequence that may apply broadly to a large group of reactors (e.g., ATWS in BWRs) or may apply in a somewhat different manner to several different classes of plants (e.g., station blackout). Note that the

plant classes do not necessarily have to have the same basic risk profile; rather, they need only to react similarly to a given accident sequence for the generic insights to be valuable.

The degree of detail necessary in establishing the classes as surrogates depends on the nature of the decision being made. In general, the decisionmaker will not rely on small differences in numerical results and will temper the insights gained from PRAs with engineering judgment. In fact, many times the qualitative insights drawn from PRAs could be more important than the quantitative insights. Sorting the reactor population into a large number of classes of plants often will not be necessary. Generic insights are gained from the examination of both the qualitative and the quantitative results of PRA, as well as from studies of specific accident sequences and probabilistic safety analyses of limited depth.

In general, these insights originate from the relative comparison of quantitative results and, therefore, must consider the uncertainty associated with the quantitative analyses on which they are based. In most situations, the uncertainties associated with relative comparisons would be less than those associated with absolute quantities. However, significant uncertainties would still remain and, in some situations, the uncertainties would not be narrowed. Therefore, an uncertainty analysis would still be necessary, even for decisions involving the relative comparisons of PRA information.

Virtually every PRA performed to date has identified some previously unrecognized deficiency in plant design or operation that has a measurable impact on either the estimated frequency of core melt or the estimated risk. These are usually associated with dependences between systems as well as man-machine interactions. It is then possible to examine these gaps and, if necessary, develop deterministic criteria

that would remove these issues from further regulatory consideration. Note that problem areas identified through risk insights do not have to be present at all other facilities; rather, it could suffice that the insights gained from PRAs have identified potential paths that reduce the level of safety at a plant in a manner that was not originally recognized by the regulator. The costs of the remedies would need to be considered.

Studies can be used to generate insights for use in developing or modifying regulatory positions. Much information can be gained from limited studies of specific issues using simplified system reliability analyses. While these limited studies are insufficient to predict accurately the absolute level of risk, they can indicate relative importance of problems, as was done in the plant-specific studies of auxiliary feedwater systems (AFWS). The AFWS study was restricted to three dominant accident sequences and the functional or systemic failures required for their occurrence. If the catalog of PRA information on various designs were considered sufficiently complete, this type of review could be expanded to review all accident sequences believed to be dominant and appropriate for a given group of plants based on generic classifications. These studies could be used to identify outlier events, without a specific need for precise quantification. Such studies might primarily examine the manner in which the plant design handles interfaces between systems and interfaces between the human and the machine.

The qualitative information on the types of failures and operational practices that have been found important in previous studies obviously identifies these as candidates for consideration for further study, regardless of the quantification uncertainties involved. If these insights reveal safety-significant gaps in the regulatory fabric, they should be closed, even if an unknown and unrecognized outlier may still

exist. If, for a particular issue, the existing catalog of core-damage or risk-significant insights is felt to be incomplete, one could perform PRAs on one or more additional plants to derive new integral insights.

The catalog of significant core-damage insights is converging but has not yet reached an asymptote. As noted earlier, experience has indicated that plant-specific design or operational differences could represent a dominant contributor to core-melt frequency. However, the existing library of PRAs should be sufficient for the reasonable evaluation of most core-melt issues. Thus, while it might be desirable to obtain additional generic insights through the performance of further PRAs on operating facilities of varying design, it would not be necessary (to support generic positions on core-melt issues) to extend this to studies of all plants because, eventually, the catalog of significant generic plant features and operational weaknesses that affect core-melt frequency should be reasonably complete.

The above discussion indicates that surrogates will have an important role to play in the generic assessment of core-melt sequences. The usefulness of surrogates for risk estimates is less clear at present. The usefulness of surrogates for evaluating containment performance cannot be judged until the ongoing source-term work is completed.

One clear conclusion can be drawn with regard to the use of PRA in generic regulatory applications: such applications normally should be at the system, component, function, or accident-sequence level. Few applications should exist where estimates of core-melt frequency or risk would be directly applied in a generic fashion in the regulatory process.

6.6 Plant-Specific Applications

The stated uncertainties in a plant-specific PRA, compounded by the inability to quantify modeling uncertainties in any but a subjective manner, make it very difficult to determine formally whether a specific safety limit (in terms of public risk or the frequency of core melt) is met with any high degree of confidence. However, PRA results can be used in a more general fashion to see whether the basic design and operational principles of the plant allow it to approach or reach the goal of a given level of safety (or risk). Thus, the "bottom-line" numerical results of a PRA are not useless, provided they are generated with care, considered with a clear understanding of their uncertainties and biases, not used in a sense of compliance versus noncompliance, and used in conjunction with other conventional regulatory tools. These quantitative descriptions of core-melt frequency or risk and the constituents thereof can be useful tools for the direction of regulatory attention and allocation of resources. While the level of reliability or risk associated with a given nuclear plant cannot be described precisely, neither can the threshold of unacceptable risk or unreliability.

Experience has indicated (Indian Point, Shoreham) that PRAs can be reasonably adjudicated in licensing hearings. Such use could well improve the hearing process, if properly controlled. An adjudicatory process is one where decisions are based on a weighing of evidence. PRAs provide evidence through their ability to portray the importance of various plant operations and failure modes and could be used to help establish the significance of potential safety issues.

Virtually every PRA has identified previously unrecognized design or operational deficiencies. In some cases, these deficiencies were rectified not primarily because of the calculated frequency values, but simply because the plant owner and operator recognized that a specific portion of the plant

(or of the operating practices) did not function in the way it was intended. Thus, the qualitative knowledge can be used to improve the operational performance of the facility without a high degree of reliance on the numerical estimates of probability and consequence.

A plant-specific PRA, performed early in the design process, can yield a tremendous number of insights about the integral performance of the plant to the designers. Because of the lack of specific design details in some areas, as well as the lack of plant-specific data, the results of such an analysis cannot be considered a true prediction of plant risk or of the frequency of core melt. Rather, such an analysis generates useful information on potential weaknesses in the design, and it allows an evaluation of the efficacy of design modifications. Also, this analysis could be used to focus quality-assurance activities on those areas with the highest potential for reducing risk (e.g., operating procedures). Again, the real significance of such an analysis lies not in the numerical estimates but, rather, in the insights into important design features and critical man-machine interfaces, which can be carefully considered in the detailed design process.

As previously noted, one use for a plant-specific PRA is in the evaluation of proposed generic solutions to unresolved safety issues and other generic items. This is not to suggest that the generic resolution of items is improper; rather, it indicates that, because of plant-specific differences, particularly in the balance of the plant, a plant-specific PRA may be able to identify a regulatory action for that plant that is more efficient than the generic solution. The regulatory decisionmaker must consider this when generic requirements are set and evaluate whether unwarranted inequity will be introduced if plant-specific resolutions are not permitted.

In the same manner, the availability of plant-specific qualitative insights into risk and quantitative results allows the regulator to assign priorities, on a plant-specific basis, to the various licensing issues and inspection activities associated with a given plant. An important factor in this process is that the regulator not rely primarily on the quantitative results, but on the qualitative insights and the associated stated and unstated uncertainties. Of particular importance is the detailed knowledge of system performance and the variety of interactions between systems and components and between the operators and the various plant systems and subsystems. A plant-specific PRA can be used to evaluate the importance of operating events and to assess the safety of the plant when equipment is not operable. Also, a catalog of accident sequences and the estimates of their frequencies can be used to train emergency-response personnel in what to expect. This could lead, for example, to improving the set of symptoms to be used as trigger points for the declaration of site or general emergencies and to developing guides on the diagnosis and prognosis of accidents as they progress. The models generated also provide the tools with which to optimize allowable outage times and surveillance intervals and can be used in evaluating the advisability of plant shutdown when equipment is out of service beyond the specified allowed outage times in current technical specifications.

Given that a plant-specific PRA has been performed, steps should be taken to track the performance of the plant to ensure that the level of safety identified in the study is not degraded with time. Thus, the PRA should be, in effect, a living document that is used and appropriately updated. The PRA should be used in the context of a safety or reliability assurance program as discussed in Chapter 3 to evaluate operational occurrences and to check the significance of experience data as they are acquired.

6.7 Other Uses

In addition to its potential uses in the regulatory process, a PRA offers many advantages to the owner and operator of the reactor. These advantages are beyond the immediate scope of this document, but a few examples are given because any actions taken by the utility to improve the operating knowledge or to modify system design or operations as a result of a PRA are likely to have a positive effect on the overall safety of the plant.

The catalog of plant-specific severe accident sequences with estimates of their frequencies, consequences, and root causes could be included in operator training and simulator design. It could also be used as a starting point for further studies intended to assess the similarity of the symptom profiles among accidents requiring different operator responses and to survey the hazards associated with misdiagnosis, compared with less-than-optimum recovery actions. Obviously, such information is also of great interest to the regulator.

A PRA will produce estimates of system reliability. It can assign quantitative measures to the importance of system components and determine the more likely failure modes that are believed to dominate the unavailability of the systems. With this information, an operator can assist and develop a strategy for repairing a given system or component within the time required in an accident situation. Operators can be trained in fault diagnosis and in effective repairs. The adequacy of diagnostic instrumentation and status monitoring can also be assessed.

Even for those safety issues that are not well handled by the quantitative aspects of PRAs (certain external events, sabotage, design or installation errors, and dependences that are not revealed by explicit hard-wired functional dependences among systems), the logic models generated in a PRA study can

be used to put concerns in perspective. For example, the PRA can identify which accident sequences might be affected by a given postulated safety issue and estimate the conditions under which the issue might emerge from the background of minor contributors to risk (or core-melt frequency) into one of the dominant concerns. Thus, the PRA can be useful even when the predictive power is poor.

Note that none of the uses suggested for consideration by the utility depends heavily on the bottom-line predictions of risk. They all depend on the more trustworthy comparative measures of importance and on the catalog of accident sequences to which a subject plant is susceptible. While some of the applications are sensitive to the limitations, particularly incompleteness and quantitative accuracy, nevertheless, the applications can be tailored to the known limitations and the models generated can provide a coherent framework to address the "what if" questions concerning its accuracies in these applications.

6.8 Conclusion

PRA as practiced today provides imprecise quantitative results, yet it has proved valuable in providing greater insights into the relative importance to safety of specific plant characteristics, regulatory issues, and alternative regulatory actions. Thus, it is recommended that the use of PRA in the regulatory arena focus on applications where issues or alternatives are placed in fairly broad categories reflecting their relative importance. These applications can include plant-specific as well as generic actions. The categories should be broad enough to be appropriate even after considering the range of uncertainties. Plant-specific applications of PRA results are not recommended where the results are to be used for a compliance-type of comparison against some numerical standard of acceptability.

The various ways regulators can use PRA techniques, results, and insights to supplement and augment the information derived from traditional analytical techniques have been discussed in detail in this chapter. The more important conclusions are presented below.

1. Assignment of prioritization to regulatory issues. The issues should be assigned to broad categories and should not require much precision from the PRA input. These assignments can aid significantly in allocating limited resources to risk-significant issues. Regulatory areas amenable to this use include generic safety issues, inspection procedures, enforcement actions, and regulatory research. Value-impact analyses would be useful in reaching a decision. However, some issues would not be amenable to reasonable quantification and thus would still require a more subjective assignment of priorities.

2. Generic regulatory applications. Such uses focus on areas where additions to existing regulatory requirements appear necessary, as well as on regulated areas that appear to be unimportant to risk. The scope and depth of the PRAs, the degree to which differences between plant classes need to be considered, and the role that uncertainty and sensitivity analyses would play would depend on the particular issue under review.

3. Plant-specific applications. Many plant-specific uses of PRA have evolved besides the strict comparison of "bottom-line" numbers with numerical criteria as a licensing or compliance exercise, and such usage is recommended. Examples are to provide information for--

   a. Plant-specific decisions on exemptions from existing requirements or the imposition of additional requirements.

b. Development of plant-specific limiting conditions for operation and surveillance testing requirements.
c. Development of plant-specific operating, testing, and maintenance procedures.
d. Development of requirements for training and quality-assurance programs.
e. Development of emergency response and operating procedures.
f. Assessment of operating experience to gain plant-specific insights.
g. Development of plant-specific inspection programs.
h. Development of reliability-based design requirements for any new plants that are not well into the design phase.

For such plant-specific applications, the PRAs would either have to be plant specific or would have to draw on information that was sufficiently plant specific to be a reasonable surrogate for that plant class.

One question that must be resolved is whether the usefulness of plant-specific applications is sufficient to warrant a regulatory requirement for the performance of such analyses by the industry. Such PRAs could be useful in integrating and assigning priorities to all identified safety issues applicable to that plant, in addition to searching out any risk outliers that would not be identified from the risk insights gleaned from PRAs of similar plants.

4. Many of the generic and plant-specific applications listed above can draw from the relative insights provided by PRAs. In many situations, these qualitative insights would be more important than the quantitative results. However, where the quantitative results are given significant weight, an analyst must be careful to consider

whether the results of sensitivity analyses conducted over reasonable uncertainty bounds (including alternative modeling assumptions) would affect the decision significantly, compared to the use of point estimates.

REFERENCES

NUREG REPORTS

NUREG-75/014, WASH-1400, "Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants," October 1975.

NUREG/BR-0058, "Regulatory Analysis Guidelines of the U.S. Nuclear Regulatory Commission," January 1983.

NUREG-0340, "Overview of the Reactor Safety Study Consequence Model," 1977.

NUREG/CR-0165, "A Value-Impact Assessment of Alternative Containment Concepts," 1978.

NUREG/CR-0400, "Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission," September 1978.

NUREG-0772, "The Technical Bases for Estimating Fission Product Behavior During LWR Accidents," 1981.

NUREG/CR-1131, "Examination of Offsite Radiological Emergency Measurements for Nuclear Reactor Accidents Involving Core Melt," 1979.

NUREG/CR-1250, "Three Mile Island: A Report to the Commissioners and to the Public," January 1980.

NUREG/CR-1278, "Handbook for Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications," 1980 (revised 1983).

NUREG/CR-1433, "Examination of the Use of Potassium Iodide (KI) as an Emergency Protective Measure for Nuclear Reactor Accidents," 1980.

NUREG/CR-1659, Volume 1, "Reactor Safety Study Methodology Applications Program: Sequoyah 1 PWR Power Plant," 1981.

NUREG/CR-1659, Volume 2, "Reactor Safety Study Methodology Applications Program: Oconee PWR Plant," 1981.

NUREG/CR-1659, Volume 3, "Reactor Safety Study Methodology Applications Program: Calvert Cliffs 1 PWR Power Plant," 1982.

NUREG/CR-1659, Volume 4, "Reactor Safety Study Methodology Applications Program: Grand Gulf 1 BWR Power Plant," 1981.

NUREG/CR-1989, "Analysis of Hypothetical Severe Core Damage Accidents for the Zion Pressurized Water Reactor," 1982.

NUREG/CR-2326, "Calculations of Reactor Accident Consequences, Version 2: Computer Code User's Guide," April 1983.

NUREG/CR-2239, "Technical Guidance for Siting Criteria Development," 1982.

NUREG/CR-2254, "A Procedure for Conducting a Human Reliability Analysis for Nuclear Power Plants," December 1981.

NUREG/CR-2255, "Expert Estimation of Human Error Probability in Nuclear Power Plant Operations: A Review of Probability Assessment and Scaling," August 1982.

NUREG/CR-2300, "PRA Procedures Guide," January 1983.

NUREG/CR-2497, "Precursors to Potential Severe Core Damage Accidents: 1969-1979 Status Report," June 1982.

NUREG/CR-2515, "Crystal River 3 Safety Study," 1982.

NUREG/CR-2723, "Estimates of the Financial Consequences of Nuclear Power Reactor Accidents," 1982.

NUREG/CR-2728, "Interim Reliability Evaluation Program Procedures Guide," 1983.

NUREG/CR-2744, "Human Reliability Data Bank for Nuclear Power Plant Operations," March 1983.

NUREG/CR-2787, "Interim Reliability Evaluation Program: Analysis of the Arkansas Nuclear One-Unit 1 Nuclear Power Plant," August 1982.

NUREG/CR-2802, "Interim Reliability Evaluation Program: Analysis of the Browns Ferry Unit 1 Nuclear Power Plant," 1982.

NUREG/CR-2906, "Sensitivity of Risk Parameters of Component Unavailability in Reactor Safety Study," May 1983.

NUREG/CR-2934, "Review and Evaluation of the Indian Point Probabilistic Safety Study," 1982.

NUREG/CR-3010, "Post Event Human Decision Errors: Operator Action Tree/Time Reliability Correlation," March 1983.

NUREG/CR-3085, "Interim Reliability Evaluation Program: Analysis of the Millstone Plant Unit 1 Nuclear Power Plant," 1983.

NUREG/CR-3226, "Station Blackout Accident Analysis (part of NRC Task Action Plan A-44)," May 1983.

Camp, A. L., et al., "MARCH-HECTR Analysis of an Ice-Condenser Containment," Proceedings, International Meeting on Light Water Reactor Severe Accident Evaluation, American Nuclear Society, September 1983.

Embrey, D. E., "Modeling and Quantifying Human Reliability in Abnormal Conditions," National Reliability Conference, Birmingham, England, July 1983.

Evans, M. G. K. and O'Reilly, P. D., "Prevention of Core Damage by Alternate Use of Existing Plant Systems, the Modification of Existing Systems, or the Addition of New Systems," Proceedings, International Meeting on Light Water Reactor Severe Accident Evaluation, American Nuclear Society, September 1983.

Joschek, H. I., "Risk Assessment in the Chemical Industry," International ANS/ENS Topical Meeting in Probabilistic Risk Assessment, Port Chester, New York, September 1982.

Potash, L., Stewart, M., Dietz, P. E., Lewis, C. M., and Dougherty, E. M., Jr., "Experience in Integrating the Operator Contribution in the PRA of Actual Operating Plants," Proceedings ANS/ENS Topical Meeting on Probabilistic Risk Assessment, Port Chester, New York, Am. Nuclear Soc., LaGrange Park, Illinois, 1981.

AIF, "A Proposed Approach to the Establishment and Use of Quantitative Safety Goals in the Nuclear Regulatory Process," Atomic Industrial Forum, 1981.

Aldrich, D. C., et al., "Recent Developments in Reactor Accident Offsite Consequence Modeling," Nuc. Saf., 23, 643-652, 1982.

Benjamin, A. S. and Harper, F. T., "Value-Impact Investigation of Filtered-Vented Containment Systems and Other Safety Options for a BWR Mark I Containment," Sandia National Laboratories, Albuquerque, New Mexico, to be published.

Benjamin, A. S., et al., "Severe Accident Risk Reduction Program (SARRP) Phase I Report," Sandia National Laboratories, to be published around June 1984.

Big Rock, "Big Rock Point Nuclear Power Plant Probabilistic Risk Assessment," Consumers Power Company, 1981.

Brune, R. L., Weinstein, M., and Fitzwater, M. E., "Peer-Review Study of the Draft Handbook for Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications, NUREG/CR-1278," SAND82-7056, Sandia National Laboratory, Albuquerque, New Mexico, 1983.

Burke, R. P., "Economic Risks of Nuclear Reactor Accidents," Ph.D. Thesis, MIT, Cambridge, Massachusetts, 1983.

CRBR, 1977, Clinch River Breeder Reactor Plant, Draft--1983.

DRS, Deutsche Risikostudie Kernkraftwerke, "Eine Untersuchung zu dem durch Storfaalle in Kernkraftwerken verursachten Risiko," Der Bundeminister fur Forschung and Technologie, Verlag TUV-Rheinland, 1979.

GAO, Letter GAO Ref. B-211642, from J. Dexter Peach, Director, Resources, Community and Economic Development Division, United States General Accounting Office, to the Honorable Richard L. Ottinger, Chairman, Subcommittee on Energy Conservation and Power Committee on Energy and Commerce, House of Representatives, 1983.

Hannaman, G. W., Spurgin, A. J., and Wreathall, J., "Systematic Human Action Reliability Procedure," EPRI Project 2170-3, September 1983. Draft Report, 1983.

Indian Point 2, 3, "Indian Point Probabilistic Safety Study," Power Authority of the State of New York and Consolidated Edison Company of New York, 1982.

Joksimovich, V., Frank, M., Hannaman, G. W., and Orvis, D. D., "A Review of Some Early Large Scale Probabilistic Risk Assessments," EPRI NP-3265, Project 2171-1, October 1983.

Joksimovich, V., et al., "HTGR Accident Initiation and Progression Analysis," GA-A-15000, April 1978.

Kemeny, T. G., et al., "Report of the President's Commission on the Accident at Three Mile Island," U.S. Government Printing Office, Washington, D.C., 1979.

Kolaczkowski, et al., "Interim Report on Accident Sequence Likelihood Reassessment (Accident Sequence Evaluation Program)," August 1983 (unpublished).

Kozinsky, E. J. and Pack, R. W., "Performance Measurement System for Training Simulators," EPRI NP-783, Electric Power Research Institute, Palo Alto, CA, 1982.

Limerick, "Probabilistic Risk Assessment, Limerick Generating Station," Philadelphia Electric Company, 1981.

Limerick, "Severe Accident Risk Assessment, Limerick Generating Station," Philadelphia Electric Company, 1983.

Oconee, Oconee PRA, EPRI/NSAC (to be published) 1983.

OECD, to be published, "Comparison of Reactor Accident Consequence Models," Summary Report of the NEA/CSNI International Comparison Study on Reactor Accident Consequence Modeling, Paris.

Rickard, C. L., "Proposed policy statement on safety goals for nuclear power plants," letter from President, American Nuclear Society, to U.S. Nuclear Regulatory Commission, May 18, 1982.

Ringhalls, "Executive Summary, Ringhals 2 Probabilistic Safety Study," Swedish State Power Board, 1983.

Shoreham, "Probabilistic Risk Assessment--Shoreham Nuclear Power Station--Unit 1--Long Island Lighting Company," Science Applications, Inc., 1983.

Sizewell-B, "Sizewell B Probabilistic Safety Study," WCAP-9991, Westinghouse Electric Corporation, 1982.

APPENDIX A

TASKS ASSOCIATED WITH PROBABILISTIC RISK ANALYSES OF VARIOUS SCOPES

The tasks associated with PRAs of various scopes are presented below. Each task is briefly described, and the relationships between tasks are discussed. The steps involved in the analysis are shown in Figure A-1.

PRAs are broad, integrated studies requiring large amounts of information, therefore, the first step is information collection. The information that is required depends on the scope of the analysis and falls into three broad categories:

1. Plant design, site, and operation information,
2. Generic and plant-specific data,
3. Documents on PRA methods.

The next task is systems analysis, which involves the definition of accident sequences; an analysis of plant systems and their operation; the development of a data base for initiating events, component failures, and human errors; and an assessment of accident-sequence frequencies. It constitutes a major portion of the PRA and hence is divided into the several subtasks discussed below.

The event-tree development subtask delineates the various accident sequences to be analyzed, combinations of initiating events and the successes or failures of systems. This activity includes an identification of initiating events and the systems that respond to each initiating event.

The system modeling subtask involves the construction of models for the plant systems covered in the PRA. The systems

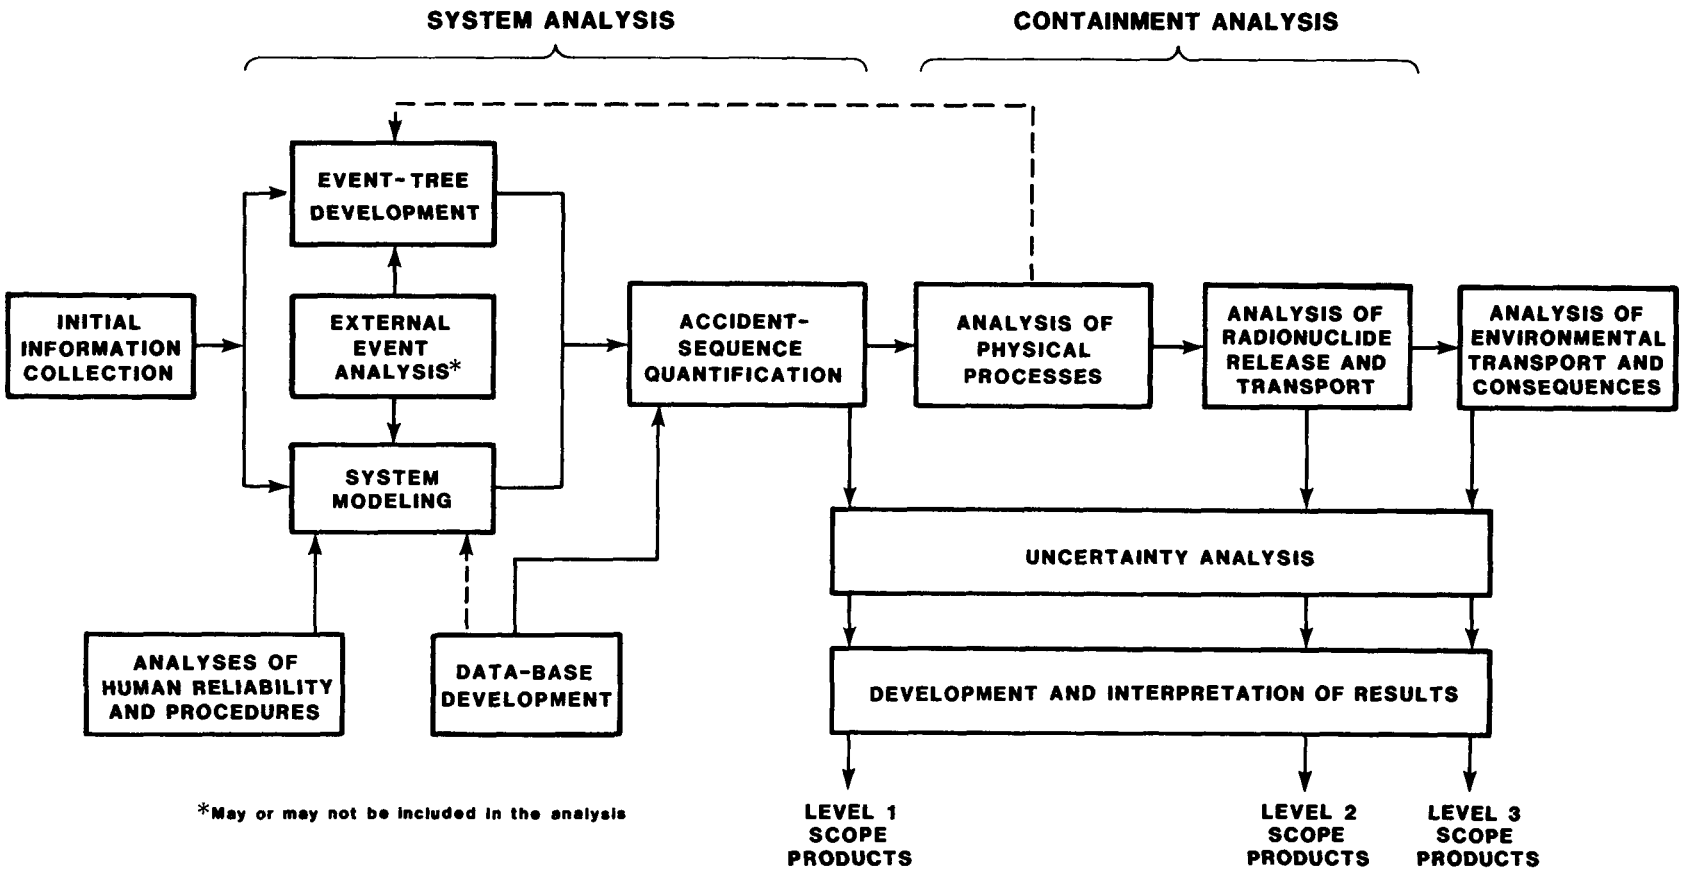

Figure A-1. Risk Assessment Procedure

to be analyzed and their success criteria are identified in conjunction with event-tree development in an iterative process. Assistance from thermal-hydraulics and containment analyses may be needed to derive realistic system-success criteria. The system models generally consist of fault trees developed to a level of detail consistent with available information and data.

Past PRAs have shown the importance of operator error. These human errors are included in the plant-system models. The analysis performed in the human reliability subtask involves a review of testing, maintenance, and operating procedures to identify the potential human errors to be included in the analysis. A review of the plant's administrative controls and procedures and the design of the control room is also performed to establish a foundation for the assignment of failure rates to the human errors found to be significant.

The next major task involves the quantification of accident sequences, which requires a component-data base, developed by compiling data, selecting appropriate reliability models, establishing the parameters for those models, and then estimating the probabilities of component failures and the frequencies of initiating events. The data used in the data base development subtask may be generic industry data or plant-specific data, or a combination of both.

In order to quantify the frequencies of the accident sequences delineated in the event trees, failure rates are assigned to each plant-system model and frequencies are assigned to each initiating event. Combining the appropriate system success and failure models with each class of initiating events yields a logical representation of each accident sequence.

The containment analysis task is important for differentiating among the consequences of various core-melt accident sequences

and consists of two subtasks. The result of this analysis is an identification of containment-failure modes and a prediction of the radionuclide inventory released to the environment for each accident sequence.

A core melt accident would induce a variety of physical processes in the reactor core, the pressure vessel, the reactor coolant system, and the containment. Computer codes have been developed to assist in the analysis of these processes. The results are insights into the phenomena associated with the accident sequence and a prediction of whether the containment fails.

A containment event tree is developed for each sequence of interest. If the containment is predicted to fail, the analysis predicts the time at which it will fail, where it will fail (i.e., whether radionuclides are released directly to the atmosphere through the containment building or to the ground through the basemat), and the energy associated with the release.

For each core melt accident that is postulated to breach the containment, it is necessary to estimate the inventory of radionuclides that would be available for release to the environment. In this subtask the analyst analyzes the radionuclides released from the reactor fuel during the accident and to assess their transport and deposition inside the reactor coolant system and the containment before containment failure. The results of this analysis are a prediction of the radionuclide inventory released into the environment at the time of containment failure for each accident sequence.

The final task is consequence analysis. To assess the risk associated with the plant, it is necessary to calculate the consequences of the release in addition to the frequency of the accident and the inventory of released radionuclides.

Consequences are generally expressed in terms of early fatalities, latent-cancer fatalities, and property damage. To perform this task, the analyst uses a computer model that begins with the inventory of radionuclides released from the containment and analyzes their transport through the environment, using site-specific meteorological data and, in some cases, information on the local terrain as well. Data on population density are then used to calculate the radiation doses delivered to the population, and a health-effects model is used to estimate health effects. The economic consequences that are estimated are those resulting from a relocation of the population and the interdiction or decontamination of the land. The results of the analysis are usually consequence distributions (i.e., plots of the predicted frequency for consequences of varying magnitudes) for each accident release category.

External initiators, frequently excluded from earlier PRAs, include winds, fires, earthquakes, and floods. This task uses the models developed in the plant-system analysis. The models are either analyzed independently from the perspective of external events or else they are modified to reflect external events explicitly. Additional event trees are sometimes developed to delineate the external event sequences to be analyzed.

The results of the external initiator analysis are incorporated into the accident-sequence analysis. In addition, external events may influence the containment analysis. The subsequent steps of the PRA are the same as those discussed above. The final result is a more complete PRA.

The final step in performing PRAs of various scopes is to integrate the data obtained in the various tasks of the analysis and to interpret the results. This integration includes, among other things, the tabulation of frequencies for accident sequences important to risk, the development of complementary

cumulative distribution functions for the plant, and the development of distributions reflecting the uncertainties associated with accident-sequence frequencies.

To provide focus for the assessment, the results are analyzed to determine which plant features are the most important contributors to risk. These engineering insights constitute a major product of the analysis. Insight into the relative importance of various components and the relative importance of various assumptions to the results may be developed from the uncertainty and sensitivity analyses. A discussion of these insights provides additional perspective to the analysis.

APPENDIX B

LEVEL OF MATURITY OF PROBABILISTIC RISK ANALYSIS (PRA)

PRA is a multi-disciplinary methodology requiring data and analyses from system engineers, plant personnel, data analysts, human behavioral scientists, experts in accident phenomenology, and geologists, to name a few. Not all of the methods have reached the same level of maturity. Some, such as reliability analysis, have been practiced in some form since World War II. Others, such as probabilistic analysis of core melt progression, are new and unique to reactor technology.

PRAs need not always include a full risk assessment. Thus, the question of how to use PRA in regulation must consider the strengths and weaknesses of PRA, in other words, the level of maturity reached by the different methods used in PRA.

The level of maturity must consider several characteristics, stability of the method, degree of realism, degree of uncertainty, desirability of major progress to improve the method, and feasibility of achieving that progress, especially in the near future. Using these indicators, the overall level of maturity of each methodological element can be gauged. Decisions based on more mature methodologies would, in general, be expected to have a higher degree of confidence than decisions based on less mature approaches.

The stability of a method is an index of the rate of change of the methodology. A methodology undergoing rapid development is described as unstable and must be treated with caution. This does not imply that stable methodologies are

necessarily more satisfactory or that they can be used without caution.

The degree of realism is the extent to which approximations or conservatisms may have been knowingly or unknowingly introduced into some parts of the PRA because of unknowns, attempts to simplify the models, or error. Whether the ultimate "result" is accurate within its stated uncertainties or is conservative or nonconservative will depend on the degree of realism.

Most of the uncertainties associated with PRA are not unique to PRA, but reflect a lack of data, experience or knowledge about system response, human behavior, or accident phenomenology. These uncertainties exist whether the decisionmaker uses PRA, deterministic modeling, or so-called engineering judgment when making decisions. They reflect the experience, the knowledge, and state of the overall technology. Since its beginning, PRA has tried to display its uncertainties explicitly and, in so doing, has focused attention on those uncertainties. Thus, analyses using PRA techniques display the uncertainties more explicitly than do other analytical approaches, even though these same sources of uncertainty often also affect the single-estimate deterministic analyses and the qualitative judgments.

