NUREG/CR-5392
SCIE-NRC-373-98

Elements of an Approach to Performance-Based Regulatory Oversight

Manuscript Completed: December 1998
Date Published: January 1999

Prepared by
R.W. Youngblood, R.N.M. Hunt, E.R. Schmidt,
J. Bolin, F. Dombek, D. Prochnow

SCIENTECH, Inc.
11140 Rockville Pike
Rockville, MD 20852

N.P. Kadambi, NRC Project Manager

Prepared for
Division of Regulatory Applications
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
NRC Job Code J6040

AVAILABILITY NOTICE

Availability of Reference Materials Cited in NRC Publications

NRC publications in the NUREG series, NRC regulations, and *Title 10, Energy,* of the *Code of Federal Regulations,* may be purchased from one of the following sources:

1. The Superintendent of Documents
   U.S. Government Printing Office
   P.O. Box 37082
   Washington, DC 20402-9328
   <http://www.access.gpo.gov/su_docs>
   202-512-1800

2. The National Technical Information Service
   Springfield, VA 22161-0002
   <http://www.ntis.gov/ordernow>
   703-487-4650

The NUREG series comprises (1) technical and administrative reports, including those prepared for international agreements, (2) brochures, (3) proceedings of conferences and workshops, (4) adjudications and other issuances of the Commission and Atomic Safety and Licensing Boards, and (5) books.

A single copy of each NRC draft report is available free, to the extent of supply, upon written request as follows:

Address: Office of the Chief Information Officer
Reproduction and Distribution
Services Section
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
E-mail: <GRW1@NRC.GOV>
Facsimile: 301-415-2289

A portion of NRC regulatory and technical information is available at NRC's World Wide Web site:

<http://www.nrc.gov>

All NRC documents released to the public are available for inspection or copying for a fee, in paper, microfiche, or, in some cases, diskette, from the Public Document Room (PDR):

NRC Public Document Room
2121 L Street, N.W., Lower Level
Washington, DC 20555-0001
<http://www.nrc.gov/NRC/PDR/pdr1.htm>
1-800-397-4209 or locally 202-634-3273

Microfiche of most NRC documents made publicly available since January 1981 may be found in the Local Public Document Rooms (LPDRs) located in the vicinity of nuclear power plants. The locations of the LPDRs may be obtained from the PDR (see previous paragraph) or through:

<http://www.nrc.gov/NRC/NUREGS/SR1350/V9/lpdr/html>

Publicly released documents include, to name a few, NUREG-series reports; *Federal Register* notices; applicant, licensee, and vendor documents and correspondence; NRC correspondence and internal memoranda; bulletins and information notices; inspection and investigation reports; licensee event reports; and Commission papers and their attachments.

Documents available from public and special technical libraries include all open literature items, such as books, journal articles, and transactions, *Federal Register* notices, Federal and State legislation, and congressional reports. Such documents as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings may be purchased from their sponsoring organization.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, Two White Flint North, 11545 Rockville Pike, Rockville, MD 20852-2738. These standards are available in the library 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—

American National Standards Institute
11 West 42nd Street
New York, NY 10036-8002
<http://www.ansi.org>
212-642-4900

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, of any information, apparatus, product, or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights.

ABSTRACT

This report discusses an approach to performance-based regulatory oversight. One key issue in developing a performance-based approach is choosing a collection of performance measures that is highly results-oriented, and will support the capability to detect and act upon emerging performance problems before they lead to adverse consequences. A related issue is the role of institutional factors, and how to reflect institutional factors in a results-oriented, performance-based approach. These issues are explored through discussion of examples. Based on these discussions, an approach is recommended. The approach entails (1) careful formulation of a safety case, which shows what the challenges are to plant safety and what the plant capability is for responding to those challenges, (2) allocation of performance goals over elements of the safety case, (3) formulation of a “diamond tree,” which is an integrated, hierarchical presentation of hardware, human, and institutional performance areas that indicates how institutional performance supports the safety case, and (4) application of the diamond tree to select a set of performance measures that is as results-oriented as possible, given the levels and kinds of performance needed in order to support the safety case, and the need to respond to emergent problems before adverse consequences develop.

CONTENTS

FIGURES

FIGURES (cont.)

TABLES

Table **Page**

EXECUTIVE SUMMARY

This report discusses the development of an approach to performance-based regulation. The term "performance-based" has been defined in the following way:

> A performance-based approach is an approach that establishes performance and results as the primary basis for regulatory decision-making, and incorporates the following attributes: (1) measurable parameters to monitor, with clearly defined, objective criteria against which to assess plant and licensee performance; (2) licensee flexibility in determining how to meet the established performance criteria that will encourage and reward improved operations; and (3) a framework in which the failure to meet a performance criterion, while undesirable, will not in and of itself constitute or result in adverse consequences.

Within an approach described some time ago by NEI, the NRC would review certain performance areas on an ongoing basis, and at any given time, the level of NRC involvement with details of plant operation would depend on the current assessment of licensee performance. Evidence of ongoing satisfactory safety performance would indicate that there is no current need for additional involvement of NRC staff with a licensee's operations. Evidence of performance problems would be grounds for increased NRC involvement with a licensee, perhaps beginning with briefings and proceeding thereafter to detailed inspections.

Implied in the above discussion is the idea that a performance basis for regulatory decision-making would supplant some current requirements, some of which are more invasive or more prescriptive than performance-based regulation would be. The hope is that a performance-based approach would be more efficient for licensees and for regulators than some of today's prescriptive regulations are seen to be.

One major focus of this report is the search for a workable balance between the third attribute and the first attribute called out in the above definition. Suppose that we tried to base a monitoring scheme at the level of performance of a particular safety function, such as removal of decay heat. An objective criterion on decay heat removal is easy to articulate and easy in principle to measure and report. (See the NEI paper [4] for an example.) Therefore, it is easy to satisfy attribute 1 (measurable parameters, objective criteria) with such an approach. Attribute 3 (no adverse consequences) is another matter. The question is whether successful removal of decay heat is good evidence of satisfactory safety performance, or conversely, whether it is a good regulatory approach to wait for an actual failure of decay heat removal before inquiring about steps to restore satisfactory safety performance.

Arguably it is not a good regulatory approach to presume satisfactory safety performance until safety functions are not only compromised but failed. Where, then, does one draw the line?

After years of work applying risk models to regulatory decision-making, it is fair to say that some important elements of the answer to that question are available. But before we simply begin applying risk models, a closely related but distinct issue needs to be addressed. This issue is suggested by examination of reports addressing concerns associated with various operational events, ranging from information notices (INs) to augmented inspection team (AIT) reports. A consistent theme in many of these reports is deficient performance in programmatic areas, or perhaps even perception of institutional weaknesses. Programmatic weaknesses can behave analogously to common cause failure mechanisms. This makes them important for two reasons: they are potentially important causes of failure, and their capacity to affect multiple trains creates the possibility that when they do manifest themselves, they do so in events having high conditional probability of adverse consequences.

While common cause failure (CCF) is certainly addressed in typical PRAs, the CCF potential of institutional factors is not normally addressed in any detail in PRA, not because of "PRA Quality" issues but because addressing issues in these areas has not been an emphasis of PRA technology. For purposes of performance-based regulation, whether risk-informed or not, we need a more comprehensive model of plant performance than PRA has generally provided. Steps towards such a model have been taken previously and are presented in the report.

It is argued that a useful litmus test for development in this area is that a model is comprehensive enough if it supports discussion of the significant findings of a comprehensive AIT report. To put it another way, a model is "comprehensive" if it addresses the performance areas that are typically called out in such a report, which generally go well beyond hardware issues. Given a significant precursor event, a typical PRA model will contain a high-level representation of an accident sequence that the event foreshadowed; but the model may not even contain all of the important components that fail in a significant event, particularly if instrumentation is involved, and those components that are modeled might be found after the fact to have had significantly enhanced failure probabilities relative to those assigned in the PRA, as a result of institutional factors.

The outline of the report is the following:

In order to sharpen our focus on certain issues, we begin in Section 2 with an example discussed previously by NEI, in a paper that presented NEI's proposed approach to performance-based regulatory decision-making and illustrated NEI's proposed approach in terms of an example based on reduced-inventory operation at shutdown. Specific examples, such as that provided by NEI, are considered an excellent basis for discussion of issues such as those mentioned above. Accordingly, a logic model was developed to support a more detailed discussion of the NEI example. This led to the conclusion that function-level monitoring of the kind originally proposed would not provide adequate ongoing assurance of satisfactory safety performance.

Next, in Section 3, a brief discussion of "diamond trees" is presented. As part of the current effort, a brief literature survey was conducted (summarized in Appendix A), and among the more immediately promising ideas to emerge from the search was the idea of the diamond tree. The diamond tree is a tool for organizing key performance elements into a hierarchy that turns out to be useful for developing a performance-based scheme. Whether or not the specific notion of a diamond tree carries forward into future work on performance-based regulation, the core idea of discussing performance in a hierarchical framework seems to be very useful, for discussion of regulatory approaches in general and even for analysis of significant operational events.

In Section 4, a process is outlined for developing a regulatory approach that is as performance-based as possible, consistent with resolution of the important issues discussed above. That is, the process is biased towards a performance-based approach, but in areas where the performance-based approach does not meet certain criteria, a need for other approaches is signaled. The general approach is supplemented with a few recommended rules of thumb for making particular choices. The approach is based on the idea that performance-based regulation is most naturally viewed as the implementation of a carefully-thought-out safety case, in which the important plant capabilities are identified and performance objectives are identified with them. This idea is not new, but is not characteristic of recent approaches to risk-informed changes to regulatory practice.

At this point in the report, having presented an approach based on the diamond tree, we return in Section 5 to the issue of the need to supplement typical PRA information in a satisfactory regulatory approach. We observe that many extant discussions of significant events dwell on performance aspects that are implicit, but seldom satisfactorily addressed, in PRA. A fairly recent AIT report is discussed, and its significant findings placed in correspondence with a diamond tree framework. A discussion is provided of the path by which the recommended approach would arguably have addressed this particular event before the fact.

A significant feature of the recommended approach is that it refrains from developing direct regulatory oversight at the level of institutional factors. This choice is made because it is perceived that regulatory oversight at this level would tend to eliminate whatever appeal performance-based regulation might have to licensees. It also seems clear that monitoring at this level would be inconsistent with the results-oriented emphasis of performance-based regulation called out in the definition given above. On the other hand, it seems that this area of performance is too important to ignore entirely in a performance-based scheme; indeed, it could be asked whether this area of performance is addressed adequately in the existing approach. It is arguable that this performance area would manifest itself adequately in the train-level figures of merit towards which the proposed scheme tends to drive; that appears to be the implication of the NEI approach (not that NEI recommended train-level monitoring, but that the NEI proposal assumes that institutional factors manifest themselves adequately in a scheme that monitors functional performance). It is also arguable that adverse institutional factors could lead to a significant event before the underlying situation is revealed by high-level functional indicators. This potential argues against neglecting these performance nodes just because they are not results-oriented enough to be considered “performance-based.”

It may be possible to address institutional factors in a non-invasive way through a scheme in which a *licensee process* would systematically address institutional factors, and the regulator’s focus of attention would normally be that licensee process, rather than observing such things as actual performance of maintenance actions. This notion is discussed in the concluding section of the report.

FOREWORD

The NRC Office of Nuclear Regulatory Research sponsored the work described in this report to explore ideas that could support the development of a performance-based regulatory process. The report is intended to be consistent with the discussion contained in SECY-98-144, “White Paper on Risk-Informed, Performance-Based Regulation,” June 22, 1998. A peer review within NRC staff has been conducted before this report was released for publication.

The focus of this report is on performance-based approaches and did not include risk-informed approaches which are considered separately. The report therefore does not take up such matters as the proper use of importance measures. Rather, it focuses on identifying which aspects of performance can be used to draw conclusions about plant safety, especially what kind of conclusions can be drawn from occurrence (or non-occurrence) of certain kinds of precursor events. The report also intends to explore the question of whether probabilistic risk assessments (PRAs), as conventionally carried out, address the kinds of performance issues that emerge from discussions of significant operational events.

The research (which began in July 1997) called for a literature search, with case studies to follow if promising approaches were found. The report concludes that widely-known analytical tools are available to deal with many aspects of the problem, but that more work would be needed in order to address performance issues. Accordingly, one of the ideas from the literature search (the diamond-tree approach) has been subjected to more detailed examination. This idea has been examined in the hope that it would help to sort out the issue of whether performance issues are adequately captured in conventional PRAs. The diamond tree’s characteristic ability to portray more transparently dependencies not normally modeled in a PRA appeared to provide the needed conceptual framework for identifying and exploring issues related to implementation of a performance-based regulatory approach. NUREG/CR-5392 is intended to be a scoping study. It is recognized that many important issues may not be addressed or that all relevant references may not be incorporated into the study. The following types of issues may need to be addressed if practical applications are to result:

- Estimating resources required to prepare and review diamond trees, or developing a more focused kind of diamond tree.

- Specifying the actions to implement performance-based concepts and the transitioning of such actions into the existing regulatory framework.

- Identifying which aspects of preparing or using diamond trees can be quantified.

- Developing the appropriate metrics that can combine quantitative and qualitative information to maximize the objectivity of the decisionmaking.

- Relating this work to other projects underway, such as the risk-based performance indicators project or the projects directed toward modifying or replacing NRC's SALP (systematic assessment of licensee performance) process.

- Differentiating among various groups of NRC licensees regarding the degree to which it would be appropriate to apply performance-based concepts.

We believe this report will be useful in the continued effort of developing a performance-based regulatory process.

John W. Craig, Director
Division of Regulatory Applications
Office of Nuclear Regulatory Research

ACKNOWLEDGMENTS

In addition to the primary authors, several colleagues contributed to the material summarized in Appendix A: John Bolin, Frank Dombek, and Dan Prochnow.

The authors also wish to acknowledge the guidance of the NRC Technical Monitor, N. Prasad Kadambi, whose vision in this area led us to consider a broader range of issues than might otherwise have been examined in this report.

ACRONYMS

ADS/SRV	Automatic Depressurization System/Safety Relief Valve
AFWS	Auxiliary Feedwater System
AIT	Augmented Inspection Team
ANS	American Nuclear Society
ASME	American Society of Mechanical Engineers
BNL	Brookhaven National Laboratory
CAM	Compliance Assurance Monitoring
CCF	Common Cause Failure
CCW	Component Cooling Water
CDF	Core Damage Frequency
CFR	Code of Federal Regulations
CRAD	Criteria Review and Approach Documents
CVCS	Chemical & Volume Control System
DHR	Decay Heat Removal
DNFSB	Defense Nuclear Facilities Safety Board
DPL	Daily Power Level
DSI	Direction Setting Issue
ES&H	Environment, Safety, & Health
ESF	Engineered Safety Feature
FMEA	Failure Modes & Effects Analysis
HAZOP	HAZard and OPerability study
HEU	High-Enriched Uranium
HPI	High-Pressure Injection
IPE	Individual Plant Examination
LERF	Large Early Release Frequency
LOCA	Loss of Coolant Accident
LR	Large Release
MPFF	Maintenance-Preventable Functional Failure
MPLD	Master Plant Logic Diagram
NEI	Nuclear Energy Institute
NPP	Nuclear Power Plant
ORNL	Oak Ridge National Laboratory
ORR	Operational Readiness Reviews
OSHA	Occupational Safety and Health Administration
PB	Performance-Based
POC	Performance Objectives and Criteria
PORV	Pilot (sometimes Power) Operated Relief Valve
PRA	Probabilistic Risk Assessment
PSA	Probabilistic Safety Assessment
PSAM	Probabilistic Safety Assessment and Management
PSM	Process Safety Management
PWR	Pressurized Water Reactor
QA	Quality Assurance

ACRONYMS (cont.)

QIP	Quality Improvement Plan
RCM	Reliability-Centered Maintenance
RCP	Reactor Coolant Pump
RCS	Reactor Coolant System
RHR	Residual Heat Removal
RMP	Risk Management Programs
RPV	Reactor Pressure Vessel
RVLIS	Reactor Vessel Level Indication System
SALP	Systematic Assessment of Licensee Performance
SBO	Station Blackout
SSC	System, Structure, or Component
TMI	Three Mile Island
USNRC	United States Nuclear Regulatory Commission

1 INTRODUCTION

1.1 Purpose of This Report

This report summarizes the results of a project to support development of a performance-based regulatory approach. The charter of the work reported here has been to explore performance-based regulatory approaches specifically, and not to perform a high-level review of regulatory practice in general, or review risk-informed aspects of recent work. The project began with a brief literature survey, identified some key issues, analyzed a specific example to illustrate the significance of those issues, and culminated in the formulation of steps in an approach to performance-based regulation.

In a recent paper [1], the staff identified certain characteristics of a performance-based regulatory approach:[*]

- There are measurable parameters to monitor acceptable plant and licensee performance.
- Objective performance criteria are established to assess performance.
- There is licensee flexibility to determine how to meet established performance criteria.
- Failure to meet a performance criterion must not result in unacceptable consequences.

That is, “performance” is assessed in terms of monitored parameters. Based on this, the present project is seen as a search for such parameters that actually bear on the current level of plant safety, and concurrently a search for a decision process that accepts these parameters as inputs and furnishes appropriate recommendations to the regulator as outputs.

In DSI-12 [1], most of the discussion is carried out with respect to *risk-informed*, performance-based regulation (as opposed to merely performance-based). The agency has a stated policy of applying risk insights to regulatory decision-making to the extent that this is practical. Correspondingly, a great deal of work is already being done both by industry and by NRC staff in development of risk-informed approaches. This project is not intended to replicate those efforts, insofar as they are aimed at reflecting risk insights in regulatory priorities. Rather, the present project is aimed at the “performance-based” portion: understanding how regulatory oversight

[*] The following excerpts are from definitions presented in a more recent NRC memorandum [2]. These are generally consistent with DSI-12 [1].

A “risk-informed” approach to regulatory decisionmaking represents a philosophy to be used in all regulatory matters whereby risk insights are considered together with other factors, such as the basis for current regulations, engineering analysis and judgment, the defense-in-depth philosophy, and preserving adequate safety margins. These integrated elements are used to establish requirements that focus licensee and regulatory attention on design and operational issues commensurate with their importance to public health and safety.…Where appropriate, a risk-informed regulatory approach can be used to reduce unnecessary conservatism in deterministic approaches, or can be used to identify areas with insufficient conservatism and provide the bases for additional requirements or regulatory action.

A performance-based approach is an approach that establishes performance and results as the primary basis for regulatory decision-making, and incorporates the following attributes: (1) measurable parameters to monitor, with clearly defined, objective criteria against which to assess plant and licensee performance; (2) licensee flexibility in determining how to meet the established performance criteria that will encourage and reward improved operations; and (3) a framework in which the failure to meet a performance criterion, while undesirable, will not in and of itself constitute or result in adverse consequences.

might appropriately be based on licensee performance in particular areas, rather than verification of licensee compliance with prescriptive regulation. Specifically, this project responds to Commission guidance to include, where practical, performance-based strategies in the implementation of risk-informed regulatory decision-making processes. The issues in this project have to do with choosing performance measures that actually support the findings that need to be made, while permitting achievement of the improvements in efficiency that are the motivation for the undertaking.

Since the Commission also directed the staff to consider ideas from other industries, a certain amount of this has been done, as summarized later on.

It is not the purpose of the work reported here to argue for or against performance-based regulation. In this report, we essentially assume that performance-based regulation has major benefits, and have tried to show how to identify approaches that realize those benefits while satisfying certain basic requirements that regulatory oversight should satisfy.

1.2 Key Issues

1.2.1 Leading Indicator Issue

One desired characteristic is that failure to meet a performance criterion ought not to lead to "adverse consequences." Given a collection of performance indicators, it is necessary to ask whether it is "leading" enough in general: whether its signal of underlying problems is reasonably sure to arrive in time to prevent adverse consequences. A trivial example of an unsatisfactory performance criterion would be a cladding temperature of 2100 degrees F; if this occurs, the plant is in a dangerous condition, and one would hope that a flag would be raised before so dangerous a precursor as this occurred.

This issue is closely related to the concept of "margin." The desired characteristic is to have margin between the failure to meet a performance criterion and adverse consequences, where margin is understood to refer both to conditional probability and to key physical parameters such as RCS temperature. Later, we will discuss an example in which a performance criterion has considerable margin as measured by RCS temperature, but perhaps not enough margin as measured by conditional probability of adverse consequences.

Support systems may be shared between "before" systems and "after" systems. This is precisely the kind of link that corresponds to the prospect of getting the adverse consequence along with tripping the monitoring threshold.

A facet of this issue is that regardless of the efficacy of the proposed monitoring of systems whose failure leads to failure to meet the monitoring criterion ("before" systems), systems that come into play after reaching the monitoring criterion ("after" systems) are not addressed directly by the monitoring scheme, essentially by definition.

In this report, we will refer to the need to resolve this issue as the "leading-indicator rule."

1.2.2 Scope of Conventional Plant Models

The focus of this project is not risk-informed in the sense of setting relative safety priorities based on risk analysis, but we will be obliged to consider models of plant safety, including risk models, in order to assess whether a given performance-based scheme is appropriate. However, most widely-used models do not address many issues that are important for this application.

It is well known that PRAs do not typically model everything that is important to safety. Developing this point in its most general form is beyond the scope of this brief report, but many discussions have been given previously. It is emphasized that this comment is not intended to disparage PRA "quality" or to insinuate that better PRAs would actually model everything important. Rather, the point is that PRAs are not ordinarily intended to identify everything "important," but are simply aimed at identifying the most likely (probabilistically dominant) causes of failure; in doing this, they tacitly credit performance in many areas in which it is reasonable to hope for good performance, and important to achieve it, but in which it is not guaranteed that the desired levels of performance will be achieved.

This idea is perhaps best illustrated by examples. One example is piping failures. A typical PRA model of a fluid system train will not explicitly display piping failure as an important event. This is not to say that piping failure would be unimportant, but only that piping failure is not expected to dominate functional failure, unless consequential effects cause additional significant failures. (If there appears to be a significant potential for this, then a focused examination of selected instances is typically undertaken, if this is within the scope of the PRA.) A more immediately relevant example is emergency lighting. It is a rare PRA that will turn out a quantified cut set explicitly showing multiple failures of emergency lighting fixtures. On the other hand, multiple failures of emergency lighting fixtures would be considered significant if they were found to have occurred; and in fact they have been found to occur. [3]

Let us pursue the example of multiple failures of emergency lighting. Suppose that in a particular plant, this occurs because of a significant weakness in implementation of the required program of surveillances of these units. Two questions come immediately to mind: how safety-significant is this multiple failure (e.g., what conditional frequency of core damage would the Accident Sequence Precursor program derive for this event at a particular plant), and perhaps more fundamentally, is this programmatic weakness linked to other programmatic weaknesses through a shared cause in the form of some institutional factor? And if so, are there accident sequences whose probability would increase significantly as a result of the coexistence of these various programmatic weaknesses? Arguably, programmatic weaknesses have to be considered potentially very significant; but PRA, although it is formally capable of generating a fairly complete set of functional accident sequences, does not convincingly model the potential effects of programmatic weaknesses on the frequencies of these functional accident sequences.