Displaying the uncertainties, as PRA can, provides important information to the decisionmaker. A proper uncertainty analysis can estimate the impact of this lack of experience and/or knowledge on engineering insights drawn from PRA by propagating uncertainties through the analysis or by performing sensitivity analyses within the PRA. Thus the treatment of uncertainties should be considered a strength of PRA rather than a limitation.

The remaining sections of Appendix B address the level of maturity of the various elements of PRA methodology and practice.

B.1 Plant Modeling and Model Evaluation

The term "plant modeling and model evaluation" is generally applied to those aspects of PRA analysis that deal with identifying and quantifying the occurrence frequency of accident sequences and system failures where the sequences are combinations of system failures and successes. The basic approach to this activity is much the same as it was for the Reactor Safety Study (RSS) (NUREG-75/014), study. The RSS identified the accident sequences using the event tree analysis technique and the system failures using the fault tree analysis approach. These models were evaluated and quantified with a combination of computer and manual techniques. While the basic approach to plant modeling activities has continued to include models for both accident sequences and system failures, many refinements in technique, especially in evaluation and quantifying models, have occurred since the RSS was published. Many of these changes have broadened the scope of the modeling activity and other changes have resulted from attempts to make the modeling activity more comprehensive.

B.1.1 State of the Art In Plant Modeling and Modeling Evaluation

Two aspects to the techniques for plant modeling are apparent; developing models, and evaluating and quantifying models. Although each of these aspects has its own unique characteristics, the activities in each depend strongly on the state of the art in the other. For example, the complexity of the models depends somewhat on the ability to evaluate and quantify them.

The development of the plant system models consists of applying techniques for postulating potential events associated with plant equipment and operation and displaying these events graphically. Also associated with model development is the discipline of each applicable technique, which provides a systematic basis for postulating events. The most common methods for plant modeling in PRA are event tree and fault tree techniques.

Since publication of the RSS, applying these techniques to PRAs of nuclear power plants has resulted in the development of various descriptions and procedures for using these methods. Among the most prominent are the PRA Procedures Guide (NUREG/CR-2300) and the Interim Reliability Evaluation Program Procedures Guide (NUREG/CR-2728). These provide some consistency in the development of plant models. However, although both methods generally emphasize the perceived problems, which are defined during the analysis process or are the subject of the study, the level of resolution of the final models can vary with the method selected. Some variability may enter based on the experience and interest of the individual analyst.

Some techniques have been developed to allow the analyst to spend less time in actual model construction so more effort can be spent on the investigative aspects of system modeling.* These techniques, developed for system failure modeling, consist of either abbreviating the model graphics or using preconstructed fault tree logic modules, as appropriate, for fault tree construction. While these refinements remove some of the drudgery from the modeling effort, they do not necessarily reduce the likelihood that the inexperienced analyst will build a model that may not accurately represent the system being analyzed and, if not used properly, could increase the likelihood of modeling error.

The investigation of common-cause failures has also been refined. Indeed, Section B.6 addresses an entire class of potential common-cause failures (external events). Some other types of common-cause failures, such as common environmental dependencies, are not usually treated explicitly in the models or addressed during model quantification. Some of the potential failures in this class are excluded primarily because of the inability to quantify their occurrence frequency (e.g., manufacturing errors, installation errors).

In addition to common-cause failures, the modeling activities examine a broader spectrum of dependencies, including those between initiating events and mitigating systems, multiple mitigating systems, and support systems and mitigating systems. The dependencies included are functional relationships, human error, shared hardware, and shared service.

The treatment of dependencies is determined somewhat by the level of resolution of the modeling activity. For initiating events that are treated statistically, their potential interactions with mitigating systems are treated implicitly by identifying the minimum set of potentially available mitigating systems for successful response.

Improvements have occurred in the treatment of initiating events. Developments in the analysis of the initiating event/mitigating system dependency now accommodate a more explicit treatment of the dependencies. Modeling techniques, primarily the failure mode and effects analysis, are being used to identify plant faults that can be accident initiators and can also degrade mitigating systems or cause their failure. Additional techniques such as constructing a master logic diagram assist in identifying a more complete set of initiators. Finally, external events are now considered as special initiators. This recognition has led to an improved treatment of

and perspective on the importance and impact of these events (Section B.5) although they are not always included in a PRA.

Identifying accident sequences has undergone some refinements primarily because of two items. The first is the changing state of knowledge of accident phenomena, which impacts the structure of the event trees and the designated outcome of some sequences. Examples of this impact are changes in perception of the importance of in-vessel steam explosions, the ability of centrifugal pumps to pump saturated fluid and the ability to cool the core after some melting occurs. Second, previous PRAs have raised questions about realistic success criteria for the various plant systems under accident conditions. Changes in definition here can impact not only the system-failure models but also the accident-sequence delineation activities by identifying new event tree headings or changes in the outcome of previously identified sequences. Some PRAs now include best-estimate thermal-hydraulic calculations to support the plant modeling effort in this regard.

One objective of plant modeling is to identify faults in the context of an accident. This includes the time of fault occurrence as well as the type of fault. However, the level of resolution of plant system models is partly determined by the data available to quantify the models. Because one of the objectives of a PRA is to produce quantitative measures of safety, events for which data are not available are usually excluded. Therefore, events like design, fabrication, and installation errors are usually not explicitly included in the models.

Model development is somewhat constrained by the ability to represent the dynamics of a postulated accident sequence with models which consist of sequences of events which either occur or do not occur. Accident timing is somewhat replicated by arrangement of event tree headings in a temporal fashion,

listing first those events expected to occur first. However, other considerations such as modeling system functional dependencies and the assumed time of system failure also impact the order of event headings.

Since the RSS, considerable activity has been devoted to developing computerized techniques to evaluate and quantify plant models. The PRA Procedures Guide, (NUREG/CR-2300) identifies the plethora of codes now available for this purpose. The primary motivations behind this activity are to handle larger models and to accommodate the interest in both the qualitative and quantitative aspects of the modeling activity. In addition, analysis activity shows a trend toward evaluation of plant models on the accident sequence level, resulting in the desire to manipulate the system models in groups to support such diverse interests as accident phenomenology, system success criteria, and identification of recovery actions.

The models can now be manipulated to provide qualitative information for evaluating accident sequences and/or system failure, depending upon the focus of the analysis. One important outcome in the evaluation of the plant system models is core damage. Therefore, core-damage accidents can be described in terms of combinations of classes of initiating events and mitigating system faults. This qualitative evaluation can be accomplished on the accident-sequence level so that each sequence can be expressed in these terms at the level of resolution of the initiating-event and system-failure models.

Evaluating quantitative models yields the predicted probability of system failures and/or the occurrence frequency of accident sequences. Quantifying accident sequences includes considering system success states, as well as failure states, when those successes are important to the postulated outcome of a sequence. Combining the accident sequences in classes

based on similar outcomes allows quantification of those outcomes.

The primary techniques used for evaluating and quantifying models are computer codes that manipulate the models to find minimum-cut sets. For an accident sequence this is a minimum set of faults that can result in the accident sequence including consideration of the initiating event and all specified system successes and failures. These codes also compute the probability of system failure or occurrence frequency of the accident sequence based on the logic of the models and the input data.

Although many of the computational techniques have been developed to accommodate large models, PRAs generally produce models which require some effort to reduce their size before and during the evaluation process. Some of this reduction activity is manual and some employs the computational technique in use. The model-reduction activity is based on the objective of retaining only the important and/or numerically significant information in the reduced model. The two major techniques associated with reduction are: (1) coalescing events independent from other models into a single event, and (2) truncating models and/or cut-set results based on numerical criteria. These reduction efforts are now aided by computational techniques that can identify independent submodels or can keep track of the number of minimum-cut sets which are not numerically included in the quantitative result.

The PRA plant-modeling activity is subject to the same limitations as most modeling activities. The major limitation is that analysts have to make assumptions and compromises to produce a workable model. This generally results in models which allow judgments regarding the best perception of reality, but do not allow absolute judgments regarding reality. The resultant limitations and uncertainties associated with

system modeling fall into the two broad categories, completeness, and representativeness.

The concern over completeness of the modeling activity is basic to the subsequent use of the results to support judgments as to the level of safety of nuclear plants. Indeed, PRAs would be very useful as tools to "verify" the level of safety. At the same time, obviously the plant models cannot include all possible occurrences.

This lack of absolute completeness, however, is only a relative problem. Judgments based upon the results can be made even though the results have some completeness limitations. Some areas of modeling seem somewhat less complete than others at present. Understanding these limitations can assist the user in drawing appropriate conclusions from the results.

Identifying initiating events is one particularly difficult area in which to establish completeness. Techniques for identifying these events include reviewing general nuclear-plant operating experience, developing master logic diagrams, and analyzing initiating event/mitigating-system interactions. Although these techniques do provide a spectrum of potential functional challenges to the plant, defining the numerous ways and contexts in which these events may occur is more difficult. This is particularly true of the class of initiators called "external initiators". These difficulties are addressed, though not in detail, by classifying groups of initiating events by potential mitigating actions. However, these classes provide only representative cases and not the particulars of each event.

Thus, PRA results represent the level of safety only in relation to the postulated initiating events. This situation is more realistic than it may seem. For example, this approach is better than the traditional design basis approach because

events are assessed in an integrated context to produce a broader view about the impact of the number of faults, operation of systems, and potential operator actions.

The plant-system models are constrained by considerations of scope, failure importance, size, and the ability to postulate what might happen. This results in excluding some potential faults from the models. However, the basis for excluding items from the models is generally either that the fault likelihood is not believed to be important to the final quantitative result or that there is no reasonable way to quantify it.

PRAs do not treat some types of accidents in detail because of the limited ability to place them in the appropriate quantitative context. These include such issues as pressurized thermal shock, reactor vessel rupture and certain containment failures (e.g., failure at less than burst pressure).

The importance of these items to the ultimate PRA results now depends upon opinion. Only technical resolution of the likelihood of these events will allow them to be included in a plant PRA. However, PRA can be applied to these and similar issues to learn more about them, individually, and to help identify the sources of the uncertainty involved.

The ability of the plant-modeling activity to model postulated accidents accurately is limited. First, it is limited by the usual requirement to provide a Yes/No statement of failure. This means that system models treat component events as two-state situations. For example, pumps are generally treated as either delivering required flow or not; reduced flow is not treated. System failures are defined in terms of less-than-minimum equipment requirements. This is a required transition from the functional definition of system failure to the operability definition.

System modeling is partially a process of constantly generating large amounts of information and classifying that information into a small number of classes, in each of which all the information shares similar characteristics. Therefore, the results represent the manipulation of classes of information sharing similar qualities, somewhat reducing the precision of the results. Further, assumptions must be made about the timing of failures and the impact of partial failures. This tends to make the results specific only to those assumptions.

Accuracy of plant system modeling is also affected by the available data. Data determine the level of resolution of the understanding of system operation and thus influence the level at which faults in the models are postulated. Further, the data are associated with some degree of uncertainty. This, coupled with the questions regarding the degree of completeness of the models, leads to uncertainty in the results.

Modeling accuracy is affected by its relationship to accident processes. On one hand, the postulation of accident sequences provides data on plant-system states for input to accident process analysis. On the other hand, the accident process analysis provides input for postulating and classifying accident sequences. The ability to postulate accident sequences accurately depends upon the ability to understand the accident processes and their impact on system operation and function. The uncertainty in the knowledge of the accident process contributes to the difficulties of accurate plant modeling.

Validation of PRA results through experience is not now readily available, because by its very nature PRA deals with events which are predicted to be very rare. This raises questions about the correctness of the results, and is especially important when considering the utilization of results in an

absolute sense. If PRA results are to be used to make decisions about regulation, the validity issue in relation to plant design and operation are also important. This returns the issue of correctness to an examination of the proper application of the methods, improvements in the completeness and accuracy of the PRA plant-modeling techniques, and comparison of plant experience with predictions to assist in methodological development.

The greatest strength of plant modeling lies in the process itself. Following the patterns of investigation dictated by application of the modeling techniques results in a rigorous look at plant design and operation. This provides the analyst with additional insight over traditional design review processes. This insight results in improved ability to identify potential design or operational problems or weaknesses and provides the basis for suggesting optimum remedies or solutions.

Another advantage of plant modeling over other examinations of plant safety is that it constructs an integrated framework for examining the importance of perceptions of individual items associated with plant design and operation. This means that the results represent the synthesis of knowledge about such diverse items as perceptions of human error rates and thermal-hydraulic conditions in containment. Not only are these diverse items treated in a combined model, but their importance to the results can be compared, individually or in groups. Also, the resulting model provides a description of the context in which to address the importance of various issues from steam explosions to system failures.

A significant strength of plant modeling lies in the use of the models to perform sensitivity studies. This effort consists of changing some aspect of the modeling input such as probability data, basic assumptions, plant design, system

success criteria, or physical process understanding, followed by redoing the model development and/or evaluation and comparing the results to the previous iterations. The analyst, designer, decisionmaker, or regulator can then identify the impact on the results of differences in perceptions of various issues. Sensitivity studies also provide insights into the impact of potential design or operational changes. Finally, sensitivity studies yield results that can assist in indicating the numerical magnitude of selected modeling and assumption uncertainties.

B.1.2 Impacts of Limitations and Uncertainties

As stated above, the primary uncertainties and limitations associated with system modeling fall into the two broad categories of completeness and representativeness. The degree to which these uncertainties and limitations affect the results of system modeling is somewhat difficult to determine. Therefore, the inherent uncertainties and limitations associated with system modeling must be recognized when the results of these modeling activities are being interpreted.

With respect to completeness, identifying all possible occurrences which affect the initiation and ultimate course of accident sequences is not possible. Thus, some potentially significant events might not be included in the models. Some events are specifically excluded from the models on probabilistic grounds. In addition, the possibility always exists that certain kinds of events might occur which have never been identified or defined. This incompleteness usually arises from a lack of knowledge of all the detailed interactions and dependencies which exist among basic failure events and failure causes. Analyses of operating history should identify those dependencies and interactions which are likely to occur. This knowledge is being fed back to improve the modeling in PRAs. Thus, the completeness issue is of concern and is a motivation for continued improvements in

knowledge of the dependencies that exist. This concern, of course, applies to all safety analyses, and PRA methods can be used to help identify where additional knowledge can have the most impact on risk characterization.

Many assumptions must be made during of plant modeling. In addition, the constraints imposed by the modeling techniques themselves require a number of compromises. Coupled with a limited understanding of some physical processes, these facts constitute the issue of representativeness of plant modeling. Modeling, by its nature, represents an abstraction, a compromised representation of physical reality. Inevitably, representativeness becomes an issue under these circumstances. The degree to which this is a problem is difficult to measure because few references are available with which to gauge accuracy. Possibly, as the accuracy of plant modeling improves, the estimated probabilities of some accident scenarios will decrease slightly because the modeling process may be somewhat conservative. As understanding of some phenomena, particularly those which are now poorly understood, increase some reduction in the estimated likelihoods of some of these events is expected. In any case, a reasonable suggestion appears to be that limitations imposed by considerations of representativeness, like completeness limitations, do not negate the usefulness of model results as long as these limitations are properly recognized when interpreting the modeling results.

B.1.3 Potential Improvements in Plant Modeling

The single greatest source of improvement to plant modeling techniques is likely to be increased experience in both actual nuclear power generation and model usage. For example, over a period of time the question of completeness will continue to be addressed. If previously unidentified events should occur, the completeness of the models can be enhanced by considering these new kinds of occurrences. If no new events

should occur, even greater levels of confidence in the completeness of the models will be warranted. Likewise, increased operating experience and continuing data collection will provide empirically derived values for initiating-event frequencies and component-unavailability data.

Scaled thermal-hydraulic experiments may provide a better understanding of some accident processes, which may, in turn, provide an enhanced ability to model various phenomena. With increased experience, some of the modeling techniques themselves may be modified to improve and extend their capabilities. An example might be the development of a model for partial failures.

Although improvements will undoubtedly result, some limitations seem to be inherent in the modeling techniques and will probably continue to exist. The need for formulating assumptions and constraints with respect to plant models is a direct outgrowth of the inability to consider everything or know everything. Assumptions and constraints, by definition, introduce uncertainty and exclude information from the analysis.

B.2 Human Interactions

Experience in many industries has shown the importance of human-plant interactions in the operation and safety of various types of plants. PRA methodology has emerged as a promising tool for prospectively assessing the impact of humans on the plant risk and understanding the importance of many man-machine interface issues.

Human-plant interactions constitute an important link in the operation, control, maintenance, and testing of equipment in virtually all industrial activities. These beneficial interactions, including repair and recovery operations, often enable various systems to achieve high availability. However,

a dichotomy exists. While such human interactions are largely responsible for maintaining high availability; the human contribution to accidents that do occur has been estimated as high as 90% in the cases of the airline (NUREG/CR-2744) and chemical industries (Joschek, 1982).

Published PRAs for nuclear power plants have yielded similar findings (Joksimovich, et al.). Past PRAs have indicated that both beneficial and detrimental contributions of the human influence impact the order of dominant sequences and, hence, the risk profile of the plant. For example, the studies invariably have included human actions that can cause initiating events or result in the unavailability of plant systems before an initiating event. In some studies, human interactions that compensate for accident causes include the diagnosis of and recovery from an accident sequence. Clearly, PRA techniques provide a framework for assessing the importance of human interactions in a spectrum of accident sequences.

The definition of specific accident sequences in PRA studies provides the analysts with a tool for determining where the human might affect the risk estimates. For example, the uncertainties in the quantitative impact can be assessed, the ways in which humans impact the course of an accident can be described, and the importance of humans in a particular sequence can be quantified.

B.2.1 Background

The basic methodology for considering human interactions stems from the techniques first developed in the RSS (NUREG-75/014).

The RSS, in addition to representing the first full-scale application of PRA techniques to LWRs, also used a human reliability technique called THERP (Technique for Human Reliability Error rate Prediction). THERP, initially developed in

1961, has undergone a number of improvements since; for example, the Human Reliability Analysis Handbook (NUREG/CR-1278) has been developed to make the technique useful to a wide spectrum of analysts. The application of the THERP methodology was documented by examples in NUREG/CR-2254.

The limitations in the early methodology were recognized and have stimulated numerous discussions with human behavioral experts outside the PRA field to gain suggestions for improvements. As a result, a revised version of the Human Reliability Analysis Handbook was prepared (NUREG/CR-1278). The revised version attempts to improve the consistency of trained analysts and to expand upon models for the diagnosis function in accident response. The proposed diagnosis model is generic for all events; therefore, considerable judgment is required in applying it because data collection has not been directed toward diagnosis. A need remains for the diagnosis models to consider the different thought processes associated with specific accident conditions. Human decisions have been key factors in several actual events and misjudgments in thought processes may result in greater impact on the systems than the processes required in following the steps of a procedure.

The techniques used to model the human errors in the PRA studies vary considerably. For example, some current PRAs try to account for cognitive human behaviors with techniques such as the operator action tree (Wreathall), time-reliability correlations (NUREG/CR-3010), confusion matrices (Potash, et al.), and specific recovery models (NUREG/CR-2787). A review of five recent PRA studies (Joksimovich, et al.) showed that the modeled human interactions have a major impact on the core-melt frequency and ordering of the dominant sequences in many studies. Furthermore, some of these methods appear to be very study-specific, and have been integrated differently in different studies, making reviews and comparisons difficult. Thus, while the development of techniques was expanding

rapidly , approaches for integrating the techniques into PRAs lagged behind.

B.2.2 State of the Art

The most recently published PRAs benefited from the groundwork established in the RSS. In general, they have recognized the importance of human interactions, although this is not always stated quantitatively. The types of human interactions identified in the recent PRA studies included some of the categories listed below.

Type 1. Prior to an initiating event, plant personnel can compromise equipment and availability by inadvertently disabling it during normal operation or when the plant is down for repair or testing.

Type 2. By committing an error, plant personnel can initiate an accident.

Type 3. By following procedures during the course of an accident, plant personnel can operate standby equipment that would terminate an accident.

Type 4. Plant personnel, attempting to follow procedures, can make a mistake that aggravates the situation or fails to terminate an accident.

Type 5. By improvising, plant personnel can restore and operate initially unavailable equipment to terminate an accident.

For each type of human interaction, the three important questions are: (a) how are the human interactions incorporated, (b) which techniques for modeling human interactions are used, and (c) what type of data are available? These are addressed below.

Type 1 interactions are modeled by selections of system unavailability data or by modeling explicitly the procedures for performing tests or maintenance with the THERP technique. Type 1 interactions are generally included in all PRA studies and are easily incorporated into the standard fault trees. The data in NUREG/CR-1278 are generally applicable, but the human error probabilities may not include all aspects of decision making.

Type 2 interactions are generally implicit in the selection of initiating events. They usually include human impacts already in the outage-frequency data base which, because of agglomeration, may not identify specific human interaction causes. A few studies have identified specific human-caused initiating events through failure modes and effects search methods.

Type 3 interactions involve the success and failure in following preestablished procedures and the ability to select the correct procedures given the information available to the operator. The THERP technique has been used to develop the framework for quantifying the reliability of following a procedure. More recent developments such as OATS (NUREG/CR-3010) and the improved THERP diagnosis model (NUREG/CR-1278) account for correctly selecting the appropriate procedure. Such interactions may be incorporated into the logic of the fault trees and event trees by the system analysts or factored into the analysis during accident sequence quantification. The data used generally are derived from NUREG/CR-1278 or expert opinion techniques such as paired comparisons or psychological scaling (NUREG/CR-2255).

Type 4 interactions are the most difficult to identify and model. Modeling requires iterations between the human reliability analysis and the system analysis to help identify the

important human interactions which could aggravate the situation. A technique has been developed to help identify these human actions (Potash). A confusion matrix is constructed to help the analysts identify cases where the operators mental image of the plant differs from actual state and thus the operator's actions become "the right actions for the wrong event". Quantification is carried out by expert opinion. Only a few PRAs have attempted to include this type of interaction, and only to a limited degree. Once the actions are identified, they can be incorporated into the logic structure of an event tree or fault tree or factored into the quantification process. Very few data are available for predicting these types of human interactions. However, retrospective analysis can usually identify these kinds of causes.

Type 5 recovery actions are generally included in the evaluation of accident sequences which dominate the risk profile. These actions may include the recovery of previously unavailable equipment or the use of nonstandard procedures to ameliorate the accident conditions. They can be incorporated into the PRAs as recovery factors on the frequency of the accident sequences. Quantification has often been based on estimates of the probability included by curves of recovery versus time without considering the many additional parameters which may be important. In most cases these estimates have been developed by expert opinion.

The incorporation of human interactions into PRAs was often left to the judgment of the systems analysts. The need for detailed interactions between the systems analysts and human factors specialists was identified in the PRA Procedures Guide and further guidance appears in a draft report which has been issued for review and comment (Hannaman, et al.).

B.2.3 Issues Contributing to Current Limitations

The usefulness of human interaction analysis in PRA can be enhanced significantly by addressing the issues that currently contribute to limitations:

1. Human behavior has been recognized as a complex subject for centuries and does not lend itself to simple models such as those for component reliability. Thus, the analysis of human interactions is the area of systems analysis most dependent on the judgment of experienced analysts. For example, in one study the assessment of recoveries may be bounded by a single parameter, such as a failure probability of 0.2 (NUREG/CR-1659) whereas, in another study the recoveries are given more realistic assessments on a judgmental basis (NUREG/CR-2787).

2. The description of the human impact has not been fully developed since human impacts have often been classified as either success or failure to match equipment failure logic. In most PRA, the use of techniques such as THERP or OATS results in the assignment of successes and failures to each branch of the tree. The possibility of other operator conditions which might affect the system in other ways has not always been considered. A framework for helping analysts make such considerations is available (Hannaman, et al.) and further improvements are anticipated in this area.

3. Generic human failure data has been applied on a judgmental basis, because a simplified mode of the various parameters which affect human performance has not yet been fully developed. The current techniques for selecting data from NUREG/CR-1278 and applying them to a human reliability analysis (HRA) tree requires considerable judgment and may not be completely reproducible by other

analysts (Brune, et al.). A simple model of human behavior such as the OAR model, or a model based on recovery time alone improves the reproducibility but does not provide information on how the likelihood of recovery varies with other parameters. Structured use of expert judgment is one way for assessing this quantitative impact (Embrey).

4. Human dependencies have been assessed primarily among the humans rather than as part of human-plant interactions. The modeling of human dependencies from the HRA viewpoint is described in NUREG/CR-1278. Although the technique provides for quantitatively assessing these dependencies between humans, the human-system dependencies may need to be addressed in greater detail. This is an area where the diagnostic models need to address multiple options. The multiple-option concept can be addressed in the confusion matrix (Potash, et al.) and in the structuring of expert judgment (NUREG/CR-2255).

5. Techniques for considering sensitivities and uncertainties currently address the quantification uncertainty in the data as opposed to alternate logic for incorporating the human. One of the weaknesses in the quantification of human-interaction uncertainties is that the modeling techniques introduce a structure for incorporating a single human link to the system reliability model. The current techniques for assessing the uncertainty of the quantitative impact of the error rates are based on changing the parameter of interest (e.g., the human failure probability) and comparing this changed condition to the unperturbed condition (NUREG/CR-2906). An improvement would be to examine changes in the logic structure also. This requires an improved ability to state the assumptions involved in incorporating the human interaction.

While the impact of these current limitations on the risk profiles assessed in published PRAs is difficult to state, many analysts feel that they are within the stated uncertainty bounds. The major impact on risk is felt to be on the probability of accidents as opposed to the consequence.

B.2.4 Developing Art

The analysis of human interactions in a PRA is clearly a developing art. Improved areas for the analysis of human-system interactions in future PRAs are likely to include:

1. development of interim methods for considering the importance of operator desicionmaking under accident conditions (NUREG/CR-1278),

2. development of certain representations of the time dependent impact of human interactions on the success or failure of a system or safety function, e.g., OATS (NUREG/CR-3010),

3. use of a more structured technique for developing data from expert opinion (Embrey),

4. development of more systematic approaches for incorporating human interactions into the PRA framework (Hannaman, et al.) and better integration of the systems and human reliability analyses, and

5. collection of training simulator data to verify some of the judgmental data and support the development of simple models of human behavior (Kozinsky and Pack).

Improved consideration of these factors in PRA should lead to substantially greater understanding of possible human behavior under accident conditions.

However, limitations to the detailed description of human interactions will still exist and they should be recognized. Both the qualitative description of the human-plant interaction logic and the quantitative assessment of those actions relies upon virtually untested judgments of experts. One area needing additional work is the development of simple mathematical human-impact models that are adequate for PRAs. Such correlations as OATS are simple to apply, but give little improved information about the behavior characteristics of the operation. This then requires little judgment during application, because all the judgment is introduced into the time-reliability curve. Factors which lead to variations in the curves should be identified from the information in the simulator studies. Thus, a more detailed model is needed so that the collection of data is directed toward identifying parameters that are likely to influence the probabilities of human interactions in an accident situation. EPRI and others are developing improved models.

Future Outlook. The future outlook for the subject of human-plant interaction modeling in PRA is excellent. The critical review of PRAs by experts in the area of human interactions has generated remarkable advances. However, the depth of the techniques must be expanded so that the impact of changes in design, procedures, operations, and training, etc., can be measured in terms of a change in a risk parameter such as the core-melt frequency. Then trade-offs or options for changing the risk profile can be identified. To do this, the methods for identifying the key human interactions, for developing logic structures to integrate human interactions with the system failure logic, and for collecting data suitable for their quantification must be strengthened. These items remain to be accomplished before the associated uncertainties can be substantially narrowed.

B.3 System Model Data Evaluation

The data that are used in a PRA consist of constants (parameters) that need to be supplied to the PRA models. The data can be divided into system model data, which are used in the accident frequency evaluations, and consequence data, which are used to evaluate the consequences associated with each accident sequence. The different types of system-model data generally used in a PRA consist of:

1. Initiating Event Data
   Example: Transient frequencies, loss-of-coolant frequencies

2. Component Failure Data
   Example: Valve failure rates, pump failure rates

3. Test Data
   Example: Surveillance test intervals, surveillance test durations for pump tests

4. Maintenance Data
   Example: Unscheduled and scheduled maintenance intervals and durations for pump maintenance

5. Common-Cause Data
   Example: Fractions of failure causes which result in multiple valve failures

6. Human Error Data
   Example: Human error rates, human recovery probabilities

7. Uncertainties Associated with the Above Data
   Example: Error factors representing approximate 95 percent bounds on the data

The PRA Procedures Guide, Chapter 5, (NUREG/CR-2300), discusses these different types of data in more detail.

The data which are used in a PRA can be either generic or plant specific in nature. Generic data represent a class of reactors or class of components. The class can have any nature, encompassing individuals with similar specifications or with dissimilar specifications. Examples of generic data are failure rates for emergency core-coolant pumps for Westinghouse reactors and transient occurrence frequencies for General Electric reactors. At a minimum, the generic data values consist of a central data value for the class (e.g., a median value) and a characterization of the spread of individual data values in the class (e.g., the difference between the maximum value and minimum value). A probability distribution describing the variation of individual data values in the class is sometimes provided. Various generic data sources are available for PRAs and are described in the PRA Procedures Guide. Generic data sources include RSS data, licensee event evaluations, and compilations of plant maintenance logs.

Plant-specific failure data are for the specific plant being analyzed by the PRA. The plant-specific failure data are obtained from the plant's records and reflect the peculiarities for the particular plant. Even for plant-specific data, failure histories of similar type components are usually aggregated. This aggregation thus involves an averaging of individual component failures. The statistical treatments used to analyze plant-specific data are described in the PRA Procedures Guide. Obtaining plant-specific data can be a significant effort in a PRA.

Because of the difficulty of obtaining plant-specific data, generic data are used in most PRAs to supplement the available

plant-specific data. The approaches which are used to integrate generic data with plant-specific data vary. Whenever generic data are used for a component, a value representing a characteristic member of the generic class is assigned to the component. Sometimes this may not be very near the actual data value for that specific component. For example, an average failure rate for a motor-operated valve in the industry may be assigned to a particular valve in the plant; however, that specific valve's failure rate may differ from the average. Therefore an important consideration is whether possible variations from the average can impact the risk results. Uncertainty analyses or sensitivity analyses accommodate this; the PRA Procedures Guide describes particular methods used for these analyses.

In some areas insufficient recorded historical experience is available to allow meaningful generic or plant-specific data to be obtained. For these areas, subjective data are used. Subjective data in general come from assessing data values that are not explicitly based on statistical analyses of past history. Subjective data thus represent the opinions of the analyst, or individuals involved in the PRA, about appropriate values for the parameters. These judgments are based on the analyst's feelings, experience, and knowledge. Subjective data are used particularly for human error rates and common-cause failure probabilities. Sensitivity studies or uncertainty analyses are important for subjective data because of the uncertainties involved.

B.3.1 State of the Art

The RSS based its estimates on approximately 17 reactor years of experience. At the present time, approximately 320 reactor years have been accumulated in the United States. The analysis of this experience to obtain required PRA data has been spotty. A brief review of the different data areas is given below.

Initiating-event data for transients have significantly improved since the RSS. Data have been tabulated for a wide spectrum of transients for both PWRs and BWRs, and both generic and plant-specific values have been tabulated.* However, because of the paucity of data, LOCA-initiating-event data have experienced only marginal improvement since the RSS, with many current PRAs using RSS numbers for LOCA or modifying them to include valve or pump rupture or leakage contributions.

A significant amount of plant-specific data have been generated for those plants which have been subjected to PRAs. In general, these sources of data for individual plants have not been combined into an industry-wide data base or used to upgrade industry-wide data bases. The generic data used in current PRAs have generally not improved over RSS. In fact, many current risk analyses use RSS data, or some variant as the generic data base. The improvements that have occurred have not exerted a major impact on either the numerical results of the PRAs or on the insights obtained.

Current data bases, including plant-specific data bases, generally do not relate component failure rates to root causes of failure; corrective actions required, if the failure rates give high risks are, therefore, seldom clear.

Test and maintenance data, including limiting conditions for operations, are obtained generally from plant technical specifications. Plant maintenance logs also sometimes provide more precise values. Corrective maintenance intervals and durations are among the more difficult data to obtain and are occasionally subjectively estimated after discussions with plant personnel.

Common-cause data, which describe the likelihood of multiple failures from common causes, have improved only marginally

since RSS. Common-cause probabilities remain largely subjectively estimated and generally are not tailored to specific plant environments and maintenance and operation policies.

Since the RSS, human-error analyses for routine procedural errors have become more codified; however, human-error data are still largely subjective with little validation from experience. Human error data for cognitive (decision) errors and that for errors occurring under accident conditions are still generally unavailable. Also, data indicating the likelihood of the operator correcting or mitigating accident situations do not exist.

B.3.2 Sizes of Uncertainties

Uncertainties in data are generally expressed as error factors where the error factor is the ratio of the upper bound estimate to the median value or best-estimate value. The upper bound estimate is usually an approximate 95 percentile value (i.e., a 95% probability, or confidence, that the true value is less than this upper bound). The range, which is the ratio of the upper bound estimate to the lower bound estimate, is the square of the error factor. The lower bound is usually an approximate 5 percentile value. The range thus represents approximately a 90% confidence range or probability for the value.

For present transient-initiating-event data, the error factors are considered to be small, only about 2 to 3. For loss-of-coolant-initiating-event data, the error factors are thought to be 3 to 10.

Component failure-rate data have error factors of approximately 3 for active components (pumps, valves, etc.). For passive components (pipes, wires, etc.), the error factors are generally in the range 10 to 20; passive-component-failure

rates are generally substantially lower than active-component-failure rates, even considering the uncertainties.

Data on test and maintenance intervals and durations generally have associated error factors on the order of 2 or less. Corrective-maintenance intervals and maintenance durations have the larger errors.

Error factors for common-cause probabilities involving two or more coupled failures usually increase as the probability values decrease. For common-cause probabilities in the vicinity of 10^{-5} the error factor is of the order of 3 to 10. These error factors apply to the probability of multiple failures occurring (i.e., unavailabilities).

The above error factor values represent gross averages over PRAs which have been performed, and can vary from PRA to PRA. PRAs which use plant-specific data generally have smaller uncertainties (error factors) than those that use only generic data. Uncertainties arising from data uncertainties are generally the only ones that are explicitly quantified in a PRA. Uncertainties in dependent failure data and human error data often dominate the data uncertainties associated with calculated system unavailabilities and accident frequencies. The specific effect of data uncertainties, however, depends on application. When a single contributor dominates the risk, then the uncertainties in data for that contributor will have significant impacts on the results. When many unrelated contributors contribute equally to the risk, then the data uncertainties on any one contributor will not have a large impact.

B.3.3 Potential Improvements in Data

A significant amount of plant-specific data have been generated by the PRAs already performed. However, these data have not been assembled together for plant comparisons and for developing a generic data base. A data system like NPRDS*

could be the vehicle for this assembling, comparing, and summarizing of plant experience. Such a data system could be particularly valuable for identifying time trends, for showing outlier component and system behavior, and for providing information on frequencies of failure causes.

Because of their importance and their present large uncertainties, additional plant-experience data on dependent failures and human error represent areas where data collection can effect the greatest improvement. The realism of models and data for test and maintenance can also benefit from improvements. This would allow more realistic analysis of plant technical specifications and would allow evaluation of the reliability assurance impacts of testing and maintenance.

B.4 Accident Progression, Containment Response and Source Terms

This section describes the status of modeling of in-plant accident processes. The principal products of the in-plant consequence analyses performed for a PRA are called the environmental source terms. These source terms describe, for accident sequences or groups of accident sequences, the characteristics of the release of fission products to the environment. The following characteristics of the release are described:

- The conditional occurrence frequency for the source term for each accident sequence or sequence category
- The time of release (with respect to reactor shutdown)
- The duration of release
- The warning time available for emergency actions
- The elevation (location) of release
- The thermal energy release rate into the environment
- The quantity of radionuclides released into the environment.

In-plant consequence analysis includes the consideration of a wide range of phenomena, some of which are not fully understood. This phase of analysis begins by determining whether sequences would cause severe fuel damage. Sequences that do not result in significant fuel damage are of little further interest to health risk because of their comparatively small consequences. Thermal-hydraulic codes are used to describe the progression of a severe accident from the time of the initiating event through core uncovering, fuel heatup, clad oxidation, fuel slumping, vessel failure, and fuel-concrete interactions. The defense-in-depth approach to the assurance of reactor safety provides a number of barriers to the release of fission products. These barriers include the fuel matrix, the cladding, the reactor coolant system boundary, and the containment building. A core meltdown accident eliminates the first three barriers and produces threats that challenge the final barrier, the containment building. The consequences of a core meltdown accident are influenced strongly by whether the containment fails and, if it fails, by the timing and mode of failure. The principal threats to containment that must be evaluated in a PRA are overpressurization by rapid steam generation caused by a molten-fuel-coolant interaction; shock loading from hydrogen detonation; rapid pressurization from hydrogen deflagration; thermal loading from hydrogen burning, hot gases or thermal radiation from the core; missile production in a steam explosion; and base mat penetration. The capability of the containment structure and penetrations to withstand these loads must then be determined. Two other potentially important containment failure modes involve inability to isolate the containment at the time of an accident and the direct bypass of the containment if the valves connecting high pressure and low pressure piping systems fail.

The potential public hazard in a severe LWR accident is the release of radioactive fission products to the environment. The quantities of fission products available for release from

the plant depend on the processes by which fission products are released from fuel and transported in the reactor coolant system and containment, and the processes acting during possible transport in buildings outside the containment, before the fission products reach the environment. In the past few years, questions have been raised about the realism of the methods used in the RSS and subsequent PRAs to analyze these processes. The uncertainties in these methods will be discussed in Section B.4.2 and areas for improvements will be described in Section B.4.3.

In a PRA, consequence analyses are performed for a discrete set of accident sequences or for conditions that are selected as characteristic of groups of accident sequences. In the same way that system event trees are used to organize the systems-analysis aspects of a PRA, containment event trees have been used to organize the consideration of the containment aspects of accident consequences. The development of containment event trees and the grouping of accident sequences by consequences will be described in this appendix as will quantification of the probabilities of the branch points on the containment event tree.

B.4.1 State of the Art

Methods for analyzing severe accident processes and fission-product release are described in some detail in Chapters 7 and 8 of the PRA Procedures Guide. Figures B-1 and B-2 from this guide illustrate the steps taken in these analyses. Although these methods are changing rapidly, the guide provides a good introduction to the methodology developed in the RSS and subsequent improvements in more recent PRAs.

B.4.1.1 Analysis of Severe Accident Behavior -- The BOIL code was written for the RSS to describe the boil-off and heatup of the fuel in the reactor vessel for accidents initiated by large breaks in primary system piping. Hand calculations

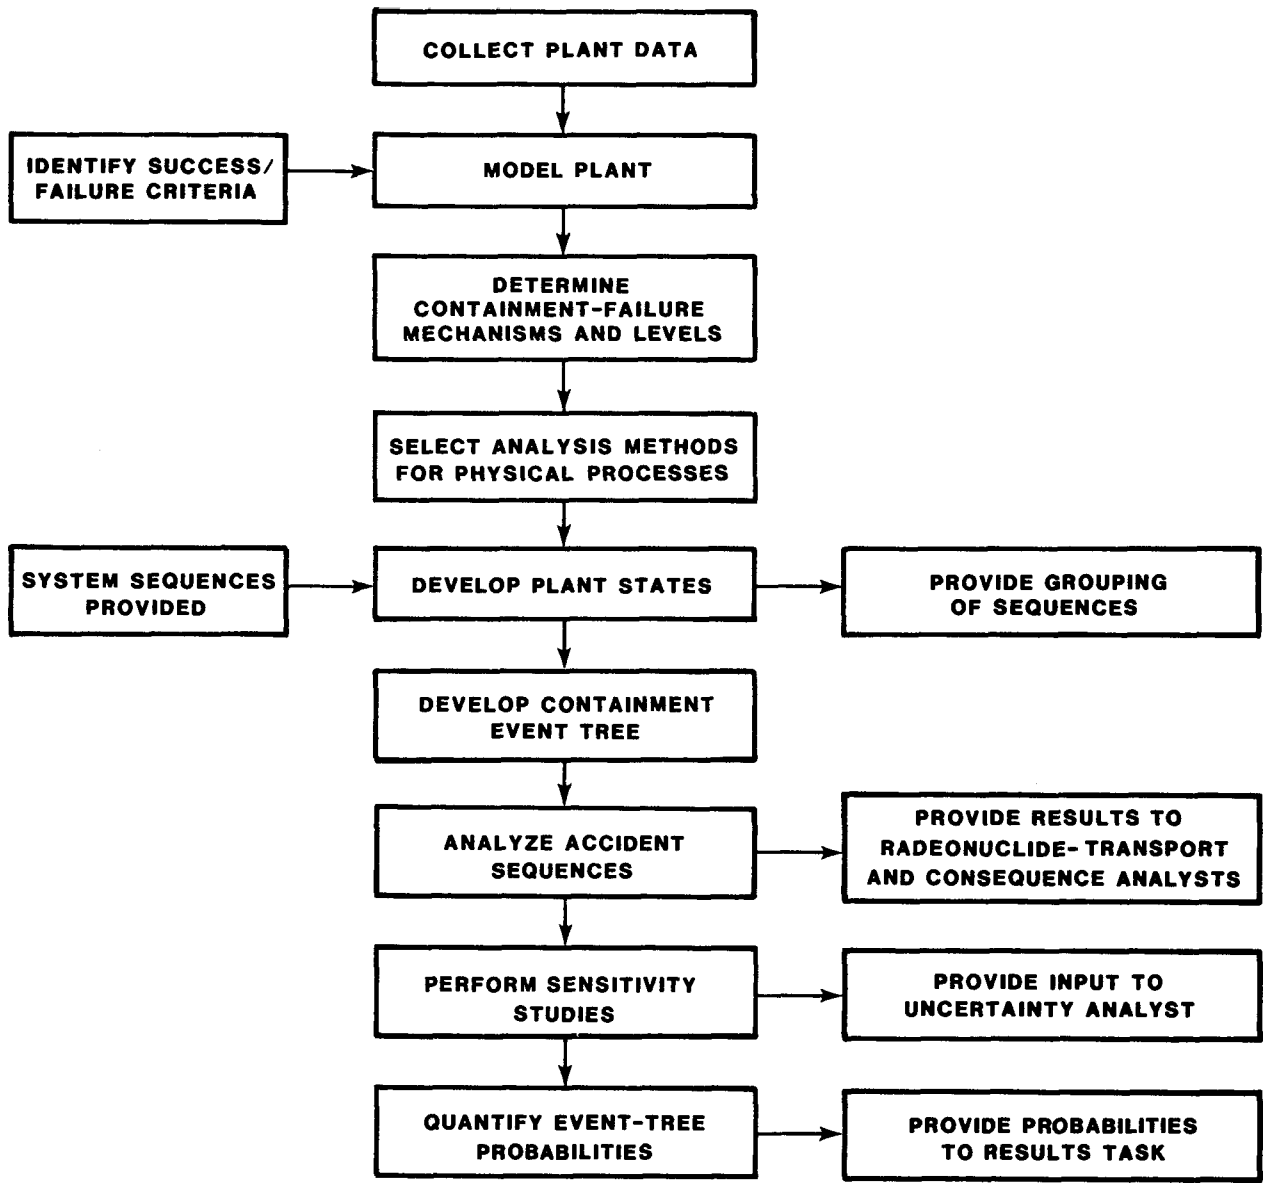

Figure B-1. Activities Diagram for the Analysis of Physical Processes

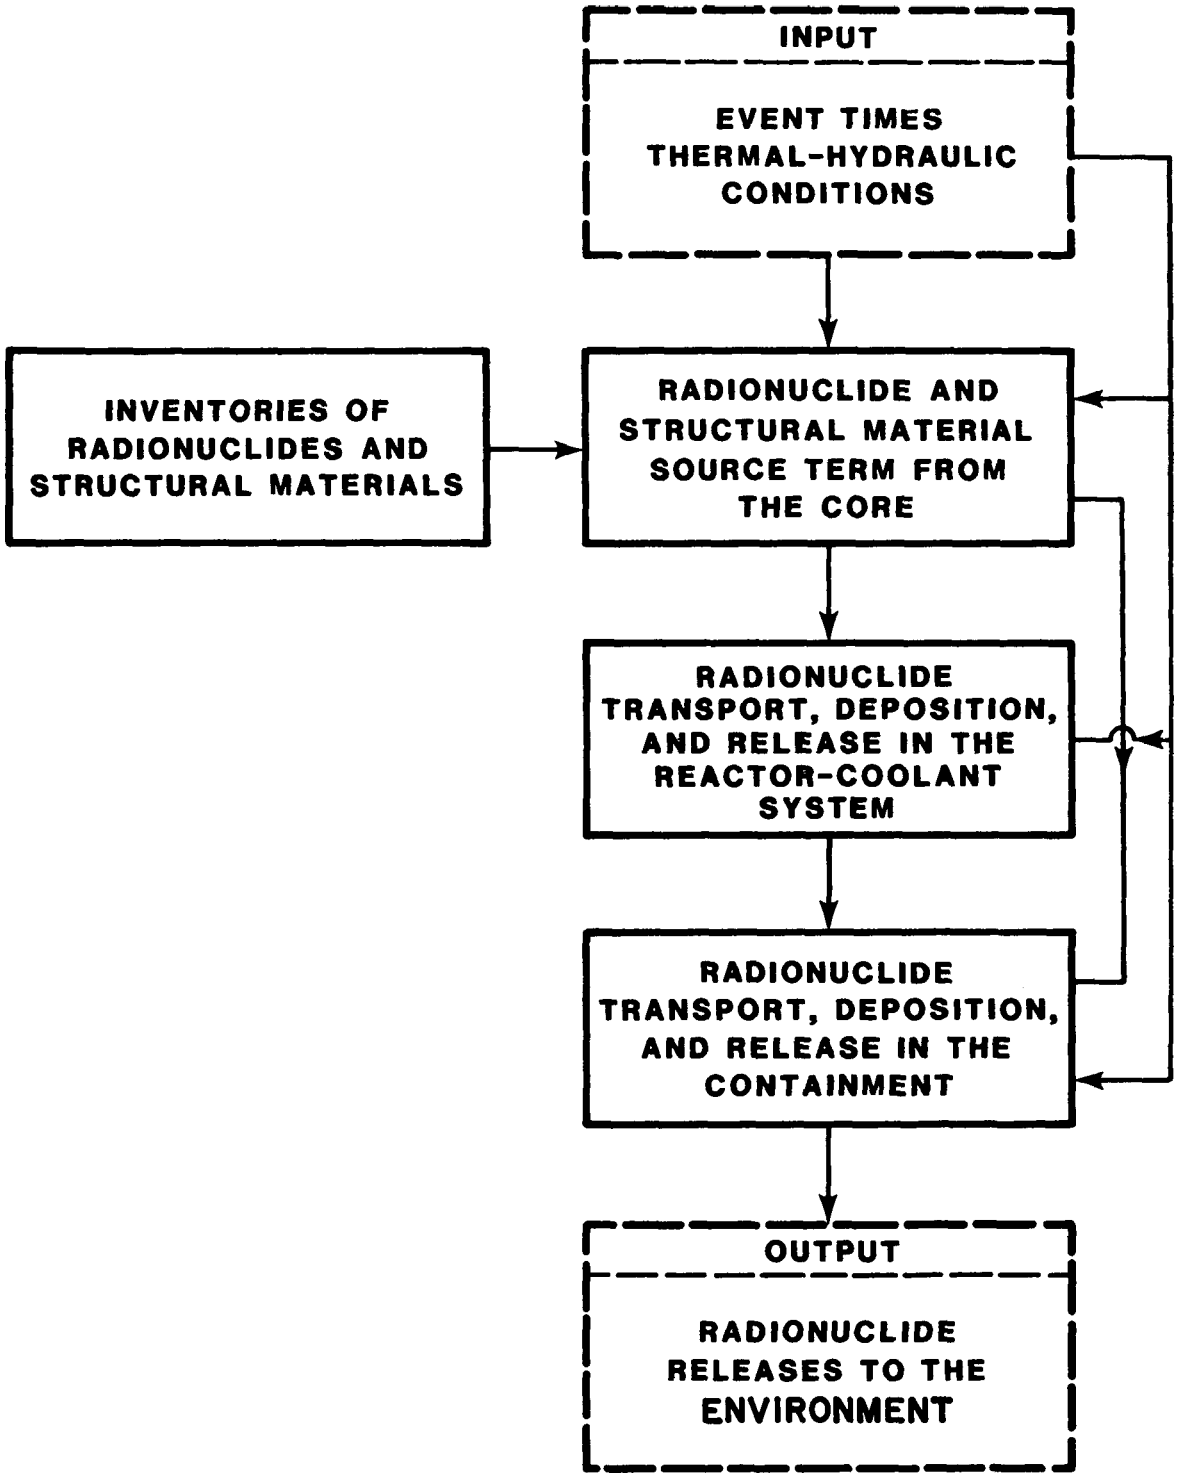

Figure B-2. Elements in the Analysis of Radionuclide Behavior in the Reactor

were used to estimate the other accident phenomena. After the RSS, the MARCH computer code* was written to enable a more consistent treatment of the physical processes in a severe accident. Some improvements were included in the modeling of processes, such as in the area of molten core-concrete interactions, but, in general, the MARCH models are simple representations of complex processes. The MARCH code very much facilitated the performance of PRAs because it is a fast running code and allows the user flexibiltiy in performing sensitivity studies. Almost all PRAs performed since the RSS have been performed with MARCH or at the same level of physical modeling as MARCH. Exceptions are the Zion and Indian Point Probabilistic Safety Studies, in which more detailed analyses were performed for important separate effects.

In the period after TMI-2 , severe accident research expanded much more broadly, focusing not so much on improving PRA but in a better understanding of severe accident behavior for possible use in plant regulation. Experimental programs in the LOFT, Semiscale, and TLTA facilities were redirected to examine the conditions of the reactor coolant system preceding severe core damage in scenarios that would eventually lead to damage. These data have provided a basis for validating codes such as TRAC and RELAP that can be used to predict if and when the core uncovers. Modeling efforts were also initiated to describe the progress of fuel degradation in more detail than in the simple models in MARCH. A more mechanistic description of the initial stages of core damage will be provided by the SCDAP severe fuel damage code under development at INEL. The MELPROG code will analyze the behavior of the degraded core from the period of slumping within the core region through failure of the vessel. The two codes are being interfaced to provide a complete description of in-vessel behavior. Validation of the models is to be provided by experiments in the PBF, NRU, and ACRR facilities. The MEDICI and CORCON codes will describe the ex-vessel behavior of

molten fuel. MEDICI will predict fuel-coolant interactions in the reactor cavity. It will be validated against experiments at SNL and BNL. The CORCON code predicts the interaction between molten fuel debris and concrete. Validation experiments are being performed in the Large-Scale Melt Facility at SNL and in the BETA facility in West Germany. These two codes will be integrated into the CONTAIN code, which analyzes the thermal-hydraulic response of the containment.

Considerable research has also been undertaken to develop an improved understanding of steam explosions and hydrogen combustion. A broad range of experiments has been performed at SNL to investigate the mechanisms for the initiation and propagation of steam explosions in mixtures of molten corium (mixtures of fuel and structural materials) and water.* The conditions under which steam explosions will occur and the energetics of the reaction are now known reasonably well. PRAs have been performed at SNL to reestimate the probabilities of energetic steam explosions leading to containment failure in core meltdown accidents.

Most of the hydrogen behavior research program has been performed since the TMI-2 accident. This program has also involved model development and extensive experimentation. Data have been collected to determine directional flammability limits as a function of composition, ignition requirements, and conditions leading to flame acceleration. The effects of engineered safety features on controlling hydrogen combustion have also been investigated. In cooperation with EPRI, large-scale tests are in progress at the Nevada Test Site. The HECTR code is being developed to predict the magnitude of loads generated in multi-compartment containments in hydrogen burning events.

The mechanistic physical-process-analysis codes that NRC contractors are developing are not intended for direct use in

PRAs. Many of these codes are in an early stage of development, will have long running times, and will be difficult to exercise in a production mode involving many sequences and sensitivity studies. They will, however, provide a capability to explore important phenomena in detail.

The codes that track the progression of fuel degradation provide input to the fission-product transport models that predict the airborne source term in the containment potentially available for release. In addition, they are used to estimate the loads on the containment. The response of the containment to these loads is also an area being studied by the NRC and EPRI. Research programs are currently investigating a number of possible containment-failure modes including localized failures at penetration seals, valve failures, overheating of electrical penetrations, and gross structural failure under quasi-steady state pressurization. The largest of these programs involves scale-model tests of steel and reinforced concrete containments. These programs will support the development and validation of models for predicting the magnitude of leakage associated with the pressure and temperature conditions in the containment.

In the more recent PRAs sponsored directly by utilities (e.g. Zion, Indian Point, Limerick, Oconee, Big Rock Point, Midland, Seabrook and the four IDCOR study plants) advances have been made in developing and quantifying containment event trees.

The containment event trees employed in the RSS constituted a delineation of different containment-failure causes such as steam explosion failures or failure caused by hydrogen burn. The probabilities assigned to these failure causes represented an integral judgment of their likelihood of occurrence for key accident sequences. Since the RSS, containment event

trees that explicitly address the underlying phenomena contributing to containment failure have evolved, such that possible combined effects as well as mutually exclusive effects can be considered. For example, both hydrogen burning and a steam pressure spike at vessel failure could, under certain circumstances, contribute to early containment failure because of overpressurization. In other cases the steam spike could render the containment atmosphere inert and prevent hydrogen burning.

Such advances allow the judging of branching probabilities on the containment event tree at a level where individual phenomena are addressed and where dependencies are explicitly considered. Radionuclide transport and release phenomena are only beginning to be considered on containment event trees.

B.4.1.2 Analysis of Fission-Product Release and Transport -- A significant effort was undertaken in the RSS to ensure that fission-product behavior was treated consistently with the existing data base and level of understanding of accident processes. The analysis was intended to be realistic. The extent to which the predicted releases of fission products were not realistic (and in general were conservatively overestimated in the RSS) resulted from limits in the existing ability to model the phenomena and by the method used to characterize the grouping of sequences in release categories by bounding the release fractions. In the RSS, fractional releases of fission products were developed for four release periods: gap, melt, core-concrete interaction (vaporization), and steam explosion (oxidation). Retention of fission products on surfaces of the reactor coolant system was not analyzed in the RSS. At the time, iodine was generally believed to be transported as I_2, which was not expected to be greatly attenuated in the reactor coolant system.

The modeling of iodine behavior in the CORRAL code (developed for use in the RSS) is largely empirical, based on the behavior of elemental iodine (I_2) in the Containment Systems Experiments (CSE). A simple aerosol model in CORRAL, which was used to predict the behavior of the less volatile fission products, was also based on the CSE tests. This model is quite primitive in comparison with existing aerosol transport codes. Recently, the CORRAL models have been shown to underpredict the removal of aerosols from the containment atmosphere in accident sequences in which the containment safety features are inoperable.* Credit for fission-product retention in buildings outside of the containment was provided only for the V sequence in the RSS PWR and some containment isolation failure sequences in the RSS BWR.

Following the RSS, the NRC undertook several research programs to improve the ability to model fission-product release and transport in severe accidents. Fuel heatup and release experiments were performed on actual irradiated fuel segments at ORNL. These experiments complemented experiments performed with simulant materials in the SASCHA facility in West Germany. The initial version of the TRAP code* was also written in this time period to predict the retention of vapors and aerosols in the reactor coolant system.

Shortly after the TMI-2 accident, questions arose about the magnitude of the possible conservatism in the RSS fission product source terms and the lack of realism in the source terms prescribed in 10 CFR 100 as a basis for regulation. In 1981, the NRC published an evaluation of the "Technical Bases for Estimating Fission Product Behavior During LWR Accidents" (NUREG-0772). As a result of deficiencies identified in that review, several new research programs have been undertaken and existing programs augmented. The temperature range of the ORNL release tests has been extended and new release tests with simulants have been initiated. Basic data have

been collected by SNL on the high temperature properties of fission-product species (e.g., CsI, CsOH, Te) and the reaction rates of these species with reactor-coolant-system surfaces. The chemistry of iodine-water systems has been extensively explored. Integral experiments for the validation of primary-system transport codes are also proceeding on an intermediate scale at ORNL and on a large scale at the Marviken facility in Sweden. Similarly, validation experiments for containment transport models have been performed at the NSPP facility and are being performed at Battelle-Institute in West Germany.*

Additional model development is also in progress. The FAST-GRASS code is being extended to provide a mechanistic prediction of the release of both volatile fission products and noble gases from overheated fuel. The VANESA code has code has been developed to describe release during core-concrete attack. TRAP-MELT has been upgraded in its ability to model aerosol agglomeration and reactions between vapor species and surfaces. The MATADOR computer code has been written as a replacement for the CORRAL-2 code for PRAs. The CONTAIN (with MAEROS routine),* TRAP-CONT,* and NAUA-4* codes have been developed to perform more detailed analyses of fission-product transport in the containment building.

Thus, a whole new arsenal of analysis capability is under development. To examine the impact of the advanced methods of analysis on the predicted release of fission products to the environment in severe accidents, the NRC has undertaken the Source Term Reassessment Study, which includes specific plant analyses. Analyses of sequences in five different plant designs will be completed by the end of 1983.

B.4.2 Uncertainties

Uncertainties in the analysis of in-plant consequences in a PRA can be subdivided into four areas:

1. the accuracy of the methods of analysis;
2. the data required by the analytical models;
3. the characterization of sequences; and
4. the estimation of branching probabilities.

The accuracy of methods for analyzing the source term (release of fission products to the environment) and the adequacy of their supporting data base are questions receiving considerable attention at the NRC. A separate program office has been established to oversee the development and application of improved methods and to modify regulations as appropriate.

In considering the uncertainty in the release of fission products from the fuel some differentiation should be made between the RSS models and current models. The fixed fractional releases during the gap and melt-release phases in the RSS approach do not account for the differences in the timing of release of different elements, which can have an important impact on their subsequent retention in the reactor coolant system and containment. The timing of release of material and quantities of inert materials released during the attack on concrete, differ substantially from the RSS analyses.

Significant gaps in knowledge exist in the current level of understanding of release phenomena as well. Substantial advances have been made in the understanding of the chemical forms of fission products. There is general agreement that, during transport in the reactor coolant system, iodine is transported primarily in the form of CsI or HI. Uncertainty still exists as to the chemical forms of many of the fission products in the fuel and the mechanisms by which they are released. Only limited aspects of the release of fission products from fuel are treated mechanistically. Currently, empirical correlations for the release rates of fission products as a function of temperature provide the best means for estimating the release from fuel. These correlations are

based, however, on small-scale and simulant experiments. In general, changes in surface-to-volume ratios during melting and local chemical reactions are not taken into account. Evidence exists that enhanced release of fission products occurs as fuel liquefies during heatup or fractures during quenching (e.g., the TMI-2 accident and PBF test SFD 1-0).* The START code models these conditions but only inadequate experimental data is available to support use of the code.

One of the greatest sources of uncertainty in predicting the release of fission products from fuel is the estimation of the time-temperature history of the fuel. Release of fission products is very sensitive to time at temperature. The MARCH code treats the melting of fuel, clad, and core internal structures very simplistically. Although more sophisticated models are under development, they have a very limited experimental basis.

The state of knowledge of fission-product transport in the reactor coolant system is changing very rapidly. The TRAP-MELT code models the transport and deposition of three important vapor species, CsI, CsOH, and Te, as well as aerosols. Basic data on the properties and deposition velocities of vapors are being provided by Sandia. Validation experiments are underway at ORNL and in the Marviken facility. Integral information on reactor-coolant-system deposition will also be obtained in the PBF experiments. However, many potentially important phenomena are not modeled in TRAP-MELT or are modeled simplistically. For example, aerosol nucleation, chemical transformations, nuclear transformations, and chemical reactions with surfaces are not currently modeled. In addition, the prediction of fission-product behavior is sensitive to the thermal-hydraulic conditions in the reactor coolant system, which are not well understood. Over the past year, the MERGE code has been developed as an extension of the MARCH code, providing an improved thermal-hydraulic model

of the reactor coolant sytem specifically for use with TRAP-MELT. That is, MERGE performs multiple-volume thermal-hydraulic calculations for severe accident conditions, using compartments consistent with TRAP-MELT specifications. However, because few experimental data are now available on, for example, flow paths and conditions in the upper internals of the reactor vessel, significant uncertainties remain despite the improved modeling capability of MERGE. In summary, the development and application of methods for predicting retention in the reactor coolant system are too formative to permit a good appreciation of their accuracy.

The status of modeling of fission-product transport in the containment is more advanced. As discussed earlier, the CORRAL code used in the RSS is not representative of the current state of the art. Although the airborne concentrations of aerosols predicted by CORRAL can differ by more than an order of magnitude from the results of the more mechanistic codes, estimates of the integral quantities released to the environment have typically agreed within a factor of two. Among the major sources of uncertainty that affect containment transport are:

1. the amount of steam condensation on aerosols;
2. the magnitude of diffusiophoresis;
3. chemical changes in hydrogen burns;
4. partitioning of fission-product vapors between air and water;
5. formation of organic iodides;
6. nuclear transmutation;
7. aerosol scrubbing in saturated water pools;
8. aerosol scrubbing in ice beds;
9. the effects of multiple compartments in containment.

Only uncertainties in the last two processes are believed to have the potential to vary the estimated consequences of accident sequences by an order of magnitude.

As described above, several aspects of fission-product behavior are not well understood or modeled, which can affect the magnitude of the source term. The most dramatic influence on the source term, however, is determined by whether or not the containment fails or by the timing of containment failure. Loads on the containment can be produced by steam spikes resulting from molten fuel-coolant interaction, hydrogen combusion, noncondensable gas generation, thermal radiation, missile generation, and steam explosions. Predicting each of these phenomena requires a detailed understanding of the progress of core meltdown accidents. How the containment will respond to a given load is also quite uncertain. The ultimate strength of the shell of a containment structure can be estimated with finite element structural codes. At some pressure less than this ultimate, however, the containment will undergo substantial leakage. Because of the uncertainties associated with the loading and response of containments, the NRC has established working groups in both of these areas to assist in evaluating the potential for containment failure.

In summary, considerable effort is being undertaken by the NRC to develop and validate methods of analysis for in-plant consequences. At the current time, however, these methods are still being developed, their sensitivities are largely unexplored, and the extent of validation is extremely limited. Note also that the cost of greatly narrowing some of the uncertainties in these methods may be prohibitive. As a minimum, the ongoing research can be expected to characterize accident source term uncertainties better and in some important areas may result in reducing the conservatisms in the analyses employed in PRAs to date.

The third area of uncertainty is the characterization of sequences. In a typical PRA, only a few sequences are identified for detailed analysis. These sequences must be considered representative of a range of related sequences involving variations in ESF performance and operator response. When the sequence is analyzed, a single set of initial and boundary conditions is selected. Little investigation has been devoted to date to the variation in consequences that can occur in the analysis of a sequence by making different assumptions regarding level of ESF performance and operator response.

The probabilistic quantification of in-plant consequences in the RSS relied on expert judgment to assess the probabilities of the different containment failure mechanisms. A probabilistic methodology for analyzing in-plant consequences has not been developed because in-plant consequence phenomena are, in general, deterministic rather than random statistical processes. Recent studies have begun to quantify explicitly the uncertainties in these processes to improve the basis for determining containment event tree branching probabilities. Considerable judgment must still be used in quantifying the branching probabilities, however, because they represent an evaluation of the state of knowledge rather than a measure of the frequency of a statistical process.

B.4.3 Potential for Improvements

In-plant consequence analysis for severe accidents is in a period of rapid transition. Indeed, developments are occurring so rapidly that, for a PRA being undertaken today, a set of computer codes is difficult to recommend. A key issue is the depth of analysis of fission-product behavior that will be required in PRAs. The decision will depend on the extent to which uncertainties are reduced through the use of complex models and the degree of potential biases associated with the simpler models.

Future PRAs, therefore are expected to employ a set of codes which have similar scope to the codes used in the Source Term Reassessment Study (i.e., MARCH2, MERGE, CORSOR, VANESA, TRAP, NAUA). These capabilities would include a method for relating fission-product release from the fuel to the local condition of the fuel (as opposed to the RSS approach) and treatment of transport and deposition in the reactor coolant system. Use of the CORRAL2 code for containment analysis in a PRA cannot currently be justified. The MATADOR code was written to replace CORRAL2. This code has had limited use to date, however, and the relative merits/demerits of MATADOR versus CONTAIN (MAEROS) and NAUA are not clear. Some of the basic modeling simplifications in MARCH2 (e.g., a single fuel melting temperature) limit its ability to support a detailed mechanistic treatment of fission-product behavior. Thus, if this is the direction of the future, MARCH2 will likely be either replaced or modified.

NRC is funding development of the MELCOR code system to perform both in-plant and ex-plant consequence analyses in a PRA. The level of detail of MELCOR analyses is still under consideration. The first version of MELCOR is scheduled for release in September 1984. It would, therefore, not be possible to use this code in a PRA prior to 1985.

The MAAP (BWR and PWR versions) and RETAIN codes have recently been developed by IDCOR for use in severe accident analysis. At the time of this writing, little information had been made publicly available on these codes. The availability of these codes will apparently also be somewhat restricted.

Ability to treat BWR plant features has lagged the ability to analyze PWRs. Recently, special consideration has been given to improving the ability to model the physical processes of severe accidents in BWRs in the ORNL Severe Accident Sequence Analysis program and in the IDCOR model development effort.

Model development for analyzing the effectiveness of suppression pools in fission-product scrubbing has been supported by the NRC and EPRI. Validation experiments are also in progress.

Advances can be anticipated in the analysis of containment failure in terms of failure pressure, failure location, and leakage rates upon failure. More realistic and plant-specific containment failure analyses will be required to evaluate the source terms for sequences involving late containment failures.

One of the most important advances to be fostered and anticipated over the next few years is the treatment of analysis uncertainties as an integral part of the accident source term analysis and the probabilistic propagation of the uncertainties. An objective of the MELCOR development effort is to provide this capability.

Uncertainties in the consequences of severe accidents are often identified as a limitation to the use of PRA in regulatory decisionmaking. Most of these uncertainties relate to the level of understanding of severe accident processes, rather than being an inherent problem of PRA. By using sensitivity studies and uncertainty analyses, the impact of these uncertainties on the results of a PRA can be examined explicitly. Often, bounding analyses can be performed. Because the effect of the uncertainties in accident source terms can be illustrated, they do not necessarily impede decisionmaking. The problem arises when the uncertainties are so large that the preferred decision is not apparent. Additional research and model development will reduce some but not all of these uncertainties. Major advances are currently being made in the understanding of processes controlling fission-product release and transport. Processes that are closely coupled to the progress of extensive fuel damage, such as the release of

the less volatile fission products from fuel, or the generation of hydrogen during core slumping, will always have a large uncertainty because of the difficulties associated with experimental validation.

B.5 LWR Offsite Consequence Analysis

B.5.1 Background and State of the Art

Offsite consequence analyses attempt to predict the frequency distribution of possible consequences for potential accidents at nuclear power plants. Accident consequences can include early fatalities and injuries, latent cancer fatalities, genetic effects, land contamination, and economic impacts. Chapter 9 of the PRA Procedures Guide (NUREG/CR-2300) contains a recent discussion of the important elements of offsite consequence analysis.