The point is that for purposes of performance-based regulation, whether risk-informed or not, we need a more comprehensive model of plant performance than PRA has generally provided. Steps towards such a model are presented later. It is argued that a useful litmus test for development in this area is that a model is comprehensive enough if it supports discussion of the significant findings of a comprehensive AIT report. Given a significant precursor event, a typical PRA model will contain a high-level representation of an accident sequence that the event foreshadowed; but the model may not even contain all of the important components that fail in a significant event, particularly if instrumentation is involved, and those components that are modeled might be found after the fact to have had significantly enhanced failure probabilities relative to those assigned in the PRA, as a result of institutional factors.

1.3 Outline of the Report

As mentioned above, the current project began with a literature search. Appendix A of the present report is adapted from a letter report documenting that search. Appendix A also raises some issues in a preliminary way, including those mentioned above in 1.2. In order for any approach to performance-based regulation to be considered successful, it arguably needs to address these issues. The discussions presented in Appendix A are the starting point for the main body of this report.

Section 2, “Development of Issues,” presents a specific and detailed discussion of certain issues that were raised in Appendix A and summarized above. This is done with reference to an example. As explained in Appendix A, NEI [4] chose to illustrate their recommendation with an example based on a particular safety function needed during shutdown. A simple model is developed partially analyzing that situation, and that model is then used to illustrate certain issues in Section 2. The model itself is discussed briefly in Appendix B.

Section 3, “The Diamond Tree,” summarizes previous discussions of diamond trees and provides specific examples. This is done as a prelude to discussion of an actual approach to performance-based regulation. Whether or not all of the specific aspects of the diamond tree carry forward into future work on performance-based regulation, the core idea of discussing performance in a hierarchical framework seems to be very useful for discussion of regulatory approaches in general and even for analysis of significant operational events. In Section 4, a process is outlined for developing a regulatory approach that is as performance-based as possible, consistent with resolution of the important issues discussed in Section 2. That is, the process is biased towards a performance-based approach, but in areas where the performance-based approach does not meet certain criteria, a need for other approaches is signaled. The general approach is supplemented with a few recommended rules of thumb for making particular choices.

Before concluding the report, we return in Section 5 to the issue of the need to supplement PRA information in a satisfactory regulatory approach. We observe that many extant discussions of significant events dwell on performance aspects that are implicit, but seldom satisfactorily addressed, in PRA. A fairly recent AIT report [5] is discussed, and its significant findings placed in correspondence with a diamond tree framework. Arguably, any approach to performance-based regulation should be able to confront significant operational events, and preferably be able to show how it serves to prevent recurrences by working against at least some of the significant failures occurring in the events; the recommended approach passes this test.

2 FACETS OF THE “LEADING INDICATOR” ISSUE

The overall purpose of the report is to drive towards an approach to regulatory oversight that is as performance-based as it can be in light of certain issues. The purpose of this particular section is to develop some of those issues by discussing a realistic example in enough detail to provide food for thought. The process recommended in a later section for identification of performance-based measures is essentially a formalization of the discussion in this section.

2.1 Shutdown Example

In order to discuss certain issues, a simple model has been developed, keyed to the example provided by NEI as an attachment to their paper. This section uses the simple model to illustrate key points. The NEI example considered the safety function “heat removal during reduced-inventory operations”; the NEI proposal is essentially for the regulator to monitor the satisfactory performance of this function according to how often mishaps occur in which the RCS heats up to a temperature appreciably (but not dangerously) greater than the desired upper operating temperature (140 F in these conditions).

This is a very pure form of “results-oriented” performance monitoring: monitoring the kind of operational events that a plant experiences. In Appendix A, it is pointed out that keying on particular kinds of operational events (such as loss of shutdown heat removal) is essentially keying on precursor events, where the notion of “precursor” is broadened a bit relative to its ASP meaning to include events having enhanced, but not necessarily high, conditional probability of core damage. Note that for purposes of this formulation, institutional failures count as precursors. Not experiencing precursor events is argued to correlate with good licensee performance, while experiencing one or more precursor events may correlate with performance problems.

Although the NEI proposal did not explicitly couch the argument in terms of a quantified frequency of observed events, it is natural to discuss their proposal in those terms, and it is believed that doing this does not significantly modify the intent of their proposal. Rather, making the frequency explicit simply helps us clarify the risk implications of one strategy or another.

As in all PRA work, the numbers should not be over-interpreted. However, for present purposes, we need to be able to interpret the numbers, at least on an order-of-magnitude basis. Therefore, in the example provided below, the frequency units are referenced to essentially perpetual reduced-inventory operation. This makes it convenient to think about what the numbers mean (compared to full-power operation), at the cost of precluding a simple practical approach to avoiding heatup in this mode, namely, limiting the dwell time. That is, the CDF contribution coming out of this model is on the order of E-4 per year, meaning that on a per-hour basis, it is loosely equivalent to a full-power level of risk.

The present purpose is to illustrate the process and the issues with a quasi-realistic model, and not to present an actual case for a real plant. The present development was accordingly carried out by adopting portions of a plant model for which initiating event fault trees had been developed, significantly modifying the fault trees, and finally embedding them into an event tree tailored to address the present issues. In other words, this model, while having details and features that we believe are appropriately illustrative for present purposes, does not present a complete picture of risk during reduced inventory operation, and does not closely correspond to any particular plant. Many simplifying assumptions were made, including an assumption that decay heat is fixed at a level corresponding to that obtaining soon after shutdown. The most significant assumptions are presented in Appendix B.

The model has the following properties.

Not all events leading to loss of RHR were modeled, but an illustrative range was considered:

- loss of all AC
- loss of component cooling water
- loss of the front-line RHR function itself

Logic expressions were used not only for the responding systems, as they are in essentially all PRAs, but also for the initiating events themselves, and these expressions were linked with the expressions for the responding systems. The point of doing this was to support a detailed examination of such questions as the following:

- what kinds of events dominate loss of RHR,
- what kinds of events are more likely to lead to core damage,
- what actual performance elements are required to respond to each initiating event in order to prevent core damage
- what the conditional probability of core damage is for each initiating event, and how this varies with the character of the initiating event.

The results of the model are as follows:

Table 1. Top-Level Results

	Initiating Event Frequency (Events/Yr)	**Core Damage Frequency (Events/Yr)**	**P(CD|IE)**
Total	4.81E-02	1.59E-04	3.31E-03

The top-level result of the model is that the overall CDF is 1.59E-4. The sum of the initiating event frequencies, defined to correspond essentially to loss of RHR for more than a few minutes (so as to cause the RCS to heat up appreciably), is 4.81E-2. This gives an overall conditional probability of core damage (conditional on having gotten one of these initiating events) of 3.31E-3, determined by dividing the CDF by the IE frequency.

This conditional probability is an averaged quantity, subject to vagaries of definition, modeling, and level of aggregation. Consider the following more detailed display of results:

Table 2. Intermediate-Level Results

		Initiating Events (IE) (events/year)	**Core Damage (CD) (events/year)**	**P(CD/IE)**
	Blackout Sequences	1.34E-04	1.00E-06	7.50E-03
	Loss of CCW Sequences	3.77E-04	1.20E-05	3.18E-02
	Loss of Front-Line RHR Function Sequences	4.76E-02	1.46E-04	3.07E-03
Total		4.81E-02	1.59E-04	3.31E-03

We see that the conditional probability of core damage is a strong function of what kind of initiator occurred. For the particular model considered, we also see that the category of event that seems to dominate the initiating event frequency has the lowest average conditional probability of leading to core damage.

This suggests that a monitoring scheme pegged at the level of loss of RHR will spend more time looking at events that have a lower conditional probability of core damage (a situation reminiscent of the false-positive problem in medical diagnosis). With this in mind, go down one more level in detail, and consider the dominant minimal cut set emerging from the sequences modeled. That cut set contains two pump failures (yielding the initiating event) and a series of non-recovery events (i.e., failures to recover) that led to the overall cut set frequency. The frequency of the initiating event portion is ≈2E-3, and the conditional probability associated with the rest is ≈5%. On a single cut-set basis, this instance of the loss of RHR initiator has as high a conditional probability as any other initiating event cut set (≈5%) and additionally dominates CDF in this limited-scope model.

In fact, in this model, for those initiating event cut sets that actually drive CDF, much of the conditional probability of core damage corresponds essentially to this 5% combination of recovery factors, the initiating event having taken down enough hardware to foreclose options other than recovery.

The SBO event’s overall conditional probability of 7E-3 derives from two possibilities: one is that the various recoveries will fail in the long run, and the other is that gravity feed will fail. Within the model used here, these two possibilities apply to essentially all SBO initiating event cut sets. (It was assumed here that the vessel is closed only a tiny fraction of the time, so that steam generator cooling never came into play. A different assumption would of course lead to a different result. The point here is the process.)

The intent of the NEI scheme is to focus the performance-based part of regulatory oversight in this area on the contribution modeled here as the “initiating event.” With the above table as an example, such a scheme can be discussed from several points of view:

- Whether observing no instances of initiating events is good evidence of satisfactory prevention of initiating events
- Whether the as-modeled conditional probability of CD provides sufficient margin to address the leading-indicator rule (that is, whether there is sufficient margin to core damage, given that an initiating event has occurred)
- Since regulatory monitoring in this area addresses events leading up to, but not following, exceedance of the threshold: what, if anything, provides regulatory assurance that the presumed conditional probability actually obtains?

These points are addressed below. Figures 1 and 2 are provided to illustrate the issues being addressed in the discussion.

2.2 Specific Issues

Initiating Events Model

As discussed in Appendix B, the present initiating event model generally requires several functional failures before an initiating event occurs. In our model, we tacitly allow a few moments for a standby pump to come on in the event of a running pump failure, without assuming that the RCS immediately heats above threshold. This means that when we get the heatup, it is as a result of more than one failure. Therefore, it is infrequent. It is theoretically possible to have a moderately significant running failure rate for individual pumps without triggering the top event, if the standby pump at least starts reliably and the failed pump is restored expeditiously.

Consider the following example. Suppose that we have identified an event whose frequency we wish to maintain below E-2 per year. How do we go about deriving assurance that measures to prevent this event are performing satisfactorily?

In particular, can we be satisfied simply by observing whether the event actually occurs in any given year? For a rate of E-2, the answer is “no.” Quite generally, if we want to derive assurance, based solely on *observations,* that a rate parameter is «1/(period) (i.e., 1/100 years for the present example), we need to observe for a time at least of order “period” (i.e., 100 years for the present example). Refer to Table 3 below, showing the probability of observing n events as a function of n, given an underlying average event rate (e.g., a “probability per year”). If the underlying event rate is E-1 per year, then our chances of observing the event in any given year are less than 10% (obtained as the difference between the n=0 entry and unity, or alternatively as the sum over entries for n greater than or equal to 1). Even if the underlying rate parameter is 1/yr, we have approximately a 37% chance of not seeing it in any given year. If we are looking for assurance that the underlying rate parameter is «1/yr, then we need more of a basis than can be derived from any single year of observation. We must either monitor events at a lower level having higher frequencies, or develop assurance from a less results-oriented approach.

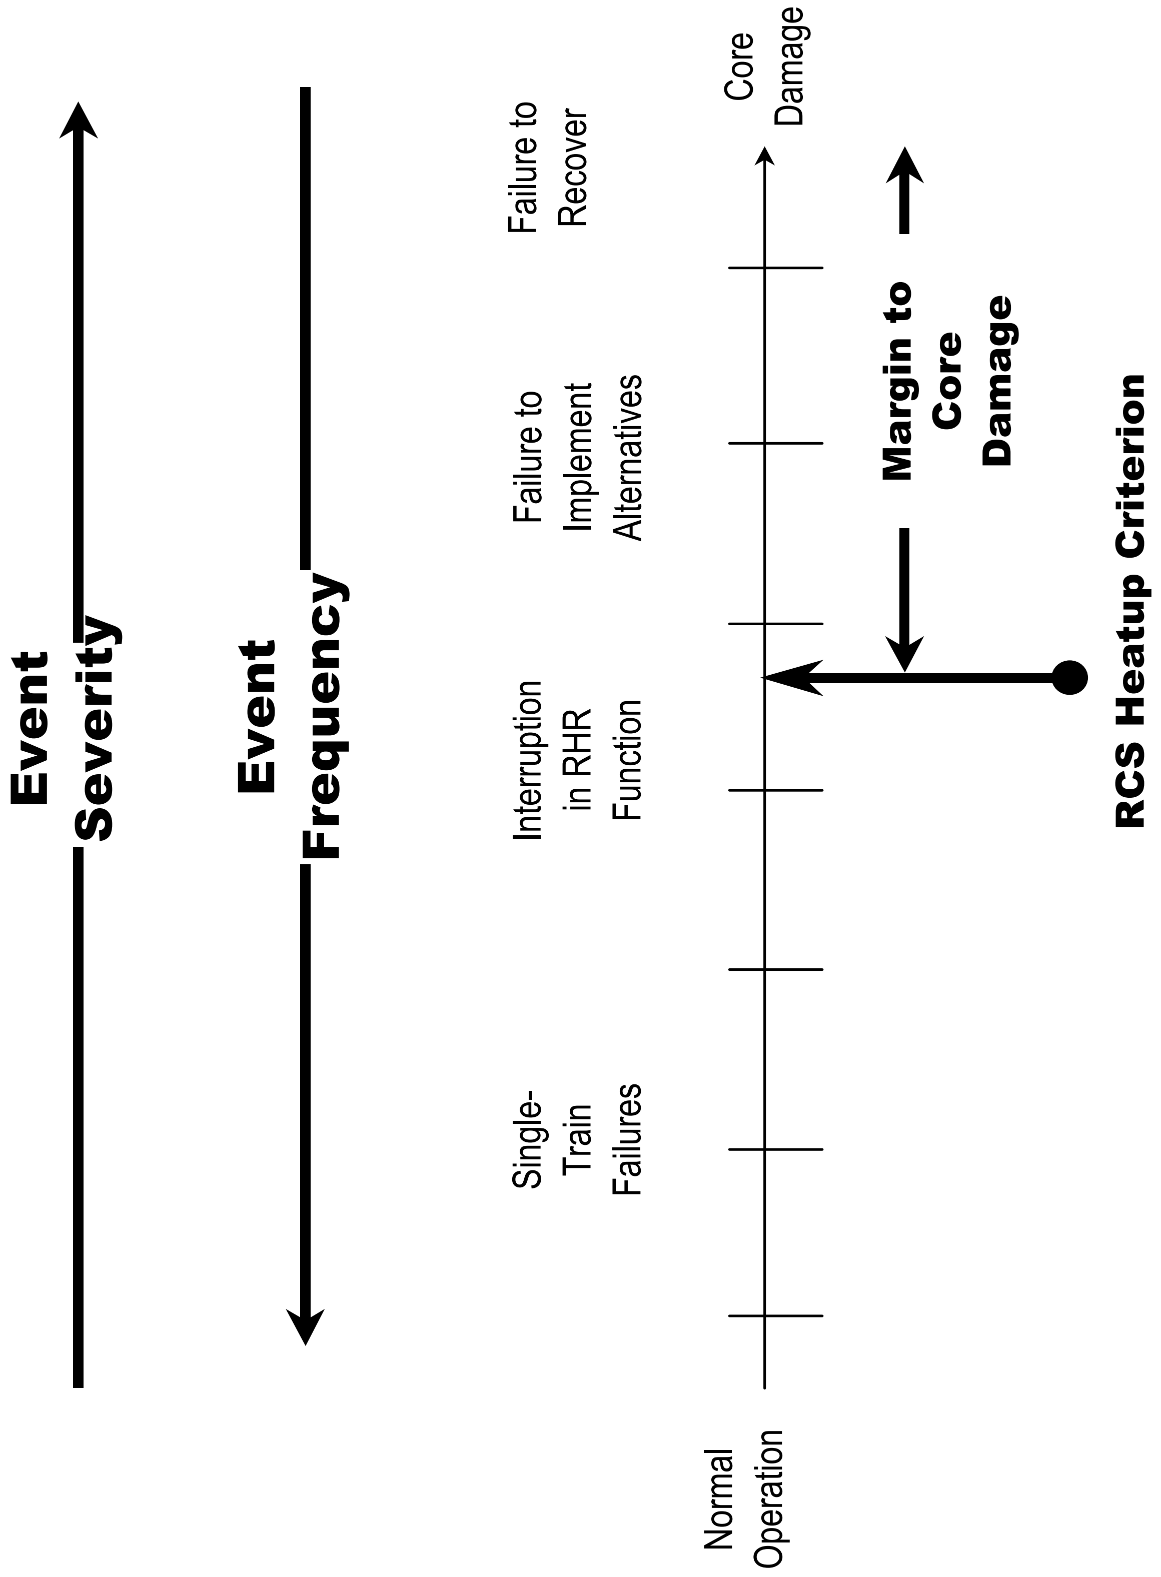

Figure 1. Considerations in Defining Monitoring Threshold

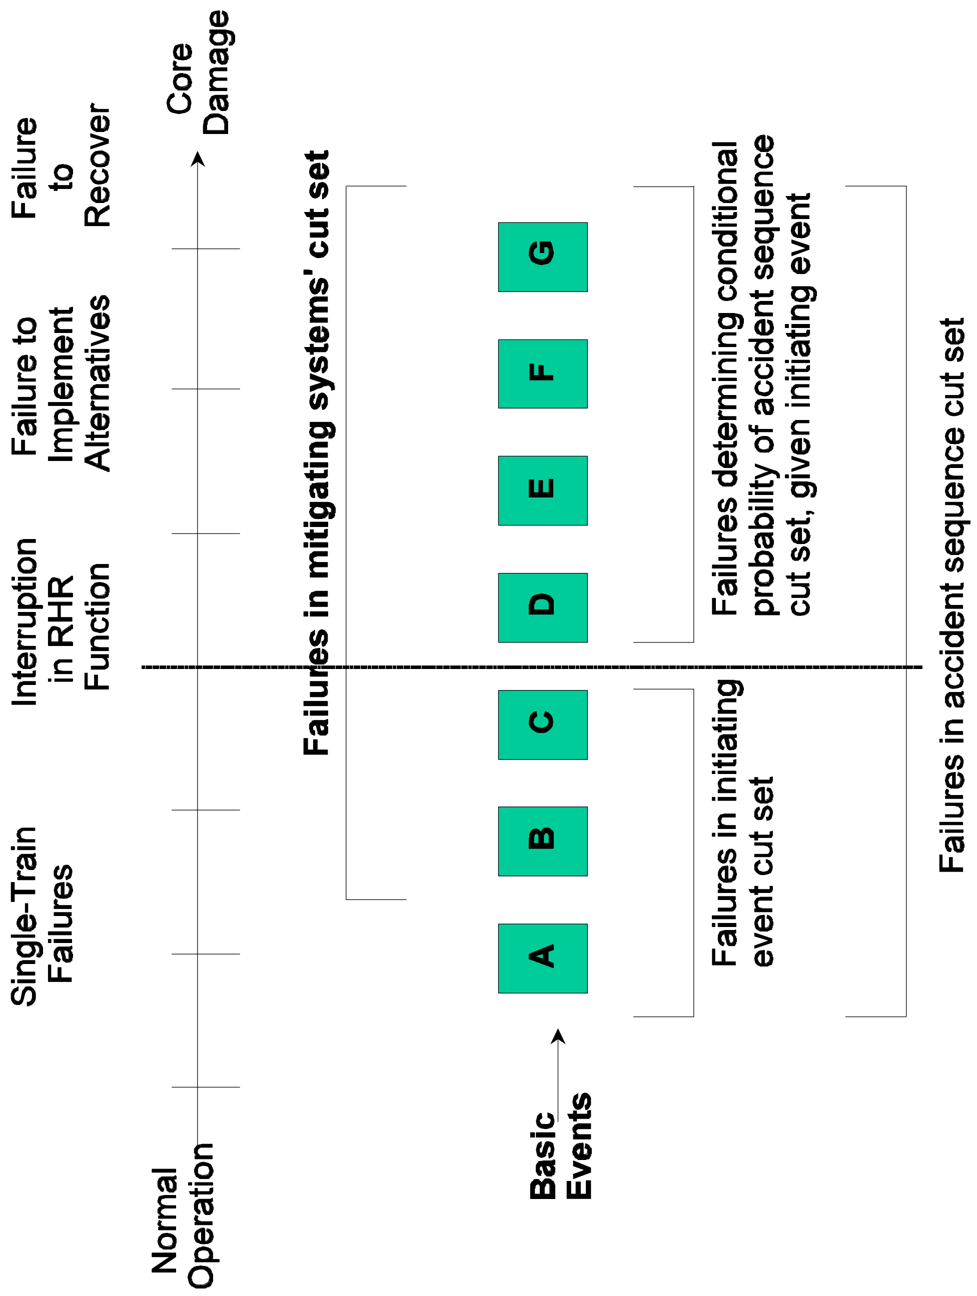
Normal Operation
Single-Train Failures
Interruption in RHR Function
Failure to Implement Alternatives
Failure to Recover
Core Damage
Failures in mitigating systems' cut set
Basic Events
A
B
C
D
E
F
G
Failures in initiating event cut set
Failures determining conditional probability of accident sequence cut set, given initiating event
Failures in accident sequence cut set

Figure 2. Conditional Probability of Core Damage Given Initiating Event

Table 3. Probability of Observing n Events in a Given Year [P(n)] as a Function of Underlying Event Rate

	Underlying Event Rate (events/year)			
	1	**0.1**	**0.01**	**0.001**
n	P(n)	P(n)	P(n)	P(n)
0	0.36787944	0.9048374	0.9900498	0.9990005
1	0.36787944	0.0904837	0.0099005	0.000999
2	0.18393972	0.0045242	4.95E-05	4.995E-07
3	0.06131324	0.0001508	1.65E-07	1.665E-10
4	0.01532831	3.77E-06	4.125E-10	4.163E-14
5	0.00306566	7.54E-08	8.25E-13	8.325E-18
6	0.00051094	1.257E-09	1.375E-15	1.388E-21
7	7.2992E-05	1.795E-11	1.964E-18	1.982E-25
8	9.124E-06	2.244E-13	2.455E-21	2.478E-29
9	1.0138E-06	2.493E-15	2.728E-24	2.753E-33
10	1.0138E-07	2.493E-17	2.728E-27	2.753E-37

The above example does not tell the whole story for the purposes of performance-based regulation. Refer to Figure 1. In the present application, the safety burden is not carried just by minimizing the frequency of exceeding the threshold; safety is also promoted by having a low probability of damage, given that the monitoring threshold was exceeded (the leading-indicator rule). This says that there should be some distance on Figure 1 between the monitoring point and core damage. To put it another way, if we are sure that the event has a low probability of going to CD, we do not need to be too concerned about the possibility that the event is more likely than we think. But this requires that we have confidence that the conditional probability is indeed low. This is the subject of the following paragraph.

Significance of Modeled Values of Conditional Probability of CD

In the shutdown model discussed above, partly because of the relative severity of our "initiating events," some of the initiating events are associated with a significant conditional probability of going to an adverse consequence. Recall that events with conditional probabilities > E-4 have "traditionally been considered important in the ASP program." [6] That is, if what we are calling an initiating event cut set having a P(CD|IE) > E-4 actually occurred, the event would be considered significant by the NRC staff. In this model, we have many such; indeed, as modeled here, each category (loss of RHR, loss of CCW and SBO) is "significant" by this test. This is not surprising; blackout events would of course be considered significant, and since we are considering a real loss of RHR, as opposed to a momentary loss of RHR, the loss of RHR initiator would be considered significant as well. As a group, loss of CCW sequences would be considered highly significant, although as we have seen, individual events within each category can vary widely in significance.

Figure 2 illustrates one reason why some initiating events are particularly significant. This figure shows a horizontal axis similar to that of Figure 1, corresponding essentially to event severity. The darkened squares below this axis represent independent failure events. The first three on the left correspond to a combination that yields an initiating event. The six on the right are failures that yield a failure of the mitigating systems. But two of these events are common to the initiator and the mitigator. These could be support system failures, for example. In this situation, the conditional probability of core damage—determined by the events to the right of the initiating event—is less than it would appear to be if one simply quantified the total probability of mitigating system failure. This concept is well known from the point of view of PRA modeling. The point of remarking it here is that in choosing monitoring points, we need to be sure that we understand the events leading up to the threshold, and whether they short-circuit some of the defense in depth that we need to have to the right of the initiating event on this figure.

A detailed review of agency reaction to previous precursor events has not been undertaken for purposes of this report, but it is suggested that events having actual P(CD|IE)s in the range of E-2 have attracted significant agency attention in the past, and that setting the bar in this range is probably leaving too little remaining margin to core damage. That is, quite apart from the above issues of inferences based on null observations of the monitored event, the leading-indicator rule is not satisfied by a scheme that accepts conditional probabilities as high as those modeled here.

Regulatory Assurance of Actual Performance of Barriers Credited as Determining the Conditional Probability of Core Damage

Above, it was suggested that there is too little apparent margin to core damage, given certain initiating events. It also needs to be asked whether there is reasonable assurance of actual achievement of the claimed performance of such barriers as there are. Of course, the barriers that come into play after an initiating event are not to be monitored at the barrier level within a performance-based approach, because their challenge frequency should be low, and their failure would be unacceptable. Therefore, unless they are challenged in the context of some other event, we must look elsewhere (e.g., at a lower level on the diamond tree, as discussed in Section 3) to find a point at which assurance can be obtained that these barriers will perform.

In the present example, because of certain modeling assumptions, a great deal of the assessed conditional probability of core damage is essentially failure to recover the RHR function: either failure to recover supports, or failure to recover the front-line function itself. The “barriers,” in other words, are presumed successes of recovery actions. How do we “assure” the credit associated with these events? Another contribution to the conditional probability of core damage arises from failure to inject. Some of the conditional probability of failure to inject arises because these initiating events tend to interfere with at least one injection path (in the case of loss of RHR, low pressure injection) and possibly more than one injection path (SBO and loss of CCW, which both interfere with everything except gravity injection). Consider a CCF of the running RHR pumps (not likely, but possible). This has a low conditional probability of going to CD, because

(a) the failure may be recoverable, (b) the nominally available train of HPI may work, (c) gravity feed may work. These possibilities combine to yield a conditional probability somewhat less than E-5. (In this sequence, again, the vessel is open and SG cooling is not an option.)

But what measures are offered to support this low number? The regulator, in approving this approach, would like to believe in a low number; what performance-based approach can the regulator adopt in order to justify that belief?

In the case of gravity feed, the important performance aspects include: the physical feasibility (is there really enough head to inject), the proper functioning of valves needed to achieve the lineup, and the correct performance of the humans. For HPI, the performance aspects include hardware functioning of the HPI itself, including necessary flowpaths, hardware functioning of support systems, and correct performance of the humans. Recovery depends on reparability (a design attribute) and, again, human performance.

At the risk of excessive repetition, it is noted that we do not wish to take a pure results-based approach to assuring the performance of these barriers, because they come into play so late that the consequences of their collective failure in this particular context are unacceptable. For HPI, we may be able to derive some results-based evidence of its generally satisfactory performance in responding to other challenges (or, at some plants, a makeup function); we may also be able to trend such things as train-level reliability, unavailability, and so on. The supports are, of course, challenged regularly (except perhaps for the diesels), and are indeed normally operating; as discussed later, there is an argument for adopting support system function as one of the "results" to be monitored in a performance-based approach. For gravity feed, the comparable figures of merit are more difficult to assess (unless the plant makes routine use of this mode of makeup, at times when it is not a last-ditch response to loss of RHR).

All of these barriers depend on the human. It is widely agreed that the performance of the human is an important aspect of this; it is difficult to imagine a convincing justification of any low conditional probability that does not take significant credit for human action. On the other hand, the "results-oriented" point of the present effort is to reduce regulatory oversight wherever it is intrusive, and monitoring human performance in such detail by the regulator would be intrusive indeed. The hybrid approach discussed later begins to address this issue.

The above discussion basically raises the question of how to optimize credit for barriers in the context of a single accident sequence. If one accident sequence obliges us to achieve high HPI performance (say), then credit for this will be helpful in many contexts; given that we have it, we can relax credit elsewhere. That is, it may be possible to choose a collection of performance elements that combine to drive down all initiating event frequencies and all conditional probabilities simultaneously, while minimizing "cost" or perhaps "regulatory invasiveness." Work has previously been done on the general form of this problem. (See [12] and references cited therein.) It turns out that achieving a reasonably optimal allocation of credit over performance elements for all sequences at once is computationally difficult.

2.3 Summary Observations from Shutdown Example

Placement of the Threshold

Refer to Figure 1. The above example suggests that for front-line functions, the threshold should be placed such that:

- the monitored event would be expected to have a frequency on the order of E-1 or more per year
- the DID barriers have a collective failure probability on the order of E-4 or less

If this placement can be achieved, then:

- the monitored events do not create a regulatory issue when they occur
- observing non-occurrence of the monitored events has at least some statistical significance
- even if the underlying rate of the monitored events is higher than E-1, the other barriers are picking up the slack, provided that their actual performance level is at least commensurable with the quantified performance level

The effort allocated to the present modeling effort has not been intended to draw a firm conclusion regarding the suitability of NEI's proposed high-level, results-oriented monitoring point (appreciable heatup). As modeled here, however, it appears that many of the complex events that trip this criterion comprise sufficient hardware failures to have eroded much of the margin to CD. This means that the frequencies of the dangerous events are high enough to matter but too low to measure. In other words, the proposed monitoring level is too far to the right on Figure 1, or equivalently (as we will see in the next section) too high on the diamond tree.

Conditional Failure Probability Associated With Monitoring Points

The shutdown example illustrates clearly that even for a particular functional loss, such as "loss of heat removal," particular initiating event cut sets vary tremendously in the associated conditional probability of going on to core damage. This variation is worse within classes of events defined at a higher level (functional versus train). Care is needed in order to ensure that the monitoring scheme appropriately addresses those events that are actually the most dangerous (those that have the highest conditional probabilities of going on to core damage). The scheme needs to reflect the expected frequency of the initiating events and the expected conditional probability of the DID barriers.

Monitoring of Key Support Functions

A pure results-oriented performance-based approach would monitor front-line functions (as in the NEI example), where the top-level "results" are achieved, and might not monitor support functions. However, in the case of the shutdown example, there is a clear argument for monitoring certain key support functions as if they were effectively front-line functions. We wish to pick up emerging problems in support systems without waiting for these problems to show up at the level of front-line performance.

This probably generalizes beyond shutdown operations. The process of working through the diamond tree will clearly drive towards such a formulation anyhow. For support functions, the above comment regarding placement of the bar would be modified. There is an argument for treating support-system performance itself as an objective having a frequency/availability target, and systematically allocating from there. That is, one might have a target frequency for loss of support, established in some way based on target frequency for loss of front-line function; the monitoring event would have some expected frequency ≈E-1 per year, and the target conditional probability would then be computed accordingly.

Reliability and Availability of Functions

Figure 3 displays important influences on functional reliability and availability.

This figure shows several levels of performance: the system level, the train level, the component level, and human actions supporting the component level. At the top, we see system or functional performance characterized in terms of frequency of functional loss (Λ) and functional unavailability ($\Lambda * T$). The frequency of functional loss depends not only on the effective in-service failure rate of each train (λ_{eff}), but also on the effective redundancy (how many failures can be tolerated before the function is lost), and on what fraction of the time each train is unavailable, either due to test/maintenance (T/M) or due to failure and repair time. By "effective in-service failure rate," we mean the rate at which a basic failure event is experienced in actual service, as distinct from the rate at which an item is pre-emptively repaired or sent off to the shop before it completely gives out. The difference is that if predictive maintenance is effective, an item is seen to be degraded and sent off for repair before it fails in service, and this replacement rate may be significantly higher than the effective in-service failure rate.

The T parameter—the time to restore lost function—can also be critical. In the shutdown example, we saw that some of the conditional probabilities were strongly affected by recovery time, which is in general influenced both by human factors and by reparability characteristics of the system. Another factor in this parameter is the possibility of consequent damage (e.g., damage to pumps as a result of cavitation or loss of cooling) due to loss of some function.

The replacement/needed repair rate (λ) (Figure 3) depends on many things. To begin with, it depends on what kind of component we are talking about (pump, valve,…), what the important failure modes are, and so on. To some extent, this rate may be influenced by code or QA requirements, although it is frequently claimed that the benefits of QA are minimal for many components. For some components, the rate is likely to depend on in-service conditions (load) and environmental conditions (temperature). It may also depend in various ways on preventive maintenance. Preventive maintenance is supposed to be beneficial; it also costs unavailability, and depending on human factors, preventive maintenance acts can contribute in other ways to component failure. The figure shows the true failure rate being affected by preventive maintenance, and the effective in-service failure rate being affected by the

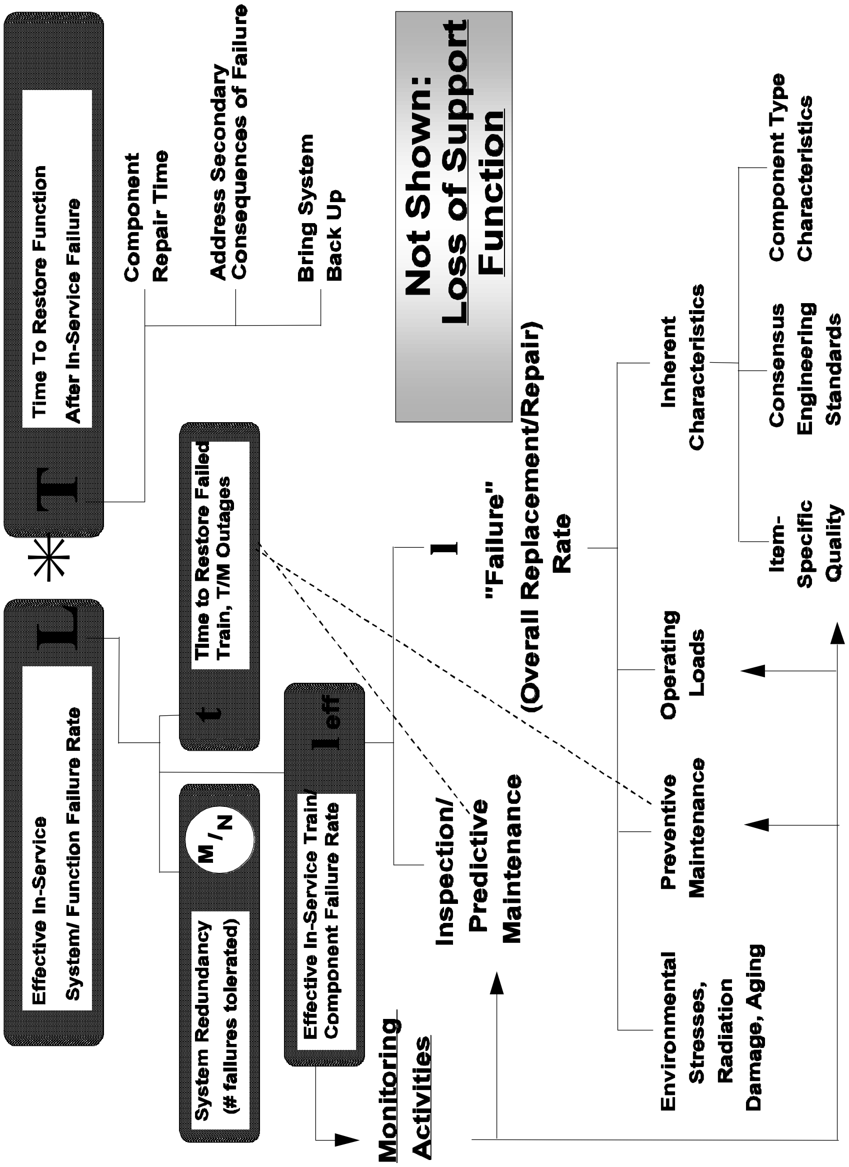

Figure 3. System/Function Level Reliability and Availability.

inspection/testing/predictive maintenance activities. These activities are connected by dotted lines to the average outage time parameter t, to show that train unavailability is also contributed by outages due to causes other than outright failure.

Consider a simple two-train system having a one-out-of-two success criterion. In this discussion, assume that testing or challenges occur often enough that repair begins very soon after a failure occurs. (This assumption is appropriate for a normally operating system.) Assume that the failure rate λ is 0.1 per year, and the repair rate is $\mu=100$, meaning that the average repair time t is 1/100 of a year. This failure rate (roughly E-5 per hour, in more familiar units) is on the order of a typical rate of failure-to-continue-running for a pump.

Within a simple reliability treatment and making a simple assumption about the availability of repair crews, the frequency of loss of this function is approximately $2 * \lambda^2/\mu$. For the stated values of λ and μ, this gives 2E-4 per year for loss of function. There is little point in trying to monitor this directly; even a shift from 2E-4 per year to 2E-2 per year will be difficult to see.

On the other hand, we experience a loss of each train at an average frequency of 0.1 per year, for a combined total of 0.2 per year if both are normally running. An order of magnitude upward shift in this quantity (i.e., 2 per year) is much easier to see than a two-order-of-magnitude shift in the frequency of loss of function. We can even more easily measure μ, and inquire after the causes of longer-than-expected train outage times. These may signal difficulties that might also cause an elevated failure rate, even if direct measurement of the elevated failure rate is still difficult. If human errors are contributing significantly to degraded performance, there is some hope of seeing it in the train level performance measures.

Predictive maintenance, and credit for predictive maintenance, is being analyzed under other programs. Progress is reported in improved testing practices, and in dealing with degraded failures. This activity has the potential to benefit performance-based approaches. It is not necessarily an area for regulatory oversight, but might be credited by regulators in their oversight of licensee programs.

Also shown on this figure is a path labeled “monitoring activities,” shown deriving information from operating experience and feeding this information back into the maintenance programs (predictive and preventive), and perhaps into operating practices or even procurement.

At every level of Figure 3, human performance is important. Working from the top down, we have the human affecting recovery of lost function, repair of failed components, inspection and predictive maintenance, preventive maintenance, and operating stresses. It is clear that these aspects can be important, but it is not typical to see them reflected in a risk model. In principle, deficiencies in these areas eventually show up in plant-specific data, including common cause failure data, which are then plugged into the models; in that sense, these aspects can be said to be reflected in risk models. But that seems an unsatisfactory foundation for a performance-based approach. As should be clear from discussions provided above, it is difficult to draw conclusions about these influences from failure statistics. It should also be clear that common cause failure potential creates a special kind of leading-indicator issue. The next section takes a more systematic look at this problem, and begins to address this area through a construct called the “diamond tree.”

3 THE DIAMOND TREE

Although diamond trees have not been widely applied, the concept appears to have something to offer in performance-based regulation. This section discusses the ideas behind diamond trees. The original point of bringing diamond trees into the present discussion was to show the effects of institutional factors on performance. However, their usefulness goes beyond that. Even without a fully-developed tree structure, a hierarchy of kinds of performance is a useful tool for comparing regulatory approaches, and even for discussing significant operating events.

3.1 Overview

3.1.1 Early Presentations of the Diamond Tree Concept

It appears that the diamond tree concept originated with Hunt and Modarres [7][*]. The diamond tree concept is an outgrowth of numerous discussions of "goal trees," and the subject of diamond trees is best approached via the subject of goal trees. For an introduction to that subject, the reader is referred to a simple introductory discussion in Modarres's book on reliability and risk analysis [8] as well as reports cited in Appendix A. Following is a brief overview.

A goal tree presents a high-level "goal," such as "achieve high level of safety," decomposed into various sub-goals, such as "achieve high functional reliability of various safety functions," which are decomposed in turn into sub-sub-goals or functions, and so on. Although diamond-tree terminology distinguishes these levels for some purposes, many of these elements could be considered as monitoring points in a performance-based scheme. For the sake of brevity, then, we introduce a bit of jargon in the following discussion: we refer generically to all elements (goals, sub-goals, etc.) at all levels as "performance nodes," and if one of these is selected for monitoring, we call it a measuring point or a monitoring point.

By virtue of the formulation, each node of a goal tree deals with some aspect of performance: reliability, efficiency, maintainability, or capability at a particular functional level. For present purposes, the important point of the development is to relate kinds of performance to a hierarchy. An important property of a goal tree is that from any given performance node below the topmost, one can look upwards at higher-level nodes to see *why* the given node needs to be accomplished, and downwards at lower-level nodes to see *how* the given node is being accomplished. Near the bottom of a goal tree, one has success paths, systems, trains, and components.

The "diamond tree" is an extension of the goal tree concept, in which the supporting role of institutional and human factors is displayed. When the goal tree display of physical functionality is essentially complete, and hardware components are specified at the bottom of the goal tree, one

[*] A paper by Hunt and Modarres, cited as reference 1 in chapter 1 in Modarres's book [8], is mentioned as the origin of the diamond tree, but the citation appears to be in error; this paper does not exist in the location given. It is possible that the citation was intended to refer to [9] or [10], but they do not address the diamond tree by name. [9] all but does so, describing what we today call the bottom half of the diamond tree as "a structure somewhat like an inverted pyramid with multiple top events, each of which represents the performance goal for a specific piece of hardware…." The first public mention of the diamond tree by that name that we can identify was in a short course taught by Hunt (1987).

begins to complete the diamond by adding additional levels below components, to display the influences of the operations staff on the components, the influence of programs and supervision on the operations staff, and so on. Institutional influences on the human performance nodes can then be displayed at a still-lower level; these include various programmatic activities and supervision. The policies that shape these programs can be displayed at a level below that. Missions can be displayed below that. And finally, a single node at the bottom (corresponding perhaps to the CEO) can be displayed below the missions.

A goal tree dealing with functional hardware performance of a complex facility naturally tends to spread out horizontally as one proceeds to lower levels, because each item is being decomposed into multiple sub-items. The goal tree is broadest at the bottom, where all relevant components appear. In making a diamond tree out of this, one adds successive layers of institutional influence below the bottom layer of the hardware goal tree. These successive layers become increasingly more general, so there are fewer nodes in each. The tree therefore narrows down again and converges to a single point at the CEO. This structure—pointed at the top and bottom, and broadest in the middle—can be said to look something like a “diamond.”

Figure 4, after Hunt and Modarres, shows this general tendency towards a diamond shape and illustrates one particular definition of the levels in the hierarchy.

As mentioned above, the diamond-tree concept is discussed here because it is a useful framework within which to compare attributes of regulatory approaches. In the case of a reactor, for example, many current testing and surveillance requirements can be placed into correspondence with component-level performance nodes; many QA requirements, especially record-keeping, would have to be binned several levels down from there. In contrast to this, “[a] performance-based approach is an approach that establishes performance and results as the primary basis for regulatory decisionmaking…a performance-based approach focuses on a licensee’s actual performance results (i.e., desired outcomes), rather than on predicted improvements or self-assessments (i.e., outputs).” [2] In the language of the diamond tree, one would say that a performance-based approach is one in which regulatory decision-making, currently focused at the middle of the diamond tree or below, focuses more attention on upper levels of the diamond tree and less attention on lower levels of the diamond tree. This point is addressed more fully in a later subsection motivating the proposed approach to development of a performance-based regulatory scheme.

The development of the diamond tree was not initiated in the context of performance-based regulation, but rather as part of an attempt to improve overall plant performance. The high-level qualitative objectives of reliability, efficiency, maintainability, and capability were aimed at production as well as safety. In fact, a comprehensive diamond tree would consist chiefly of nodes addressing priorities other than safety considerations. The original developers evidently hoped that thinking more clearly along these lines would lead to improved plant economics.

Far from being irrelevant to performance-based regulation, focusing on overall plant performance may suggest a way in which certain difficulties with regulatory oversight might be addressed. This discussion will be initiated in the following subsection, and resumed in a later section.

3.1.2 Plant Information and its Value to Performance-Based Regulation

First, recall that the diamond tree idea came out of work aimed at improving overall plant performance. [9,10] Let us set aside for a moment the question of whether the regulator or the licensee is assessing performance, and focus on the process of learning from experience. For purposes of discussion, it is useful to distinguish two kinds of thought processes. One kind of process is the evaluation of events that are “significant.” The other kind of process is based on trending behaviors that bear on the safety situation even in the absence of seemingly significant events: things like maintenance backlogs, individual component failures, adverse trends in unavailability, error rates, and so on.

Event-Based Evaluation

This involves the collection, analysis and use of actuarial evidence for lower order “precursor” events and its use to infer whether degraded levels of plant safety are present. This is based on the idea that the occurrence of “significant” events says something about the performance. Definitive conclusions about overall plant performance are not drawn simply from the occurrence of a significant event, but a significant event that violates expectations can fairly be called a warning; much has been learned from studying significant events in order to understand and address their causes.

The NEI proposal for regulatory oversight falls into this category. The idea is to characterize a class of “significant” events whose repeated occurrence would be considered to violate expectation, and whose nonoccurrence would be considered evidence of satisfactory performance.

Evaluation Based on Lower-Level Performance Measures

This kind of evaluation involves the analysis of more mundane plant failures, such as:

- analysis of failure and repair information, including excessive delays in restoring equipment, to identify statistically important hardware or human failure rates which are outside the bounds of normality and which may provide direct evidence of less than adequate plant management systems.
- analysis of any observed institutional failures to determine whether they reflect directly upon the quality of plant management systems which are important to plant safety, e.g. procedural violations or other evidence of inadequate programmatic controls.