The first comprehensive assessment of the consequences from potential accidents at nuclear power plants was performed in the RSS. The RSS, published in 1975, examined the aggregate risk posed by commercial nuclear power plants in the United States. As part of the study, a computer model (CRAC, for Calculation of Reactor Accident Consequences) was developed to predict the offsite consequences of releases of radioactive material to the atmosphere for typical (i.e., "generic") sites (NUREG-0340). CRAC describes the atmospheric transport, dispersion, and deposition of released radioactive materials, and predicts the resulting interaction with and influence on the environment and man. The computation steps in the model are shown schematically in Figure B-3. Other computer models developed for offsite consequence analysis consist of these same basic steps. Given a description of the release of radioactive material (source term) and a range of possible weather conditions as input, submodels for atmospheric transport and dispersion, radiation dosimetry, population location and behavior, offsite protective measures, radiological health

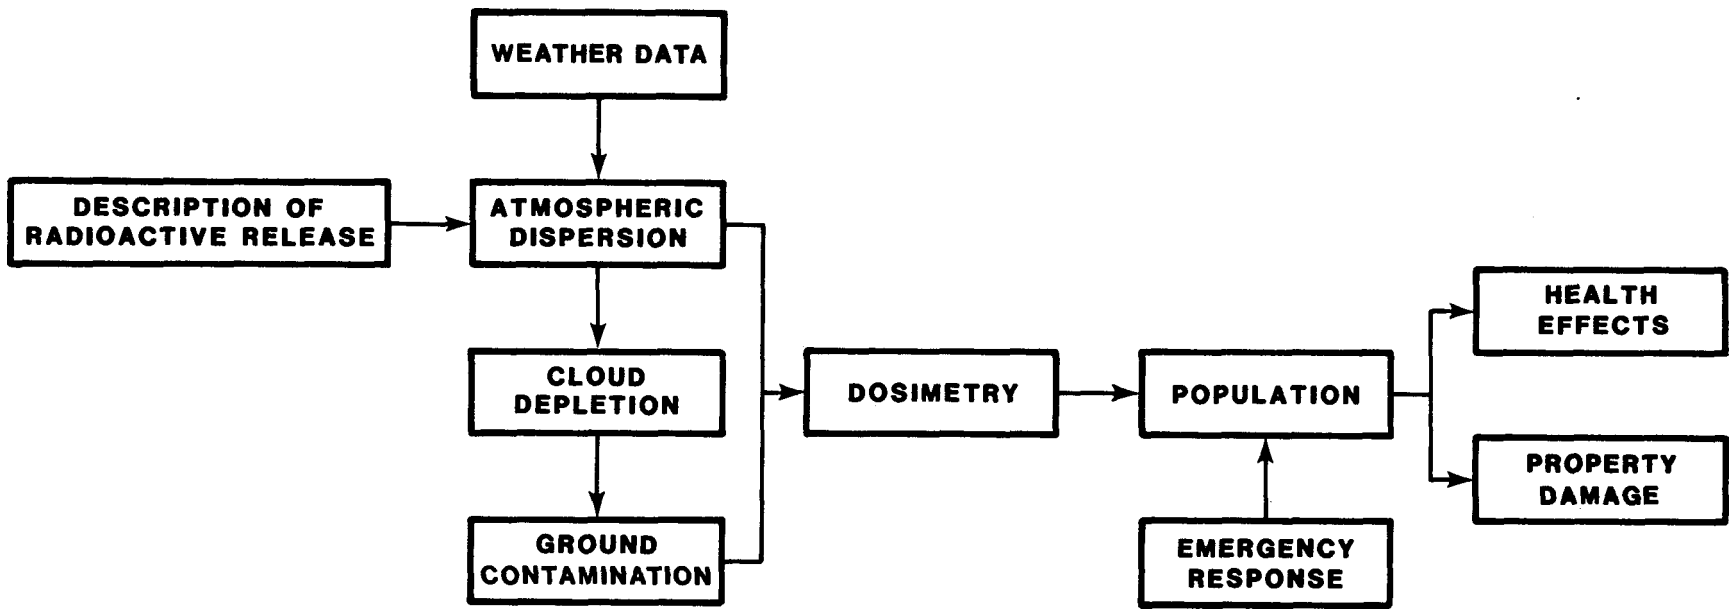

Figure B-3. Schematic Outline of Reactor Safety Study Offsite Model, CRAC

health, and property damage are used in turn to estimate the resulting frequency distribution of potential consequences. The distributions of results obtained are normally displayed in the form of complementary cumulative distribution functions (CCDFs) and expected (mean) values.

Since the completion of the RSS, several improvements have been made in the field of offsite consequence analysis (Aldrich et al.). Two improved versions of CRAC have been developed and are currently in use in the United States: CRAC2* and CRACIT.* CRAC2, developed under an NRC-sponsored research program, incorporates significant improvements in the areas of weather-sequence sampling and emergency response. CRACIT (CRAC Including Trajectories), developed by Pickard, Lowe, and Garrick, Inc., includes modifications to the atmospheric dispersion and evacuation models that permit some of the unique features of a specific site (e.g., terrain, evacuation routes) to be considered. CRAC2 is widely used by utilities, National Laboratories, and NRC; CRACIT, which is proprietary, was used in the Zion and Indian Point PRAs. In addition to the U.S. models, offsite consequence models have been developed for use in risk evaluations in other countries; examples include the Sizewell PWR Inquiry in the United Kingdom and the German Reactor Safety Study.

To understand better the influence of different consequence modeling techniques, the International Comparison Study of Reactor Accident Consequence Models was organized in 1981 under the auspices of the Organization for Economic Cooperation and Development/Nuclear Energy Agency's Committee on the Safety of Nuclear Installations (CSNI). Approximately 30 organizations, representing 16 countries and the Commission of the European Communities, participated to some degree in the comparison study. As part of the study, a series of standard problems was specified to allow a step-by-step comparison of individual models as well as consequence and risk

estimates. The study showed that a great many modeling techniques and assumptions are being used to estimate the consequences of potential reactor accidents. The estimates of consequences made by the participants, however, were generally in fairly close agreement. In most cases where significant differences did occur, they could be explained readily by differences in modeling techniques or assumptions. A detailed comparison and evaluation of the results of the study, along with important observations and conclusions, will be presented in the forthcoming "Summary Report" (OECD, to be published).

The potential consequences resulting from accidental releases of radioactive material to water pathways have not been examined with the same degree of detail as those resulting from releases to the atmosphere. Risks from the atmospheric pathway are generally considered to be dominant for two interrelated reasons. First, the initial time scale on which radioactive contaminants could reach the human population would probably be significantly shorter for the atmospheric pathway, and therefore hydrospheric transport of contaminants would allow for greater decay of radionuclides. Second, initial atmospheric exposure would usually be involuntary, whereas in most cases exposure to hydrospheric contamination could be largely avoided by the implementation of appropriate protective measures.

Several generic studies of the potential effects of radioactive releases to water pathways have been performed. In general, these studies concluded that short-term radiation doses to individuals via liquid pathways would probably never be large enough to cause early health effects, and that for most sites the public risks posed by core-melt accidents will be dominated by releases to the atmosphere rather than by direct releases (e.g., basemat melt-through) to liquid pathways. However, sites may exist with characteristics such that risks via the liquid pathways could be important relative

to those from terrestrial pathways, and further analyses should be performed to identify those characteristics and to select appropriate models for evaluating risks at such sites. A review of previous liquid-pathway studies and a discussion of methods appropriate for performing site-specific liquid-pathway analyses are included in the PRA Procedures Guide (NUREG/CR-2300).

Potential economic consequences and risks from commercial nuclear power plants are receiving considerable attention because of their importance for cost/benefit analyses. Several recent studies have examined the economic consequences and risk from nuclear power plants (Starr and Whipple; NUREG/CR-2723; Burke). These studies have pointed out the overriding importance of potential onsite costs (e.g., cleanup or repair and replacement power) to the overall consequences of reactor accidents. For example, the TMI-2 accident resulted in minimal offsite consequences but major onsite damage to the plant. Moreover, relatively high-frequency "routine" forced-outage events have been shown to dominate the aggregate economic risks from reactor operation. Only for severe core-melt accidents will offsite costs (decontamination, land-use denial, health effects, etc.) equal or exceed the onsite costs, and then only for densely populated sites or extremely adverse weather conditions.

As already mentioned, the RSS was performed to assess the aggregate risk from commercial nuclear power plants in the U.S. Since the completion of the RSS, the capabilities of offsite consequence analysis have been extended to provide assessments of the risk posed by reactors at specific sites and to provide guidance for planning and decisionmaking. Examples of site-specific applications of offsite consequence analysis include the Limerick, Zion, and Indian Point PRAs and the recent environmental statements for Susquehanna and

Fermi. In addition to use in risk evaluation, offsite consequence analysis has been used to aid decisionmaking in several. Examples include evaluation of alternative design features (Benjamin and Harper, to be published; NUREG/CR-0165), emergency planning and response (NUREG/CR-1131; NUREG/CR-1433), reactor siting recommendations (NUREG/CR-2239) and determinations of risk acceptability.

Even in the presence of large uncertainties, which are discussed in Section B.5.2, offsite consequence analysis can provide (and has provided) several useful insights and perspectives on severe reactor accidents. For example, analyses have shown that the extremely low-probability, high-consequence events ("tails") do not contribute significantly to the mean (or expected) consequences (their significance for decisionmaking, therefore, is not clear). However, clearly, if large releases of radioactive material are possible, the potential health effects could be extremely severe and economic damage could exceed tens or even hundreds of billions of dollars. A summary of some of the important insights for nuclear reactor safety gained from these applications of offsite consequence analysis is provided in Chapter 5 of the main report.

Currently, the capabilities for performing offsite consequence analyses are more mature than those for the evaluation of accident progression, containment behavior, and source terms; and a fair evaluation is that offsite consequence analysis is not the weakest link in PRA. In the development of offsite consequence models, a large pool of supporting data is available from which to draw. Such data include air pollution studies (atmospheric transport, dispersion, and deposition), nuclear weapon fallout studies (behavior of radionuclides in the environment), radiation therapy and Hiroshima and Nagasaki data (radiation, health effects), and so forth. In general, relatively simple empirical relationships (i.e., "fits") can

be derived from these data to model phenomena that are extremely complex in nature. In contrast, the phenomena associated with the progression of severe reactor accidents are much less well understood. Nevertheless, offsite consequence analysis is not without significant uncertainties. Moreover, because of the potential applications of offsite consequence analysis to decisionmaking, especially in cost-benefit analyses, these uncertainties should be minimized to the extent possible and quantified.

B.5.2 Uncertainties in Offsite Consequence Models

Uncertainties in offsite consequence predictions stem principally from two types of uncertainties: modeling uncertainty and uncertainty in the input data to the models. Modeling uncertainty arises from (1) an incomplete understanding of the phenomena involved in the transport of released radionuclides to man and of the health, environmental, and economic effects that result and (2) simplifications made in the modeling process to reduce costs, complexity, and requirements for input data. Uncertainty in the data used as input to the models arises from the quality or appropriateness of the data and from statistical fluctuations. In addition to uncertainty in the models and data, the weather conditions following a release can have a very large impact on predicted consequences. Uncertainty in the meteorological conditions is usually addressed by treating weather as a stochastic parameter.

A comprehensive assessment of the uncertainties in offsite consequence predictions has not been performed yet. A large body of parametric (or sensitivity) analyses, however, does exist, that is, studies in which consequences are calculated for a range of plausible values of a key parameter or model. The PRA Procedures Guide (NUREG/CR-2300) made a tentative listing of the relative contribution to total uncertainty of the major parameters and models in an offsite consequence

analysis. The contributions of the factors to uncertainty were ranked as "major," moderate," or "low." This list, which was based on past parametric/sensitivity studies and the subjective judgment of the authors, contains 51 factors, 14 of which were deemed to be major contributors to uncertainty in at least one type of consequence. Among the "major" contributors are: (1) the magnitude of the source term, which strongly influences all consequences, (2) the form and effectiveness of emergency response, which can make a large difference in predicted early health effects, (3) the dry deposition rate of particulate matter from the plume, which impacts early health effects and the distances to which land-use restrictions or crop impoundment may be required, (4) the modeling of wet deposition caused by rainfall, which impacts the low-probability, high-consequence end (tails) of the distributions of all consequences, and (5) the dose-response relationships for somatic and genetic effects. In addition, questions have been raised over the importance of modeling atmospheric transport and dispersion with a "straightline" versus a trajectory model, particularly for sites with significant topographic features. Though a complete evaluation of the importance of trajectory models on predicted risk has not been performed, the results of the CSNI International Comparison Study indicate that the impact of trajectory modeling is less than that of other major modeling assumptions. Efforts to quantify better the uncertainties in the estimates of offsite consequences are currently underway and are described in Section B.5.3.

Even though a thorough examination of uncertainties in offsite consequence analyses has not been performed, the magnitude of these uncertainties may be inferred from the results of the large number of existing sensitivity studies. Major contributors to uncertainty, and the magnitude of uncertainties, depend strongly on assumptions about accident source terms and site characteristics. For estimates of the consequences

resulting from very large source terms at a highly populated site (Aldrich and Sprung, 1982), the following crude estimates of uncertainties can be made. Mean early fatalities could range from approximately a factor of ten above present "best" estimates to nearly zero. This broad range is in large part due to uncertainty in the effectiveness of short-term emergency response near the plant. The uncertainty in mean predicted population dose (person-rem) is estimated to be a factor of three or four, while the uncertainty in the predicted mean number of latent cancer deaths (which is a function of population dose) is approximately a factor of ten. In general, the uncertainties are somewhat larger in the magnitude of the extremely low-probability, high-consequence portion ("tails") of predicted consequence frequency curves.

One effect which has not been considered to date in LWR offsite consequence analyses is the possibility that condensation of moisture in the released plume could result in a significant fraction of the radioactive material being deposited in the immediate vicinity of the reactor. Were this to occur, the "rainout" of radioactive material could have a dramatic influence on risk, depending on the extent and location of the enhanced deposition. The likelihood of this occurring would depend on the nature of the containment breach and on the physical characteristics of the plume (e.g., temperature, momentum, moisture content, etc.). An NRC-sponsored program is currently examining this effect.

Clearly, all consequences are sensitive to the amount of radioactive material that could be released (the source term). However, early fatalities and injuries are particularly sensitive because of the existence of dose-thresholds for these effects. If potential source terms are found to be substantially smaller (at least one order of magnitude), then the risk of early health effects would generally no longer be a principal concern. Nonetheless, the consequences of such

accidents could still be large; the nature of the risk, however, would be different. Focus would be directed on latent health effects and on the more localized problem of land contamination which is roughly proportional to the amount of long-lived radionuclides released (mainly cesium). Tradeoffs between decontaminating an area, barring its use (interdiction), and a possible increased risk of cancer would need to be considered. In the limit, releases of just the inventory of noble gases (krypton and xenon) could still result in significant offsite radiation exposures and resulting population dose.

In addition to the source term magnitude, the numbers of estimated early health effects are very sensitive to assumptions about the nature and effectiveness of potential emergency protective measures. Studies have shown that for large releases of radioactive material, prompt evacuation and sheltering are potentially effective means of reducing the numbers of early health effects.* Latent cancer fatalities are not as sensitive to emergency response assumptions because of the larger areas and longer time frame involved.

The weather at the time of the accident can have a very large impact on offsite consequences (e.g., low or high windspeed, dry or precipitating). However, the variation in weather from site to site does not appear to impact total risk appreciably because the probabilities of weather types which contribute the most to variation in consequences are not significantly different in different climates. However, total risk is very dependent on the characteristics of a site, such as population density and land use, and these considerations are important for reactor siting.

B.5.3 Potential for Improvements

In Section B.5.2, five major sources of uncertainty in offsite consequence analysis were described. The single largest contributor to the uncertainty in the offsite consequence estimates is uncertainty in the magnitude of the source term. Ongoing research in the United States and throughout the world is directed toward providing improved quantification of potential source terms for severe reactor accidents. In addition, three important aspects of offsite consequence analysis exist where significant improvements are possible: (1) the capability to consider the specific characteristics of a particular site; (2) the capability to consider more detailed descriptions of the source term such as release duration, moisture content, and particle size; and (3) the tools used to provide estimates of uncertainties. Efforts are underway in these areas as part of the NRC-sponsored MELCOR program. Specifically, improvements are being made in the atmospheric dispersion and transport model which include developing a multi-puff model which will permit the analysis of site-specific terrain and plume trajectories and provide an improved treatment of long-duration releases and precipitation modeling. Other improvements in modeling capabilities will include the incorporation of more detailed land-use characteristics, especially the differentiation of urban and rural areas, and a reevaluation of the available emergency response data, which will provide improved estimates of the risk of early health effects. Improved models for radiological health effects and potential economic impacts are also being developed. In addition, a key objective of the MELCOR offsite consequence modeling effort is to develop tools which can provide estimates of the uncertainties in the predicted consequences. Although uncertainties are likely to remain quite large, a thorough examination of their origin and magnitude will provide both a firmer basis for applying offsite consequence analysis and a better understanding of its limitations. Finally, as mentioned above, a program is currently underway to assess the

potential impact on offsite consequences of localized "rain-out" from a moist plume.

These efforts, which will be completed in about one to two years, should provide improved estimates of offsite consequences, quantitative estimates of uncertainties, and increased confidence in the results, thus expanding the usefulness of offsite consequence analysis for the decisionmaker.

B.6 Accidents from External Initiators

PRA analysts have conveniently grouped accidents resulting from various 'external initiators' into a separate category of analysis, principally because the method for treating them differs from the method for treating so-called internal initators. The external initiators** are:

Earthquakes
Internally initiated fires
Floods
High winds--tornadoes, hurricanes
Other: aircraft
barge and ship collisions
truck, train, pipeline accidents
external fires
volcanoes
turbine missiles
lightning

The unifying characteristic of all these initiators is the potential for the initiator not only to start an accident, but also to compromise simultaneously the efficacy of the

---

*"Both 'internal initiators' and 'external initiators' are misnomers, since the former category is usually taken to include accidents starting with loss of offsite electric power, while the latter usually includes internally initiated fires and floods.

safety systems designed and needed to halt or mitigate the accident.

Each of the first four initiators listed above (earthquakes, internal fires, floods, high winds) has been the subject of one or more comprehensive PRAs, leading to calculation of core-damage frequency and offsite risk. So far, none of the initiators listed in the last group has been subjected to such a comprehensive analysis, typically because simpler analysis has shown the risk from these to be acceptably small.

B.6.1 Introduction to PRA Methodology for External Events

The basic approach to the probabilistic analysis of external initiating events is similar for all such types of events, and consists of four different types of analyses which are then combined. The sequence of these analysis steps is not necessarily in the following order, but all must be performed in a full analysis.

1. The expected frequency (events per year, or, say, events per millenium) that the plants might experience the particular external event must be quantified. Considering floods, for example, the likelihood of floods of various sizes must be determined, recognizing that the very largest floods are much less likely than the somewhat smaller (though still quite large) floods.

2. The analyst must determine the effects that various external initiating events will have on specific pieces of equipment (components, systems, operator command and control functions, etc.). This includes determining the coupled likelihood of what are known as "common cause" failures, in which several systems or functions experience failure or degradation together in a correlated way.

3. The analyst must determine the effect of degraded or inoperative systems, components, or functions on the ability of the entire plant to reach a safe, stable shutdown state. This typically involves event-tree/fault-tree methods.

4. An analysis must be done on the phenomena and consequences associated with those rare accident sequences that might lead to undesirable outcomes. This part of the analysis is nearly identical to that performed for other types of (internal) initiating events.

The methodology of PRA that has been developed for the various external initiating events differs in detail, of course, from one category to another. However, all external events are analyzed using the general approach just described, except in cases where the first part of the analysis shows that the expected probability of the initiating events is so low that the overall contribution to risk would be very small. In such cases, the analysis can stop after step one without considering the other parts of the full PRA approach.

Through methodological advances and some recent highly successful applications, the PRA methods for analyzing externally initiated accidents have matured enormously since the pioneering studies of the RSS. Major engineering insights are now available, even though large uncertainties persist in the numerical results of the analyses. The methodology for seismic analysis, for example, has reached a stage where the insights gained from recent PRA studies can be applied to specific, detailed pieces of hardware (such as the fragilities of specific structures and equipment). This was certainly not true of the RSS analysis, which used a relatively primitive approach to conclude that these initiators were probably unimportant. Most of the PRAs that followed the RSS in the

first two or three years after its publication did not treat any external events at all.

This comment is not meant to imply that the methodologies for analyzing these external initiators are mature enough to be relied upon in a quantitative way; the general consensus in the PRA community is that they are not, and that the insights gained today are by-and-large qualitative in nature, even though the results are quantitative. Much active methodological development is now in progress, and the major advances of the past two or three years are sure to be only a precursor to even greater advances in the coming years. Important insights have already resulted from some of these analyses. Although some active disagreements exist among experts as to the validities of specific approximations or judgments used in these analyses, the disagreements should be viewed as a healthy sign of the great maturation process being experienced today in this field. They are not a sign of the invalidity of the present insights, only of their incompleteness or uncertainty.

The limitations occur in each of the various parts of the analysis. Thus for some external events, the likelihood of a major initiator (say, a very large earthquake or an extreme flood) is often neither known from the historical record nor reliably inferred from analysis based on extrapolations of that record. Also, the effect of some of these events on plant components, systems, and functions is in some cases not well understood. The "fragility" values used for equipment and structures are often based on incomplete data or approximate analysis. Finally, the ability to analyze well some of the correlations among failures is still limited.

For some of the categories of externally initiated accidents, the overall risk can be acceptably bounded only by analysis

of the frequency of the initiator in comparison with the calculated core-damage frequency of other accident sequences. This is typical of the initiators listed in the 'other' category above: aircraft, barges, ships, trucks, trains, pipelines, volcanoes, turbine missiles, lightning, and external fires. Indeed, these have almost never been shown to be important contributors to overall plant risk. More important, the methods used for calculating the frequency of an initiator sufficiently large to compromise the plant (e.g., a large enough external fire, or a serious aircraft or barge collision) are typically adequate to determine a very small upper bound with high confidence.

The categories where special discussion is needed are earthquakes, floods, internal fires, and high winds.

B.6.2 Earthquakes

B.6.2.1 State of the Art and Discussion of Uncertainties

Of all the various external initiators, earthquakes are the ones for which the PRA methodology is the most mature: several comprehensive seismic PRAs have been completed, a significant body of research has been completed to develop the methods and explore their sensitivity and data uncertainties, and symposia/conferences/workshops have been held to advance the state of the art. Results of a few recent seismic PRAs have been used to modify the plants studied, to affect regulatory decisionmaking, and to upgrade the quality of the methodology itself. Also, several specialized, limited studies of narrow issues have been performed using elements of the broader methodology.

Despite this relatively advanced state of maturity, the results of the recent seismic PRAs are still highly uncertain, especially in a quantitative sense: indeed, while the results are expressed in quantitative form, the key insights are

widely accepted as qualitative. Two different but similar approaches are described in the PRA Procedures Guide, (NUREG/CR-2300) which complement each other.

The state of the art of seismic hazard analysis has progressed to the point that several different groups are now carrying out such analysis, using somewhat different methods that can easily be differentiated from one another but among which choosing a 'preferable' approach is sometimes difficult. The different methods can yield quite different results. The intrinsic problem with seismic hazard analysis is that the dominant contributors to reactor risk come from earthquakes significantly larger than the earthquakes used as the design basis, and these large earthquakes have return periods that are difficult to estimate well, there being no historical record. This is especially true in the eastern U.S. but is even true in California. Thus there are major uncertainties in the likelihood estimates. These uncertainties propagate through to large uncertainties in the final results. The extrapolations to larger earthquakes, with very long recurrence intervals, present another problem, because for any given site or seismic province a maximum earthquake motion is believed to exist that can be sustained by the specific geological features--thus the extrapolations are usually cut off in one fashion or another. It turns out that the numerical results of some recent PRAs may be more sensitive to the nature of the cut-off procedure than to any other of the several assumptions!

Still another uncertainty comes from efforts to characterize the motion associated with these large earthquakes in terms of physical parameters, such as frequency-dependent acceleration, velocity, displacement; the shape of the motion in time; the energy dispersion; and so on. Sometimes the historical record can be valuable in characterizing the expected motion of future earthquakes; but unfortunately for many of the

important historical earthquakes, records are limited to known structural damage and area over which the motion was felt, and even these can be less than reliable.

The PRA requires a probabilistic description of the propagation of the earthquake motion through the rockbed and soil to the building substructure and the interaction of the motion with that substructure, with the goal of characterizing the motion felt by safety-related structures and equipment. Regarding substructure interaction, advances have been made under the auspices of the NRC's research program; and, although much remains to be accomplished, this area is now in better shape than other elements of the analysis, contributing rather less to the overall uncertainties.

The fragility of structures and equipment is the other major element of seismic PRA. Here the progress has been quite significant in recent years, with a large number of reactors now studied. The fragility analysis relies partly on a data base that is neither strong nor extensive, and partly on analytical methods designed to overcome many of the weaknesses in the data base. One roadblock to precise fragility analysis is the difficulty of characterizing the input motion adequately. Various surrogate accelerations have been proposed and used parametrically, each with limitations that offset some of the advantages. Furthermore, 'failure' in the context of seismic PRA means failure to perform a safety function, not necessarily structural collapse or physical distortion.

Another issue is the statistical nature of the very problem addressed: obviously not all 'identical' components (say, identical valves) will behave identically, even in a statistical sense. Yet the issue of most concern is that identical or similarly configured components might all fail together in an earthquake, defeating the redundancy of safety systems.

This difficulty has not been adequately solved methodologically. Of course, the problem can be bounded by assuming complete independence and complete dependence as the extreme cases, and sometimes this approach is adequate.

While the state of the art of fragility analysis is becoming more advanced as each new reactor is studied, there are still not many practitioners of this part of the PRA art. Also, while the number of reactors studied is increasing (it is now well over a dozen), many of the analyses are still underway at this time and have not been either published or reviewed. The rapid increase in the number of applications will surely produce increased confidence in the techniques in the near future.

The final element of seismic PRA analysis is linking the failures of structures and equipment into a systems analysis of the plant. Here the event trees developed in the other parts of a comprehensive PRA usually provide the basis for the analysis. Unfortunately, this part of the problem is not as easy to accomplish in practice as an analogous internal-events analysis, because a large earthquake can cause numerous failures that compromise redundant systems. The analyses completed to date, however, have found that the seismic risks are often dominated by accident sequences involving only a few important seismic failure modes, often failures of structures. This simplifies the analysis somewhat, but only to the extent that the simplifying assumption can be made and defended that failure of the single or few components or structures actually brings the plant to core melt. This has often been assumed with the full knowledge that it is conservative, sometimes highly so. The human factors aspects have also not been treated with the depth ultimately desired, for example, recovery by the operators is not well treated in the studies to date. For all of these reasons, the analyses tend

to be, in the end, conservative in character rather than realistic in a numerical sense.

Overall, a consensus prevails that the uncertainties in bottom-line numbers (core melt frequency, offsite risks) remain quite large for seismic PRAs. Error factors of 10 to 30 (implying spreads of factors of about 100 to 1000 in the 5%-95% confidence interval) might reasonably characterize this aspect today. Of course, these large spreads on the numerical conclusions for bottom-line numbers do not compromise the highly significant engineering insights obtained using these PRA methods, and many of these insights are new or are unavailable with traditional methods. In particular, the analyses generate system vulnerabilities and common-cause dependencies that can point the decisionmaker to the need for careful consideration using other approaches and other decision criteria.

B.6.2.2 Methodological Insights (Earthquakes)

Several important insights have emerged from the completed PRAs that include seismic analysis:

1. With a few exceptions, the PRAs completed to date indicate that most major equipment items are calculated to be less fragile than structures. However, equipment failure in the PRA context means failure to perform a safety function, which is hard to analyze well. Thus the problems associated with equipment fragility analysis are sometimes great, especially concerning dependencies among failures.

2. The largest uncertainty in the numerical quantification of risk arises from the difficulty that the hazard analysis encounters in quantifying the probability of the very large earthquakes that seem to dominate the calculated risk. However, the engineering insights are not highly

dependent on the actual numerical values of these probabilities.

3. In contrast, the largest uncertainty in the engineering insights results from the difficulty in quantifying fragilities, especially considering dependencies and correlations among failures, human factors issues, and the problem of understanding how equipment within or dependent on a structure will fail when the structure fails.

B.6.2.3 Potential for Improvements and Areas That Will Remain Uncertain (Earthquakes)

1. Major improvement in the generally applicable insights from seismic PRAs may occur as more and more of these PRA studies are performed and reviewed. Up to now, the number of studies completed has been too small and the idiosyncratic issues that have dominated the risk have been too plant-specific to allow much generalization.

2. In contrast, major improvement in the quantification of seismic sequence probabilities, from better understanding of the probabilities of the large earthquakes that seem to dominate seismic risk, will not progress very fast, because the intrinsic limitations (from the shortness of the historical record) are not likely to be dramatically overcome soon by analytical advances. However, progress may occur in understanding the "cut-off" issue, which is the issue of how to cut off the extrapolation of earthquake motion at the high end to account for the limitations in the ground's ability to sustain the very largest motions.

3. Integrated systems analyses that incorporate the seismic PRA work with the broader 'internal initiating events' PRA work are important and have been accomplished in the

most recent applications. This will significantly improve insights.

4. The role of operators in mitigating, or contributing to, the risks from seismic-initiated sequences is still not well understood. Some progress is expected in this area, but this issue is likely to remain a difficulty for some time to come.

5. Improvements will probably occur in the ability of analysts to choose and apply various surrogate accelerations as a means or parameterizing the ground motion for the fragility analysis.

6. Soil-structure interactions are already in reasonably good shape when considered in the context of the other larger uncertainties in seismic PRA. The ability to model and quantify these interactions will improve further in coming years as more and more studies examine different sites with different physical conditions and configurations.

B.6.3 Internal Fires

B.6.3.1 State of the Art and Discussion of Uncertainties

Only recently has the probabilistic analysis of internal fires become an accepted part of a full-scale PRA: only a very few such probabilistic studies of fires have been accomplished, and the methodology used is by no means mature. Although the applications have been few so far, the literature covering both methodological improvements and applications is growing. The state of the art has now reached the stage that the methodology can be used as part of a plant-specific PRA. Also, important new research is now underway.

Those few recent applications of the methodology in the context of full-scale PRA (e.g., Big Rock, Zion, IP) have proven to be very useful in several ways. First and most important, they have demonstrated that the methodology can provide important engineering insights about plant vulnerability to fires. Secondly, they have revealed the extent to which these insights can be relied on quantitatively, the problems with application to specific plants, and areas where the qualitative lessons learned may be generic despite the continuing uncertainties in the numerical results. Finally, these applications have been important in guiding future research and methodological development.

While several different approaches are evident in the literature, they share a common framework, as discussed in the PRA Procedures Guide (NUREG/CR-2300). Analysis begins with the identification of critical areas of vulnerability, then calculates the frequency with which fires might begin in each area, and follows with analysis of the extent to which critical safety functions and equipment are disabled by the fire, attempting to account for possible detection and suppression. Finally, the disabled equipment and functions are analyzed in a systems sense using event-tree/fault-tree methodology similar to that used in other parts of the larger PRA.

The initial phase of probabilistic fire analysis is the identification of critical areas. The criterion in this step is whether a fire could compromise important safety equipment; and in practice this criterion is narrowed to emphasize areas where multiple equipment could be compromised, in particular several trains of redundant equipment to perform the same safety function.

Identification often begins with the fire zones delineated in the more classical regulatory analysis, supplemented or modified by a walk-through with attention to questions such as

potential for the cross-zone spread of fire and the likelihood that transient fuels might supplement fuels always present in a zone. While this part of the analysis remains more an art than a science, the general consensus is that it is reasonably mature in the sense that uncertainties introduced in the ultimate PRA results from this aspect are thought to be smaller than uncertainties from other aspects. Barrier identification feeds the fault tree analysis by identifying candidate physical areas of concern for fire, or fire coupled with relatively high probability random failures.

The next phase of fire PRA is the determination of the frequency of fire initiation for each critical zone. The frequency can sometimes differ as a function of location within the large zone if fuel loading conditions, cross-zone spreading potential, or other idiosyncracies require that level of detail. While a historical data base exists about fires initiated in various areas in the past, the data base is not fully adequate as an empirical basis for working out the desired initiation frequency. The analyst must bring to bear location-specific information from his walk-through and other experience, so this part of the analysis has important numerical uncertainties. Despite these numerical uncertainties, the ability of a skilled analyst to rank the important potential fire-initiating locations is generally quite good, and it is now accepted that a good analyst will usually not overlook important zones.

The more difficult aspects of the probabilistic analysis occur in the next phase of study, which is the likelihood of disabling equipment given fire initiation. Here the whole arsenal of methods used in fire analysis is brought to bear on the problem, and the process may be described in terms of four tasks:

1. analysis of fire growth and spread,

2. analysis of detection/suppression of effectiveness,

3. assessment of component 'fragility' to fire and combustion products, and

4. calculation of probability estimates (distributions) for fault tree quantification. Overall systems analysis then proceeds through the containment-challenge and consequence analysis steps.

Detection and suppression (manual and automatic) should be analyzed in conjunction with fire spread and growth as competing processes. Work to improve the methodology in this area is now underway in two areas: first, incorporating detection and suppression into the computer-based models used; and second, development of an analytical model for sprinkler effectiveness, which will provide the time of sprinkler actuation for a given fire.

This analysis problem is compounded by uncertainties concerning modeling of detection and suppressing actual fuel availability (amount and character of transient fuels, etc.), the stochastic nature of fire growth over time, the size of the affected secondary zone where hot gases can cause equipment failures or induce secondary fires or at least secondary equipment failures, and issues of access for firefighting. Several important models have been developed over the years to assist the analyst in calculating the likely progression of the phenomena, but in even the best cases the uncertainties remain large in a numerical sense.

These insights from fire PRA analysis have already proven useful in a few PRA studies in guiding actions to upgrade the fire-safety of certain installations. Whether the analysis is complete in the sense that it is capable of identifying

all of the main vulnerabilities is still an issue. The available models in even the most advanced cases are only approximate in character, and are not capable of accurately modeling the spread of fire in, say, a compartment full of crowded objects in a unique configuration. Also, analysis of failure modes for components exposed to the whole spectrum of combustion products needs more methodological development and more test data. Finally, additional methodology is required for treating intercompartmental spread of fire and combustion products. Even more important, a firm perception does not yet exist on the amount of uncertainty introduced by these incomplete capabilities. Research on this question, along with studies of how well the analyst can be expected to perform in the detection-suppression arena, may contribute much to answering the question of the achievable accuracy from these probabilistic analyses.

The analysis of systems effects from fires involves the coupling of the fire studies with the event trees used in the more traditional PRA analysis of other internal initiators. It is important to realize that simply adding the fire-induced vulnerabilities to existing event trees is not adequate: the time sequences can be different for fires, the dependence on common-cause failures can be subtle, and human intervention in fire suppression can differ sufficiently for fires that the existing event trees cannot handle the issue. Event trees for the fire sequences must be drawn in an integrated way taking into account the fire issues in parallel with the other initiators. In a proper analysis, the fire vulnerabilities would be integrated into the overall fault trees to allow a comprehensive treatment of dependent failures including secondarily induced failures. At the present level of maturity of fire PRA, this is only partly accomplished. Also, because human intervention cannot be analyzed as well for

fires as would be desirable, approximations or bounding calculations are required to determine the sensitivity of the final results.

The NRC has recently begun an integrated program of methodological development in this area, along with a program of applications intended to gain insights to guide the ongoing research. At the present time, the methods are already fully capable of identifying many important types of vulnerabilities in a qualitative sense, including a ranking of their relative importance.

It is still too early to give a useful judgment as to the achievable accuracy of the results of fire PRA with regard to core-melt frequency or offsite consequences. The uncertainties are quite large, at least as large as those for the internal events analysis. Of course, the engineering insights obtained from the few PRAs performed have been very useful already and are in no way invalidated by the large uncertainties in the bottom-line numbers. Research now underway will probably narrow these large uncertainties somewhat but it is now too soon to know how much fire PRA will improve the usefulness of bottom-line values.

B.6.3.2 Methodological Insights (Fires)

The PRAs performed to date that have incorporated fires have contributed several insights:

1. The nature of the fire phenomena is such that the conditional probability of containment failure, given a fire leading to core degradation, is likely to be significantly higher than the conditional containment-failure probability from other reactor accidents. This insight implies that great care must be taken to consider dependencies in the containment analysis.

2. The fire PRAs have generally revealed that the key vulnerabilities are located in places where multiple safety systems are collocated, or where their controls or instrumentation and support systems are collectively vulnerable. Whether this 'finding' is the result of the PRA analyses, or simply the outcome of analyses based on this 'finding' as a postulate, is not clear...probably some of each.

3. The numerical results of these analyses are greatly influenced by the modeling of detection and suppression, human as well as automatic; and the analyses are difficult in this area. Research is now addressing this problem.

4. The largest uncertainty in the numerical quantification of fire-related risk seems to arise from the difficulty with quantifying the likelihood that a fire, once initiated, will disable critical equipment.

5. In contrast, the largest uncertainty in the engineering insights arises from the question as to whether the analysis might have overlooked entirely some critical fire zones, the issue of how combustion products can induce failures, and the question as to whether the human intervention in detection and suppression and in coping with the accident sequence once initiated has been modeled correctly.

6. The need for knowing the precise layout, including questions of amount of transient fuel loading, is such an important determinant of the numerical results that a fire PRA attempted before a plant is actually designed in detail would probably not be very useful.

B.6.3.3 Potential for Improvements and Areas That Will Remain Uncertain (Fires)

1. Major improvements in the usefulness of fire PRA results will occur simply from the continuing application of these techniques to more and more plants. This is probably one of the most important modes for progress in this field.

2. Improvements in the ability to model secondary fire growth will occur soon as a result of development and application of advanced models under current support of the NRC research program.

3. Improvements in modeling of fire growth and spread, including comparison of different modeling approaches, will provide insights into the strenghts and limitations of the various modeling schemes.

4. The issue of the importance of intercompartment spreading, of both fires themselves and combusion products and hot gases, is an area where work now underway may yield important insights. However, these questions are highly configuration-dependent in actual application, and some limit probably exists as to how well this part of the overall analysis can be done. The area of boundary penetration or failure, barrier-violating pathways such as ducts and drains, and isolation devices requires careful study.

5. In the foreseeable future, the uncertainty in the PRA results will not be dominated by the ability to quantify fire initiation, but by the difficulty in working out the vulnerability of specific equipment given a type of fire, and the problem of quantifying the effects of human intervention in fire-suppression and accident-sequence mitigation. Particularly significant is the analytical difficulty in coupling suppression with fire growth (i.e., is statistical coupling adequate and can it be done, or is

the much more difficult physical and mechanistic coupling required to achieve adequate results?).

B.6.4 High Winds

B.6.4.1 State of the Art and Discussion of Uncertainties

The high winds' referred to here include both tornadoes and cyclones, which are meteorologically distinct phenomena. The cyclones include both hurricanes (tropical cyclones) and extra-tropical cyclones. In most PRAs, the cyclones and tornadoes are treated separately.

The methodology for treating high winds probabilistically in the context of a reactor PRA is similar in concept to that used for earthquakes: the similarity arises because each phenomenon can affect almost all parts of a reactor plant simultaneously, at widely dispersed physical locations. This is in contrast to an internal fire or flood, whose effect is almost always confined to a particular part of the overall plant. The similarity extends to the observation, which seems to be true of both earthquakes and high winds, that the main vulnerability occurs for large structures rather than for specific pieces of equipment, except for offsite power equipment whose loss is almost sure to occur in any sizable wind storm.

The approach to analyzing high winds involves the same steps as for earthquakes: first the likelihood of a hazard of a certain size must be worked out (i.e., an external wind field of a certain velocity and pressure). Next, the fragility of structures and equipment in the presence of the supposed hazard must be established. Finally, the nature of and probability that the overall reactor system can fail must be determined, given the failures of certain structures and equipment. Of course, the analysis is quite different in detail for high winds than for earthquakes.

The state of maturity of this part of PRA is still only modest: only a very few PRAs have included this segment of the overall analysis, and neither the methods for determining the wind-hazard potential nor the fragilities methods have been applied enough to enable an understanding of all of the problems with the analyses. However, some useful insights have already been gained, despite the rather sparse experience to date.

The phenomenological difference between hurricanes and tornadoes is due to the differing character of their winds. Hurricanes tend to produce mainly straight winds whose duration in a given location can last for many minutes or even for hours or more; however, the velocities of hurricanes are limited to about 150 mph (miles per hour) for most practical purposes, with winds exceeding about 130 mph being rare. Tornadoes, on the other hand, can produce winds much higher than 200 mph, characterized by the more familiar 'twister' wind forms, and typically lasting only a few minutes to a fraction of one hour at any location. A further important distinction is that, while either type of wind can pick up objects and throw them great distances, the likelihood of this is so much greater for tornadoes that analysis of 'tornado missiles' is almost imperative while missiles associated with hurricanes are seldom analyzed.

An additional complicating factor is that hurricanes are usually associated with torrential rain and flooding; this association is also sometimes true for tornadoes.

Several methods are available to determine the likelihood of a tornado or cyclone producing a wind speed exceeding a particular value at a reactor site. These methods rely on historical records for the most part, which exist nearly everywhere in the U.S. with enough data to provide a useful starting point. These historical records must be extended

and modified to provide a useful foundation for the analysis under discussion, first because the records may not exist at the specific site (so that some extrapolation from other nearby sites is needed); second because local topographical features will modify the wind profile at any specific site; and third because for the very highest wind speeds of possible concern, an extrapolation of the historical record is almost surely needed. This last issue is a point of contention among the experts, since more than one method is available for doing the extrapolation, and differences arise which for PRA purposes can only represent uncertainty as to which is the 'correct' value to use.

Uncertainties in the hazard analysis--the calculated likelihood of a given high wind speed at a site--are today still a significant contributor to overall uncertainty in the PRA. For example, at a given site, two different analytical methods might give answers that differ by as much as an order of magnitude to the question of the likelihood per year that a wind in excess of, say, 150 mph will occur. While the spread can sometimes be smaller than an order of magnitude among estimates, the actual wind speed to which a particular building is exposed is more uncertain, because local building shape factors modify the open-field winds by as much as several mph (sometimes even more), corresponding to as much as an order-of-magnitude probability difference. Effects of building height must also be analyzed, although these are typically easier analytical problems.

The problem of working out the likelihood of tornado missiles of various sizes is also quite difficult: the availability and number of objects of various sizes in the vicinity for the wind to pick up (e.g., telephone poles, autos, trees, even heavier objects) must be determined. The likelihood that a given missile type will attain a high enough velocity to cause harm is also known only roughly. Thus the effects of

these missiles on the integrity of a reactor are not easy to analyze.

The fragility analysis takes several forms. First and most important is the assumption, now commonly agreed to as a sure occurrence, of the loss of offsite power and any onsite power openly exposed to the high winds. Second, structures must be analyzed. Here the typical assumption is that metal-sided buildings are much more vulnerable than concrete-sided ones, with failure modes including buckling, pressure collapse, and corners tearing away. For example, in the Indian Point PRA, the vulnerability of the reactor to winds was entirely due to metal structures. The tornado missiles are also assumed to be highly hazardous to metal structures, while the concrete structures are usually assumed to be rather immune to them; again, the Indian Point analysis assumed that any tornado missile striking any metal building causes its structural failure. The third and most difficult analysis is the fragility of equipment within a building. Here the assumption is usually made that the failure of any building because of a high-speed windstorm will imply failure of all enclosed equipment.

Clearly, each of these assumptions is probably conservative in general, although not necessarily always so. For example, some concrete buildings could easily be more vulnerable at lower windspeeds than now thought because of design idiosyncracies. Also, the response of operators to the extreme weather conditions is difficult to model. Finally, the complicating presence of flood-like rain torrents is not well analyzed for hurricanes in any of the analyses done to date.

The best way to summarize the present state of the art of high-winds PRA is that, while engineering-type insights are available, the quantitative results for core-melt probability or offsite risk are highly uncertain.

B.6.4.2 Methodological Insights (High Winds)

1. The threat from tornado missiles can only be modeled in the most approximate manner at present. They could pose an important threat, but their quantitative analysis is difficult because in general the spectrum of missiles of different types and sizes is not known, and the data base is weak.

2. In locations where hurricanes and/or tornadoes are a threat to other civil structures, they probably should be included as potential accident initiators in any comprehensive PRA.

B.6.4.3 Potential for Improvements and Areas That Will Remain Uncertain (High Winds)

1. A need exists for continuing application of high-winds PRA to several plants to supplement those very few for which the analysis has been done. These new analyses will teach the extent to which lessons learned to date are generic or plant-specific.

2. There is much potential for improving the way the analysis of wind speed hazard likelihoods is done. Refinements, including improvements in calculating open-field wind speeds at a site, better ways to account for local topographic features, and better ways for working out building shape factors and wake effects, could reduce the uncertainties in this part of the PRA analysis considerably.

3. Analysis of the damage potential from tornado missiles will probably continue to be highly uncertain, mainly because of the difficulty in knowing the number and nature of missiles likely to be present as a function of tornado size.

4. For high winds themselves (in contrast to tornado missiles) analysis of the fragility of buildings, especially metal-sided buildings, is in reasonably good shape at present, and contributes rather less to the overall uncertainty than does the windspeed hazard analysis. Moreover, progress in this aspect of the analysis will undoubtedly occur as more studies are done, including the application of the extensive existing data base from non-nuclear experience if it is deemed important to do so.

B.6.5 Flooding

B.6.5.1 State of the Art and Discussion of Uncertainties

The analysis of flooding as a cause of severe reactor accidents has not been carried out in most comprehensive PRAs to date, although flooding clearly poses a potentially serious challenge to the facilities' overall safety. The methodology used in those few studies accomplished so far has been limited to internal flooding and is quite similar to the approach used in studying internal fires: the analyst must first identify critical areas, then work out the probability that a flood might initiate, then determine how long the flood might continue before it is stopped or its floodwaters drained. Finally, the effect on critical safety functions must be determined, and the harm to safety functions must be integrated into an overall systems study using the event-tree approach.

Of course, flooding can occur from either an external cause (river, lake, ocean, torrential rainstorm, etc.) or an internal cause (pipe break, tank rupture, etc.). In all PRAs to date that have considered externally initiated floods, the analysis has been limited to calculating the frequency of a flood large enough to compromise important safety equipment or structures. In every case so far, this frequency has been shown to be small enough that further analysis of plant response has been unnecessary. Thus the remainder of this

discussion will concentrate mainly on internal flooding, although much of the methodology is very similar.

Internal flooding is in practice simpler to study than internal fire for several reasons. First, the mechanisms for terminating the flooding are better understood and the effectiveness of using the mechanisms can be specifically estimated. Second, the areal or volumetric effect of the flooding is much easier to determine, so the zone-designation problem can be handled better. Third, the rate of spread of the flood is usually known and limited. Finally, most of the fragility issues are thought to be understood, although by no means all of them.

Flooding analysis is complicated by several factors, however: fragility of safety functions in the presence of a spray-type flood from a pipe break is very difficult to analyze quantitatively, for example. Also, the corrosion of equipment from the flooding can compromise the ability of a safety function to maintain its operation over the very long postaccident recovery period after a particular flood has been nominally 'controlled. Another issue is the limited ability to quantify partial blockage of drains or sumps relied on to mitigate flooding. Finally, flooding (especially from an external source) can bring solid matter such as sludge, silt, or even sizable objects into areas where they could cause problems difficult to analyze.

Because time-sequence issues are different for most flooding scenarios than for accidents initiated from other sources, the analyst usually must think through these sequences separately and draw special event trees to handle their quantification.

The analysis of the probability of an internal flooding scenario of a given size and location starts by identifying all

major piping or tanks that might be a source of water. The likelihood of a pipe break of a given size is not well known, the historical data base being sparse and not easily transferable to many important scenarios requiring study...indeed, this aspect dominates the ultimate uncertainty in these analyses at this time. Pipe leaks and breaks are not the only potential initiators of flooding, however: another initiator is a possible failure of an isolation valve while a section of pipe is being maintained on-line during reactor operation. Since on-line maintenance (of, say, a valve or an instrument) occurs commonly during operation, a significant chance exists that the isolation valving might be opened either by the error of an operator or maintenance crew, or by hardware failure.

The analyst must determine the approximate flow rate of the break, as well as the ultimate capacity of the source of water (a tank, or a large reservoir, or possibly just 'city water'). With this information he can work out how much water will fill the available volume in how quick a time, taking into account drainage, sump capacity, and sump blockage. The analyst determines, for example, that a given compartment will fill up with water at, say, one inch per minute under certain conditions. Then the analyst must determine the likelihood, in a probabilistic sense, that operator intervention will terminate the flood at a given time, before it reaches whatever height that will compromise critical equipment. While all of this seems straightforward, it poses for the analyst the need to make estimates (sometimes only postulates) about various probabilistic issues that are not well known.

Compounding the analytical problem is the issue of spray flooding from a pipe or tank leak, which could cause electrical failures at nearby locations. The data base and analytical methods for coping with this issue are not well known. There is also the possibility that unusual dependencies among

equipment, for example because of spatial collocation of electrical or support equipment, will cause additional vulnerabilities. The difficulty in modeling human intervention can also complicate the analysis.

Despite the analytical difficulties, those few PRAs in which internal flooding has been analyzed have carried out this part of PRA quite successfully. The analytical problems, while by no means easily overcome, are fully tractable if uncertainties only in the order-of-magnitude range are sought.

B.6.5.2 Methodological Insights and Potential for Improvements (Flooding)

1. Problems with analyzing the human-caused initiation of flooding (by inadvertent removal of isolation from an opened piping system) remain an important contributor to analytical uncertainty and will likely continue to be difficult for the analyst until either a better data base or better analytical methods are developed. Neither of these is now likely to happen soon.

2. Development of experimental information on the fragility of equipment exposed to spray flooding phenomena might strongly improve the analytical methodology and might not be very difficult to obtain if only modest data are sought.

3. Until more attempts are made to carry out a full probabilistic analysis of internal flooding scenarios, the various methodological difficulties and potential achievements of flooding PRA analysis will not be fully known.

4. Externally initiated flooding is not usually an important accident initiator. Analyses have typically placed acceptably small upper bounds on the core-melt probability

from external floods by calculating that the frequency of sufficiently large external floods is small enough. There has been no comprehensive PRA of external flooding so far because it has not been necessary.

B.6.6 Other External Initiators

Besides the four major initiators discussed separately (earthquakes, high winds, internal fires and floods), several other categories of accident initiators which can threaten a large power reactor are usually considered within the category of 'external events':

aircraft impacts
barge and ship collisions
truck, train, and pipeline accidents
external fires
volcanoes
turbine missiles
lightning

The state of the art of analyzing all of these in the context of PRA is adequate at least conceptually, but they are all undeveloped in actual practice...to be precise, most of these have been examined at least to some extent in various PRAs, but never with a full-scale analysis.

Typically, these external initiators are analyzed probabilistically by performing a bounding analysis on the likelihood of the initiating event. An initiating event serious enough to merit 'concern' is usually semi-quantitatively estimated. (For example, the analyst might consider how large an external fire must be to compromise important safety equipment.) The analyst next determines quantitatively the likelihood that such a large initiator might occur. If the likelihood is small enough, or can be bounded well enough, the

analysis ends with the statement that the effect under study "does not contribute significantly," or some such language.

This approach is fully adequate if the analyst performs his work well. Several pitfalls could lead an analyst astray. Among them are:

1. The analyst might not have chosen properly that initiating event he thinks is worth analyzing because it might compromise the plant. That is, a much smaller (and more likely) event might lead to undesirable consequences.

2. The analyst might overlook some coupled failure modes from the initiator.

3. The analysis of the likelihood of occurrence might be badly flawed, for example, because no historical record exists and the extrapolation procedure used is erroneous, or because the historical data base is actually erroneous itself, or inapplicable.

The main insights gained to date from the analyses performed on these initiators of lesser importance are that, generally, they have less risk significance. That is, seldom has any one of them turned out even to need further study. This insight is quite important, because it tells how good the deterministic design and operational requirements have been in assuring plant adequacy in these areas. The design and regulatory approaches seem to be adequately conservative.

The main limitations to the analyses are the possibility that some oversights, of the kinds mentioned above, could invalidate the conclusions. Given the conservatisms in the assumptions (specifically, even if the postulated external initiator is more probable than thought, the plant fragility analysis

must still be performed with a high likelihood of plant survival), the general conclusion regarding plant design adequacy is likely to be correct.

B.7 Sabotage

Treatment of sabotage as an initiating event has not been traditionally included in PRAs. The threat of sabotage has been long recognized and treated outside the PRA arena. PRA techniques have on occasion been used to analyze various vital areas and penetrations related to sabotage, but the risk of sabotage itself has never been calculated, principally because of difficulty in quantifying the threat frequency.

The use of PRA techniques to address the sabotage issue dates back to 1975 when a fault tree analysis was used to identify the combination of events which, if caused by a saboteur, could result in significant releases of radioactive material.* Sabotage vulnerability studies have shown that sabotage cannot result in higher consequences than those considered in PRAs.* A methodology was later developed which uses fault trees to aid identification of vital areas, that is, areas which warrant special attention when providing sabotage protection for the plant.* The techniques used are not probabilistic.

Some probabilistic computer modeling has been used to identify weak links in physical protection systems. These codes, highly subjective in nature, model the detection and response capabilities of physical protection systems, given an external sabotage attempt. What has not been done is to develop models which allow meaningful predictions of the probability of a sabotage attack.

In some sense, a sabotage attempt can be regarded as another initiating event. The accident sequences stemming from the attempt is not unlike those modeled in PRAs. The saboteur can be, in effect, a common cause for the failure of several

components or systems concurrently. The difficulty lies in being able to predict meaningfully the frequency of this initiating event. So called random initiating events (component failures, human errors, earthquakes, tornadoes, etc.) can be estimated meaningfully by considering past experience. The assumption inherent in looking at past experience and using that as a basis for future predictions is that the failure rates do not vary significantly from year to year. For example, if data indicates that pumps have failed on the average once out of every 100 demands, then it is assumed that in the future they will fail on the order of once in every 100 demands. Uncertainties are placed on the estimate reflecting how much data one has to base the estimate. Nevertheless, the inherent assumption is that past performance is indicative of future performance.

Such an assumption cannot be made for sabotage. The frequency of sabotage events are a function of social and political unrest among other things, which may differ significantly with time. Therefore, existing statistical methods, which use past sabotage frequency experience to predict future sabotage frequencies are not valid.

The development of PRA is not expected to improve this situation in the foreseeable future. Therefore, the evaluation of sabotage is expected to remain, appropriately, outside the formal discipline of PRA.

B.8 Concluding Remarks

As this appendix has shown, PRA consists of a multitude of different disciplines, having different levels of maturity, and different uncertainties both in the size and root cause of the uncertainty. The uncertainties must be recognized as being, for the most part, not unique to PRA, but reflecting a lack of data (or experience) or a lack of knowledge about system response, human behavior, or accident phenomenology.

These uncertainties exist whether the decisionmaker uses PRA, deterministic modeling, or so-called engineering judgment when making regulatory decisions. They reflect current experience and knowledge.

Since its beginning, PRAs have attempted to display the uncertainties and, in so doing, have focused attention on those uncertainties. Displaying the uncertainties, as is done in a PRA, provides important information to the decisionmaker. This analysis can provide an estimate of how this lack of experience and/or knowledge impacts engineering insights drawn from the PRA. This is done by propagating uncertainties through the analysis or by performing sensitivity analyses within the PRA. Thus, the treatment of uncertainties should logically be considered a strength of PRA rather than a limitation.

The current level of experience or knowledge (data base) should be recognized as differing with different parts of the PRA. Thus, the reliance the decisionmaker places on the PRA insights should likewise differ with different parts of the analysis. As the state of the art exists today, the following conclusions can be reached.

B.8.1 Systems-Analysis, Internal-Initiated Events

The methodology and information base for this type of analysis is reasonably mature and stable. A relatively high level of confidence can be placed in insights about the relative importance of plant characteristics, dominant accident sequences, and core-melt frequencies from analyses done within this scope. The weakest part is the human reliability analysis, thus confidence is best placed in conclusions that are robust in the face of human error uncertainties.

B.8.2 System-Analysis, External Events

The weakness in the analysis of analysis of the frequency of external event initiators, the lack of data on component response to these initiators, and the lack of uniformity in PRA methodologies for external-event initiators results in decreased confidence in insights derived from this part of the PRA. Little confidence should be placed in insights that are derived from relative comparison between internal and external event analyses or across different types of external events. More confidence can be placed in insights which stem from comparisons within the analysis of a specific initiator type. That is, more confidence can be placed in an estimation of which seismic-initiated accident sequence dominates the frequency of core melt from seismic initiators than any conclusions about whether a seismic, fire, or small LOCA sequence dominates the overall core-melt frequency.

B.8.3 Accident Progression, Containment Response, and Source Term Prediction

Currently, this represents the most uncertain and most unstable part of the PRA methodology, so much so that the best approach to the analysis is by no means clear in many cases. This situation can be expected to continue until the ongoing source term work has somewhat stabilized predictions. Very limited confidence should be placed on insights being derived from this part of the analysis.

B.8.4 Offsite Consequence Analysis

This type of analysis is relatively mature. Although it does not have the long experience of usage as system reliability analysis, a relatively high level of confidence can be placed on the resulting insights; however, they are only applicable to the assumed source term. Of course, the stochastic uncertainties associated with meteorology are quite large, so the

actual consequences associated with a radiological release are quite uncertain.

In a field as complex as nuclear reactor safety and PRA, generalizations are difficult to support. Nevertheless, the motivation to do so is strong. Therefore, the following chart is provided in the interest of crisp communications, despite the risk of oversimplification.

Table B-1

PRA Activities and Related Level of Confidence Levels

PRA Activities	Level of Confidence in Insights
Systems-analysis internal events	Relatively high.
System-analysis external events	Modest for insights applicable to specific external events. Very limited for insights spanning different external events or internal events.
Accident progression, containment response, source term	Very limited.
Offsite consequence	Relatively high, but applicable only for assumed source term.

Starr, C. and Whipple, C., "Coping with Nuclear Power Risks: The Electric Utility Incentives," EPRI, 1981.

Susquehanna, "Susquehanna Steam Electric Station: Probabilistic Risk Assessment," Pennsylvania Power and Light Company, prepared by NUS Corporation, NUS-4376, Draft--1983.

von Herrmann, J. L. and Wood, P. J., "Engineering Applications of Probabilistic Risk Assessment," Wood-Leaver and Associates, submitted to Progress in Nuclear Energy, July 1983.

Woods, D. D., Wise, S. A., and Hanes, L. F., "Evaluation of Safety Parameter Display Concepts," NP-2239, EPRI RP 891, Electric Power Research Institute, Palo Alto, California, February 5, 1982.

Wreathall, J. W., "Operator Action Trees, An Approach to Quantifying Operator Error Probability During Accident Sequences," NUS Report 4159, NUS Corporation, Gaithersburg, Maryland, July 1982.

Yankee Rowe, "Executive Summary, Probabilistic Safety Study, Yankee Nuclear Power Station," Yankee Atomic Electric Company, 1982.

Zion, "Zion Probabilistic Safety Study," Commonwealth Edison Company, 1981.

APPENDIX C

SUPPLEMENTAL INFORMATION SUPPORTING RISK INSIGHTS FROM PRAs

C.1 Supplemental Information Regarding Plant Risk Insights

In the course of performing PRA studies, those involved gain valuable engineering and safety insights. Conceptual insights are the most important benefits obtained from PRAs, and the most general of these is the entirely new way of thinking about reactor safety in a logic structure that transcends normal design practices and regulatory processes. PRA thought processes introduce realism into safety evaluations, to the extent possible, in contrast to deterministic thinking which may mask important matters due to its generally conservative approach.

Several studies including WASH-1400 (RSS, 1975) indicate important distinctions between contributors to different types of outcomes of potential accidents. Table C-1 presents some results from the Zion/Indian Point (ZIP) studies (Zion, 1981; Indian Point, 1982) comparing the important accident initiator contributors to core-melt with those important to public risk. This indicates that risk cannot be measured in terms of any single indicator and that changes in plant configuration that significantly affect one indicator may or may not impact the others; thus, core melt fixes may not impact public risk and vice versa.

Extremely important insights gained from some studies, benefitting from the research performed upon completion of the RSS, suggest that (a) phenomena challenging the containment are both less severe and less likely than previously believed, and (b) containments are stronger than previously assessed. As a result, the containment may be more effective than previously perceived in its ability to withstand the bulk of the

Table C-1

Results from Zion and Indian Point PRA Studies

	Major Contributors		
		Public Risk	
Reactor Unit	Core Melt	Acute Fatalities	Latent Fatalities
Zion 1 and 2	Small LOCA	Seismic	Seismic
Indian Point 2	Fire, Seismic	Seismic, Interfacing LOCA	Seismic, Fire
Indian Point 3	Small LOCA, Fire	Interfacing LOCA	Fire

early threats, i.e., events which occur in the first few hours of an accident sequence. These findings could be generally true for strong, large, dry containments like Zion/Indian Point. Energetic steam explosions, which can generate a missile that can penetrate the containment, are now judged more improbable than before; some studies suggest that such explosions are so unlikely as to be considered physically impossible.* Delayed containment failure modes are basically due to late overpressurization or basemat melt-through. Late overpressurization occurs only if containment heat removal fails, or if the quantity of hydrogen and other non-condensibles can accumulate in excess of that from 100% core/zirconium reaction. Basemat melt-through can only occur if the debris is not coolable. The evidence appears to be strong that basemat melt-through is not a failure mechanism that contributes appreciably to risk. The studies that have been made relative to liquid pathways for radioactive material have indicated relatively little risk compared to atmospheric pathways, but this might may be design and site specific.*

Changing perceptions of accident phenomenology, enhanced effectiveness of the containment, and differences in source terms rather dramatically impacted estimates of the health effects in some studies. Offsite consequences were analyzed with either CRAC2 or CRACIT. More realistic modeling of plume meander in CRACIT has only a small impact on risk as expressed by the complementary cumulative distribution function in early and latent fatalities. The use of the PRA study results to affect the rationale of emergency planning has emerged as another significant application of PRA.

The results of PRA studies are usually expressed in terms of core-melt frequencies, frequencies of release of various magnitudes, or curves presenting the frequencies of occurrence of different reactor accident consequences (e.g., early and

latent fatalities), depending on the level of the PRA. These are further discussed below.

C.1.1 Supplemental Information Regarding Core Melt Frequencies

One of the results of a PRA study is the identification of a relatively small number of accident sequences that represent the dominant contributors to core melt. Analysis of the salient features of the dominant accident sequences from eleven PRAs yielded a characterization of accident sequence categories as shown in Table C-2. The table shows the contribution (percentage) of each sequence category to the total core-melt frequency quoted in the study. In the cases of Zion, Big Rock Point and Indian Point, the total core-melt frequency includes the contribution from external events. Externally initiated accident sequences were characterized by their effect on the plant; e.g., if an earthquake caused a loss of AC power the sequence was categorized under loss of offsite power.

Figure C-1 represents a composite chart that combines the first five columns of Table C-2 (those studied in the EPRI sponsored review) for PWRs and BWRs respectively. The BWR chart does not include Big Rock Point because its design was considered atypical of other BWRs and its relatively high accident sequence frequencies would have biased the results. The RSS BWR was substituted for it because it was deemed more representative of operating BWRs. The grouping was slightly modified in order to account for negligible contributions of large and interfacing LOCAs and small LOCAs to BWRs. Recent work (Garrick, 1983)* indicates the contribution of transients to core melt frequency may be as large as 80-90% for some newer PRAs not yet published.

Several studies (e.g., ZIP, Big Rock Point, Limerick) highlighted the importance of a probabilistic treatment of such

Table C-2

Core Melt Sequence Contributions

Sequence Category	% Total Core-Melt Frequency										
	BRP	Zion	Limerick	Grand Gulf	ANO	Surry	Peach Bottom	Sequoyah	Ocones	IF-2	IF-3
Small LOCA's-Injection Failure	10	0	0	0	28	27	0	18	14	37	33
Small LOCA's-LTDHR* Failure	4	41	0	14	5	20	1	67	21	3	43
Large LOCA's-Injection Failure	0	3	0	0	0	4	1	0	0	0	1
Large LOCA's-LTDHR Failure	0	18	0	0	0	2	0	1		4	11
Transients-PCS* Hot Available											
a. Loss of Off-site power	14	18	48	27	20	7	0	0	12	26	3
b. Injection failure	36	0	34	0	23	14	2	5	15	28	2
c. LTDHR failure	5	0	0	38	0	0	0	0	21	0	0
Transients-PCS Available											
a. Injection failure	0	0	5	0	0	0	0	0	1	0	0
b. LTDHR failure	0	4	3	0	0	0	47	0	0	0	0
ATWS	0	15	2	14	4	9	47	0	11	0	1
Interfacing LOCA	9	0	0	0	0	9	0	9	5	0	0
TOTALS	78%	99%	92%	93%	80%	92%	98%	100%	100%	98%	94%

*LTDHR is long term decay heat removal which includes recirculation and RHR. FCS is power conversion system.

BRP - Big Rock Point
ANO - Arkansas Nuclear One Unit One
IP-2 - Indian Point unit No. 2
IP-3 - Indian Point Unit No. 3

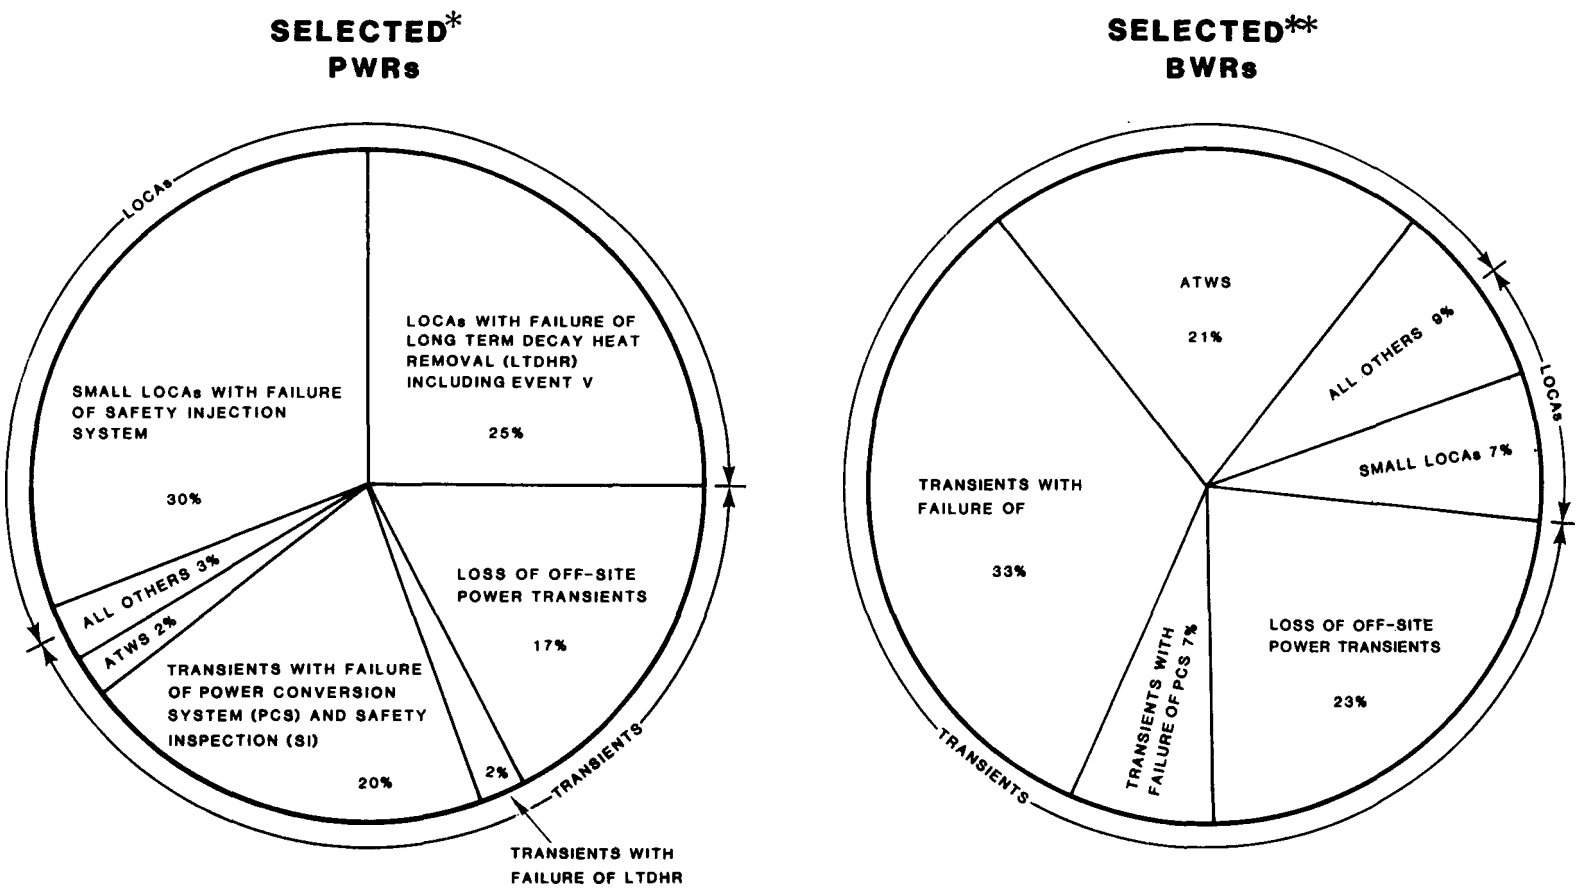

* Arkansas-1, Ocones-3, Sequoyah-1, Surry, Zion

** Grand Gulf, Limerick, Peach Bottom-2

"All Others" contain roughly equal mixes of LOCA and transients

NOMENCLATURE:

LOCA	**Loss of Coolant Accident**	**PCS**	**Power Conversion System**
SI	**Safety Injection**	**LTDHR**	**Long Term Decay Heat Removal**
ATWS	**Anticipated Transients Without Scram**	**Event V**	**Interfacing Systems LOCA**

Figure C-1. PRA Study Results - Core Melt Probabilities

external events as earthquakes, fires and floods. Part of the reason for the high external event contribution found in these studies is related to the considerable uncertainty associated with their frequency of occurrence as well as the structural and containment responses to such events. These uncertainties, by and large, are attributable to the state of knowledge and ability to model. It is believed that as both knowledge and modeling improves, their contribution to the uncertainties is likely to change.