This provides focus on the identification of management system vulnerabilities that represent “common cause” influences on plant safety, and that when improved, represent a “common cure” for enhanced safety. If important trends in management system effectiveness can be

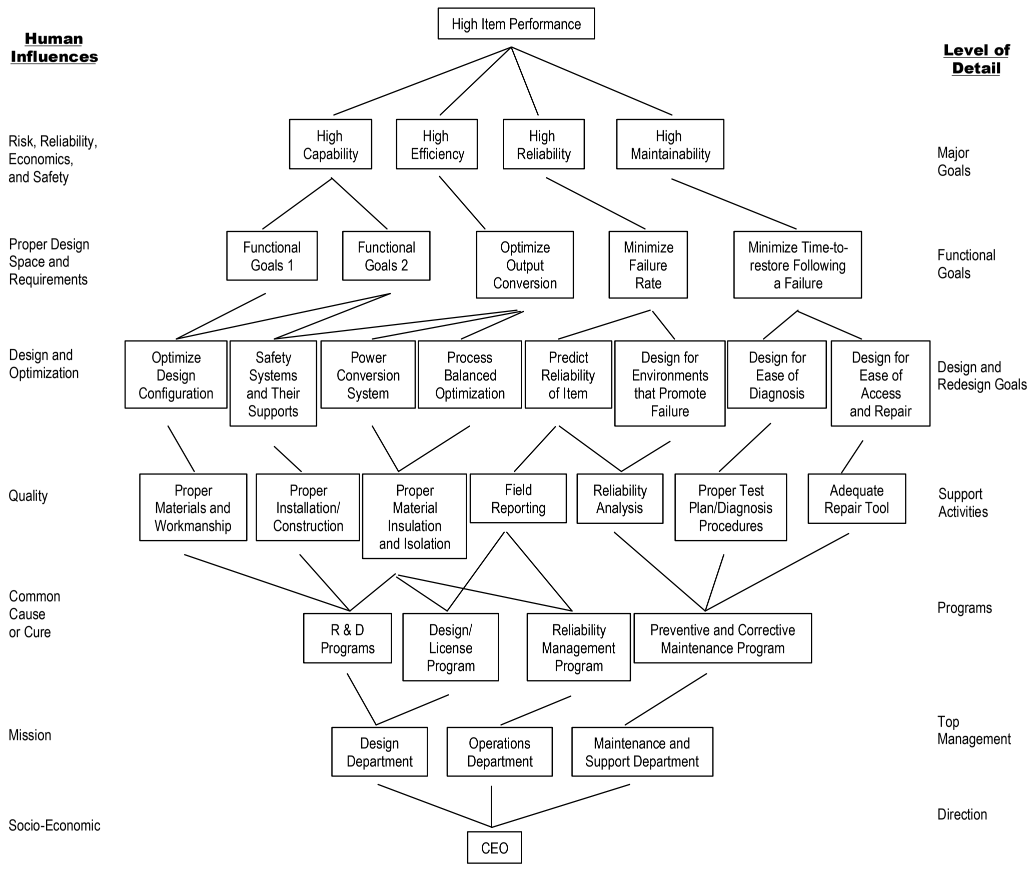

High Item Performance
Human Influences
Level of Detail
Risk, Reliability, Economics, and Safety
High Capability
High Efficiency
High Reliability
High Maintainability
Major Goals
Proper Design Space and Requirements
Functional Goals 1
Functional Goals 2
Optimize Output Conversion
Minimize Failure Rate
Minimize Time-to-restore Following a Failure
Functional Goals
Design and Optimization
Optimize Design Configuration
Safety Systems and Their Supports
Power Conversion System
Process Balanced Optimization
Predict Reliability of Item
Design for Environments that Promote Failure
Design for Ease of Diagnosis
Design for Ease of Access and Repair
Design and Redesign Goals
Quality
Proper Materials and Workmanship
Proper Installation/ Construction
Proper Material Insulation and Isolation
Field Reporting
Reliability Analysis
Proper Test Plan/Diagnosis Procedures
Adequate Repair Tool
Support Activities
Common Cause or Cure
R & D Programs
Design/ License Program
Reliability Management Program
Preventive and Corrective Maintenance Program
Programs
Mission
Design Department
Operations Department
Maintenance and Support Department
Top Management
Socio-Economic
CEO
Direction

Figure 4. A conceptual Diamond Tree Representation for Achieving High Performance (After Modarres)

identified, inferences of global plant safety levels are possible. However, because their baseline influence is implicit to the data used to quantify the PSA, the explicit importance of a degraded trend may be hard to measure.

Much potentially useful plant information is not generally reportable to outside agencies, but is collected internally in the guise of:

- QA/QC non-conformance reports,
- Reported procedural violations,
- Issues submitted to the Plant Operating and Safety Review Committee for resolution,
- Human errors, component failures and their causes,
- Repetitive failures of hardware or administrative controls.

Since the detailed information needed to guide an effective risk or safety management and improvement program is only available internally, its implication to the regulatory process is as follows:

- Primary assurance that the plant is operating within each prescribed safety envelope *must* come from the plant's use of this available information to perform critical self assessments, and make changes wherever the need for change is indicated,
- Though the regulator can, and should, retrospectively use all of the relevant information reported by the licensee to assess the adequacy of plant safety during mid-loop operation, further evaluation appears limited to an assessment of the effectiveness of the licensee's self assessment process, and whether and how it is used to effect positive changes to plant safety.

Examples of Information That May Be Important

Examples of the types of events which are of concern to the inference of adequacy of plant safety and safety management programs during mid-loop operation, and may, or may not, be reportable include:

- Routine operation in mid-loop with the minimum number of success paths available, i.e. the plant makes no attempt to maximize the functional reliability of the plant in this particular operating state.
- Events that may influence the probability of a LOCA, potentially one of the more severe initiating events that can occur during mid-loop operation. These will likely be related to human errors. Examples of these type of events are provided below:

⇒ Premature failure of flange connections on the RCS or RHR systems, which result from uneven torquing or are over-torquing to the point that bolt cracking occurs. This latter failure mode can occur when hydraulic stud or bolt tensioners are used and there is either an error in the correlation of hydraulic pressure to actual torque, or in reading the hydraulic pressure.

⇒ Freeze plugs on the RHR or RCS boundary which are less than adequate and as a result have a relatively high probability of failure,

⇒ Improper installation of nozzle dams,

  - ⇒ Relief valves returned to service after repair or adjustment without adequate testing—resulting in their failure to reclose after a challenge (as in the case of the Peach Bottom ADS/SRVs, where multiple failures were caused by improperly replaced insulation).

- Undetected, or unreported unavailability of instrumentation and control systems which result from less than adequate management and control of the many calibration and testing activities which are part of the surveillance and maintenance programs conducted during refueling outages—this may become an even more critical issue as outage lengths are further reduced.
- Events or conditions which are external to the RCS/RHR systems but influence their reliability as a result of spatial interactions (fires, dropped loads, etc,.).

The occurrence of any of these or other similar events may be very important to a plant program to improve performance , but traditionally, and under foreseeable reporting practices, much of this information would not be readily available to the regulatory authorities on a routine basis. It may be argued that the regulator has a stake in ensuring the adequacy of this process, but it is probably infeasible for the regulator to receive, process, and apply data at that level for all plants.

We will see at the end that it is difficult to base a really satisfactory regulatory scheme solely on high-level results indicators. The regulators need assurance of performance at lower levels as well. But regulatory intervention at those lower levels is undesirable. A solution worth considering is a plant information system capable of tracking performance at lower levels for plant purposes; if a means can be established by which the regulator can derive assurance that the plant is satisfactorily monitoring and actively managing performance at these lower levels, then some assurance of desired performance at higher levels might be obtainable through less prescriptive, if not truly “performance-based,” means.

3.1.3 Other Presentations of the Diamond Tree Concept

A more recent, and more safety-oriented, discussion of the diamond tree idea appears in Wreathall et al. [11] The focus of Wreathall’s work is performance indicators, a subject very closely related to the present topic.

This work was part of the large body of NRC-sponsored work on performance indicators. As with other work in that area, the focus was on diagnosing the safety status of a plant through interpretation of various performance figures of merit, and it therefore bears on the present effort. Indeed, the major difference in emphasis between that work and the current effort is that in performance-based regulation, we are not just trying to supplement existing prescriptive requirements with additional layers of interpretation, but rather looking for ways to improve the existing regulatory process itself, potentially by replacing existing prescriptive regulations with performance-based regulations. As part of that effort, we need to characterize the level of assurance that we can derive from a relatively limited set of performance figures of merit, an issue not much discussed in earlier work.

Figure C.1 of the Wreathall report is reproduced here for convenience as Figure 5. This figure shows two things. One is the “diamond tree” idea itself. The upper half of this tree is essentially a traditional goal tree: at the top of the tree is one or more high-level “objectives,” which are decomposed into “goals,” which are in turn decomposed into “sub-goals,” which are in turn decomposed into “success paths.” At this level the tree is at or near its widest point. Moving to the lower half, the next level is intended to address “activities” that support these success paths, followed by programs that support these activities, and policies that mandate and inform these programs.

Most “performance” figures of merit that one might consider can be placed into correspondence with a structure like this. Performance of a particular safety function might be a “goal” and performance of individual systems would be factored in at the “success path” level; monitoring hardware performance at the function, system, train, or even component level could be discussed in terms of a structure like this, although the intent of the development in the report appears to have been aimed more at system-level or higher. Note that in the lower portion of the tree, performance of plant staff is being addressed, including the effect of “institutional factors,” corporate policy, and so on; in the upper portion of the tree, functional performance is seen to be a blend of plant staff performance and equipment performance.

With this structure in place, the figure then classifies different rows of the tree into different kinds of performance indicators: “direct” indicators, “programmatic” indicators, and “organizational” indicators. The “direct” indicators include all of the top-level “results” figures of merit that one might consider, including the RCS temperature figure of merit used as an example by NEI. The higher a given measure appears in this structure, the more results-oriented the measure is. The programmatic and organizational figures of merit are a major focus of the Wreathall report and, as discussed below, are one reason that the diamond tree suggested itself for the present effort.

The appendices of the Wreathall report contain explicit diamond trees. Selected sheets are reproduced here for convenience as Figure 6. The upper-level sheets (exemplified in Figure 6 Sheet 1) display MPLD-like structure: a presentation, in success space, of the logical relationships between safety functions. Logic gates relating performance nodes to each other are actually shown at this level on this tree; subsequent sheets do not all show gates, and it seems that even at the higher levels, the gates are meant conceptually rather than literally. Figure 6 Sheet 2, “Effective Item Time to Restore,” does not show logic gates, but identifies very specific factors bearing on performance.

Figure 6 Sheet 3 addresses several programmatic activities: preventive maintenance, surveillance and testing, and procurement of spare parts. Some of these are the kinds of performance aspects that show up in discussions of contributors to significant events; indeed, some of these will turn up in Section 5.

3.2 Specific Examples

This section presents two partial developments of a diamond tree to demonstrate further what kind of issues are addressed and what kind of performance nodes show up in the development.

One development covers fire protection, and the other covers residual heat removal in reduced inventory conditions, essentially the example discussed in Section 2.

The performance nodes that emerge as verifiable in the fire protection example are at relatively low levels in the diamond tree hierarchy. This means that they are less outcome-oriented than might have been hoped for, in the present context. However, these are useful areas for the plant to monitor (cf. Section 3.1.2 above). The shutdown example develops many performance nodes that have potential application in performance-based regulatory oversight.

3.2.1 Fire Protection Example

Fire protection is a blend of prevention activities, which are obviously heavily influenced by institutional factors, and hardware performance of many kinds: equipment qualification, reliability of fire detection/suppression, and successful performance of surviving trains of safe shutdown equipment, potentially including operator actions to bring safe shutdown systems into play. Housekeeping in particular is an area that is important to fire protection, but arguably an instance of the sort of thing that PRA is not good at modeling explicitly. Following up on these ideas, the following discussion presents an attempt at a diamond tree addressing fire protection.

The top-level objective for this tree is the following:

> Each licensee shall provide adequate protection from the effects of internal plant fires by ensuring that fires which may possibly initiate and propagate within, and from, each individual designated in-plant fire area
>
> - have an acceptably low frequency of occurrence
> - cannot, of themselves result in damage which causes the loss of all success paths for a single plant critical safety function
> - neither result in an unacceptable threat to the safety of the nuclear power plant, nor present an unacceptable risk to the health and welfare of the general public

Figure 7 Sheets 1–7 represent the first steps in development of a Fire Protection Goal Tree. It is intended to:

- Define the basic high level requirements for an effective fire protection system which in turn can serve to define the basic functional sections of a monitoring approach
- Provide the first insights into each of the types of plant activities which must be explicitly included within the scope of the approach
- Provide the context for a single set of detailed trees which can be used to understand how information can be collected and analyzed to guide the selection, definition and use of performance based indicators or measures, in the assessment of the “adequacy” of all or parts of the Fire Protection Systems in any particular facility.

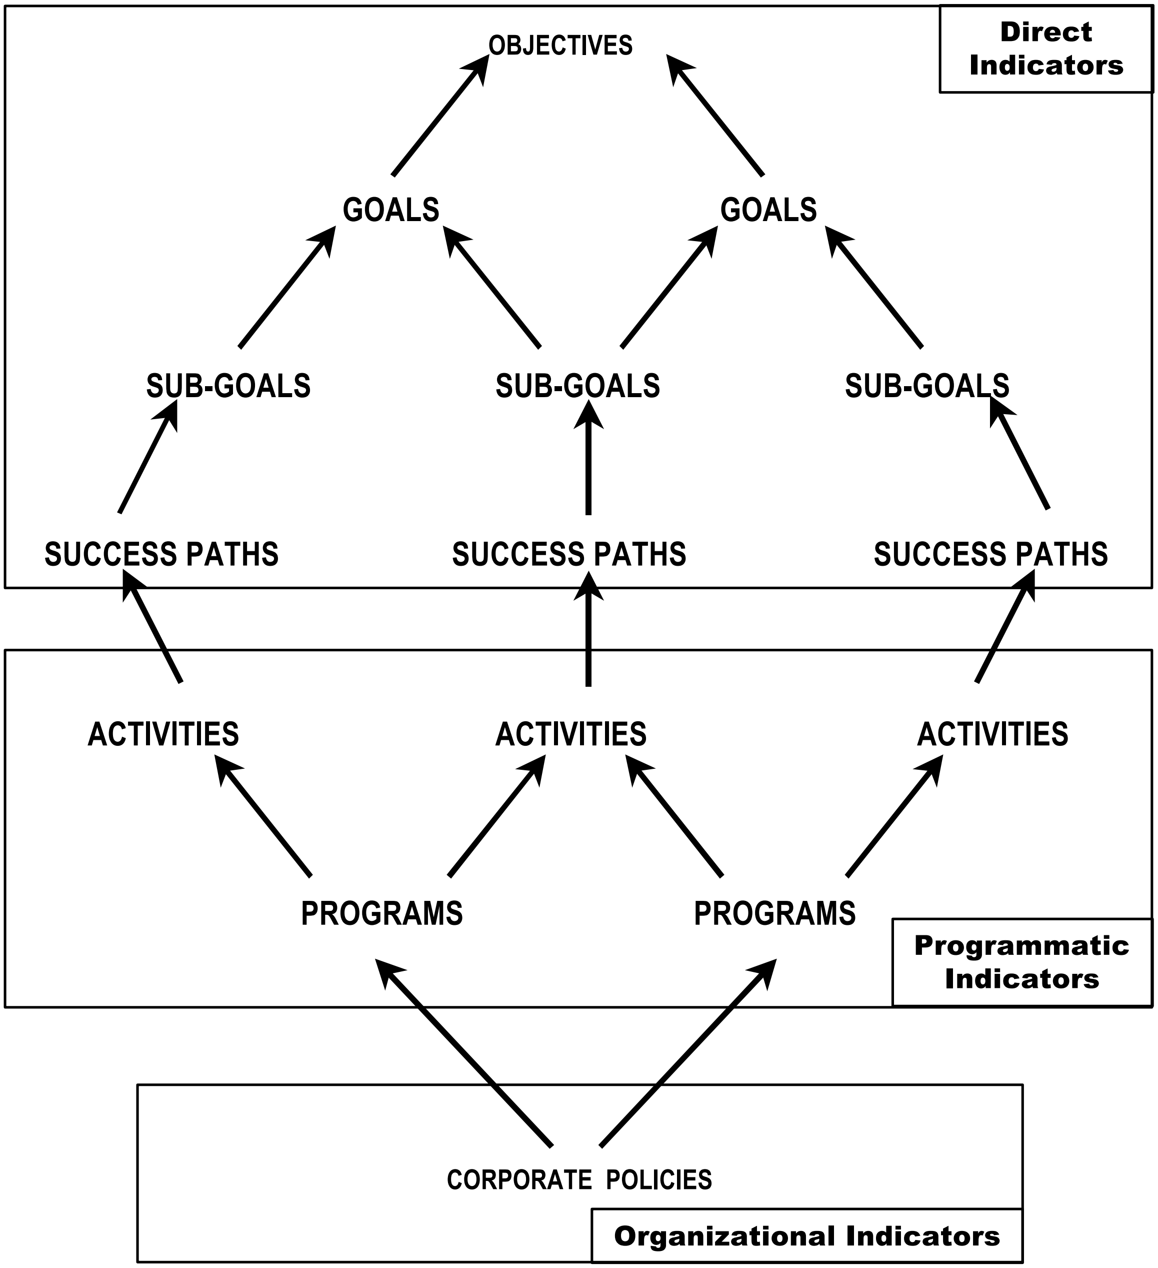

Figure 5. Correlation of Indicator Classes to Diamond Tree (After Wreathall).

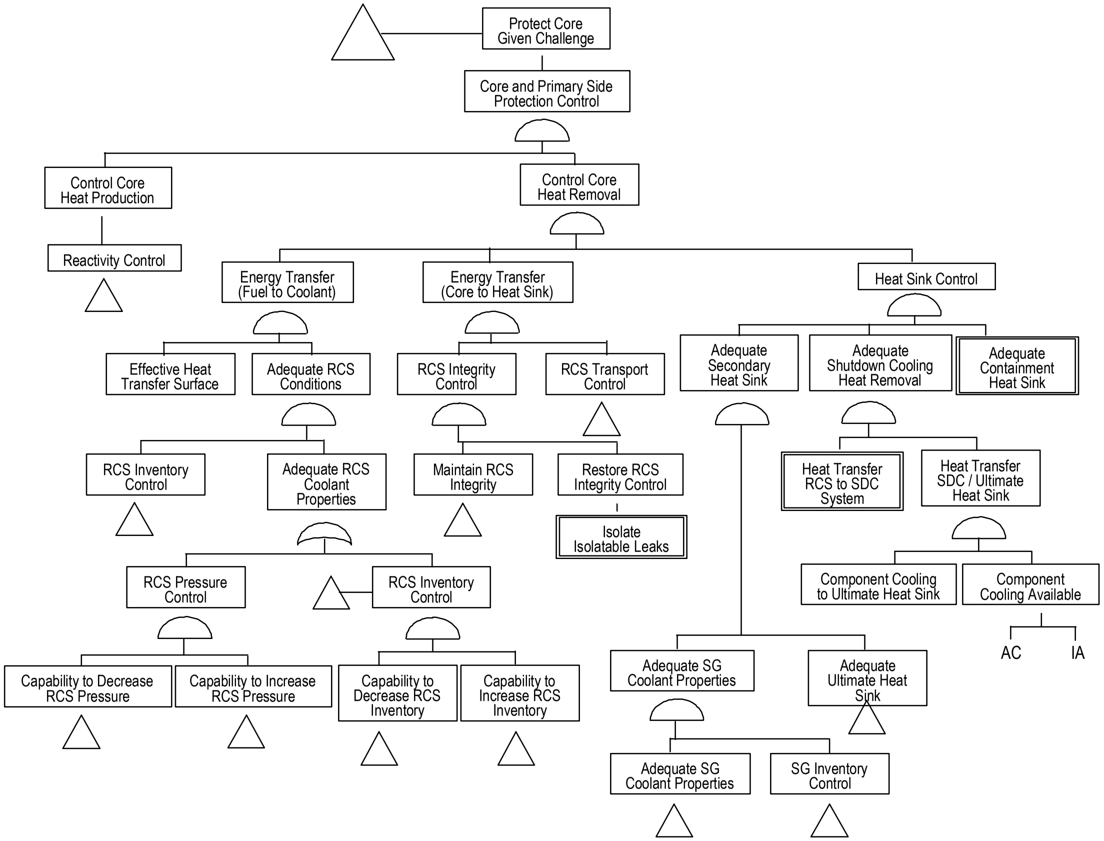
Protect Core Given Challenge
Core and Primary Side Protection Control
Control Core Heat Production
Control Core Heat Removal
Reactivity Control
Energy Transfer (Fuel to Coolant)
Energy Transfer (Core to Heat Sink)
Heat Sink Control
Effective Heat Transfer Surface
Adequate RCS Conditions
RCS Integrity Control
RCS Transport Control
Adequate Secondary Heat Sink
Adequate Shutdown Cooling Heat Removal
Adequate Containment Heat Sink
RCS Inventory Control
Adequate RCS Coolant Properties
Maintain RCS Integrity
Restore RCS Integrity Control
Isolate Isolatable Leaks
Heat Transfer RCS to SDC System
Heat Transfer SDC / Ultimate Heat Sink
Component Cooling to Ultimate Heat Sink
Component Cooling Available
AC
IA
RCS Pressure Control
RCS Inventory Control
Capability to Decrease RCS Pressure
Capability to Increase RCS Pressure
Capability to Decrease RCS Inventory
Capability to Increase RCS Inventory
Adequate SG Coolant Properties
Adequate Ultimate Heat Sink
Adequate SG Coolant Properties
SG Inventory Control

Figure 6. (Sheet 1 of 3) PWR Diamond Tree: Protect Core (After Wreathall et al. [11])