C.1.2 Supplemental Information on Releases

Figure C-2 provides a range of categorized radionuclide release fractions, their release frequencies and descriptive comments. Three illustrative cases are displayed: (a) severe containment failure modes, i.e., early overpressurization or containment bypass; (b) late containment failure; and (c) containment intact despite core melt. Figure C-3 displays the same type of information but with more detail showing the results of individual plant studies in terms of iodine release only.

C.2 Supplemental Information Regarding Dominant Accident Sequences

Shortly after the Reactor Safety Study a program was instituted by the NRC to address, among other things, the similarities of dominant accident sequences for PWRs and BWRs. The results of the program indicated that the dominant accident sequences were not consistent in detail across broad plant classes such as all PWRs or all BWRs. In fact, the plant features and operational characteristics which gave rise to specific dominant accident sequences were not of a nature that would lead to the conclusion that dominant accident sequences would be similar even for smaller classes such as "all B&W PWRs." Characteristics which determined what accidents were dominant often reflected individual utility prac-

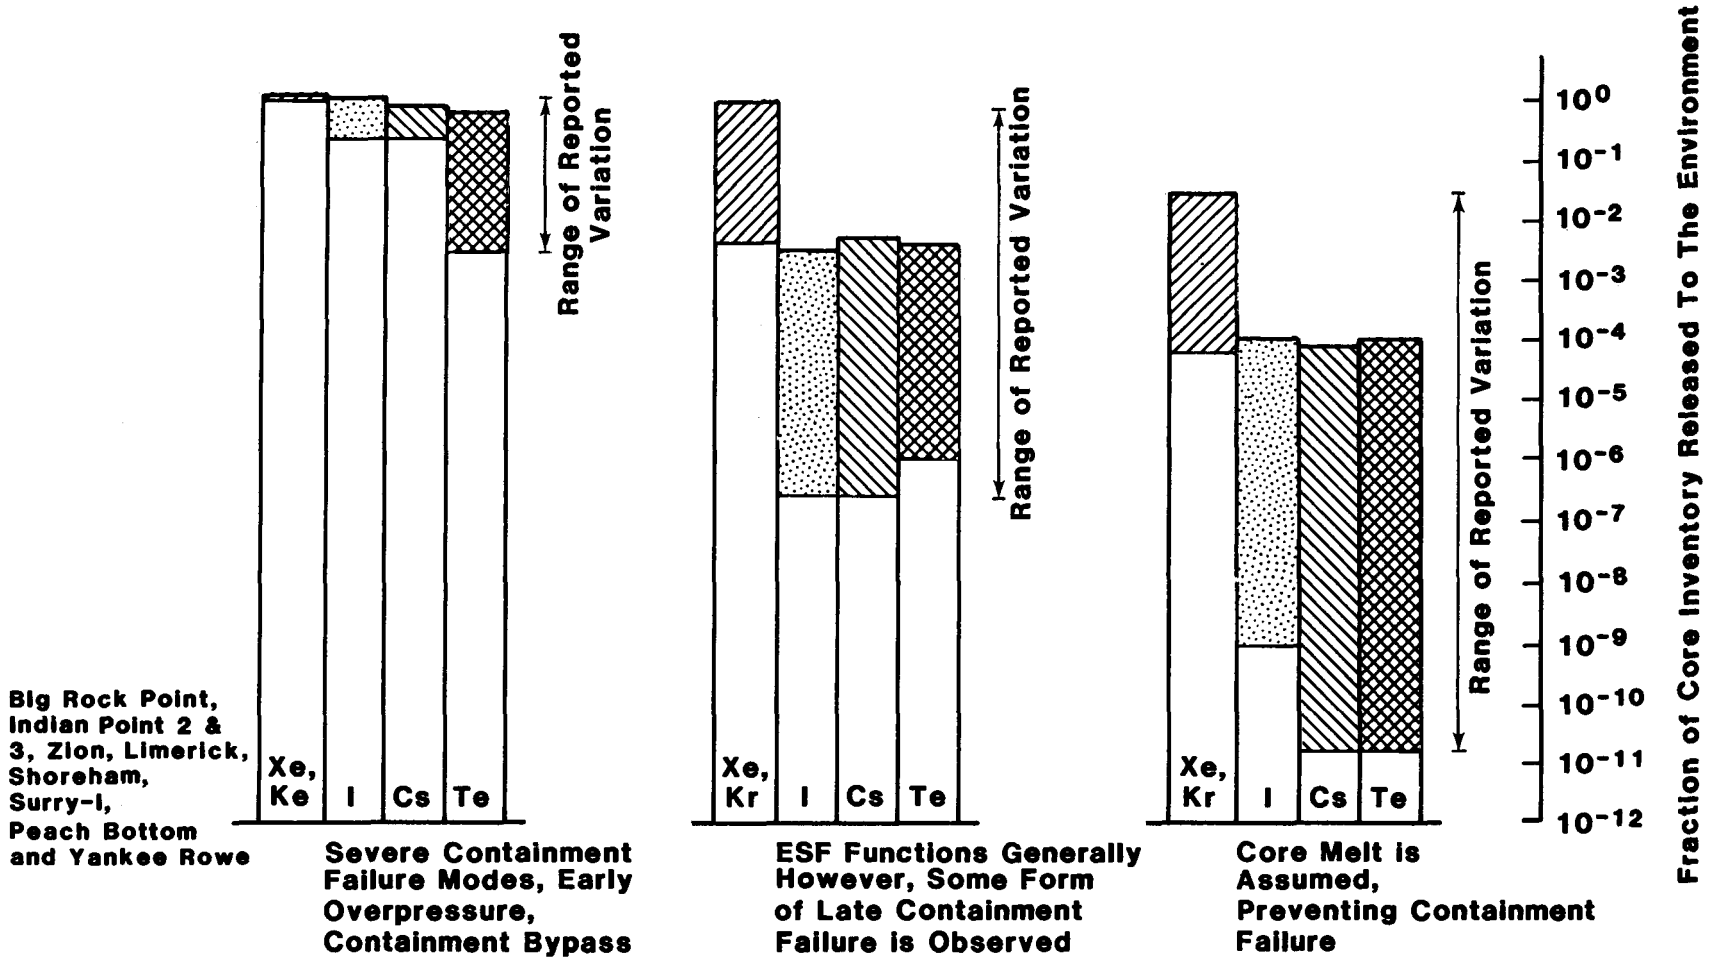

Figure C-2. Range of Radionuclide Release Fractions from Listed PRA Studies

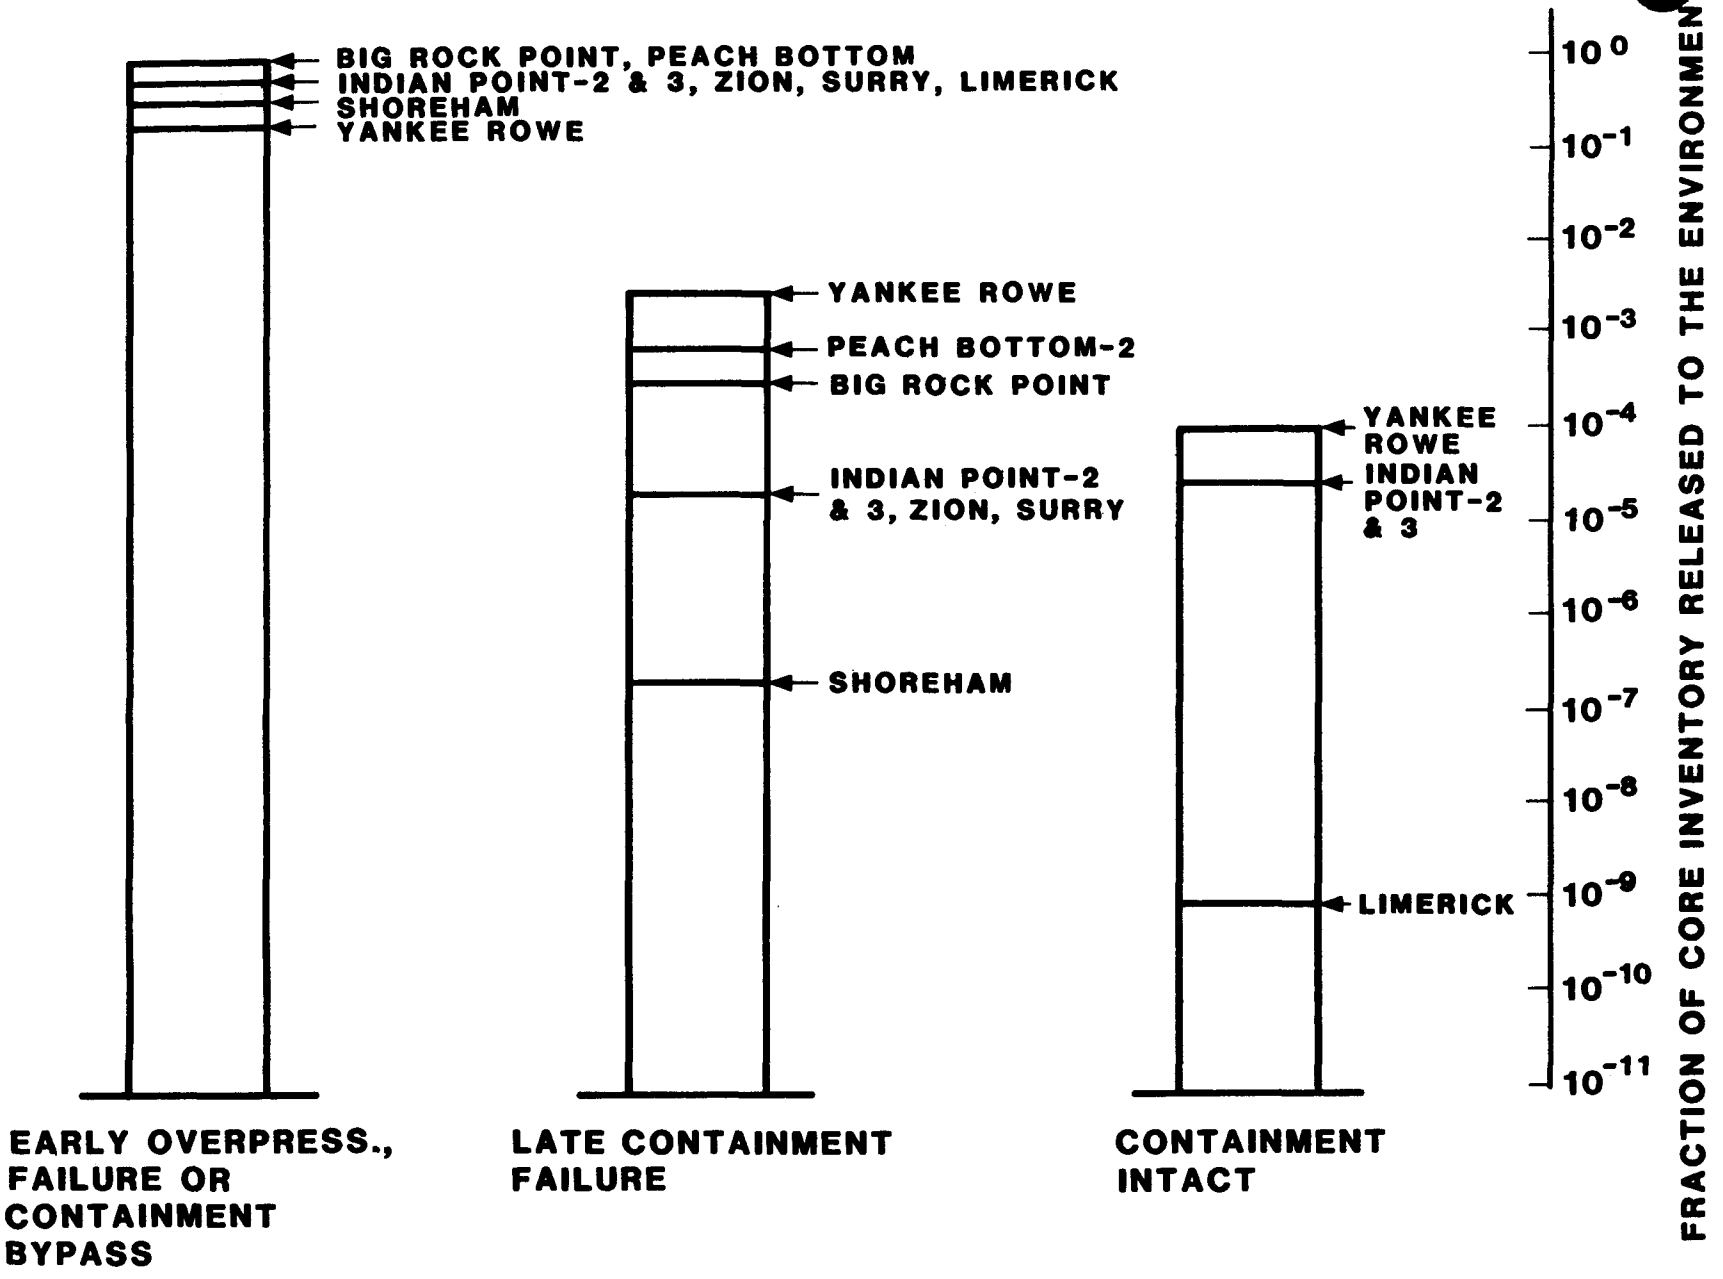

Figure C-3. Iodine Release Fractions from Listed PRA Studies.

tices such as what check valves were tested before returning to power after shutdown or, if tested, how they were tested.

Thus, the conclusion that began to form in the late 70s was that the specific dominant accident sequences for plants across the industry could be quite different in detail and therefore, the ability to reduce the existing plant risk through generic decisions (as opposed to plant specific decisions) would be more difficult than perhaps hoped.

Many more risk assessments have now been done, and the results of these studies have tended to confirm this earlier conclusion and to add new insights. Much of these new insights stemmed from the growing body of work in the area of accident phenomenology taking place within containment after a core melt. Containment failure prediction and the resultant source term to be used for consequence calculation began to appear to have wider uncertainties than originally thought, and the possibility was recognized that the source term used may be quite conservative. This evolving body of work was important to the determination of dominant accident sequences generically from two respects:

1. The assessment of accident sequences which dominate core melt frequency, and

2. The instability of the information base being used to determine containment response and source term values leads to the conclusion that dominant accident sequences with respect to core melt frequency can be predicted with greater confidence than those contributing to risk.

However, as noted in Chapter 5, though each plant is unique and may exhibit accident sequences that are specific to the design and operation of the plant, it is still possible to identify broad accident sequence categories on a functional

basis and identify those characteristics of plants which lead to the high frequencies (dominance) of specific accident sequence categories. Based on these characteristics plants can be placed into classes such that each member of the class would be expected to have dominant accident sequences in similar categories. This task is currently underway under sponsorship of the Office of Nuclear Regulatory Research in the Accident Sequence Evaluation Program (ASEP) which is concentrating on those accident sequences which dominate core melt frequency.

This program has reviewed a dozen existing PRAs, including both NRC and industry sponsored studies. The dominant accident sequences identified have been placed into broad categories representing functional accident sequences. As a group they represent a summary of the functional sequences which have been found to dominate core melt frequency in past PRAs.

The functional accident sequences evaluated in the ASEP study have been first grouped by initiating events (transients or LOCAs) and then by major functional failures.* Using this approach, the past PRA dominant functional accident sequences for PWRs are as follows:

1. Transients with loss of the reactor subcriticality function

2. Transients with induced loss of the reactor coolant system integrity and the loss of the core cooling function

3. Transients with loss of the core cooling function

4. Transients with loss of the core cooling and containment heat removal functions

5. LOCAs with loss of the core cooling function

6. LOCAs with loss of the core cooling and containment heat removal functions

7. Failure of a high-to-low pressure reactor coolant system interface in a non-mitigatable LOCA.

The sequences are shown in Table C-3 together with a range of assessed frequencies and are described below.

Accident Sequence Category 1 - Transient with loss of the reactor subcriticality function (ATWS). Anticipated transients without scram has been an unresolved safety issue for some time. One of the principal unanswered questions is whether the primary system can survive the pressure spike expected, if the transient causes the main feedwater system to trip off or run lack when the failure to scram occurs. Given the RCS does not fail due to the pressure spike, actuation of the high pressure injection system may terminate the sequence by injecting borated water to achieve reactor subcriticality and to make up RCS inventory lost via the pressurizer safety valves during the pressure transient.

Variations in sequence frequencies in past PRAs have not been large; the range has been from 1 to 60 per million reactor years (MRY). The staff's estimate of the frequency is at the high end. The recent Salem ATWS event, which has cast a shadow across past reactor protection system reliability predictions, and lingering uncertainties in the phenomenology occurring in the primary system following an ATWS suggest greater uncertainties than past PRA results reflect. The NRC staff is in the process of formulating rules intended to reduce the potential for core melt from this category of accident sequences. They cover such things as diverse scram signals, signals for tripping the main turbine and starting the auxiliary feedwater system which are separate from the scram

Table C-3

Functional Accident Sequence Categories (PWR)

Sequence Category	Freq. Range x 10^{-6}	Major Uncertainties	Comment
1) Transient Loss of Reactor Subcriticality	60 1	RPS Reliability RCS Ability to Withstand Pressure Spike	ATWS Rule Pending
2) Transient Loss of Integrity Loss of Core Cooling	30 <1	PORV Demand Rate HPIS Availability Necessity to Switch-Over to Recirc.	TMI Fixes (Raising PORV Set Point and Anticipatory AFWS Start Signal) Should Reduce Sequence Freq.
3) Transient Loss of Core Cooling	1000 0.1	Feed and Bleed Capability AFWS Availability	TMI Fixes Have Called for Many Improvements in AFWS Availability
4) Transient Loss of Core cooling Loss of Containment Heat Removal	140 0.2	Redundancy of AC Power Sources Battery, CST Depletion Times Possibility of Induced RCS Pump Seal Leak Long Term Ventilation Loss Effects AFWS Availability	NRC Position Statement-Forthcoming
5) LOCA Loss of Core Cooling	200 <0.4	LOCA Frequency ECCS Success Criteria ECCS Redundancy	Small LOCA May Be Higher Than Thought Due to RCP Seal Leaks TMI fixes Stressed Better Procedures for Small LOCA
6) LOCA Loss of Core Cooling Loss of Containment Heat Removal	6 <1	LOCA Frequency ECCS Success Criteria ECCS Redundancy	Small LOCA May Be Higher Thought Due to RCP Seal Leaks TMI Fixes Stressed Better Procedures for Small LOCA
7) Event V			

signal, and a reactor protection system reliability assurance program.

Pending final regulatory resolution of the ATWS unresolved safety issue, this sequence can be expected to have the potential of being dominant in PWR power plants.

Accident Sequence Category 2 - Transient with induced loss of reactor coolant system integrity and loss of the core cooling function. This sequence, as in Category 1, is initiated by a transient. The reactor protection system operates successfully, and subcriticality is achieved leaving only decay heat to be removed. Failure of reactor coolant system integrity occurs, which is in effect a transient-induced LOCA. Failure of core cooling then results in failure to keep the core covered and subsequent core damage. The principal cause of loss of reactor coolant system integrity has been found in past PRA work to stem from 1) failure of a primary relief valve to reclose after opening due to pressure buildup associated with the transient, or 2) a reactor coolant pump seal leak occurs as a result of seal cooling loss caused by the transient. Loss of core cooling in this case usually implies failure of the ECCS to respond to the transient-induced LOCA.

Variations in sequence frequencies for this sequence in past PRAs has ranged from 1 to 30 per MRY. These variations from plant to plant and sequence frequency uncertainties within a plant stem primarily from three sources. First, the chances of a transient-induced LOCA are dependent upon the number of times the pressurizer relief valves are opened which, in turn, is a function of the type of transient and the relief valve setting. Valve settings vary from plant to plant. Following TMI, relief valve settings have been increased and anticipatory trips are being used to start the auxiliary feedwater system earlier following a trip. Both of these improvements should aid in reducing PORV openings and thus reducing the

frequency of this sequence. A second variation across plants is the availability of the high pressure injection system (HPIS) due to differences in HPIS design and the procedures used for test and maintenance (T&M). Third, past PRAs have indicated that the switchover from the injection mode to the sump recirculation mode of operation following depletion of the refueling water storage tank (RWST) is a critical time for emergency core cooling system failure . However, additional research may indicate that for these transient-induced LOCAs, which tend to be small LOCAs, enough time may be available for many plants to achieve cold shutdown prior to the depletion of the RWST, and the need to switch over to the sump for core cooling is not required. If this is true, the frequency of this sequence could be significantly less than current predictions indicate.

The consequences of this sequence are highly dependent upon the operability of the containment pressure suppression systems (containment sprays and/or fans). The design of containment pressure suppression systems vary widely within the industry.

Pending further research, Sequence Category 2 should continue to be considered as having the potential for being a dominant accident sequence.

Sequence Category 3 - Transients with loss of the core cooling function. This sequence, as for Categories 1 and 2, is initiated by a transient. The reactor protection system operates successfully, and subcriticality is achieved leaving only decay heat to be removed. Although pressurizer relief valves may open as a result of the transient, they reclose when the pressure drops, thus a transient-induced LOCA has not occurred. However, the core cooling functions, both heat removal via the steam generator and or steaming off primary

coolant via the pressurizer relief valves, fail to keep the core covered and adequately cooled.

For those transients in which the power conversion system is lost, the preferred method of removing decay heat is the AFWS feeding the steam generator. If it fails, a backup method appears feasible in some plants. This method consists of steaming off primary fluid via the pressurizer relief valves and replenishing the fluid with the high pressure injection system.

Variations in sequence frequencies for this sequence in past PRA work has been 0.1 to 3000 per MRY. This large range stems from variations in both the preferred and backup method of removing decay heat. Following TMI, a study of AFWS was undertaken by the NRC and the variations in design resulting in variations in AFWS availability from plant to plant were found to be significant. These variations stemmed from different numbers of trains, different methods of actuation, single valves in the water source serving all trains, different limiting-conditions-of-operations, and different power dependencies to name a few. Steaming off primary coolant and replenishing reactor coolant with the high pressure injection system (called feed-and-bleed) was also not widely accepted as a viable means of removing decay heat in the absence of secondary system cooling. Thus, there exists wide variations in feed-and-bleed procedures and limited study on the actual capability of plants to successful feed and bleed.

In addition to plant-to-plant variations, the Accident Precursor Program,* sponsored by the USNRC, Office of Nuclear Regulatory Research, suggests that past PRAs may have overestimated auxiliary feedwater system availability, which would imply an even higher potential for this sequence to dominate in PRAs.

Accident Sequence Category 4 - Transient with loss of both the core cooling and the containment heat removal functions. This sequence is the same as sequence Category 3, except that the containment suppression systems are lost. These systems vary from plant to plant; but, for most PWRs, they consist of sprays and/or fans; and, although designed for LOCA mitigation, they serve to reduce containment pressure caused by steam in a transient initiated core melt accident. This sequence category is dominated by the station blackout sequence. This sequence is initiated by a loss of offsite power, then followed by failure of the onsite emergency AC power system. This sequence leaves the typical plant with only steam driven and DC powered systems operating to remove decay heat. Typically, this means one train of AFWS. Past PRAs have shown variations in this sequence from 0.2 to 100 per MRY.

Station blackout is an unresolved safety issue and is being pursued under Task Action Plan (TAP) A-44.* The research performed under TAP A-44 for station blackout found that failure of AC-independent core cooling systems due to DC power supply (battery) depletion or depletion of the condensate storage tank could lead to core melt in long periods (such as eight hours) after the loss of all offsite and emergency AC power. Since this was not considered in all past PRAs, such a scenario may make this sequence dominant.

The wide variation of frequency of this sequence stems from several factors: (1) variations in redundancy of offsite and onsite AC power sources, (2) variations in depletion times of batteries or the condensate storage tank, (3) the possibility of an induced reactor coolant pump seal failure due to the loss of seal cooling, (4) unknowns in the long-term effect of loss of ventilation systems, and (5) variations in availability of AC-independent trains of AFWS.

An NRC position is expected in the near future on the station blackout issue. It is expected that the action taken will reduce the frequency of this sequence in many plants.

Sequence Category 5 - LOCAs with loss of the core cooling functions. This sequence category differs from those preceding in that the initiating event is a LOCA stemming from a pipe rupture or component failure which releases coolant from the reactor coolant system. This is followed by failure of the ECCS.

The variation in sequence frequencies in past PRAs range from less than 0.4 to 200 per MRY. This range stems from the range of LOCA frequencies (20,000 per MRY for small LOCAs stemming from reactor coolant pump seal leaks to 100 per MRY for large LOCAs), from variations in success/failure criteria (number of ECCS trains required to mitigate various size LOCAs), and, of course, the availability of the ECCS. The predominance of existing evidence indicates that greater risk stems from the smaller LOCA sequences due to their higher frequency of occurrence.

Higher small LOCA frequencies may exist than estimated in some previous PRAs, since not all of them considered the possibility of a reactor coolant pump seal rupture as an initiating LOCA. Since the TMI accident, the need for better operating procedures for small LOCAs has been stressed which has also led to improved training of operators. This latter development may lead to a smaller chance that this sequence will dominate in the future.

Currently, small LOCAs with loss of the core cooling function should be considered as having the potential for dominating core melt or risk in some PWRs.

Accident Sequence Category 6 - LOCAs with loss of both the core cooling and the containment heat removal functions. This sequence differs from the preceding category only by the addition of the containment heat removal function.

The frequency of this sequence has historically been relatively low, 1 to 6 per MRY. However, the loss of the containment heat removal function implies a greater potential for relatively higher consequences. The variations in frequency stem from some of the general variations described for accident sequence Category 5. The same comments also apply; namely, that some past PRAs did not consider the reactor coolant pump (RCP) seal leak as an initiating event and better procedures and operator training exists to mitigate sequences in this category.

Accident Sequence Category 7 - Failure of high to low pressure reactor coolant system interface resulting in a non-mitigatable LOCA.

This sequence first received attention in the Reactor Safety Study as sequence "V". The frequency of this sequence is very dependent upon plant design and operational procedures and has been estimated to be as high as 700 per MRY. In general this sequence results from valve failures between the reactor coolant system at high pressure and interfacing emergency and/or shutdown systems which are designed for low pressure and extend outside containment. The high pressure release from the reactor coolant system is assumed to cause failure of the low pressure piping and thus a LOCA outside containment.

This sequence was identified in WASH-1400 (RSS, 1975) and industry was made aware of the sequence some years ago by the NRC. Some corrective action has taken place; however, recent PRAs indicate that this sequence continues to appear to be a

major risk contributor due to the potential release of fission products directly to the atmosphere.

For BWRs a similar set of potentially dominant functional accident sequences have been identified. These fall into 6 functional classes.

1. Transients with loss of the subcriticality function

2. Transients with induced loss of reactor coolant system integrity and loss of the core cooling function

3. Transients with induced loss of reactor coolant system integrity and loss of the containment cooling function

4. Transients with loss of the core cooling function

5. Transients with loss of the containment heat removal functions

6. LOCAs with loss of containment heat removal function.

The sequences are shown in Table C-4 and are discussed below.

Accident Sequence Category 1 - Transient with loss of the reactor subcriticality function (ATWS). Anticipated transients without scram has been an unresolved safety issue for some time. Uncertainties related to primary system response, the adequacy of high pressure cooling, operator control of water level, and the effect of other system failures are important. ATWS sequences lead to core melt by one of two principal modes. In the first mode, the resulting high primary system pressure cause steam to dump to the suppression pool at a rate greater than high pressure makeup. In the second mode, the heat dumped to the suppression pool is greater than the capacity of the suppression pool cooling

Table C-4

Functional Accident Sequences Categories for BWRs

Sequence Category	Freq. Range	Major Uncertainties	Comment
1) Transient Loss of Reactor Subcriticality	5E-5 1E-7	RPS Reliability Adequacy of ECCS Unknown Phenomenology in RCS Ability of Open to Control Water Level	ATWS Rule Pending
2) Transient Loss of RCS Integrity Loss of Core Cooling	7E-5 <2E-7	ECCS Availability Operator Procedures for ADS SRV Demand Rate	
3) Transient Loss of RCS Integrity Loss of Containment Cooling	1E-3 1E-7	RHR Availability SHV Demand Rate	Estimated Time to Core Melt Appear Longer Than Previously Expected, Thus Longer Times for Recovery
4) Transient Loss of Core Cooling	7E-4 2E-7	ECCS Availability Operator Procedure for ADS	Station Blackout Rules Pending
5) Transients Loss of Containment Cooling	1E-4 <4E-7	RHR Availability ECCS Success Criteria ECCS Redundancy	Estimated Time to Core Melt Appear Longer Than Previously Expected, Thus Longer Times for Recovery
6) LOCA Loss of Containment Cooling	5E-6 <1E-7	RHR Availability Time Available for Recovery	Estimated Time to Core Melt Appear Longer Than Previously Expected, Thus Longer Times for Recovery

system. This results in containment overpressure failure followed by core melt.

Variations in this sequence in past PRAs have ranged from 0.1 to 50 per MRY. The staff's assessment of the frequency is at the high end. As in the case of the PWR, the NRC staff has proposed a set of rules intended to reduce the potential risk from this sequence, which included diverse scram, augmentation of the standby liquid control system (SLCS), automatic actuation of SLCS for new plants, reactor pump trip, and a reliability assurance program on the reactor trip system. Pending final regulatory resolution, this sequence category continues to have the potential of being dominant in BWR power plants.

Accident Sequence Category 2 - Transients with induced loss of reactor coolant system integrity and loss of the core cooling function. This sequence, as in the above Category 1, is initiated with a transient. The reactor protection system operates successfully and subcriticality is achieved leaving only decay heat to be removed. Failure of reactor coolant system integrity is in effect a transient-induced LOCA stemming from the safety relief valves (SRVs) failing to reclose after being opened by the effects of the transient. This releases primary coolant into the suppression pool. The loss of coolant in the primary must be replenished by the ECCS. If the ECCS also fails, the core will be uncovered and core melt results. Variations in sequence frequencies in past PRA work are from less than 0.2 to 70 per MRY. BWRs having the least redundancy in primary system makeup capability, such as some of the earlier designs using isolation condensers, tend to yield results in the higher end of the range. Other factors that contribute to variations are (1) variations in operating procedures for use of the automatic depressurization system to reduce primary system pressure, to allow low pressure core cooling systems to serve as a backup to failed high

pressure core cooling systems, and (2) variations in the SRV demand rates due to transients.

Accident Sequence Category 3 - Transients with induced loss of the containment cooling function. This sequence is similar to the Category 2 in that a transient-induced LOCA has occurred via a stuck-open SRV, and subcriticality has been achieved. In this category, the core cooling system is supplying makeup coolant, and decay heat is being successfully transferred to the suppression pool. However, suppression pool cooling has failed resulting in high water temperatures and inability to pump suppression pool water back to the reactor, which results in core uncovery and core melt.

The frequency range in this category is from less than 0.1 to 20 per MRY. This range stems from differences in the residual heat removal (RHR) system availability including variations in design of electrical and service water support systems. SRV demand rates due to transients is also a contributor to the variability in sequence frequency.

These sequences are long term in the sense that a significant amount of time is expected between failure of the heat removal system and core melt. In earlier PRAs that time was placed at approximately one day; more recent information indicates two days may be a better estimate. This implies a longer period prior to core melt to recover from the loss of suppression pool cooling and therefore greater chance of recovery than considered in earlier PRAs. The level of consequences for this sequence will vary depending on whether the core cooling pumps fail due to stress from pumping hot suppression pool water, thus resulting in core melt in an intact containment, or whether the core cooling pumps do not fail until after the containment has failed due to loss of suppression pool cooling. This latter case results in core melt in an already failed containment, thus higher consequences.

Accident Sequence Category 4 - Transient with loss of the core cooling function. This sequence is similar to Category 2 in that it is initiated with a transient and subcriticality has been achieved. It differs in that the SRV has not remained open, thus a transient-induced LOCA has not occurred. The loss of the core cooling function causes failure to remove decay heat.

Variation in this sequence category frequency range from 0.2 to 70 per MRY. As before, the variations stem from ECCS availability differences from plant to plant and difference in the operator procedures for operation of the automatic depressurization system (ADS). The forthcoming NRC position on the station blackout issue could also tend to reduce this sequence frequency.

Accident Sequence Category 5 - Transients with loss of the containment heat removal function. This accident sequence is similar to Sequence 3, with the exception that no transient-induced LOCA has occurred due to a SRV failing to close. The sequence frequency as predicted in past PRAs vary from 0.1 to 100 per MRY. As in Sequence Category 3, this variation stems from differences in RHR availabilities and uncertainties in the time available for recovery.

Accident Sequence Category 6 - LOCAs with loss of the containment heat removal function. This sequence is similar to accident Category 5 with the difference that the initiating event is a LOCA directly from failed piping or a component as opposed to a transient-induced LOCA. The predominance of evidence from past PRAs indicate this sequence category will be dominated by small LOCAs.

The sequence frequency ranges from less than 0.1 to 5 per MRY. This variation stems from the same plant differences and uncertainties as described in sequence Category 5.

The above accident sequence categories represent those which PRAs have indicated have the greatest potential for dominating the core melt frequency and/or risk. Which category actually dominates for any given plant depends on features of that particular plant. Some of the more important features which tend to determine the dominant accident sequences are listed as "major uncertainties" in Tables C-3 and C-4. It should be recognized that the specific component failure modes and human errors leading to these functional accident sequences can be expected to be quite different from plant to plant. It must also be recognized that the potential exists for new dominant functional accident sequences to be found in plants for which risk assessments have not been done if PRAs are performed on these plants. The above categories only reflect our current knowledge.

This discussion has not addressed accident sequences initiated by external events such as earthquakes, fires, and high winds. Such accident sequences do have the potential for dominating core melt frequency and/or risk. They result in similar functional accident sequences as described for transients and LOCAs. However, as discussed in C.5, the uncertainties can be expected to stem from quite different unknowns and plant features. The specific component and structural failures and the specific human errors that give rise to the functional accident sequences can be expected to be quite different.

As can be seen from the above discussion and from Appendix B, most of the uncertainties which give rise to large accident sequence frequency variations are being addressed in ongoing research or are being addressed in various regulatory initiatives. It is expected that these activities will result in new knowledge or plant changes which will tend to reduce the high end of the accident sequence frequency range, particularly for those cases where the sequence frequencies may be in the 100 to 1000 per MRY range.

Plant design variations will remain, and these variations can be expected to cause plants to have different dominant accident sequences. To allow generic studies to be done, which can contribute to the regulatory decisionmaking process, it is necessary to identify which plants can be expected to have similar dominant accident sequences.

Completion of the ongoing source term work will lay the ground work for changing the classes such that the criteria will be dominant accident sequences important to risk rather than core-melt frequency.

C.3 Supplemental Information on Insights Regarding Front Line and Support Systems

Those systems important to performing the functions identified in Section 5.3 and Appendix C.2 fall into two broad groups, often referred to as the front-line systems and support systems. The front-line systems are those which are designed to directly perform the above functions. Support systems are those which provide power, control, cooling, or other supportive needs to the front line systems.

Front-line systems differ from plant to plant. Furthermore, different vendors or utilities may give very similar systems slightly different names. Sometimes, the names reflect different uses of the systems; other times, the different name reflects no more than a different naming preference. The situation is further complicated by the fact that the same system may be given different names within a given plant to reflect different functions it serves when aligned for different modes of operation. Thus, for example, the low pressure injection system and low pressure recirculation system may represent nearly the same set of components only realigned to different water sources. Tables C-5 and C-6 provide some of the front-line systems currently being used in LWRs.

Table C-5

Typical Front-Line System for PWRS

Initiating Event	Function	Front-Line System
LOCA	Render Reactor Subcritical	Reactor Protection System
	Remove Core Decay Heat	High Pressure Injection System Low Pressure Injection System High Pressure Recirculation System Low Pressure Recirculation System Core Flood Tanks Auxiliary Feedwater System Power Conversion System
	Prevent Containment Overpressure	Reactor Building Spray Injection System Reactor Building Spray Recirculation System Reactor Building Fan Coolers Ice Condensers
	Scrub Radioactive Materials	Reactor Building Spray Injection System Reactor Building Spray Recirculation System Ice Condensers
Transients	Render Reactor Subcritical	Reactor Protection System Chemical Volume and Control High Pressure Injection System
	Remove Core Decay Heat	Auxiliary Feedwater System Power Conversion System High Pressure Injection System Power-Operated Relief Valves
	Prevent Containment Overpressure	Containment Spray Injection System Containment Spray Recirculation System Containment Fan Cooling System Ice Condenser
	Scrub Radioactive Materials	Containment Spray Injection System Containment Spray Recirculation System Ice Condenser

Table C-6

Typical Front-Line System for BWRS

Initiating Event	Function	Front-Line System
LOCA	Render Reactor Subcritical	Reactor Protection System
	Remove Core Decay Heat	Main Feedwater System Low Pressure Coolant Injection System Low Pressure Core Spray System Automatic Pressure Relief System High Pressure Coolant Injection System Reactor Core Isolation System
	Prevent Containment Overpressure	Suppression Pool Residual Heat Removal System Containment Spray System
	Scrub Radioactive Materials	Suppression Pool Containment Spray System
Transients	Render Reactor Subcritical	Reactor Protection System Standby Liquid Control System
	Remove Core Decay Heat	Power Conversion System High Pressure Core Spray System High Pressure Coolant Injection System Low Pressure Core Spray System Low Pressure Coolant Injection System Reactor Core Isolation Cooling System Feedwater Coolant Injection Standby Coolant Supply System Isolation Condensers Control Rod Drive System Condensate Pumps
	Prevent Reactor Coolant System Overpressure	Safety Relief Valves Power Conversion System Isolation Condenser
	Prevent Containment Overpressure	Residual Heat Removal System Shutdown Cooling System Containment Spray System
	Scrub Radioactive Materials	Suppression Pool Containment Spray System

As can be seen from these tables, often the same systems are used to perform different functions.

Support systems, which provide power, control, cooling or other supportive needs to front-line systems, also differ considerably from plant to plant. In general, they may be electrical systems, water systems or air systems. Most front-line systems will require some form of support systems to operate. Notable exceptions are the reactor protection system and core flood tanks which require none. Table C-7 displays a typical set of front-line systems and the support systems needed for each. This table was taken from a PRA of the Arkansas Nuclear One, Unit 1 power plant.*

Table C-8, taken from the same report, reflects that support systems often require yet other support systems. Thus, a typical power plant is a complex set of interdependent systems which perform a range of functions which are important to prevention of core melt and reduction of public risk.

C.4 Relative Importance of Systems

As identified in Chapter 5, risk importance measures have been used to evaluate a feature's potential for further reducing the risk and its importance in maintaining the present risk level. These measures have been applied to four plants and the results are shown in Figures C-4 through C-7.* Although this sampling of four plants reflects a single PRA method (that used in RSSMAP), it does not provide a basis for drawing broad generic conclusions. However, it is still instructive to review these results.

First, systems providing the greatest potential for risk reduction may not be the same as those which should receive highest attention in a reliability assurance program. Not surprisingly, the most important systems by either measure are not the same for different plants, since both plants'

Table C-7

ANO-1 Front Line vs Support System Dependencies

(Read Across)

FRONT LINE SYSTEMS / SUPPORT SYSTEMS	OFFSITE AC POWER	DIESEL AC GENERATORS	125V DC POWER	ENGINEERED SAFEGUARDS ACTUATION SYSTEM	EMERGENCY FEEDWATER INITIATION AND CONTROL SYSTEM	SERVICE WATER SYSTEM	INSTRUMENT AIR SYSTEM	INTEGRATED CONTROL SYSTEM	INTERMEDINTE COOLING SYSTEM	AC SWITCHGEAR ROOM COOLING	DC SWITCHGEAR ROOM COOLING	HIGH PRESSURE PUMP ROOM COOLING	LOW PRESSURE/SPRAY PUMP ROOM COOLING	NON-NUCLEAR INSTRUMENTATION POWER
REACTOR PROTECTION SYSTEMS														
CORE FLOOD SYSTEM														
HIGH PRESSURE INJECTION/RECIRCULATION	●	●	●	●		●				●	●	●		
LOW PRESSURE INJECTION/RECIRCULATION DECAY HEAT REMOVAL	●	●	●	●		●				●	●		●	
REACTOR BUILDING SPRAY INJECTION/RECIRCULATION	●	●	●	●		●				●	●		●	
REACTOR BUILDING COOLING SYSTEM	●	●	●	●		●				●	●			
POWER CONVERSION SYSTEM	●		●		●	●	●	●	●	●	●			●
EMERGENCY FEEDWATER SYSTEM	●	●	●		●	●				●	●			
PRESSURIZER SAFETY RELIEF VALVES														

NOTE: ALL REQUIREMENTS FOR DIESEL GENERATORS ASSUME LOSS OF STATION POWER.

Table C-8

ANO-1 Support vs Support System Dependencies

(Read Across)

SUPPORT SYSTEMS	OFFSITE AC POWER	DIESEL AC GENERATORS	125V DC POWER	ENGINEERED SAFEGUARDS ACTUATION SYSTEM	EMERGENCY FEEDWATER INITIATION AND CONTROL SYSTEM	SERVICE WATER SYSTEM	INSTRUMENT AIR SYSTEM	INTEGRATED CONTROL SYSTEM	INTERMEDIATE COOLING SYSTEM	AC SWITCHGEAR ROOM COOLING	DC SWITCHGEAR ROOM COOLING	HIGH PRESSURE PUMP ROOM COOLING	PRESSURE/SPRAY PUMP ROOM COOLING	NON-NUCLEAR INSTRUMENTATION POWER
OFF SITE AC POWER	■													
DIESEL AC GENERATORS		■	●	●		●				●	●			
125V DC POWER	●	●	■					●			●			
ENGINEERED SAFEGUARDS ACTUATION SYSTEM	●	●	●	■						●				
EMERGENCY FEEDWATER INITIATION AND CONTROL SYSTEM	●	●	●		■					●	●			
SERVICE WATER SYSTEM	●	●	●	●		■				●	●			
INSTRUMENT AIR SYSTEM	●		●			●	■		●	●				
INTEGRATED CONTROL SYSTEM	●	●	●					■		●	●			
INTERMEDIATE COOLING SYSTEM	●		●	●		●			■	●	●			
AC SWITCHGEAR ROOM COOLING	●	●				●				■				
DC SWITCHGEAR ROOM COOLING	●	●				●				●	■			
HIGH PRESSURE PUMP ROOM COOLING	●	●		●		●				●		■		
LOW PRESSURE/SPRAY PUMP ROOM COOLING	●	●		●		●				●			■	
NON-NUCLEAR INSTRUMENTATION POWER	●	●	●								●			■

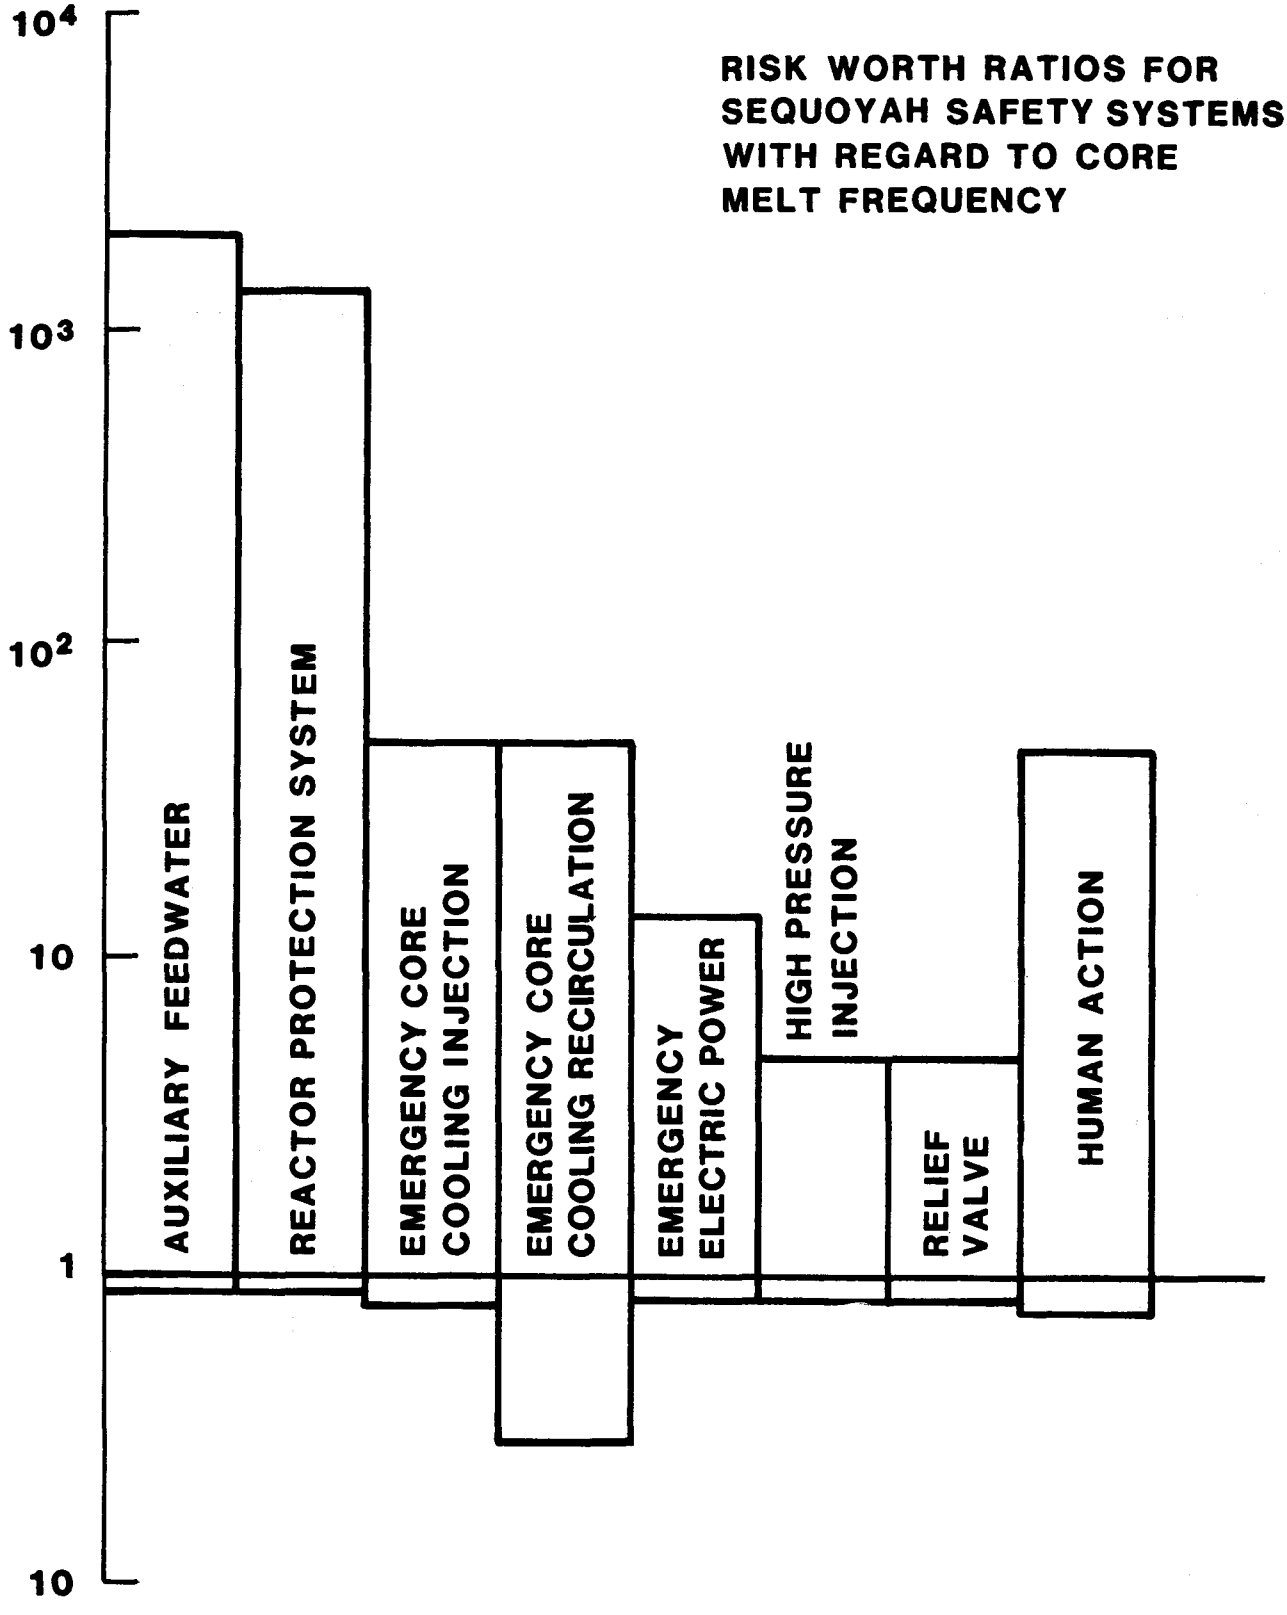

Figure C-4. Risk Worth Ratios for Sequoyah Safety Systems with Regard to Core-Melt Frequency

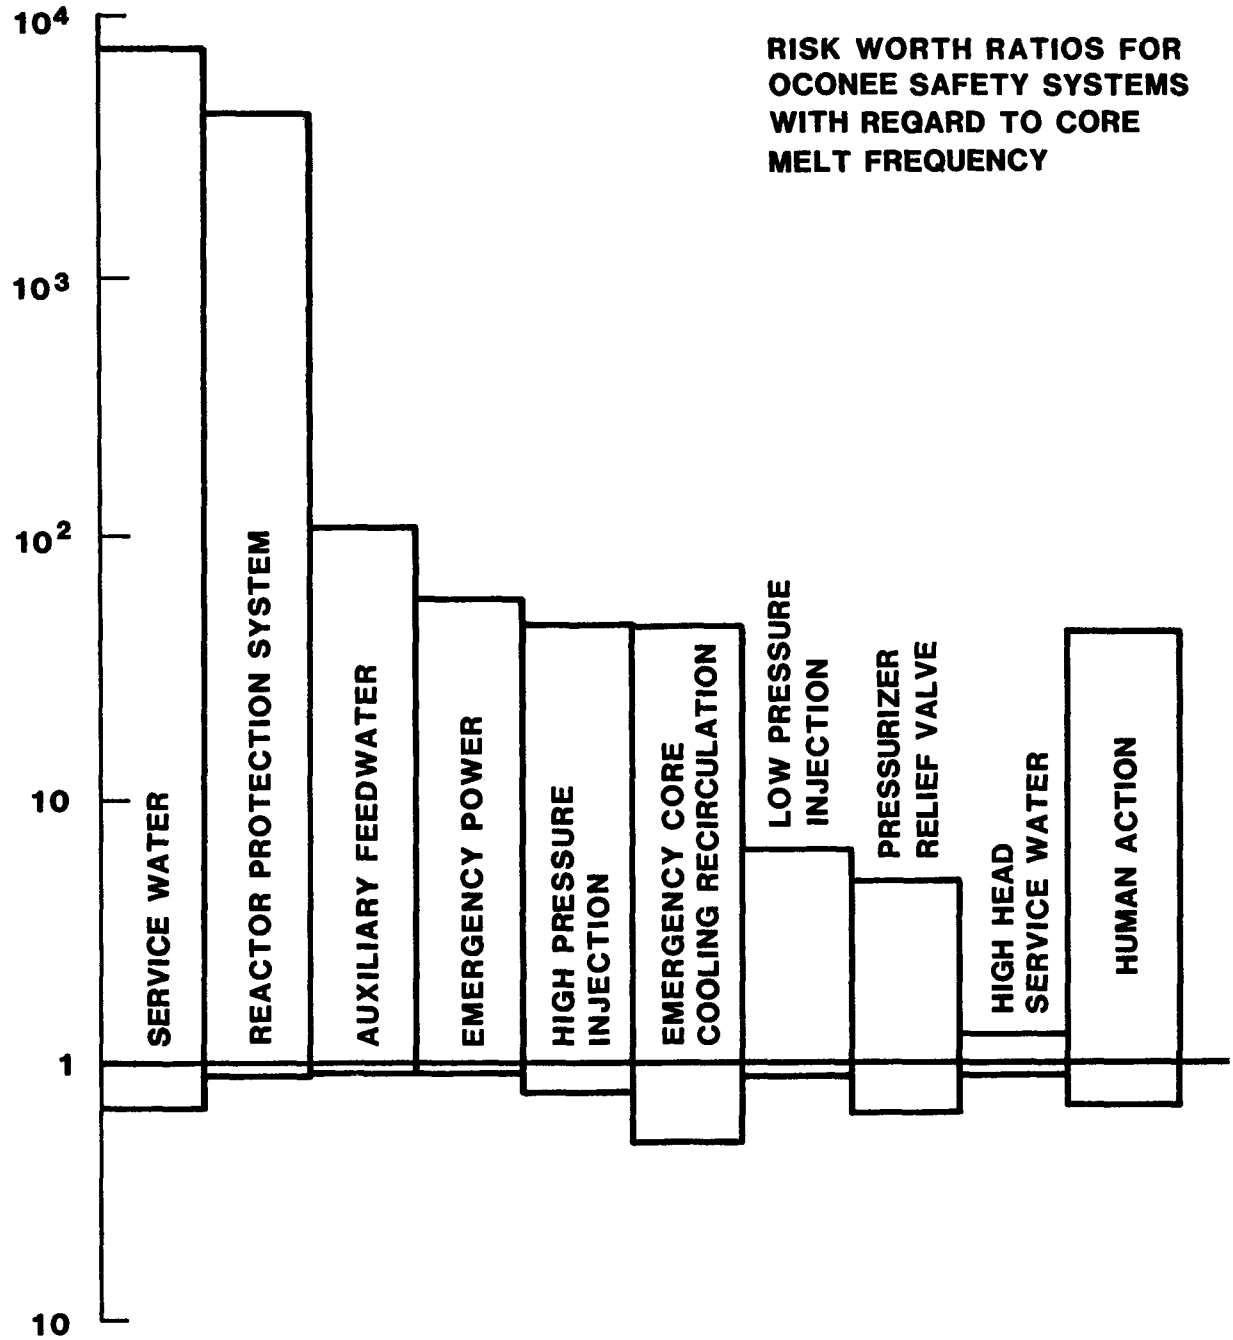

Figure C-5 . Risk Worth Ratios for Oconee Safety Systems with Regard to Core Melt-Frequency

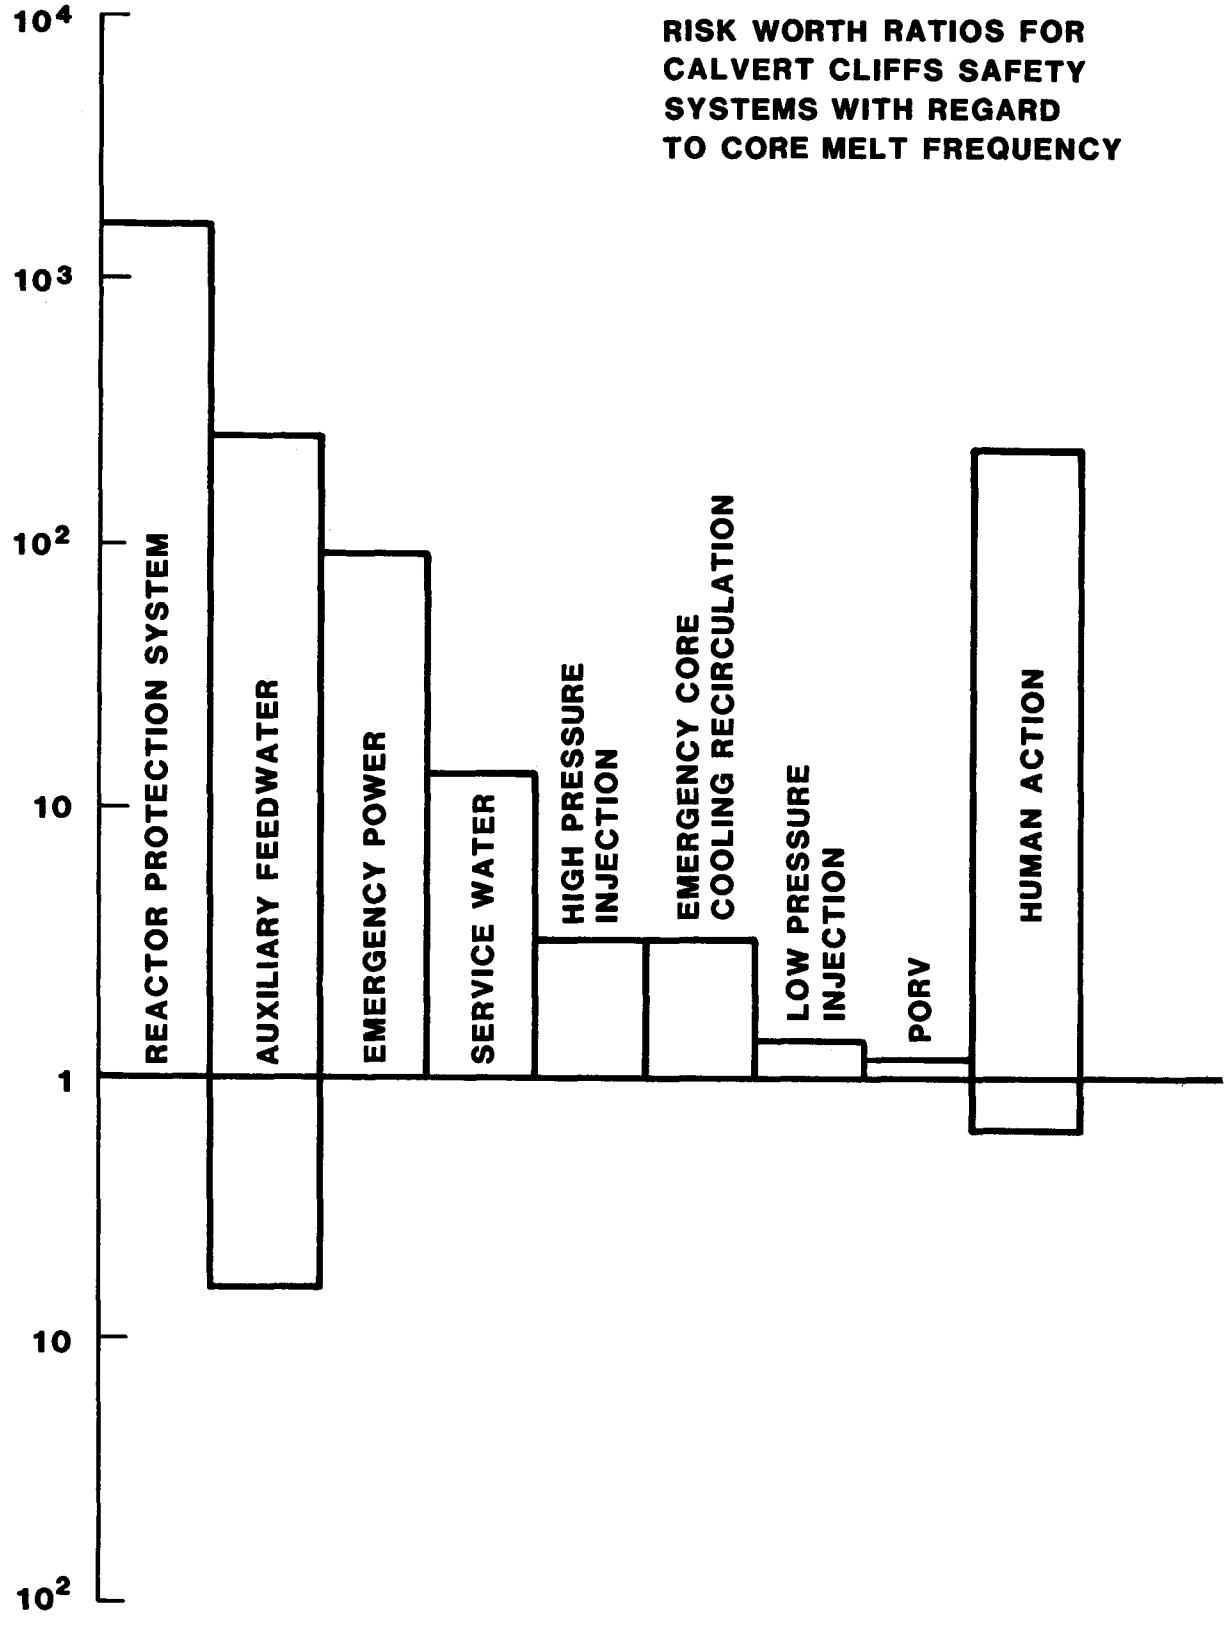

Figure C-6. Risk Worth Ratios for Calvert Cliffs Safety Systems with Regard to Core-Melt Frequency

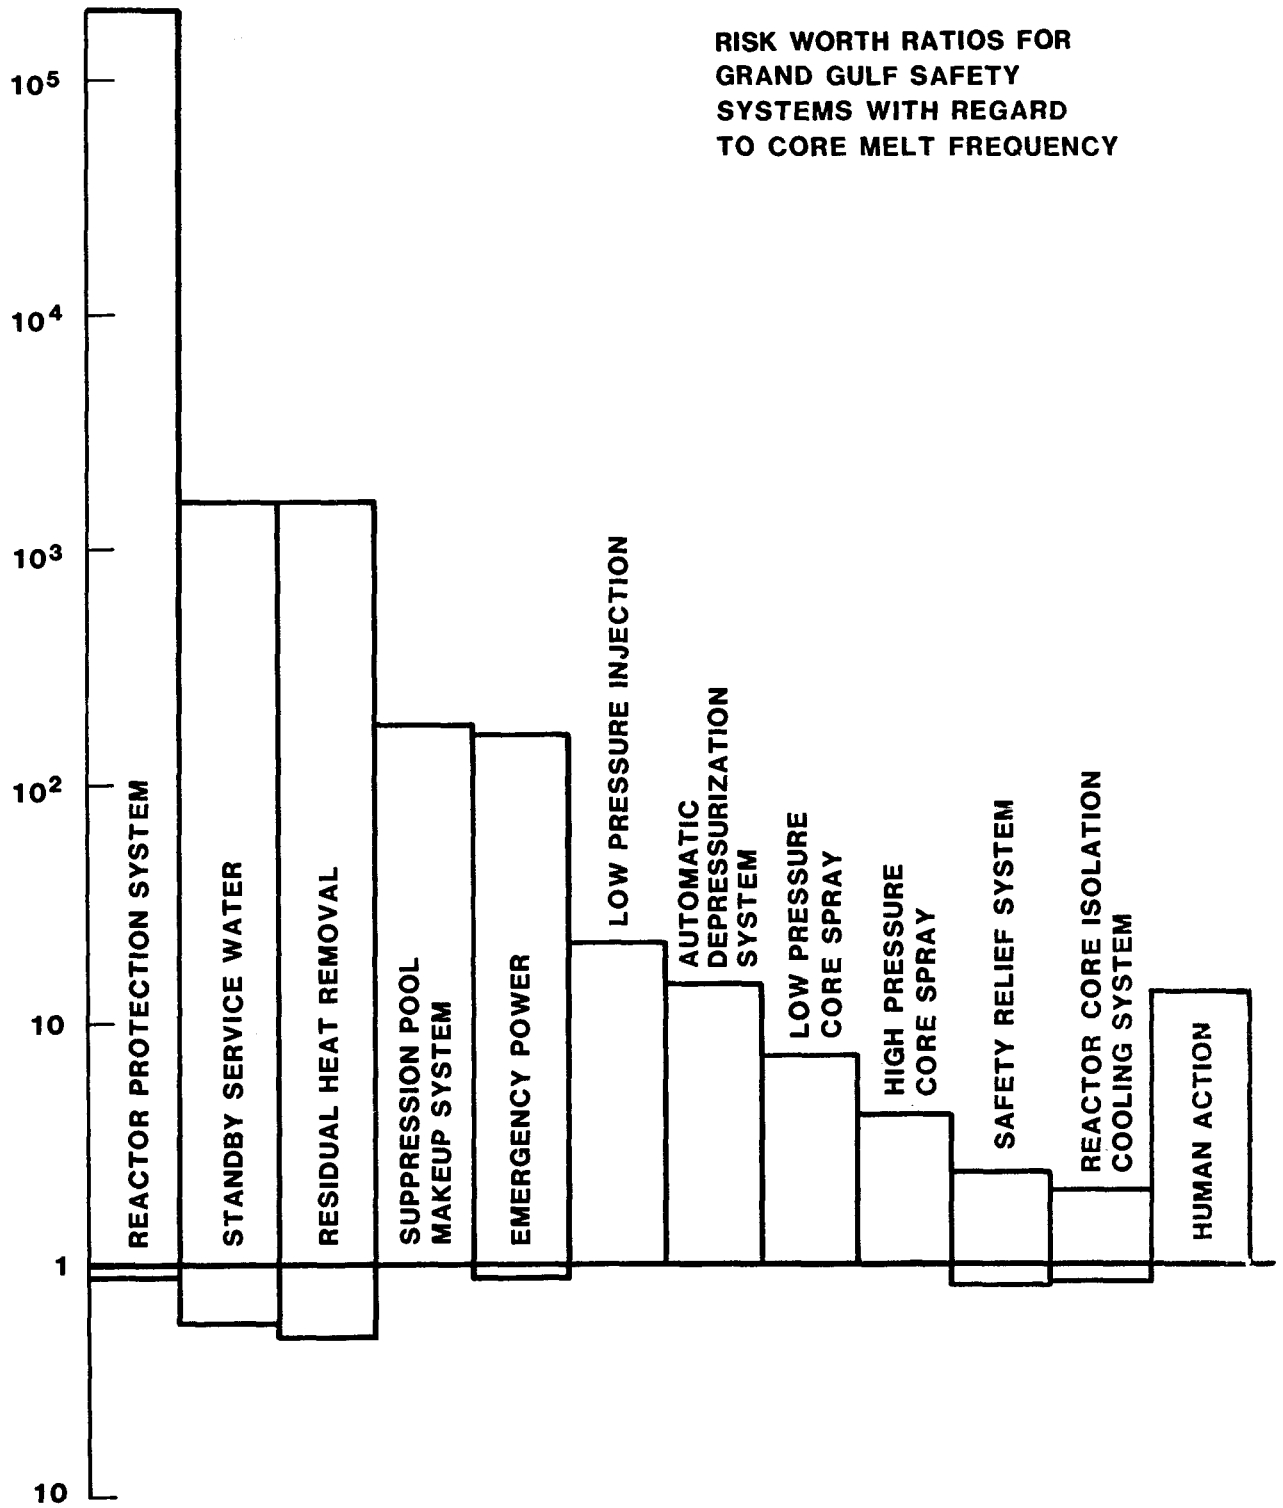

Figure C-7. Risk Worth Ratios for Grand Gulf Safety Systems with Regard to Core-Melt Frequency

dominant accident sequences and system designs vary significantly. Comparisons of three PWR plants do suggest some trends, however. From a risk achievement standpoint, the reactor protection system, auxiliary feedwater systems and high pressure injection systems do consistently appear high, suggesting these three systems would be prime candidates for safety assurance programs in PWRs. Electric power and service water appear to have rather consistent, moderately high, risk achievement worths.