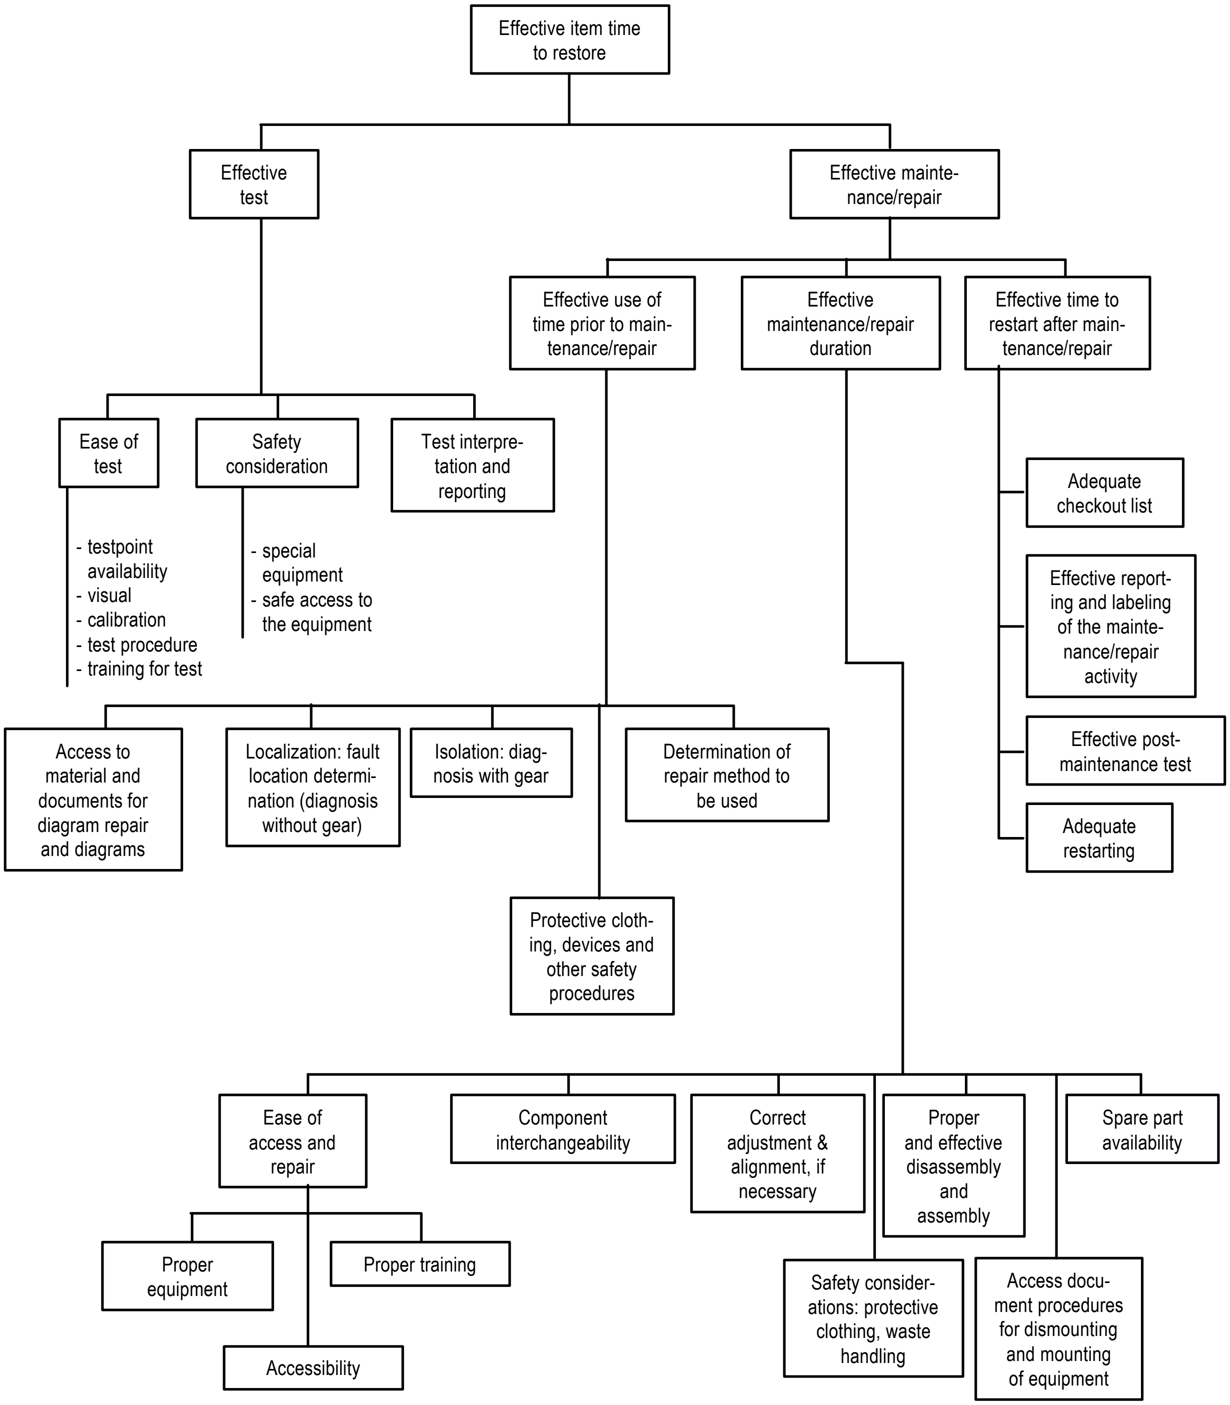
Effective item time to restore
Effective test
Effective maintenance/repair
Effective use of time prior to maintenance/repair
Effective maintenance/repair duration
Effective time to restart after maintenance/repair
Ease of test
Safety consideration
Test interpretation and reporting
- testpoint availability
- visual
- calibration
- test procedure
- training for test
- special equipment
- safe access to the equipment
Adequate checkout list
Effective reporting and labeling of the maintenance/repair activity
Effective post-maintenance test
Adequate restarting
Access to material and documents for diagram repair and diagrams
Localization: fault location determination (diagnosis without gear)
Isolation: diagnosis with gear
Determination of repair method to be used
Protective clothing, devices and other safety procedures
Ease of access and repair
Component interchangeability
Correct adjustment & alignment, if necessary
Proper and effective disassembly and assembly
Spare part availability
Proper equipment
Proper training
Accessibility
Safety considerations: protective clothing, waste handling
Access document procedures for dismounting and mounting of equipment

Figure 6. (Sheet 2 of 3) PWR Diamond Tree: Effective Item Time to Restore (After Weathall et al. [11])

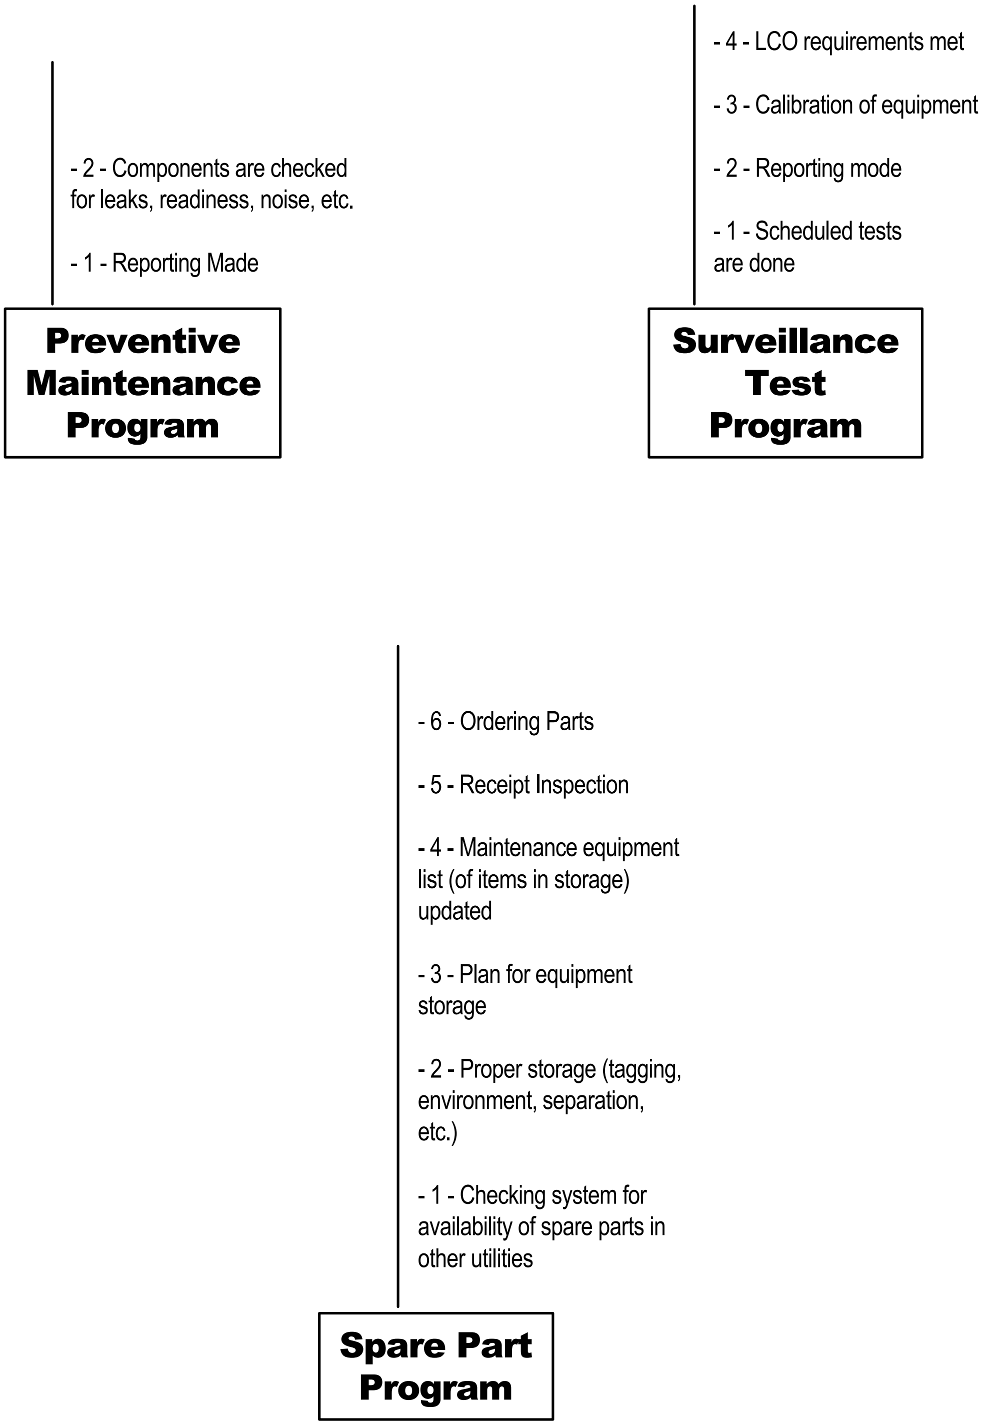
- 2 - Components are checked for leaks, readiness, noise, etc.
- 1 - Reporting Made
Preventive Maintenance Program
- 4 - LCO requirements met
- 3 - Calibration of equipment
- 2 - Reporting mode
- 1 - Scheduled tests are done
Surveillance Test Program
- 6 - Ordering Parts
- 5 - Receipt Inspection
- 4 - Maintenance equipment list (of items in storage) updated
- 3 - Plan for equipment storage
- 2 - Proper storage (tagging, environment, separation, etc.)
- 1 - Checking system for availability of spare parts in other utilities
Spare Part Program

Figure 6. (Sheet 3 of 3) PWR Diamond Tree: Programmatic Activities (After Wreathall et al. [11])

Use of the Fire Protection Model

The potential institutional contributors to an effective and adequate fire protection are shown in the lower levels of the tree structure. For each discrete institutional objective, the tree in principle shows the organizational structure, functional programs, and individual activities that must be focused on these objectives, if they are to be satisfied.

Figure 7 Sheets 1–7 provide hints about the nature of this hierarchy, from its uppermost levels to the detailed lower order activities. These lower level activities will eventually appear in each of the individual institutional goal trees, when they are fully developed. These activities are currently provided in the form of lists to provide examples in the current structure, although to be useful they must be developed in the form of a hierarchy if completeness, cause-consequence and relative importance of issues is to be inferred.

Figure 7 Sheets 8 and 9 extend the development of one particular goal *"Minimizing the Number of Potential Ignition Sources"* (Figure 7 Sheet 2, goal 1-1-1) and demonstrate how it can be expanded in detail to provide the basis for an information definition and collection process whose products can be used to fuel a performance indicator program. This same diagram has been used to demonstrate how it may be possible to score individual goals within the hierarchy.

In this particular example, weighting and rating factors have been assigned to a limited portion of the tree to demonstrate how such a system may work:

- Develop a relative weighting factor for each of each sub-goal in terms of its relative importance to the higher order goal it serves
- On a relative scale of 0–100, rate the plant's success paths which are relied upon to satisfy the parent goal, where 0 equates to the minimum procedural standards which could be accepted and 100 equates to a perfect procedural response.

Note: Success paths are physical entities or activities which are carried out to meet specific objectives or goals.

- Normalize all weighting factors to provide a cumulative value of 1.0 at each level in the goal tree and use these factors to propagate the ratings to the top goal of the tree to synthesize a figure of merit.
- This synthesized measure can represent a performance measure for the fire protection goal, i.e. how well, based on observations about the processes involved in assuring adequate fire protection, it is being achieved.

Test Application to Develop Performance Indicators

Figure 7 Sheets 8 and 9 provide an expanded level of detail to the baseline goal tree to describe how to satisfy the goal *"Identify each potential source of ignition within the perimeter of each room or fire area which has safety significance, and prevent any Unplanned Additions of Ignition Sources which have sufficient energy to ignite available Combustibles"* (paraphrased from third order goal, Figure 7 Sheet 2, Number of Potential Ignition Sources).

The implication of the use of the word “unplanned” is that the number of ignition sources has been increased without an explicit risk assessment and the conscious acceptance of any associated change in risk.

The hierarchy of lower order goals that serve the selected goals is as follows:

- *“Develop a baseline inventory of ignition sources in each important fire area”*
- *“Minimize the number of risk important sources of ignition”*
- *“Prevent the unplanned addition of any new ignition sources in any important fire areas or rooms”*

This latter goal can be decomposed into the set of lower order goals as described by the following:

⇒ *“The licensee has established explicit controls in plant design, construction and installation procedures to prevent the inadvertent addition of any permanent new ignition sources in a safety important fire area or room”*

To collect information to assess the quality of the success paths implemented to meet this goal would probably involve the following:

*   Review engineering design procedures to confirm the presence of explicit requirements relating to ignition sources in important fire areas

*   Confirm that each important fire area is identified and that engineering personnel are familiar with their listing

*   Confirm that the procedural requirements are explicit (not just a general statement)

*   Assess the review format (check-off, second check, independent review, etc.) to estimate the quality of the requirements and procedural guidance

*   Addition of new ignition sources is explicitly addressed by the plant in its 50.59 review process

At the conclusion of the review, each of the above would receive a plant specific rating which could be propagated upwards to participate in the synthesis of an overall figure of merit for “Adequacy of Fire Protection.”

⇒ *The licensee has “Implemented controls to prevent unplanned addition of transient ignition sources in risk important areas a result of plant construction, operations and maintenance activities.”*

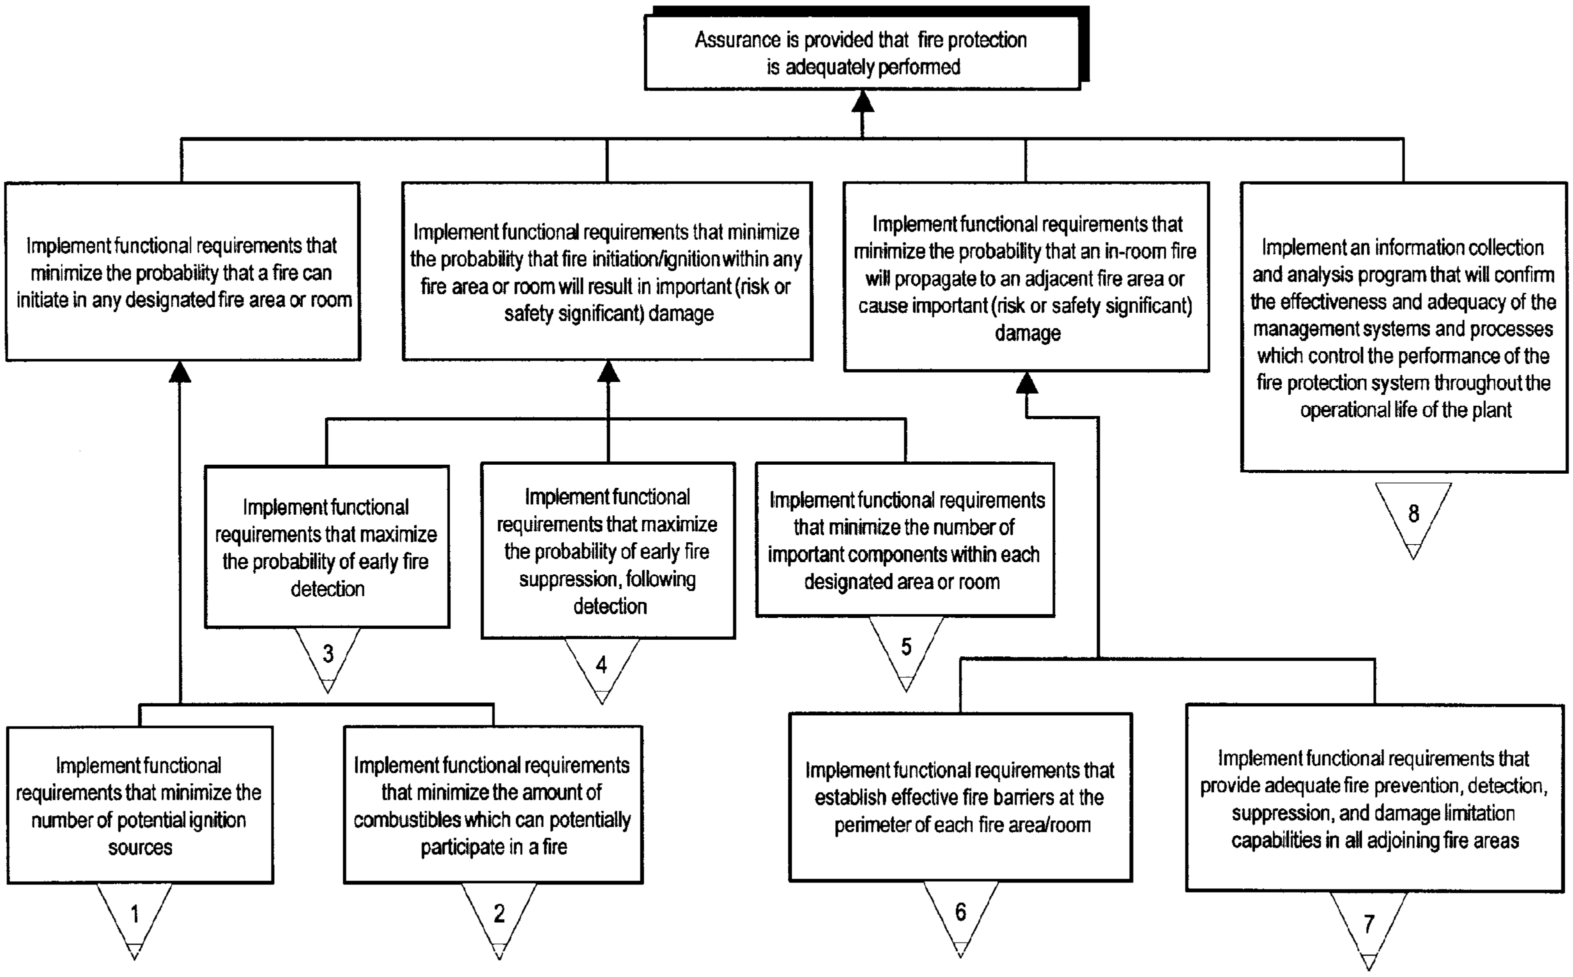
Assurance is provided that fire protection is adequately performed
Implement functional requirements that minimize the probability that a fire can initiate in any designated fire area or room
Implement functional requirements that minimize the probability that fire initiation/ignition within any fire area or room will result in important (risk or safety significant) damage
Implement functional requirements that minimize the probability that an in-room fire will propagate to an adjacent fire area or cause important (risk or safety significant) damage
Implement an information collection and analysis program that will confirm the effectiveness and adequacy of the management systems and processes which control the performance of the fire protection system throughout the operational life of the plant
Implement functional requirements that maximize the probability of early fire detection
3
Implement functional requirements that maximize the probability of early fire suppression, following detection
4
Implement functional requirements that minimize the number of important components within each designated area or room
5
8
Implement functional requirements that minimize the number of potential ignition sources
1
Implement functional requirements that minimize the amount of combustibles which can potentially participate in a fire
2
Implement functional requirements that establish effective fire barriers at the perimeter of each fire area/room
6
Implement functional requirements that provide adequate fire prevention, detection, suppression, and damage limitation capabilities in all adjoining fire areas
7

Figure 7. (Sheet 1 of 10) Overall Fire Protection Goal Tree

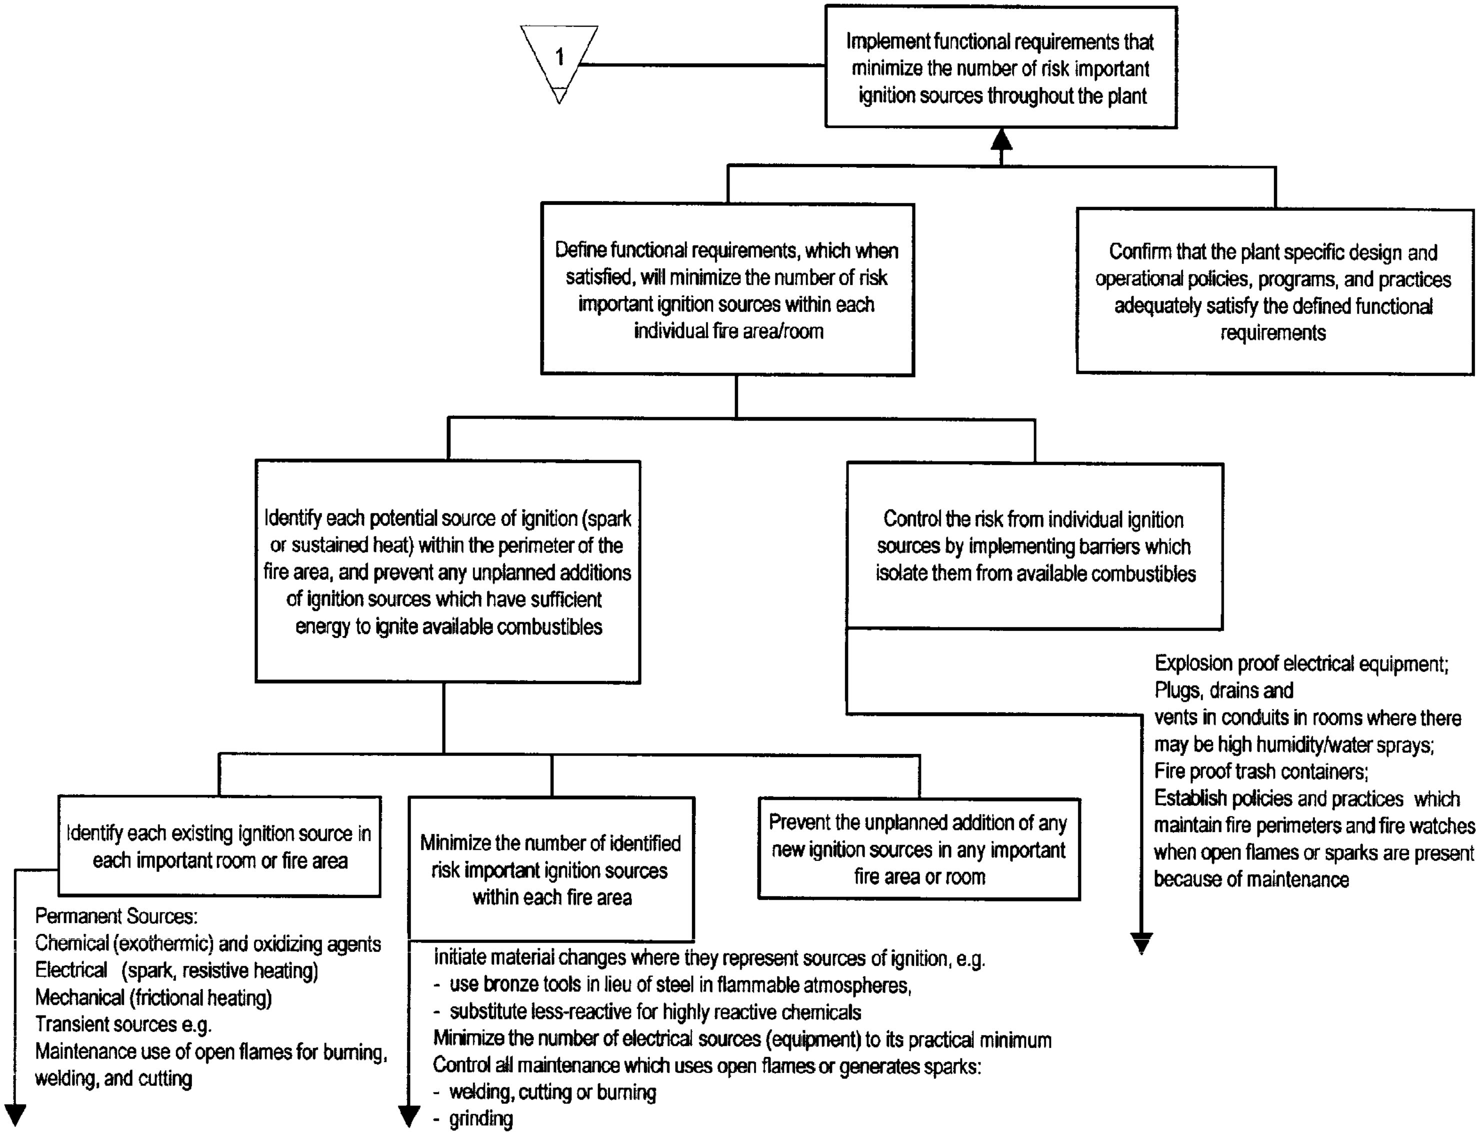
1
Implement functional requirements that minimize the number of risk important ignition sources throughout the plant
Define functional requirements, which when satisfied, will minimize the number of risk important ignition sources within each individual fire area/room
Confirm that the plant specific design and operational policies, programs, and practices adequately satisfy the defined functional requirements
Identify each potential source of ignition (spark or sustained heat) within the perimeter of the fire area, and prevent any unplanned additions of ignition sources which have sufficient energy to ignite available combustibles
Control the risk from individual ignition sources by implementing barriers which isolate them from available combustibles
Explosion proof electrical equipment;
Plugs, drains and
vents in conduits in rooms where there may be high humidity/water sprays;
Fire proof trash containers;
Establish policies and practices which maintain fire perimeters and fire watches when open flames or sparks are present because of maintenance
Identify each existing ignition source in each important room or fire area
Permanent Sources:
Chemical (exothermic) and oxidizing agents
Electrical (spark, resistive heating)
Mechanical (frictional heating)
Transient sources e.g.
Maintenance use of open flames for burning, welding, and cutting
Minimize the number of identified risk important ignition sources within each fire area
Initiate material changes where they represent sources of ignition, e.g.
- use bronze tools in lieu of steel in flammable atmospheres,
- substitute less-reactive for highly reactive chemicals
Minimize the number of electrical sources (equipment) to its practical minimum
Control all maintenance which uses open flames or generates sparks:
- welding, cutting or burning
- grinding
Prevent the unplanned addition of any new ignition sources in any important fire area or room

Figure 7. (Sheet 2 of 10) Number of Potential Ignition Sources.

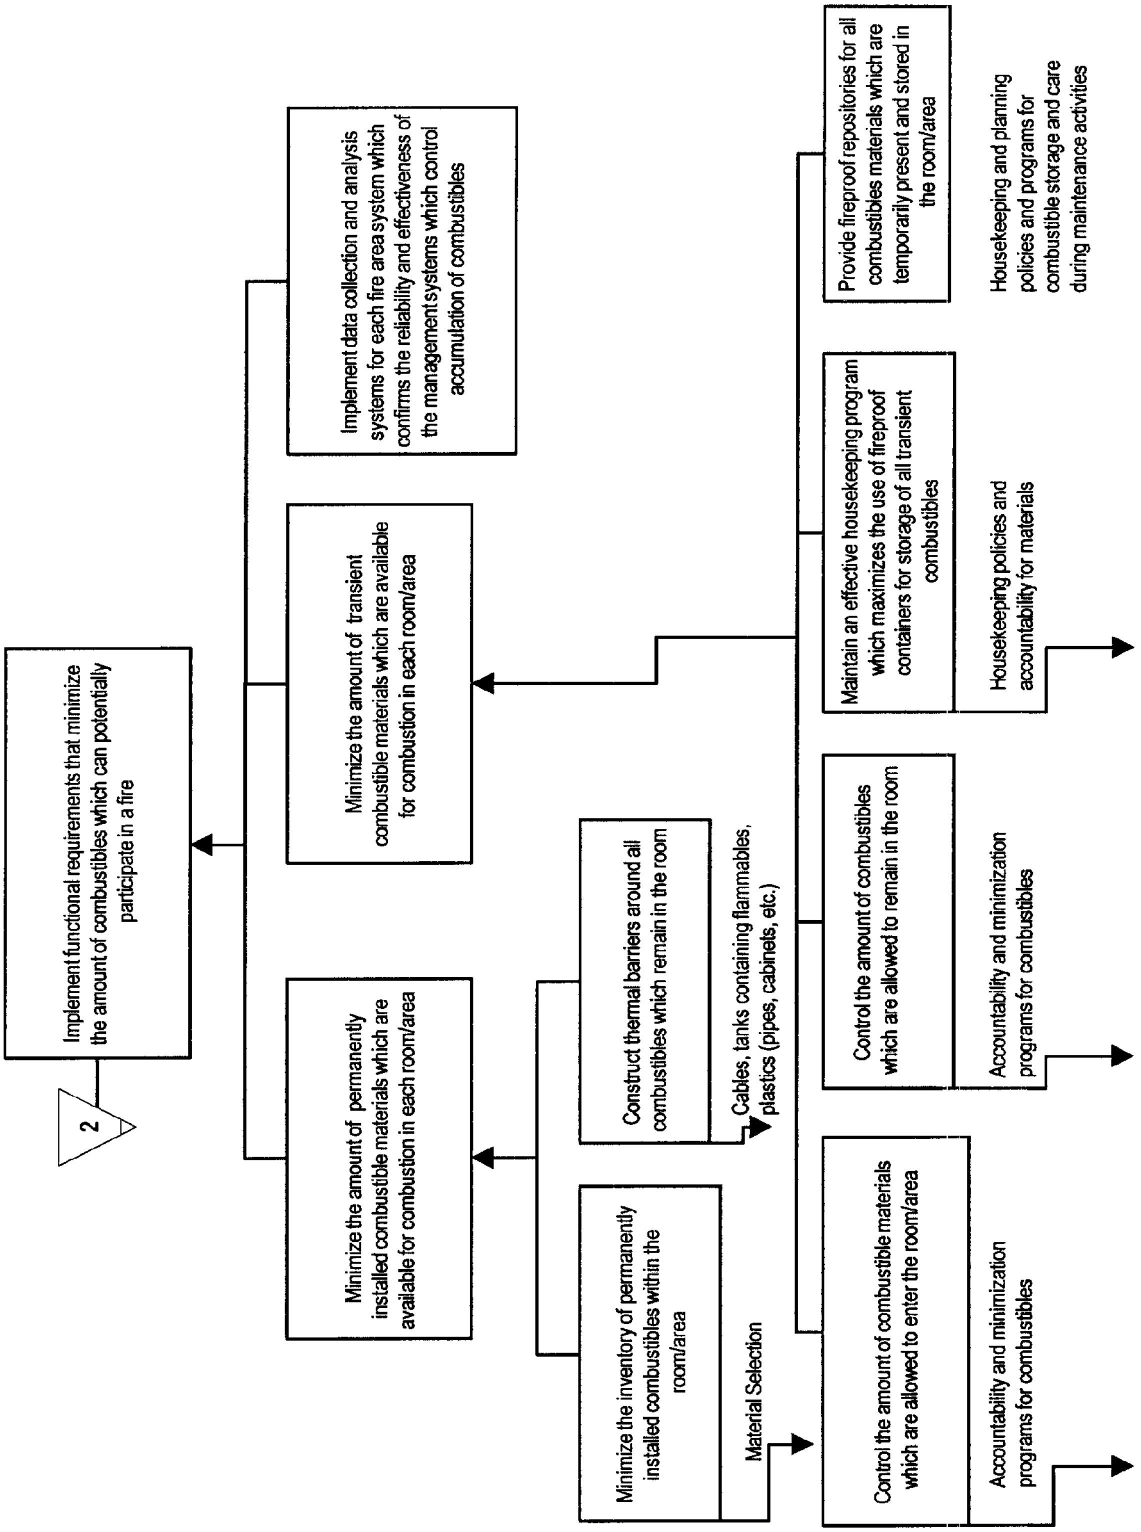
2
Implement functional requirements that minimize the amount of combustibles which can potentially participate in a fire
Minimize the amount of permanently installed combustible materials which are available for combustion in each room/area
Minimize the amount of transient combustible materials which are available for combustion in each room/area
Implement data collection and analysis systems for each fire area system which confirms the reliability and effectiveness of the management systems which control accumulation of combustibles
Minimize the inventory of permanently installed combustibles within the room/area
Material Selection
Construct thermal barriers around all combustibles which remain in the room
Cables, tanks containing flammables, plastics (pipes, cabinets, etc.)
Control the amount of combustible materials which are allowed to enter the room/area
Accountability and minimization programs for combustibles
Control the amount of combustibles which are allowed to remain in the room
Accountability and minimization programs for combustibles
Maintain an effective housekeeping program which maximizes the use of fireproof containers for storage of all transient combustibles
Housekeeping policies and accountability for materials
Provide fireproof repositories for all combustibles materials which are temporarily present and stored in the room/area
Housekeeping and planning policies and programs for combustible storage and care during maintenance activities

Figure 7. (Sheet 3 of 10) Amount of Combustibles

Note:
"Adequate" capabilities
imply that levels of effectiveness
and reliability are commensurate
with importance

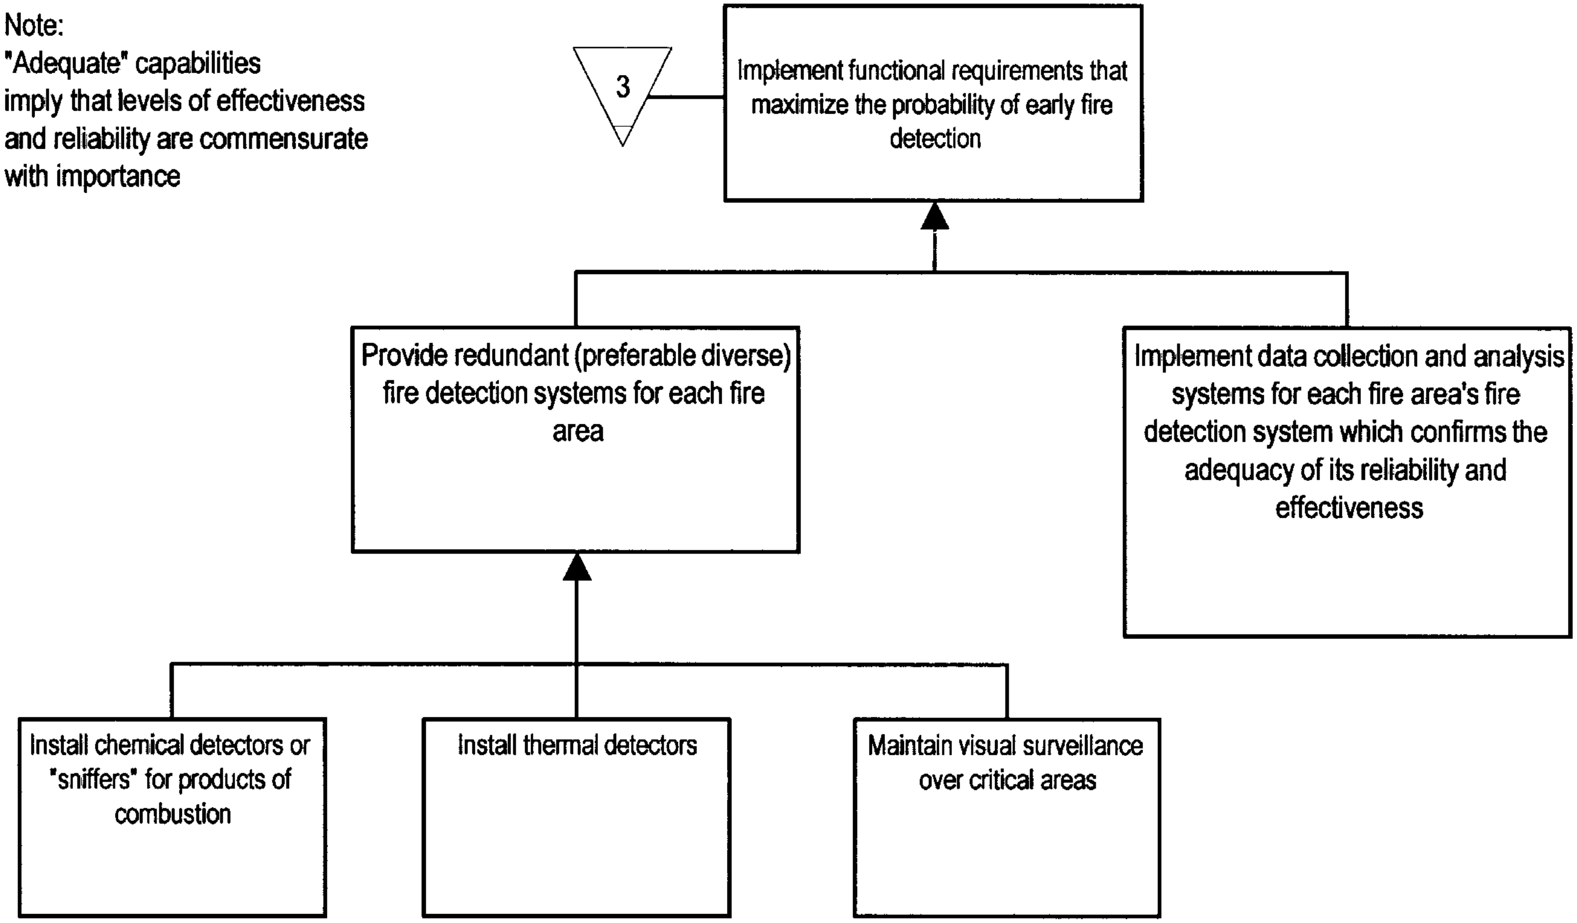
3
Implement functional requirements that maximize the probability of early fire detection
Provide redundant (preferable diverse) fire detection systems for each fire area
Implement data collection and analysis systems for each fire area's fire detection system which confirms the adequacy of its reliability and effectiveness
Install chemical detectors or "sniffers" for products of combustion
Install thermal detectors
Maintain visual surveillance over critical areas

Figure 7. (Sheet 4 of 10) Early Fire Detection.

Note:
"Adequate" capabilities
imply that levels of effectiveness
and reliability are commensurate
with importance

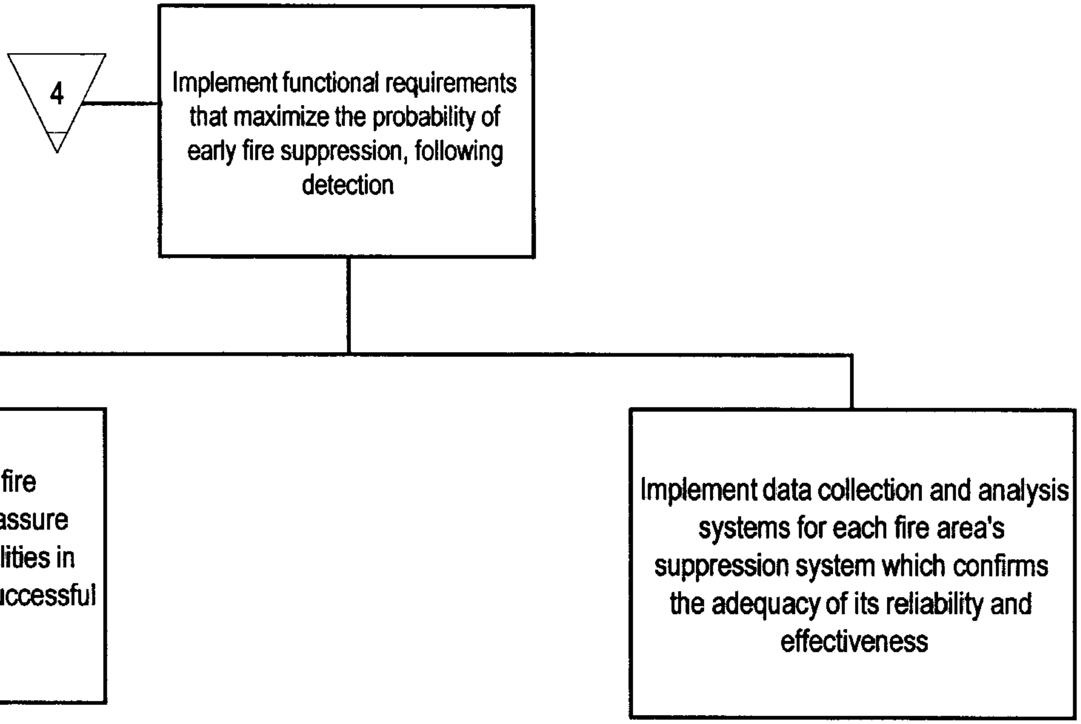
4
Implement functional requirements that maximize the probability of early fire suppression, following detection
Implement redundant/diverse fire suppression systems which will assure adequate fire suppression capabilities in each room or fire area, following successful fire detection
Implement data collection and analysis systems for each fire area's suppression system which confirms the adequacy of its reliability and effectiveness

Figure 7. (Sheet 5 of 10) Number of Important Components.

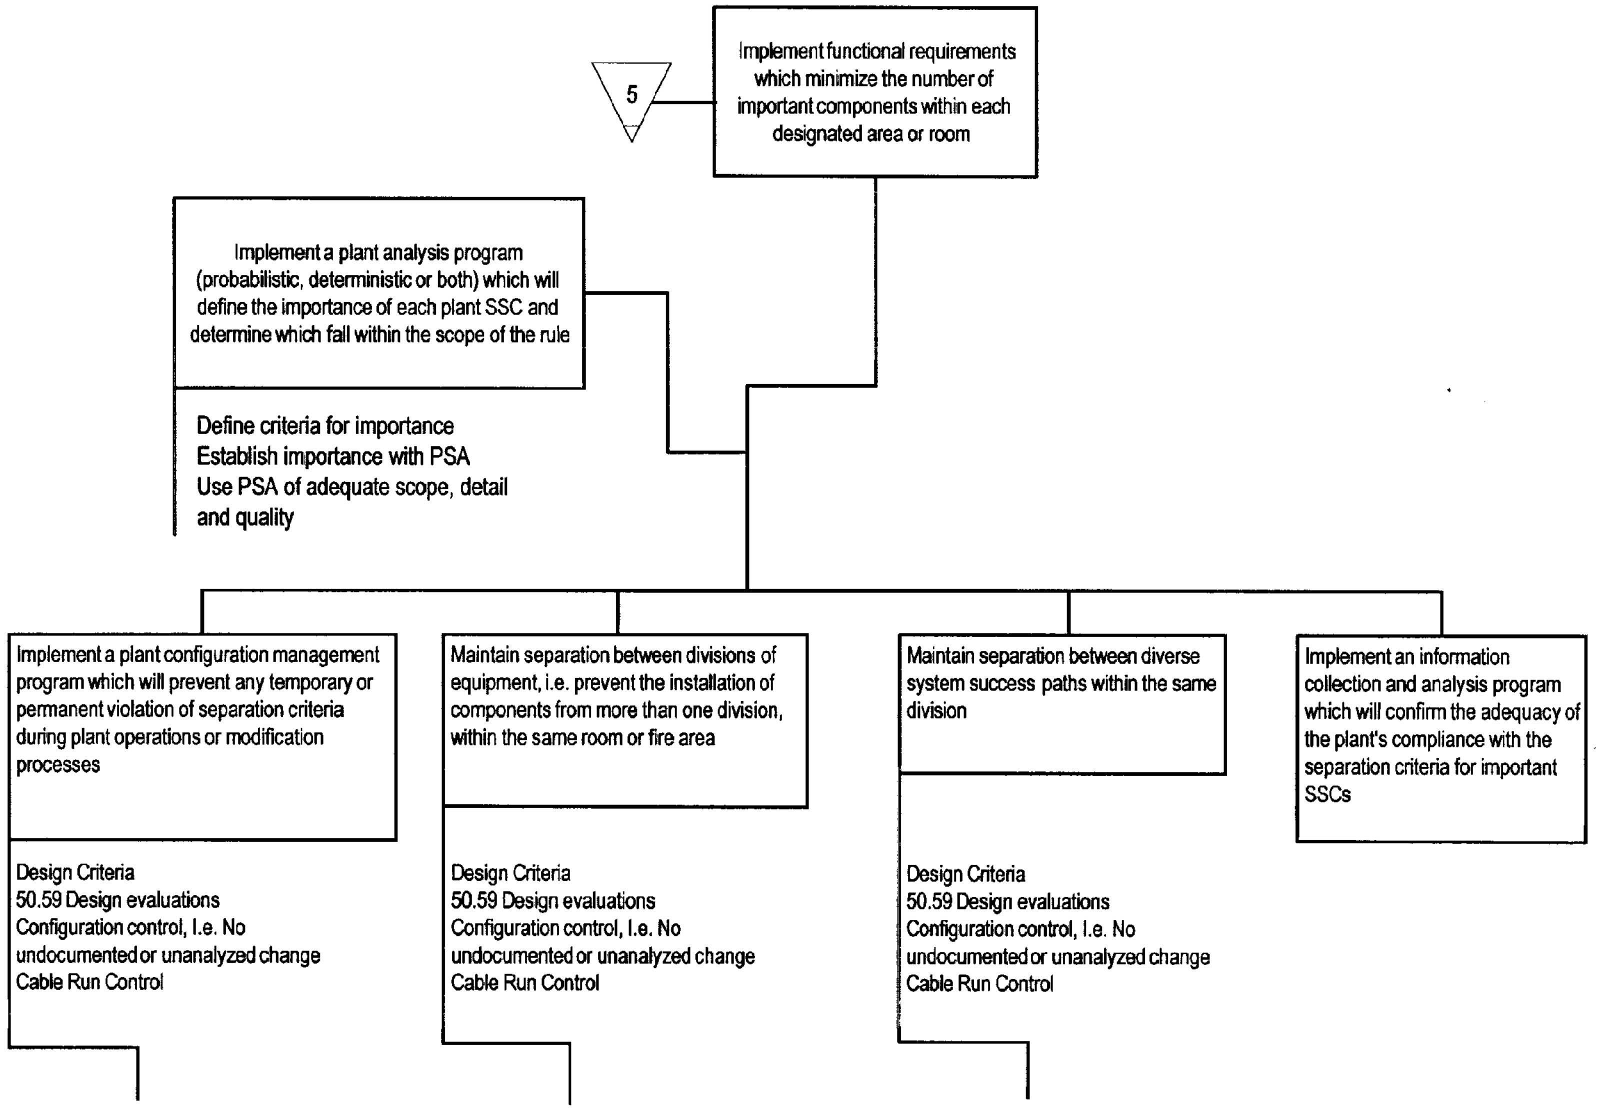
5
Implement functional requirements which minimize the number of important components within each designated area or room
Implement a plant analysis program (probabilistic, deterministic or both) which will define the importance of each plant SSC and determine which fall within the scope of the rule
Define criteria for importance
Establish importance with PSA
Use PSA of adequate scope, detail and quality
Implement a plant configuration management program which will prevent any temporary or permanent violation of separation criteria during plant operations or modification processes
Design Criteria
50.59 Design evaluations
Configuration control, I.e. No undocumented or unanalyzed change
Cable Run Control
Maintain separation between divisions of equipment, i.e. prevent the installation of components from more than one division, within the same room or fire area
Design Criteria
50.59 Design evaluations
Configuration control, I.e. No undocumented or unanalyzed change
Cable Run Control
Maintain separation between diverse system success paths within the same division
Design Criteria
50.59 Design evaluations
Configuration control, I.e. No undocumented or unanalyzed change
Cable Run Control
Implement an information collection and analysis program which will confirm the adequacy of the plant's compliance with the separation criteria for important SSCs