On the other hand, the system which would appear to have the highest potential for risk reduction stemming from system improvement, based on these PRAs, is the auxiliary feedwater system on Calvert Cliffs. The AFWS on Calvert Cliffs has been improved since the completion of the RSSMAP studies.

A second study sponsored by the Office of Nuclear Reactor Regulation which addresses system importance is reported in draft form.* In this study, the Fussel-Vesely importance measure was used which provides a measure of relative importance of modest changes in component failure probabilities and initiating event frequencies on risk or core melt probability.

In this study, fifteen PRAs were analyzed representing different scopes and different methods. Their results provide additional insights about external events, and on the problem introduced by different scopes and different methods when using past PRAs and are shown in Figures 5-6 and 5-7 in the main text.

Five PRAs of Westinghouse plants were evaluated. Two of the PRAs (Zion and Indian Point; a total of four plants) considered external events. Even across the four plants on two sites, significant differences in the importance of the external events reflect important differences in plant and

site characteristics. Location of systems become as important as the systems themselves in the case of fire initiated events. At both Zion and Indian Point cables in the spreading rooms and cable tunnels were most significant. The impact of seismic events upon plants was also very site and design-dependent. Zion and Indian Point 3 have similar core melt frequencies due to seismic events, but Indian Point 2 has a significantly higher frequency. The dominant accident sequences were all different. Site differences and system location differences which impact the seismic and fire risk calculations also directly impact the relative importance of plant systems and features.

Also, in this study, emergency AC power availability was found important on two of the five Westinghouse plants, Surry and Indian Point 2. On the remaining plants, accident sequences not involving emergency AC power were found to dominate, thus for those plants modest changes in AC emergency power would have only modest or no effect on the risk on core melt probability. Safety injection systems availabilities were found significant only at Sequoyah and Surry. Specific design or site features of the plants can be identified which give rise to most of the apparent differences in the relative importance of plant systems. The difference in study scope (only two PRAs covered external events) also contributed to apparent differences in relative system importances.

Three PRAs of Babcock and Wilcox plants were also evaluated (Crystal River-3, ANO-1, and Oconee). This set is instructive to consider since they all were done using similar methods. All were sponsored by the USNRC and none covered external events. Two (Crystal River-3 and ANO-1) were done under the Interim Reliability Evaluation Program (IREP) and had more detailed system models than the third (Oconee) which was done under the more moderately funded RSSMAP.

Even in this case, where PRA methods, scope and assumption differences were minimized, but by no means eliminated, a high degree of consistency in relative importance of systems could not be found. For example, in Oconee, the low pressure service water, power conversion system, and reactor protection system were found to be most important. For Crystal River-3, the emergency AC power system, DC power system, and emergency feedwater system were found most important and for ANO it was the emergency feedwater system and high pressure injection system. With the exception of the reactor protection system and high pressure injection systems, the systems found important are not those provided by B&W, the nuclear steam system supplier, but were designed by architectural engineering firms or the utility itself.

In summary, it appears that the relative importance of plant systems to risk are highly dependent upon plant and site characteristics. They also differ in terms of which importance measure is used. Because the importance measures are based on PRAs, the scope, methods, and assumptions used in the PRA can also have a significant impact. Industry-wide conclusions as to the importance of systems are difficult to achieve with any degree of refinement. However, several systems do appear to be important with a reasonable degree of consistency. They are:

PWR: Auxiliary Feedwater System
High Pressure Injection System

BWR: High Pressure Coolant Injection System
Reactor Coolant Isolation System
Reactor Protection System
Residual Heat Removal System

C.5 Insights Regarding Scope and Depth

PRA analyses can be applied to issues of risk and safety at varying levels of scope and depth. The term "scope" refers to the breadth of the analysis, and is primarily determined by the study objectives, user perspective, and desired form of results. "Depth" refers to the degree of detail to which each item of the analysis is examined. The depth of an analysis must be consistent with study scope and objectives, but is also determined by time and manpower con straints, and the availability of data and information.

The PRA Procedures Guide (NUREG/CR-2300) defines three levels of scope for plant-specific PRAs, all of which are plant level analyses. Level 1 consists of a quantitative analysis of plant design and operation focused on accident sequences, their basic causes, and their frequencies. A Level-2 PRA would expand the scope of the Level-1 PRA by adding analyses of physical processes, radionuclide source terms, and containment responses, while a Level-3 PRA would include the environmental transport of the radionuclides and an assessment of the potential consequences. However, the scope of a PRA study may also be limited to analysis of one or a few related systems at one or multiple plants, or to a specific set of accident initiators or sequences. Some of the most noteworthy limited scope analyses to date are the auxiliary feedwater reliability analysis accomplished by NRC shortly after the TMI-2 accident,* and the PRA study of ATWS sponsored by the Utility Group on ATWS.* The scopes of these two analyses are below the Level-1 definition of the PRA Procedures Guide.

The question is often raised about the insights that can be gained by each type of study, particularly those limited in scope. The following sections address the particular types of information that can be gained from various levels of study.

C.5.1 System Reliability Studies

Analyses of individual systems utilizing PRA techniques such as fault tree analysis or GO methods* can be accomplished. For these studies it is necessary to define the context in which system operation is required and to make assump tions regarding interfaces with other systems. In particular, the operability or degraded status of support systems such as electrical power and component cooling systems must be identified, the type or types of initiating events to which the individual system must respond must be characterized, and the criteria for system success must be defined.

The results of a system reliability study may be expressed in terms of the probability that the system will satisfactorily perform its safety function given the boundary conditions imposed. Other useful information may be obtained concerning particular system weaknesses, such as single hardware or electrical component failures that could lead to system failure, the importance of test and maintenance actions to system reliability, and the relationship of the operator and procedures to system reliability. Additional information on system performance can be obtained through system uncertainty and sensitivity analyses.

These results and insights may be used in decisions regarding optimum system design configuration, system reliability/cost analysis, establishing proper system testing and maintenance schedules, and revising operating procedures or training. If reliability criteria have been predetermined, then the analysis results can be used to verify system compliance. However, the results and uses are limited in several ways. First, a system reliability study is performed within a limited analytical context, usually with a significant set of boundary conditions. The validity of the results is therefore limited to this context. Since the study is limited to one system, the impact of the results on plant safety are not readily

discernable. The benefits of a potential system improvement is measured in terms of a system availability increase, but cannot easily be related to reduction in plant core melt frequency or public risk. Although decisions to upgrade components within the studied system can be prioritized based on the reliability impacts and cost constraints, the basis for resource allocation on a plant-wide level would depend on an assessment of the importance of the system to overall plant safety.

System reliability studies, however, have been used both for regulatory purposes and internally by utilities. The NRC assessment of auxiliary feedwater system reliability resulted in the identification of several dependencies and single component failures. A number of modifications have been made at some plants to raise system reliability up to an acceptable level. As one example of an internal utility use of a system reliability analysis, some bid specifications for safety parameter display systems have required a minimum system reliability or availability. Depending on the measurement criterion, fault tree or GO models have been used to demonstrate compliance with the bid specifications.

C.5.2 Sequence Level Issue Studies

A second group of analyses performed using PRA techniques can be termed sequence level issue studies. Examples of issues that may be examined at the sequence level include anticipated transient without scram (ATWS) and station blackout. Such studies involve a number of plant systems performing an integrated set of safety functions in response to an initiator. An event tree is required to structure the analysis and tie together the system level fault trees or other system models. Since a number of systems are involved, the investigation of system dependencies and interactions is usually an important aspect of a sequence level issue study. The extent and manner in which support systems are included can vary depending on

study objectives and other factors. One approach is to treat support systems as operational for the best-estimate analysis, and then examine the impact of degraded support systems as a sensitivity issue. If a support system is determined to be critical to the analytical results, then it may be added to the sequence level model to provide a better basis for decisions.

The depth or level of detail of the analysis is important, because it can determine the reliability of the results. There is much more information on pump or valve reliability than there is on entire systems, so an analysis which has gone down to the individual component level is more likely to give a better estimate of system performance than an analysis that has relied on information about other similar systems. The configuration of individual systems is often unique to a plant, particularly when the details of support systems are included. However, if an issue is generic to a class of plants, it may be preferable to use system reliability data from a less rigorous depth of analysis. The issue of "generic" systems, however, must then be addressed to demonstrate that the analysis is valid.

The results of sequence level issue studies may be expressed in the form of sequence core melt frequency, or plant damage state frequency, or may be extended to include accident processes, containment response, fission product transport, and consequences. Insights may include rankings of system importance to sequence outcomes, opportunities for operator intervention, and alternative success paths with balance of plant systems. Because the context and scope of analysis is broader than an individual system reliability study, the uses of a sequence level issue study are more varied. Since it is possible to relate the impact of system hardware modifications to core melt probability, or perhaps to public risk, benefit/

cost questions may be considered. In particular, more efficient decisions can be made concerning resource allocation to improve safety, although plant-level decisions are not optimized when dealing at this sequence level. Operator actions and procedures may be considered at this sequence level, particularly with respect to alternate safety system or balance of plant systems to prevent core melt or mitigate accidents. However, the perspective is still limited to the sequence level scope.

The Utility Group on ATWS submittals concerning the ATWS issue are examples of extensive sequence level issue analyses that have been used in the regulatory process. Elements within the ATWS issue, such as the comparisons between different shutdown systems and configurations can be examined, and relative benefits gauged on a broader basis than system reliability. Sequence level analysis of ATWS has also been used internally by utilities to provide information concerning alternative design of plant systems and cost-benefit improvements.

C.5.3 Plant Level PRAs - Level 1

A Level-1 PRA includes a comprehensive set of initiators, event trees, and system models; a Level-2 PRA includes the systems analysis plus analyses of accident process, fission product transport within the containment, and containment responses; and a Level-3 PRA extends the analysis to fission product transport in the environment and a consequence analysis. For any of these three levels of analysis, the PRA may be restricted to internal initiators or it may include external events as well.

The Level-1 PRA stops when the frequencies of the core damage sequences have been determined. The analysis is focused on

plant design and operation and the potential core damage accident sequences, with their causes and their frequencies. Results are usually expressed in terms of dominant accident sequence frequencies and overall plant core damage frequency, often categorized by initiating event, or by resulting plant damage state. The status of containment safety features is sometimes included to give some feeling for accident sequence/containment interactions. The importance of a front line system or support system to overall plant core damage frequency can be traced by several different methods, as can the contribution of operator errors or test and maintenance actions. Most aspects of accident prevention can be evaluated and ranked on the basis of core damage contribution. Impacts of system changes can be determined, and priorities set for resource allocation, although the perspective of public safety or risk cannot readily be included at this level.

The usefulness of a Level-1 PRA has been questioned in the past, since it provides no release information at all. However, release categories can be reasonably assigned and consequences estimated based on the results of analyses at similar plants. From the utility point of view, the Level-1 PRA provides an answer of great value: how likely is it that the large investment in this plant will become a liability through a core damage accident. It also indicates the ways in which a potential core damage accident is most likely to occur, and what systems should be improved to reduce the probability of core damage at the least cost. In many plants the core damage frequency is dominated by only a few sequences, and improving the reliability of one or two systems or removing a common dependency may significantly reduce the expected frequency of core damage. A Level-1 PRA provides utilities with the probabilistic information required to perform an expected value calculation comparing costs of a core melt accident versus costs of modifications to increase plant safety. Many of the plants that have had PRAs performed have made modifications

that resulted in large reductions in potential core melt frequency.

From the regulatory viewpoint, the Level-1 PRA is very useful in determining the basic reliability of the reactor systems, the importance of these systems to accident sequence prevention, and how plant safety in terms of core damage frequency may be most efficiently increased. Whether or not a core damage accident results in a release, these types of accidents are to be avoided. If the accident progression is terminated before the core is damaged, or the systems are improved so that the accident does not happen, then the question of how much radioactivity may be released from containment becomes moot. It is in the interest of both the regulators and the utilities that accidents such as TMI-2 are avoided in the future. A Level-1 PRA addresses the frequency of these accidents, and allows decisions concerning safety to be made at the plant rather than system level.

Level-1 PRA results may be utilized in both relative and absolute ways. If a study of several plants shows the expected core damage frequency to be much higher for one of the plants than for the others, for example, the presence of some sort of problem at this plant may be indicated. Similarly, comparing the frequencies of similar or identical sequences at similar plants may indicate differences in reliability of certain systems from which approaches for improved reliability can be gained. For this kind of use, the absolute numbers of the results are not important. As long as the basic assumptions and methodology used were the same, the relative comparison of the results is valid. There are many other uses of Level-1 PRA results on a relative basis. For example, relative comparison of the results is useful for indicating which backfit at which plant may bring the most improvement, and where inspection at frequent intervals can enhance safety.

C.5.4 Plant Level PRAs - Level 2

The Level-2 PRA continues on from the Level-1 PRA to discover how the accident defined by each sequence will actually proceed: when the reactor vessel will melt through, whether the containment will fail and in what manner, the form of the fission products, what portion of the fission products are expected to be airborne in the containment at the time it fails, and so on. Whereas the end result of the Level-1 PRA is a list of sequences with a frequency for each; the end result of the Level-2 PRA is usually information about a number of release categories or groups. Sequences which result in similar releases are usually grouped for ease of treatment. The information for each category would include the total frequency, dominant sequences included in this category, release fractions, and time and energy of release.

In this way, core melt accidents can be categorized by their potential severity, which adds more perspective on plant safety. Insight into accident process mitigation and containment safeguards is also obtained. However, information concerning consequences and public risk is still not provided at this level.

The use of release groups or release categories has been the subject of some discussion. In cases where there are only a few dominant sequences the use of release categories may not be necessary, in which case it is best to avoid them. In other cases, there may be 20 or more sequences which all make significant contributions, and some sort of grouping is desired to make the problem tractable. When release categories specific to the reactor in question have been used and there has been constraint to make the dominant sequences fit into a very few categories, the categorization of the release information has not distorted the results. When the results have been forced into too few categories, into generic categories, or into categories derived for some other plant,

there are often problems. The RSSMAP studies, for example, were constrained to place their releases into RSS release categories for comparison purposes. Each RSSMAP study commented that the RSS categories did not fit their results and considerable distortion was required to utilize the RSS categories.

C.5.5 Plant Level PRAs - Level 3

A Level-3 PRA is a Level-2 PRA for which the analysis team has gone on to determine the offsite risk. While the release definitions which the Level-2 PRA generates give a general indication of the severity of the accident, quantitative results for the offsite risk can be obtained only by continuing on to the consequence analysis. A Level-3 PRA analyzes the transport of radionuclides through the environment and assesses the public health and economic consequences of an accident (e.g., early fatalities or latent cancer fatalities) and combines this with the probabilistic results of the Level-2 PRA. There are many examples of Level-3 PRAs: Zion, Limerick, and Indian Point are among the more recent ones. The RSS was a Level-3 PRA, but differed from most others in that a series of generic sites was used with systems analyses for two specific reactors to estimate the societal risk for all the reactors in the country.

One of the main uses of the RSS results was to compare risk to the offsite public from a reactor accident to the risks from other sources. Current interest lies in determining how various reactors measure up to the proposed safety goals. Results from a Level-3 PRA can be used in a variety of other ways; for example, in studying the efficiency of different evacuation plans, or indeed, the need for any evacuation at all.

The benefits and potential uses of a Level-3 PRA are obviously greater than for a Level-2 or a Level-1 PRA, as would be expected. However, the effort involved is also considerably greater, and whether or not this additional effort is justified depends upon the reasons for which the PRA was undertaken. If the goal is to determine the effect of the reactor on the health and safety risks to the surrounding population, then a Level-3 PRA is desirable. On the other hand, if the goal is to find out where a limited amount of resources may be most efficiently applied to decrease the frequency of a core damage accident, then a Level-1 PRA will suffice. The scope of a PRA, therefore, must be consistent with the objectives of the study. When the scope and depth of a PRA are properly defined, then the PRA results can provide very useful information and insights into reactor safety, both from utility and regulator viewpoints.

C.6 General Insights Regarding Areas Amenable to Improvement

C.6.1 Introduction

Once a safety concern is identified, a variety of possible solutions have to be evaluated. The NRC is confronted each year with a variety of proposed modifications to plant design or procedures. Many proposed changes have direct or indirect impact on plant availability. Since utilities must work with a finite resource base, some evaluation of the alternatives based on either financial or safety considerations is necessary. PRA has the potential to become the model by which the safety significance of proposed changes can be evaluated and prioritized. (It could also be directed toward plant avail ability.) This application of PRA methods and results represents the most ambitious use of PRA to date. It calls upon and brings together all of the engineering applications associated with PRA. It also focuses on one of the most, if not the most, important decisions--how and when to spend a limited supply of safety-related capital dollars and labor resources.

This chapter will consider how PRA has been used in the past and is being used in the present to evaluate proposed safety improvements. We will explore the theme of improvement relative to three areas: (1) existing plant systems design and maintenance procedures, (2) plant operations, and (3) consequence mitigation add-ons. For each area, we will describe how PRA has provided useful insights, even though the uncertainties are in many cases very large. Finally, we will consider the usefulness of cost-benefit analysis for evaluating proposed safety improvements.

C.6.2 Use of PRA to Identify and Evaluate Plant Systems Modifications and Maintenance Procedures

Some of the plant-specific changes that have either been made or been committed to based on PRA insights are described in Table C-9. Many of these changes were made while the PRA was ongoing so that the resulting effects could be included in the final PRA results. It can be noted that insights which give rise to plant modifications were found for PRAs of all levels.* For example, several plant changes were made during the Reactor Safety Study Methodology Applications Program (RSSMAP). The RSSMAP PRAs were limited in both depth and scope of their analyses and the nature of their "insights" was rather broadly defined. Later studies which included more rigorous plant modeling, external events, and explicit consequence analyses have produced a wider variety of insights with finer resolution.

Although the PRAs in Table C-9 were not performed with the specific objective of defining plant modifications, the results indicated areas of potential vulnerability and appropriate changes were made. An example of a PRA specifically intended to identify safety concerns and define cost-beneficial fixes is the Big Rock Point PRA. It is worth stressing that the goal of the BRP PRA was not to demonstrate

Table C-9

Examples of Plant Modifications Made or Committed to

Based on PRA Insights

Plant	Plant Modification	PRA
Sequoyah	Procedures changed to insure upper compartment drain plugs replaced after refueling	RSSMAP*
Oconee	Procedure and hardware changes made to reduce frequency of interfacing system LOCA	RSSMAP
ANO-1	Station battery test scheduling changed to reduce common mode failure probability	IREP**
ANO-1	AC and DC switchgear room cooler actuation circuitry test procedure established	IREP
Millstone	Logic changes made to emergency AC power load sequencer to eliminate single failure	IREP
Indian Point	Upgrade of charging pump alternate shutdown power supply to reduce probability of RCP seal failure	IPPSS***
Indian Point	Replacement of manual valves with motor-operated valves in fan cooler service water lines	IPPSS

*Reactor Safety Study Methodology Applications Program
**Interim Reliability Evaluation Program
***Indian Point Probabilistic Safety Study

"acceptability" against a pre-defined standard, but to identify real plant-specific concerns and propose cost-effective fixes. The fact that Big Rock Point is one of the older, low-power commercial units undergoing systematic review by NRC to determine if the plant should make changes to meet current licensing requirements made these objectives and goals all the more noteworthy.

Table C-10 shows some results from the BRP PRA analysis of the risk outlier "limited feedwater during ATWS." Eleven different modifications were considered in the PRA; only four are shown in Table C-10. As can be seen, there is a wide range in the impact of the modifications on the core damage frequency resulting from ATWS events.

The risk-reduction potential of various plant system modifications are being evaluated also by the NRC-sponsored Severe Accident Risk Reduction Program (SARRP) and the Industry Degraded Core (IDCOR) program.* Both programs indicate that a factor of about 2 reduction in overall core-melt frequency may be possible with specific hardware and maintenance procedure modifications. These include auxiliary feedwater system modifications, improvement of emergency ac power systems, reactor protection system modifications, improved maintenance for interfacing system check valves and ice condenser floor drains, etc. The SARRP Program has found that further risk reductions require add-ons designed for a broader class of accidents.

C.6.3 Use of PRA to Investigate and Enhance Plant Operations

The recognized importance of the operations staff in maintaining plant safety has led to considerable activity among the regulators, utilities, and other industry groups to enhance the operator's ability to detect, diagnose, and respond to accident conditions. These activities have addressed a wide spectrum of issues, including instrumentation design, display

Table C-10

Big Rock Point Analysis of Potential Design Modifications to Address Limited Feedwater During ATWS

Modification	Core Damage Frequency Resulting from ATWS, Per Year (10-5)
Plant As-Is	2.7
Auto Recirc. Pump Trip	2.6
Auto Liquid Poison System (LPS)	1.9
LPS and High-Pressure Recycle	0.5
LPS, High Pressure Recycle, and Load Rejection	0.2

systems, emergency procedures, training, control room design, etc. Just as the results and models of PRA can be used to identify and evaluate plant-specific risk significant design features, they can also be used as a logical framework to systematically identify the specific operational issues of risk significance at a particular plant or group of plants.*

The NRC-sponsored Severe Accident Sequence Analysis (SASA) program considers the impact of operator actions as a means for managing the outcome of severe accidents. The accident scenarios which have been analyzed are those which have been identified by PRAs as being dominant contributors to the overall risk and in which some mitigative operator actions are possible. The analyses to date have been focused upon specific plants representative of three basic containment types, namely, a large dry PWR, an ice-condenser PWR, and a Mark I BWR. Some operator actions have been identified for these plants which will mitigate the consequences of specific accidents. Some of these actions are identified as being generic within that type of plant; however, others turn out to be plant specific due to peculiarities in a particular plant design or its operating procedures.

The following example illustrates how PRA insights can be used to define operator procedures during severe accidents. PRAs have found that for many PWRs operator inability to manually realign the ECCS from an injection to a recirculation mode is a significant contributor to risk. The need for realignment comes about when the water inventory in the refueling water storage tank (RWST) falls below a set level. Conservation of the water in the RWST provides the operator with more time to achieve cold shutdown before ECCS recirculation becomes necessary. As a result, SASA found that,* until significant fission products are released to containment, the operator should minimize the use of the containment

sprays (which draw from the RWST) and should use the containment fan coolers to control containment pressure.

A second example illustrates how insights gained from PRAs can be used to improve safety by directing the approach to operator training. PRAs have concluded that, if the reactor protection system fails during an ATWS scenario in a BWR, there is a high likelihood that the operator will fail to initiate the liquid poison injection system or to manually insert the control rods. Based on this insight, SASA program members at the Oak Ridge National Laboratory invited several operators to be tested during a simulated ATWS at the Browns Ferry simulator. In the first test, the operators made errors and the outcome was less than optimal. However, after repeating the procedure, the operators proved to be quite efficient in performing the necessary tasks.*

C.6.4 Use of PRA to Identify and Evaluate Plant Add-Ons for Consequence Mitigation

Consequence mitigation systems are intended to reduce offsite consequences (e.g., offsite health effects and property damage). Although some mitigation systems are designed to enhance the depletion of airborne fission products (e.g., containment sprays and recirculating air filtration systems), the majority are designed to protect containment from a particular type of failure mode. Thus, it is convenient to characterize the offsite risk in terms of dominant containment failure modes.

Below is a list of the ways in which severe accidents can threaten containment, presented roughly in order of the timing of containment failure relative to the state of core degradation:

1. Direct bypass--The initiating event causes containment to fail or to be bypassed. Examples include large earthquakes, steam generator tube ruptures, and check valve failures which cause primary system inventory to be released outside containment.

2. Failure to isolate--The containment isolation system fails to provide a leaktight boundary.

3. Pre-core-melt overpressurization--Failure to remove heat from containment as fast as it is being produced in the core region causes the containment to fail by steam overpressurization. The ECCS may subsequently fail because the pumps cavitate or because large structural deformations may damage the cooling lines.

4. In-vessel steam explosion--An explosive interaction between molten core materials and water in the lower plenum of the reactor pressure vessel destroys the vessel and causes containtment to be breached by a missile.

5. Ex-vessel "steam spike"--Containment fails as a result of rapid pressurization by steam when the molten core penetrates the reactor vessel and is rapidly quenched by water in the reactor cavity or on the containment floor.

6. Hydrogen burning--A widespread hydrogen deflagration, or a local detonation, causes containment failure at any time during the accident if airborne hydrogen concentrations are sufficiently high and flammability conditions are attained.

7. Post-core-melt overpressurization--Containment fails as a result of gradual overpressurization from steam and noncondensibles while the molten core is attacking the concrete basemat of the reactor cavity.

8. Thermal degradation--Thermal radiation from the hot core materials in the reactor cavity and/or hot gases from the decomposition of concrete raises the containment structural temperature beyond the point where integrity can be maintained. Leakage paths may develop through containment penetration seals.

9. Basemat meltthrough--The hot core materials melt through the concrete basemat.

The mitigation approach most likely to succeed for a given reactor depends upon which of these containment failure modes dominate the risk. This in turn depends upon both system and containment design characteristics--e.g., the reliability of containment cooling systems compared to core cooling systems, the potential for common mode failures of both, the containment structures capability (i.e., failure pressure times free volume), the potential for water ingression into the reactor cavity, the temperature capability of containment penetration seals, etc. The determination of dominant containment failure modes has higher uncertainties than the determination of dominant accident sequences. That is, because it depends not only upon an analysis of the systems but also upon most other aspects of PRA (i.e., containment loading analysis, containment response analysis, and fission product transport analysis).

In spite of the uncertainties, some general insights on containment failure modes and applicable mitigation approaches have emerged from the various PRAs and from NRC's severe accident research programs, particularly the Severe Accident Risk Reduction Program (SARRP). Within the near future, many more results are expected from the SARRP program and from the industry's IDCOR program. In the meantime, the following statements appear to be substantiated by results reported to date:

1. The majority of offsite risk for the strongest PWR containments is associated with post-core-melt containment overpressurization from steam. Most often, this occurs as a result of a loss of all AC power, and contain ment does not fail until at least 10 hours after the accident initiation. Applicable mitigation systems include a low-volume filtered vent or an AC-independent containment cooling or spray system.

2. The offsite risk profile for the less strong PWR containments, including subatmospheric containments and ice condensers, is more uncertain. Some NRC-sponsored analyses* indicate that early failures from hydrogen burning and/or ex-vessel steam spikes may be important. (Note that according to these analyses, hydrogen burn failures in ice condensers are not necessarily precluded by the glow-plug igniters now in place.) Thus, hydrogen control and containment water management may be effective mitigation approaches for these containments.

3. The dominant containment failure modes in the BWR containments are pre-core-melt overpressurization from steam and post-core-melt overpressurization from the combination of steam and noncondensibles. The majority of risk is associated with accidents in which suppression pool scrubbing of fission products is not guaranteed, and a significant fraction of that risk emanates from anticipated transients without scram (ATWS). Mitigation approaches could include a high-volume unfiltered vent (for the ATWS) together with those mentioned above for the strong PWR containments.

C.6.5 Cost-Benefit Analysis

It is often said, in nonnuclear areas of human endeavor, that an ounce of prevention is worth a pound of cure. If this

statement of value can be assumed to apply, at least qualitatively, to nuclear reactor safety, then it should be expected that a small portion of the expenditure for accident prevention should go toward consequence mitigation. Whether the appropriate ratio is 16:1 depends upon cost-benefit considerations. While the original formulator of the prevention-cure theory most likely did not perform a formal cost-benefit analysis, he or she must have relied upon some subjective feelings about costs and benefits gained from experience. In nuclear reactor safety, we have a more rigorous cost-benefit approach which relies upon PRA.

To perform a cost-benefit analysis, it is necessary to have a means for evaluating the financial worth of a reduction in risk. One approach was suggested in the original NRC safety goal formulation:

> The benefit of an incremental reduction of risk below the numerical guidelines for societal mortality risks should be compared with the associated costs on the basis of $1,000 per man-rem averted.

A difficulty with this approach is that the dollar figure is rather arbitrary and subject to dispute.* It also does not consider other benefits to society that would normally be included in a cost-benefit analysis, such as financial impacts.

Recently, analyses of the financial risk from reactor accidents have been performed by considering the actual monetary impacts that may be associated with each of the consequences.* In these analyses, the impact of health effects is inferred from the expenditure that society has traditionally been willing to make to prevent deaths (typically 10^5 to 10^7 dollars per death, based on data from traffic safety, cancer detection,

and medical treatment programs, as well as safety standards for other industries). Property damage impacts are obtained from computer analysis.* Onsite consequence impacts, such as the costs of replacement power and post-accident cleanup, are estimated from available industry data. The benefits (averted impacts) are usually time-discounted at a fixed percentage rate (e.g., 4%), except that the NRC normally does not discount health impacts.

In general, the uncertainty of the offsite impacts is much higher than that of the onsite impacts, since the former depends to a much greater extent on uncertain phenomenologies (i.e., containment response, fission product behavior, plume dispersion, dose effects, etc.). Nonetheless, even when the offsite impacts are evaluated very conservatively, they are found in almost every case to be dominated by the onsite impacts. From a purely financial point of view, therefore, a pound of prevention is worth considerably more than a pound of mitigation.

Results to date from the SARRP program and related programs indicate that the expected, or mean, financial risk from internal events is rarely more than $20 million per plant. At face value, this means that safety improvements costing more than that amount would not be cost-effective for most plants. However, the effects of external events (earthquakes, fires, etc.) have not been evaluated for most plants, and that factor could increase the ante. The relative importance of external events is, of course, highly plant-specific. Furthermore, secondary financial risks (e.g., the effect of a severe accident on the nuclear industry) have not been explicitly considered and could shift the balance toward a higher expenditure for safety.

There are large uncertainties in the analysis of costs versus benefits for safety improvements. Still, much useful information can be obtained from it. In many cases, a safety improvement can be ruled out with high confidence, regardless of the uncertainties, because the cost outweighs the benefit by a very large margin. In some cases, the benefit may outweigh the cost by such a large margin that the converse becomes true. Even if the difference between the benefit and the cost is within the range of uncertainty, it is possible to use the results of the cost-benefit analysis to rank each safety option relative to the others. Use of cost-benefit analysis in this manner appropriately places the emphasis on utilizing PRA insights rather than relying on the bottom line.

C.6.6 Closure

Identification of safety concerns and evaluation of possible solutions are well recognized and utilized applications of PRA. In several instances, potential plant vulnerabilities identified by PRA have resulted in actual plant design changes or procedural modifications. The changes implemented to date have been preventive in nature (i.e., their purpose is to reduce the likelihood of an accident involving core damage). Other application programs are using PRA results to identify and evaluate system add-ons and operating procedures that can mitigate severe accidents (i.e., reduce their consequences).

PRA has been most effective when used to provide <u>relative</u> rankings of safety issues and solution alternatives. This use of PRA places the emphasis on its strength, which is the determination of the relative importance of competing factors. Specifically, it relies upon PRA's ability to identify dominant accident sequences, dominant containment failure modes, and dominant contributors to financial risk.

C.7 Insights Regarding Reliability Assurance

A PRA presents a "snapshot" of the risk profile at a given plant at a given time. As time progresses, modifications to plant equipment or operational or maintenance practices can change the risk associated with the plant. Further, as operational data accumulates, the improved information base may suggest that generic failure rates used for some components should be modified or that there is a different potential for dependent failures than previously assessed. Thus, there is a need to update the analyses and to make the PRA essentially a "living" document that reflects the impact of plant modifications and acquired data.

Techniques are available, and discussed in Chapter 3 and Appendix B, which permit an analyst to measure the incremental effect of a degradation in a given safety function, system, or component. Analyses such as these permit the plant operator/owner and the NRC to focus inspection and quality assurance attention to those plant features which the PRA results indicate could significantly increase the calculated plant risk or core damage frequency derived by the PRA. The features identified by such an analysis may not necessarily be those that are major contributors to risk. Rather, they encompass those features which could become dominant, if they degrade significantly relative to the failure characteristics used in the analysis. Other importance measures are useful in identifying the important contributors to the assessed risk. These are useful in aiding decisions on where to improve the plant if a reduction in risk is desired.

The availability of an updated PRA would also make possible a means for inter preting the risk (or core melt frequency) significance of variations in component failure rates as determined by acquired plant-specific data. Similarly, plant models could be compared with operational occurrences to

assure that they reflect the best information on plant performance and interactions between systems and components. Use of probabilistic analysis techniques alone will not constitute an adequate reliability assurance program. PRA techniques can be applied to potential problems such as improper design, faulty installation, or improper specification of performance requirements only in a subjective manner. Thus, the use of PRA techniques must be coalesced with appropriate quality assurance and quality control measures in an integral process for a comprehensive reliability or safety assurance program.