Figure 7. (Sheet 6 of 10) Number of Components in a Room or Fire Area

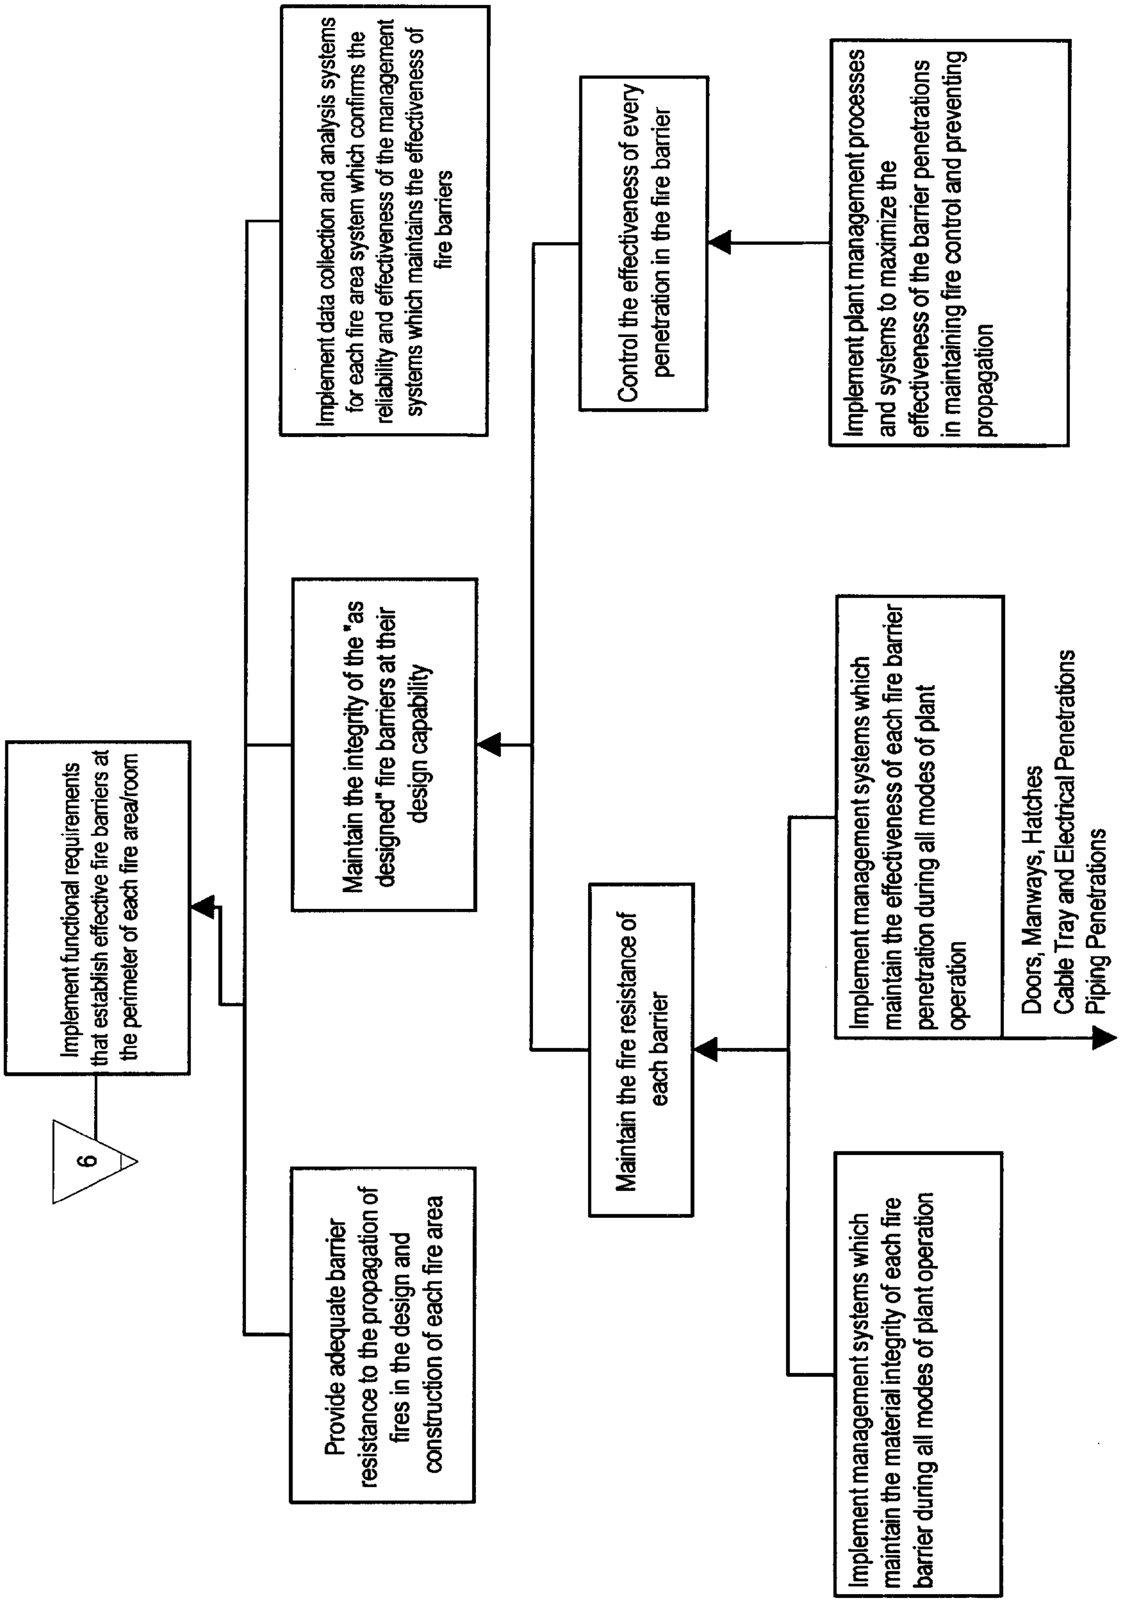
6
Implement functional requirements that establish effective fire barriers at the perimeter of each fire area/room
Provide adequate barrier resistance to the propagation of fires in the design and construction of each fire area
Maintain the integrity of the "as designed" fire barriers at their design capability
Implement data collection and analysis systems for each fire area system which confirms the reliability and effectiveness of the management systems which maintains the effectiveness of fire barriers
Maintain the fire resistance of each barrier
Control the effectiveness of every penetration in the fire barrier
Implement plant management processes and systems to maximize the effectiveness of the barrier penetrations in maintaining fire control and preventing propagation
Implement management systems which maintain the material integrity of each fire barrier during all modes of plant operation
Implement management systems which maintain the effectiveness of each fire barrier penetration during all modes of plant operation
Doors, Manways, Hatches
Cable Tray and Electrical Penetrations
Piping Penetrations

Figure 7. (Sheet 7 of 10) Fire Barriers.

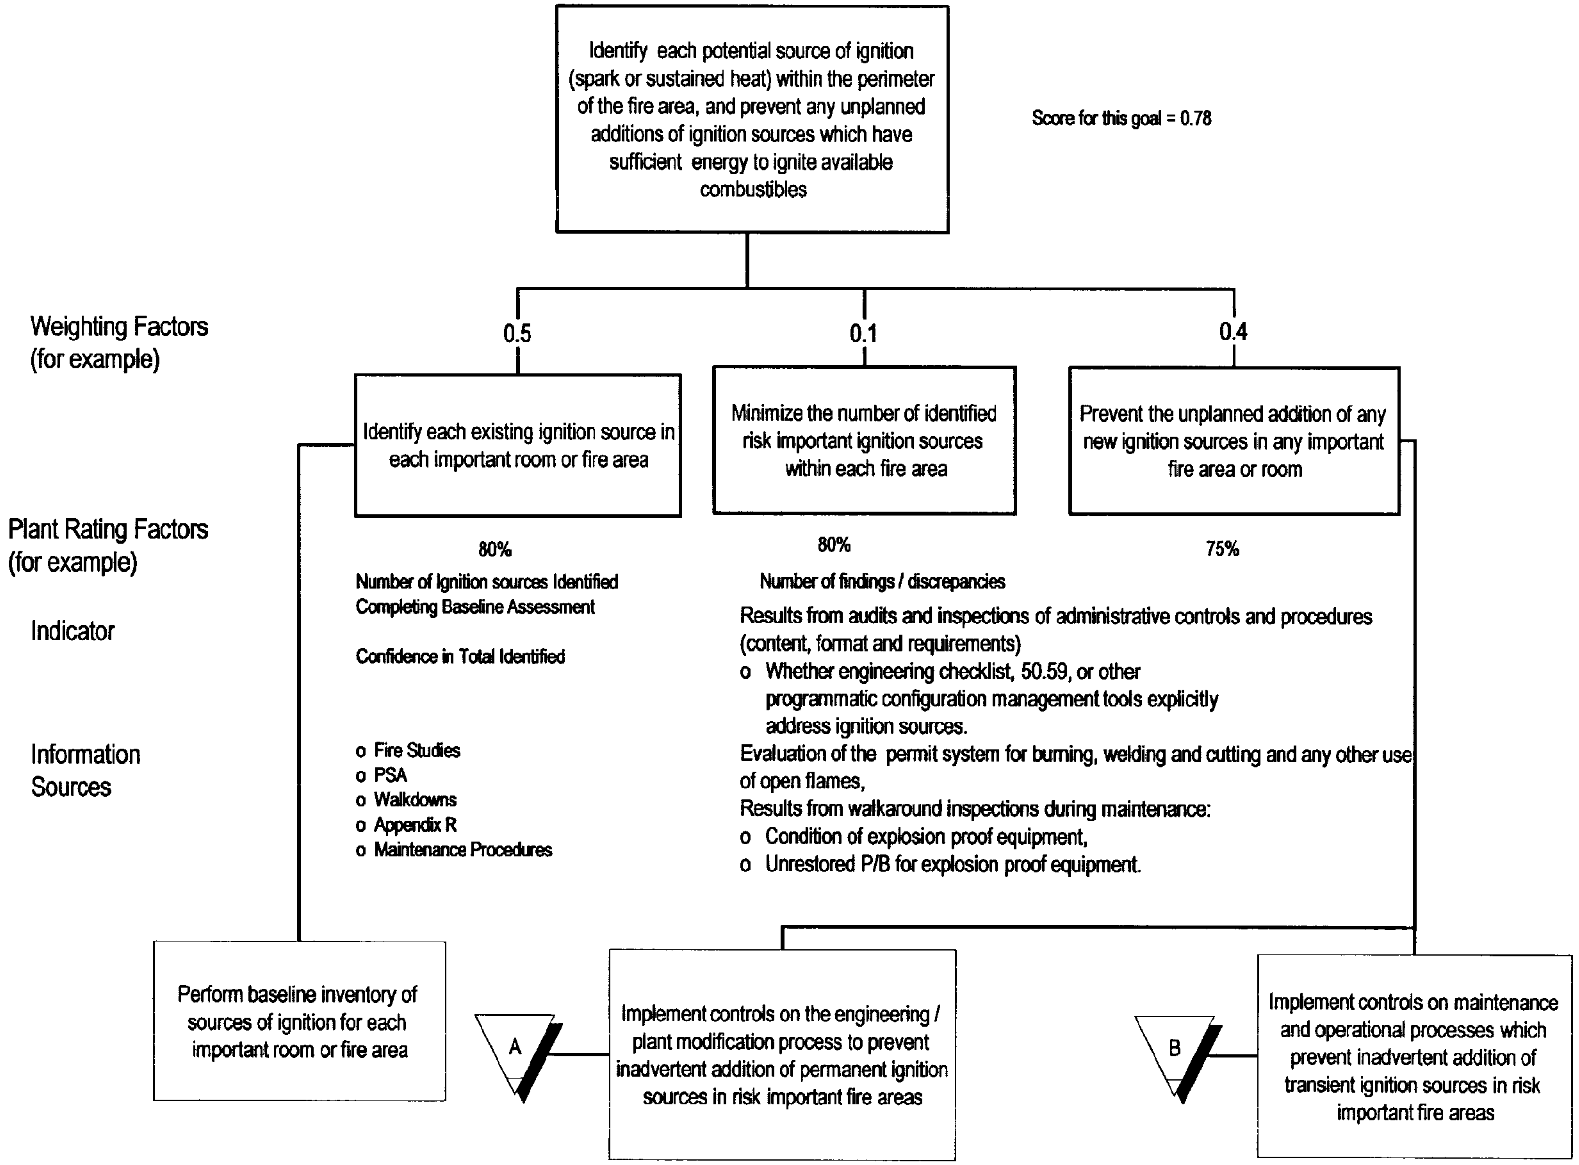
Identify each potential source of ignition (spark or sustained heat) within the perimeter of the fire area, and prevent any unplanned additions of ignition sources which have sufficient energy to ignite available combustibles
Score for this goal = 0.78
Weighting Factors (for example)
0.5
0.1
0.4
Identify each existing ignition source in each important room or fire area
Minimize the number of identified risk important ignition sources within each fire area
Prevent the unplanned addition of any new ignition sources in any important fire area or room
Plant Rating Factors (for example)
80%
80%
75%
Indicator
Number of Ignition sources Identified Completing Baseline Assessment
Confidence in Total Identified
Number of findings / discrepancies
Information Sources
o Fire Studies
o PSA
o Walkdowns
o Appendix R
o Maintenance Procedures
Results from audits and inspections of administrative controls and procedures (content, format and requirements)
o Whether engineering checklist, 50.59, or other programmatic configuration management tools explicitly address ignition sources.
Evaluation of the permit system for burning, welding and cutting and any other use of open flames,
Results from walkaround inspections during maintenance:
o Condition of explosion proof equipment,
o Unrestored P/B for explosion proof equipment.
Perform baseline inventory of sources of ignition for each important room or fire area
A
Implement controls on the engineering / plant modification process to prevent inadvertent addition of permanent ignition sources in risk important fire areas
B
Implement controls on maintenance and operational processes which prevent inadvertent addition of transient ignition sources in risk important fire areas

Figure 7. (Sheet 8 of 10) Ignition Sources (More Detailed).

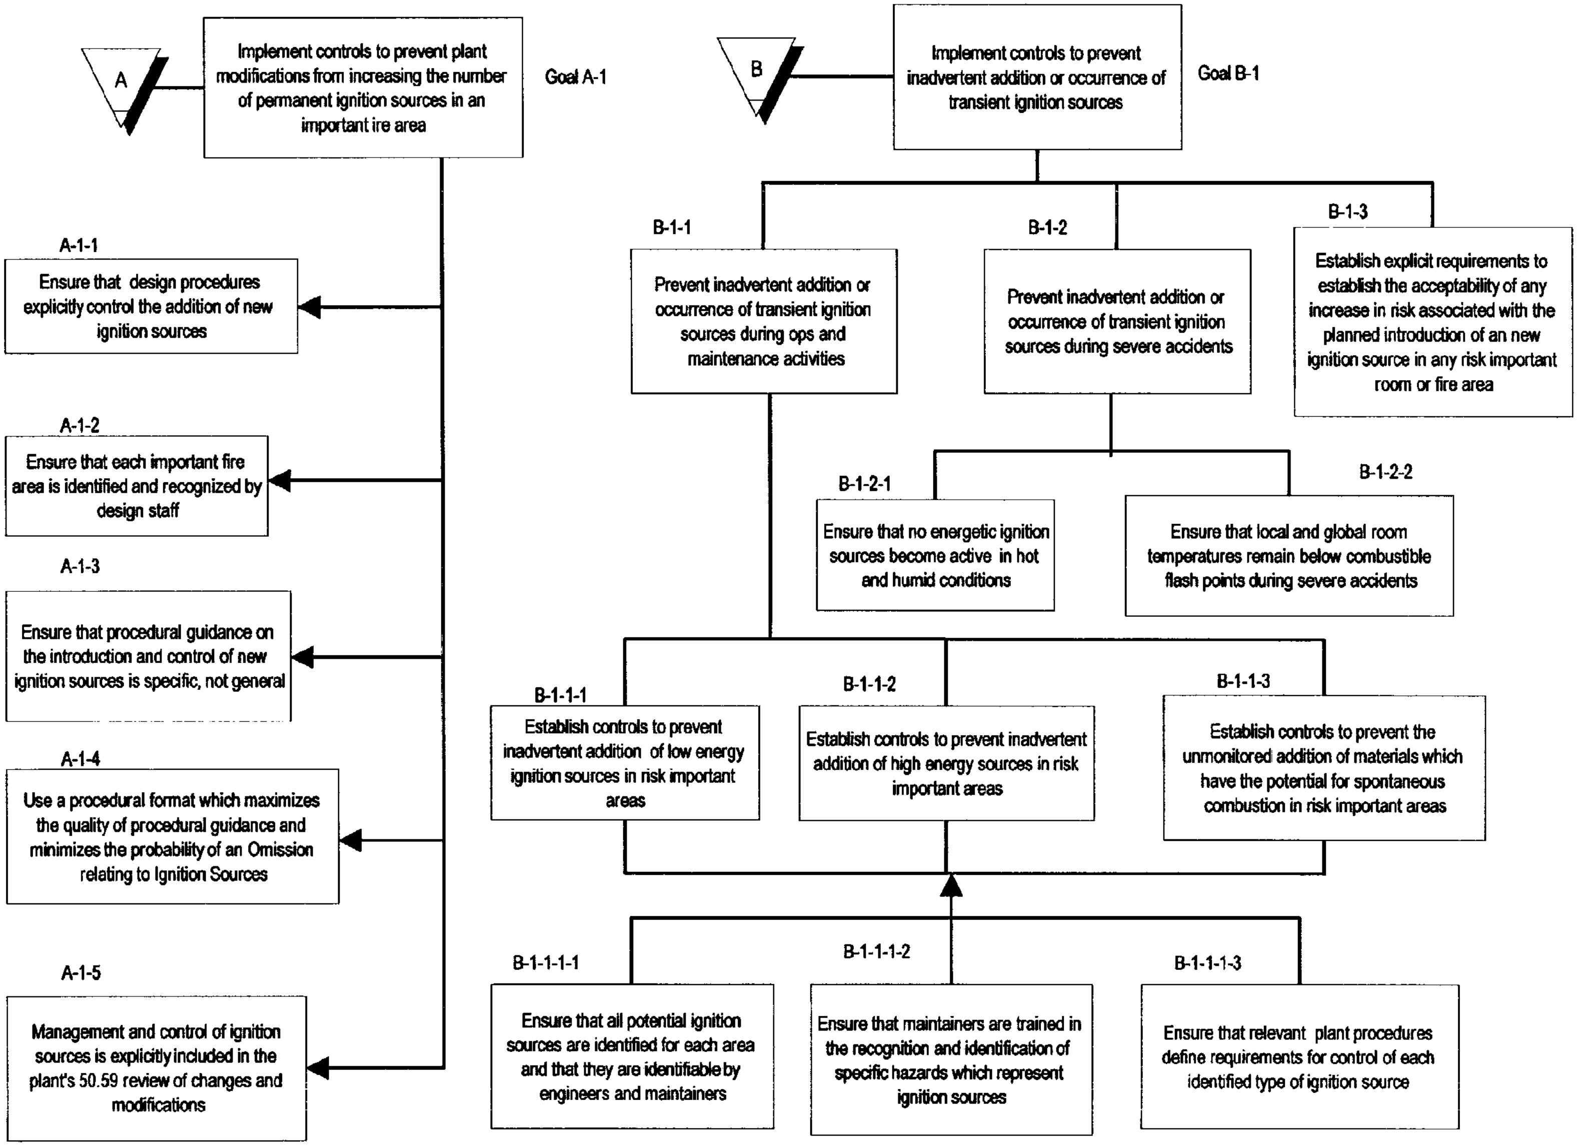
A
Implement controls to prevent plant modifications from increasing the number of permanent ignition sources in an important ire area
Goal A-1
A-1-1
Ensure that design procedures explicitly control the addition of new ignition sources
A-1-2
Ensure that each important fire area is identified and recognized by design staff
A-1-3
Ensure that procedural guidance on the introduction and control of new ignition sources is specific, not general
A-1-4
Use a procedural format which maximizes the quality of procedural guidance and minimizes the probability of an Omission relating to Ignition Sources
A-1-5
Management and control of ignition sources is explicitly included in the plant's 50.59 review of changes and modifications
B
Implement controls to prevent inadvertent addition or occurrence of transient ignition sources
Goal B-1
B-1-1
Prevent inadvertent addition or occurrence of transient ignition sources during ops and maintenance activities
B-1-2
Prevent inadvertent addition or occurrence of transient ignition sources during severe accidents
B-1-3
Establish explicit requirements to establish the acceptability of any increase in risk associated with the planned introduction of an new ignition source in any risk important room or fire area
B-1-2-1
Ensure that no energetic ignition sources become active in hot and humid conditions
B-1-2-2
Ensure that local and global room temperatures remain below combustible flash points during severe accidents
B-1-1-1
Establish controls to prevent inadvertent addition of low energy ignition sources in risk important areas
B-1-1-2
Establish controls to prevent inadvertent addition of high energy sources in risk important areas
B-1-1-3
Establish controls to prevent the unmonitored addition of materials which have the potential for spontaneous combustion in risk important areas
B-1-1-1-1
Ensure that all potential ignition sources are identified for each area and that they are identifiable by engineers and maintainers
B-1-1-1-2
Ensure that maintainers are trained in the recognition and identification of specific hazards which represent ignition sources
B-1-1-1-3
Ensure that relevant plant procedures define requirements for control of each identified type of ignition source

Figure 7. (Sheet 9 of 10) Ignition Sources (More Detailed).

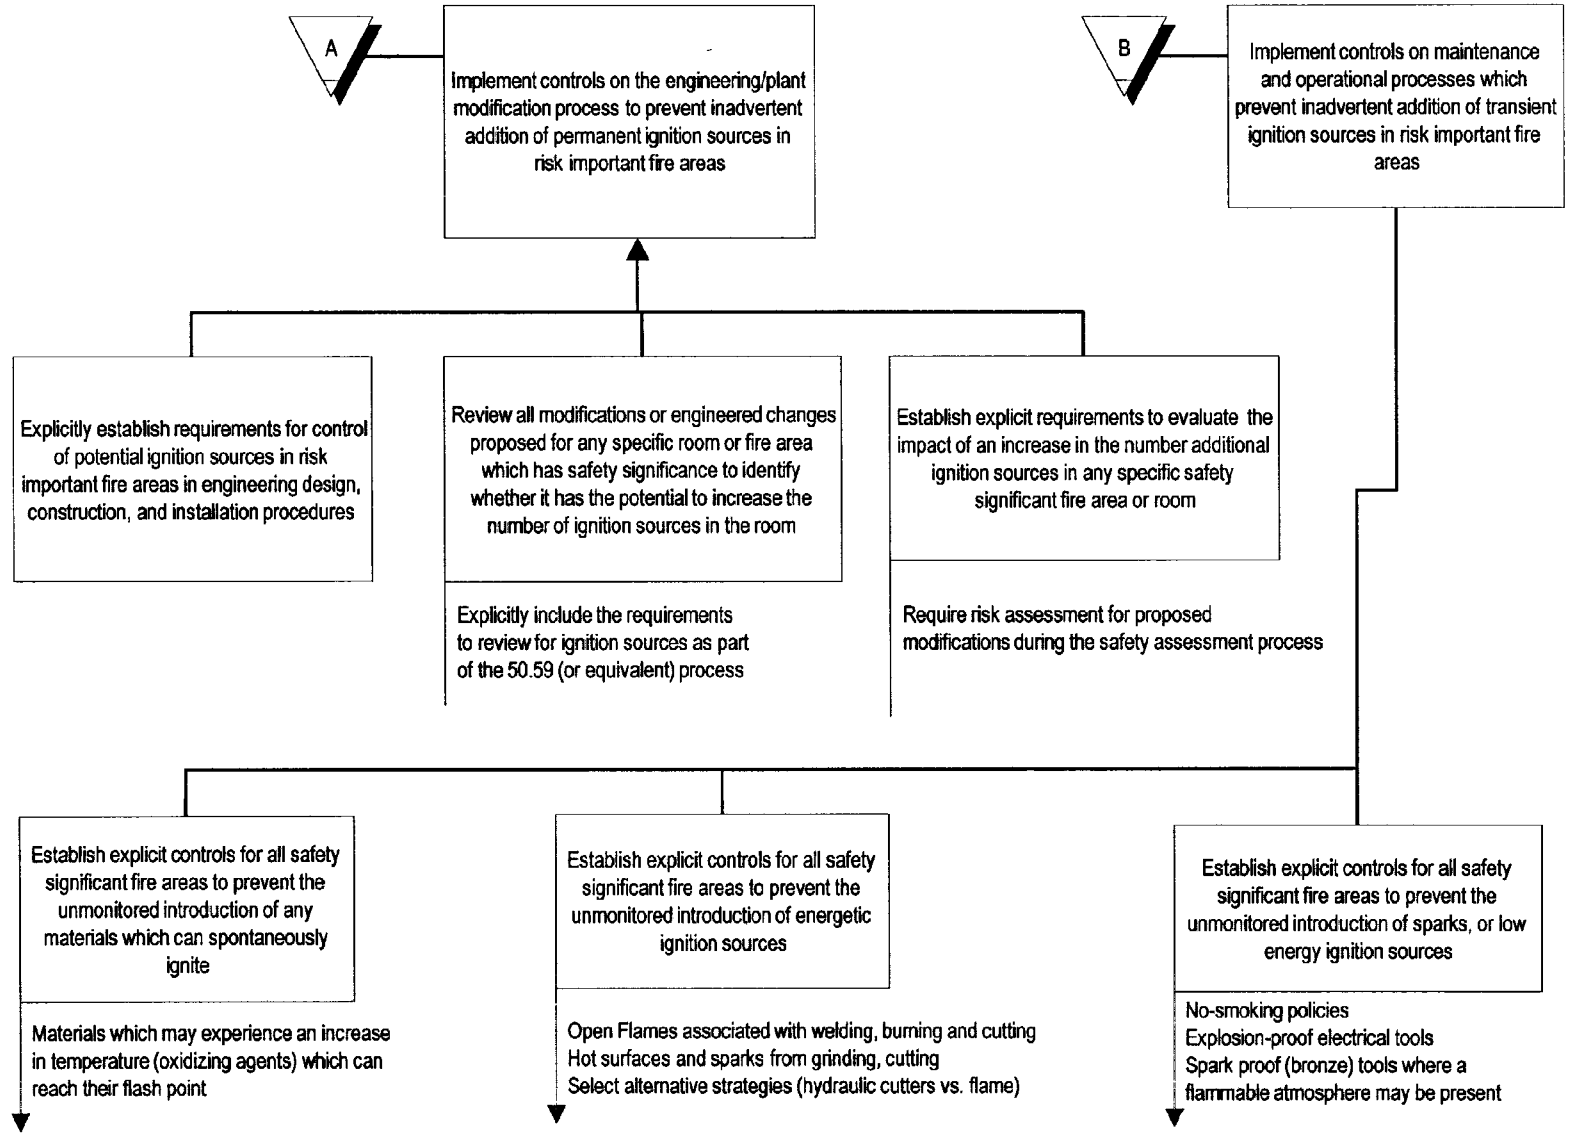
A
Implement controls on the engineering/plant modification process to prevent inadvertent addition of permanent ignition sources in risk important fire areas
B
Implement controls on maintenance and operational processes which prevent inadvertent addition of transient ignition sources in risk important fire areas
Explicitly establish requirements for control of potential ignition sources in risk important fire areas in engineering design, construction, and installation procedures
Review all modifications or engineered changes proposed for any specific room or fire area which has safety significance to identify whether it has the potential to increase the number of ignition sources in the room
Explicitly include the requirements to review for ignition sources as part of the 50.59 (or equivalent) process
Establish explicit requirements to evaluate the impact of an increase in the number additional ignition sources in any specific safety significant fire area or room
Require risk assessment for proposed modifications during the safety assessment process
Establish explicit controls for all safety significant fire areas to prevent the unmonitored introduction of any materials which can spontaneously ignite
Materials which may experience an increase in temperature (oxidizing agents) which can reach their flash point
Establish explicit controls for all safety significant fire areas to prevent the unmonitored introduction of energetic ignition sources
Open Flames associated with welding, burning and cutting
Hot surfaces and sparks from grinding, cutting
Select alternative strategies (hydraulic cutters vs. flame)
Establish explicit controls for all safety significant fire areas to prevent the unmonitored introduction of sparks, or low energy ignition sources
No-smoking policies
Explosion-proof electrical tools
Spark proof (bronze) tools where a flammable atmosphere may be present

Figure 7. (Sheet 10 of 10) Ignition Sources (More Detailed).

This goal is reduced to its constituent lower order goals as follows:

◊ *The licensee has established explicit controls which will prevent the inadvertent introduction of any low energy ignition sources in risk important areas which contain flammables:*

- * Potential ignition sources of concern are clearly identified to, and identifiable by plant staff
- * Maintainers are trained in the specific hazards associated with low energy ignition sources and the importance of not using them in an uncontrolled manner in risk important areas
- * Procedures define the exact requirements for control of each type of low energy ignition source whenever they are used in a safety important fire area or room

◊ *The licensee has established explicit controls which will prevent the inadvertent and uncontrolled introduction of any high energy ignition sources in risk important areas*

- * Ignition sources of concern are clearly identified to, and identifiable by plant staff
- * Procedures define the exact requirements for control of each type of ignition source whenever they are used in a safety important fire area or room
- * Maintainers are trained in the specific hazards associated with energetic ignition sources and the importance of not using them in an uncontrolled manner in risk important areas
- * Accountability procedures are used to track and control energetic ignition sources whenever they are used in risk important fire areas

◊ *"The licensee has established explicit controls which will prevent the inadvertent introduction of any materials which may spontaneously ignite in any room or fire area which is important to safety"*

- * Plant has identified (and eliminated?) all materials used in maintenance or operations which may ignite spontaneously
- * Plant has established accountability processes for transient combustibles needed during maintenance or operations which have the potential to ignite spontaneously. (This ensures that any materials of this type are controlled while in use and then removed. If they are to remain; the PSA will be updated to reflect the new room combustible loading and ignition frequency).

⇒ *"The licensee has established explicit controls to assess, and mitigate the effects of severe accident conditions and their potential for inducing consequential fires in individual rooms or fire areas":*

- * In humid or wet environments, confirm that there is no possibility of a highly energetic electrical fire (caused by poor insulation, inadequate electrical terminations or accumulation of electrolytes around high voltage/high current lines)

  - * In dry, hot environments, ensure that there is no possibility of local or ambient temperatures which approach the flash point of any combustibles in the area.

⇒ *"The licensee has established explicit requirements, to establish the acceptability of any risk increase which is associated with the addition of any planned new ignition sources which are identified during the design review process"*

  - * The plant PSA is used routinely to assess the risk significance of changes to the plant
  - * The plant PSA scope includes detailed assessment of the effects from fires in individual areas and rooms
  - * The plant PSA quality is adequate to make a confident prediction of the magnitude of the expected risk change which is associated with an increase in the number of ignition sources in a specific area

Summary Observations on Fire Protection

Extensive discussion of regulatory approaches will be deferred to a later section, but here, it is appropriate to observe that direct regulatory oversight of the performance aspects identified in the above development would be considered highly invasive. Monitoring these performance nodes would entail verification of performance of highly prescribed tasks, rather than judging performance based on results achieved.

This outcome should generalize to other safety functions to which challenges are rare. There may be less prescriptive regulatory approaches to be considered for these areas, but truly results-oriented performance-based regulation is faced with difficult issues in these areas. NEI's paper [4] on performance-based regulation expressed this reservation, using equipment qualification as an example:

> It is not always possible to mix risk-based and performance-based approaches. For example, although risk insights can be used to identify risk significant SSCs to which the environmental qualification rule would apply, it would be inappropriate to establish performance criteria under a performance-based approach because the equipment may never experience the conditions assumed in the rule. It would therefore not be possible to monitor performance. However, implementation of the regulations for those SSCs could still be accomplished through programmatic or prescriptive regulations.

Independently of regulatory approach, addressing low-level performance aspects arises in the context of the licensee's management of his own performance, along the lines apparently intended by the originators of the diamond tree concept. As we will see later, the fact that monitoring requirements would drive us to low-level nodes in some areas does not foreclose all less-prescriptive regulatory options; one can consider less-prescriptive approaches at this level that are less outcome-based than a strictly performance-based approach would be.

The above remarks are not intended to contradict extant discussions of performance-based approaches to fire protection (see Appendix A). Rather, it is believed that those discussions are using the phrase “performance-based” in a different way: to refer to regulations that are based on physical performance requirements derived from mechanistic analysis of realistic scenarios. This is very different from assessing licensee performance of fire protection activities based on results.

3.2.2 Shutdown Example

Following is a conceptual diamond tree covering essentially the same territory as the logic model presented in Appendix B.

The first sheet of Figure 8 is an overview. In this particular development, there is a band in the middle labeled “MPLD [for ‘master plant logic diagram’] for Mid-Loop Operation.” This band represents the system level, train level, and component levels portrayed on many other diamond-tree illustrations in this report. As will be seen on subsequent figures in this series, most of the lower-lying performance nodes connect to many of the hardware nodes, and a shortcut was taken on this figure, connecting essentially all human performance aspects to essentially all hardware in this area. It is not literally true that all human performance aspects really apply to all hardware (for example, some of the performance nodes are specifically aimed at active components, others at passive components).

The subsequent sheets are essentially self-explanatory. However, the hierarchical context of each performance node is not always clear. The following discussion is a step towards relating each performance item to a hierarchical level.

The first two sheets of Figure 8 focus on causes of “initiating events,” meaning in the present context “loss of RHR.” In Figure 8 Sheet 2, Initiating Events and Their Institutional Influences (Active Components), we see several kinds of events. At the top is a performance node that basically plugs into many RHR-related components. Below this node are diverse nodes that correspond to programs: investigate repetitive failures, maximize the reliability of active RHR components, assure the quality of maintenance practices, RCM, testing, inspection,…The details of these nodes, if shown explicitly, should logically (in the inverted diamond-tree sense) be placed physically above the program-level nodes from which they emanate. There is also a node calling for verification of the operability of the standby RHR train, which, in our model, would need to fail as part of the initiating event. This verification would arguably be procedural, because it is the kind of thing that an outage management plan would require before entry into a reduced-inventory mode was allowed. On the right, there is a node calling for reliability and availability of support systems. Later, we will argue for treating key support systems as if they were key front-line systems, and a diamond tree developed in that way would have reliability and availability of support systems effectively at the “system” level, physically several levels above this level. Of course, human and institutional performance aspects at this level would then act to support that higher-level objective. On the next sheet (Figure 8 Sheet 3, Initiating Events…Passive), we see a number of controls that would act directly on passive components. These seem essentially procedural in nature.

The following sheet (Figure 8 Sheet 4, Conditional and Subsequent Failures and Their Institutional Influences) begins to address activities that assure performance of the fall-back strategies that come into play after an initiating event occurs. By definition, the systems on this sheet are either in standby or in a kind of service that does not bear directly on residual heat removal in the current configuration, so the performance aspects necessarily relate to standby systems. On the left, we see an activity to verify operability of key standby trains, essentially the kind of verification activity that an outage management plan would likely mandate. This would be procedural in nature. Most of the rest of the nodes on this sheet are programmatic in nature.

On Figure 8 Sheet 5, Recovery Actions and Institutional Influences, we see different kinds of nodes. On the left, under "provide as much time as possible," we see nodes that pertain to some extent to a design phase and to some extent to a procedural node. By "design phase," we refer to the argument given elsewhere that performance-based oversight of a plant should be based on a systematically developed safety case. That is, for clarity, one should articulate a safety case including a set of performance objectives whose ongoing verification is the subject of performance-based regulation. A node that says "maximize success paths" or "maximize available time" belongs properly to that previous "design" or "certification" phase, and not to the implementation phase. Finally, under "maximize the probability that the operator will quickly detect…" we have a node calling for cuing and instrumentation, both for diagnosis and for event management. This node is developed on a later sheet (Figure 8 Sheet 8, Information Needed for Detection and Diagnostic Processes), with examples of desirable kinds of information. As one would expect, level information is called for. Later, based in part on a discussion of a loss of inventory event, we will argue that instrumentation, too, should be treated at the "system" level for present purposes.

On the next sheet, we see two areas treated. The first (Figure 8 Sheet 6, Recovery Actions and Institutional Influences [Errors of Omission and Commission]) contains nodes that refer to procedures, to training, and to human factors engineering. The second (Figure 8 Sheet 7, Recovery from Errors of Omission and Commission), refers again to instrumentation and procedural guidance, and this time adds a reference to simulator training.

These sheets mention many kinds of performance nodes. Based on the above summary review, it appears to be possible in retrospect to correlate most of them to a reasonably clear hierarchy, even though these sheets were prepared in a relatively free-associating way.

3.3 Summary Comments on Diamond Trees

Based on the examples discussed above, a level hierarchy that seems to span the variations encountered above is given in Figure 9.

The important thing about diamond trees for present purposes is the hierarchical presentation of performance nodes. This makes the diamond tree a useful conceptual tool for discussing degrees to which a given regulatory approach is performance-based. It will be argued later that for essentially the same reason, the tree is useful for discussing significant events.

The diamond tree concept has not been applied widely, so that there seems to be no standard format for developing diamond trees. The various levels are given slightly different names in different illustrations. Some diamond trees have logic gates displayed on them (cf. some portions of the appendix in the Wreathall report), while some do not (cf. other portions of the appendix in the Wreathall report, the illustration in Chapter 1 of Modarres's book, and the examples given here). It presently appears that the main usefulness of the idea is as a thought tool, rather than as a basis for a calculation.

The diamond tree seems to lend itself to identification of very specific operations tasks that strongly influence hardware performance. But these tasks seem very far down in the tree to be appropriate as targets of regulatory attention directly. They may well be very useful in the plant's program of monitoring its own performance.

Diamond trees seem to have the potential to be quite voluminous, relative to logic models. A fault tree is essentially an abstraction of logical relationships into graphical form, with event definitions stripped down to bare essentials. A disciplined approach to fault-tree development and event definition can keep the size of the tree under control. On the other hand, the surveyed examples of diamond trees indicate that a diamond tree tends to be more of a catalog of influences, and that the spirit of the development is to reflect whatever there is to reflect. A plant activity worth carrying out, for example, probably needs to show up on a diamond tree, but very few such activities would show up on a fault tree. Moreover, there are potentially very many connections between nodes on different levels.

In short, diamond trees seem to be more of an abstract thought tool than a basis for a calculational algorithm. Partly for this reason, a later recommendation to "develop a diamond tree" will be tempered to say something like "the analyst should develop those portions of a diamond tree that are required before proceeding to the next step." However, it will be argued later that the core of the concept—the hierarchy of performance types—is a very useful thought tool for comparing regulatory approaches and even for discussing the significance of operating events. This is the reason for driving towards a standard version of Figure 9.

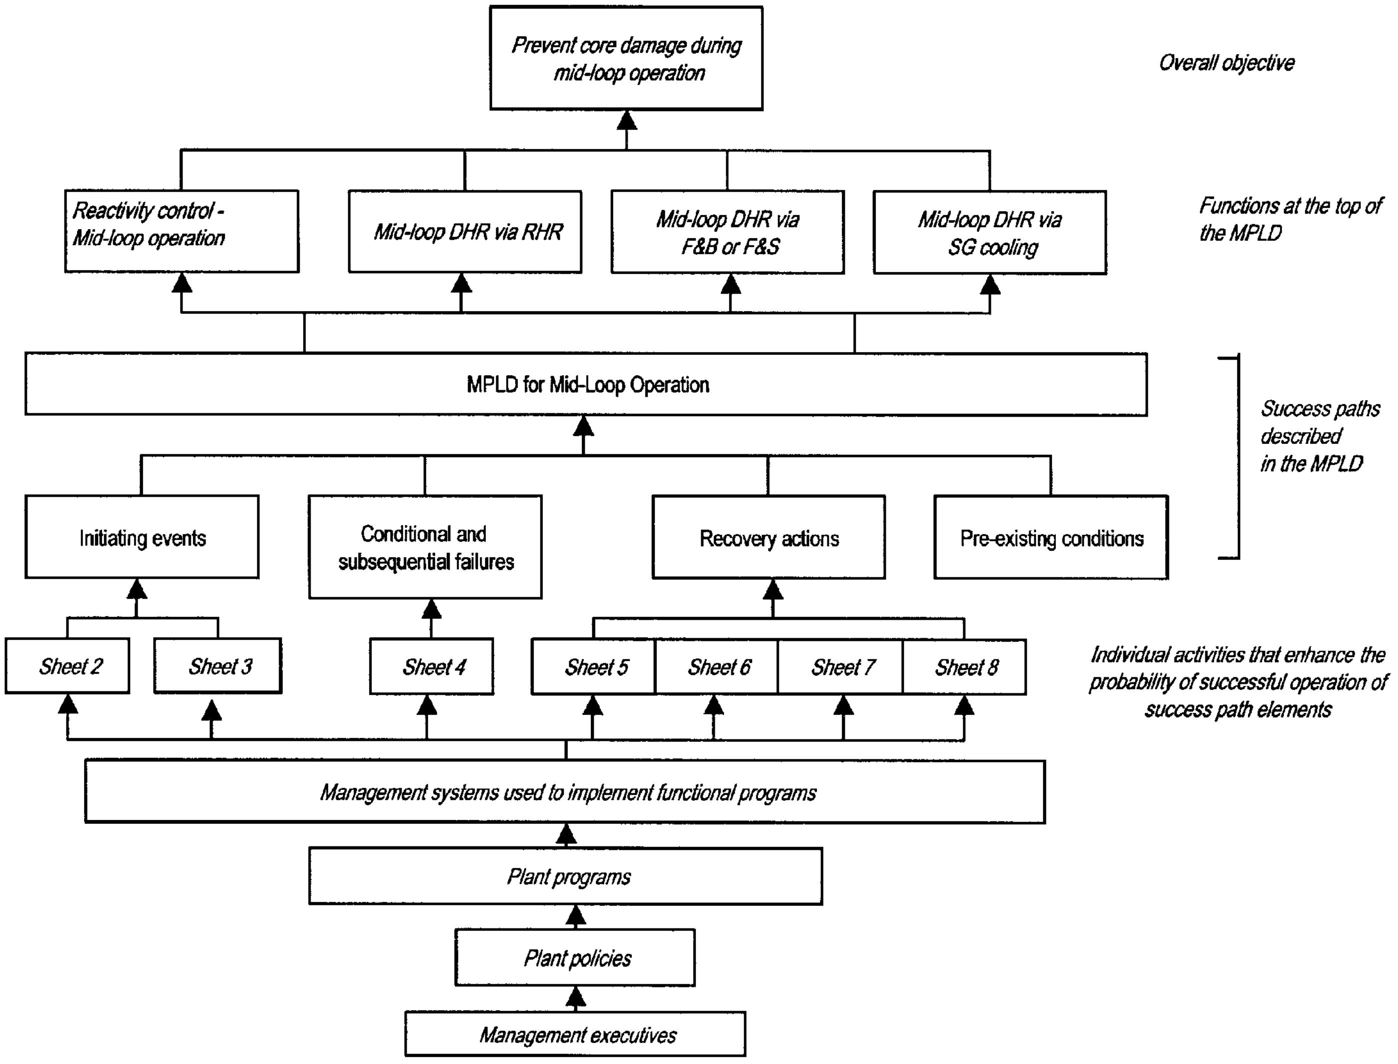
Prevent core damage during mid-loop operation
Overall objective
Reactivity control - Mid-loop operation
Mid-loop DHR via RHR
Mid-loop DHR via F&B or F&S
Mid-loop DHR via SG cooling
Functions at the top of the MPLD
MPLD for Mid-Loop Operation
Success paths described in the MPLD
Initiating events
Conditional and subsequential failures
Recovery actions
Pre-existing conditions
Sheet 2
Sheet 3
Sheet 4
Sheet 5
Sheet 6
Sheet 7
Sheet 8
Individual activities that enhance the probability of successful operation of success path elements
Management systems used to implement functional programs
Plant programs
Plant policies
Management executives

Figure 8. (Sheet 1 of 8) The Diamond Tree for Reduced Inventory Operation.

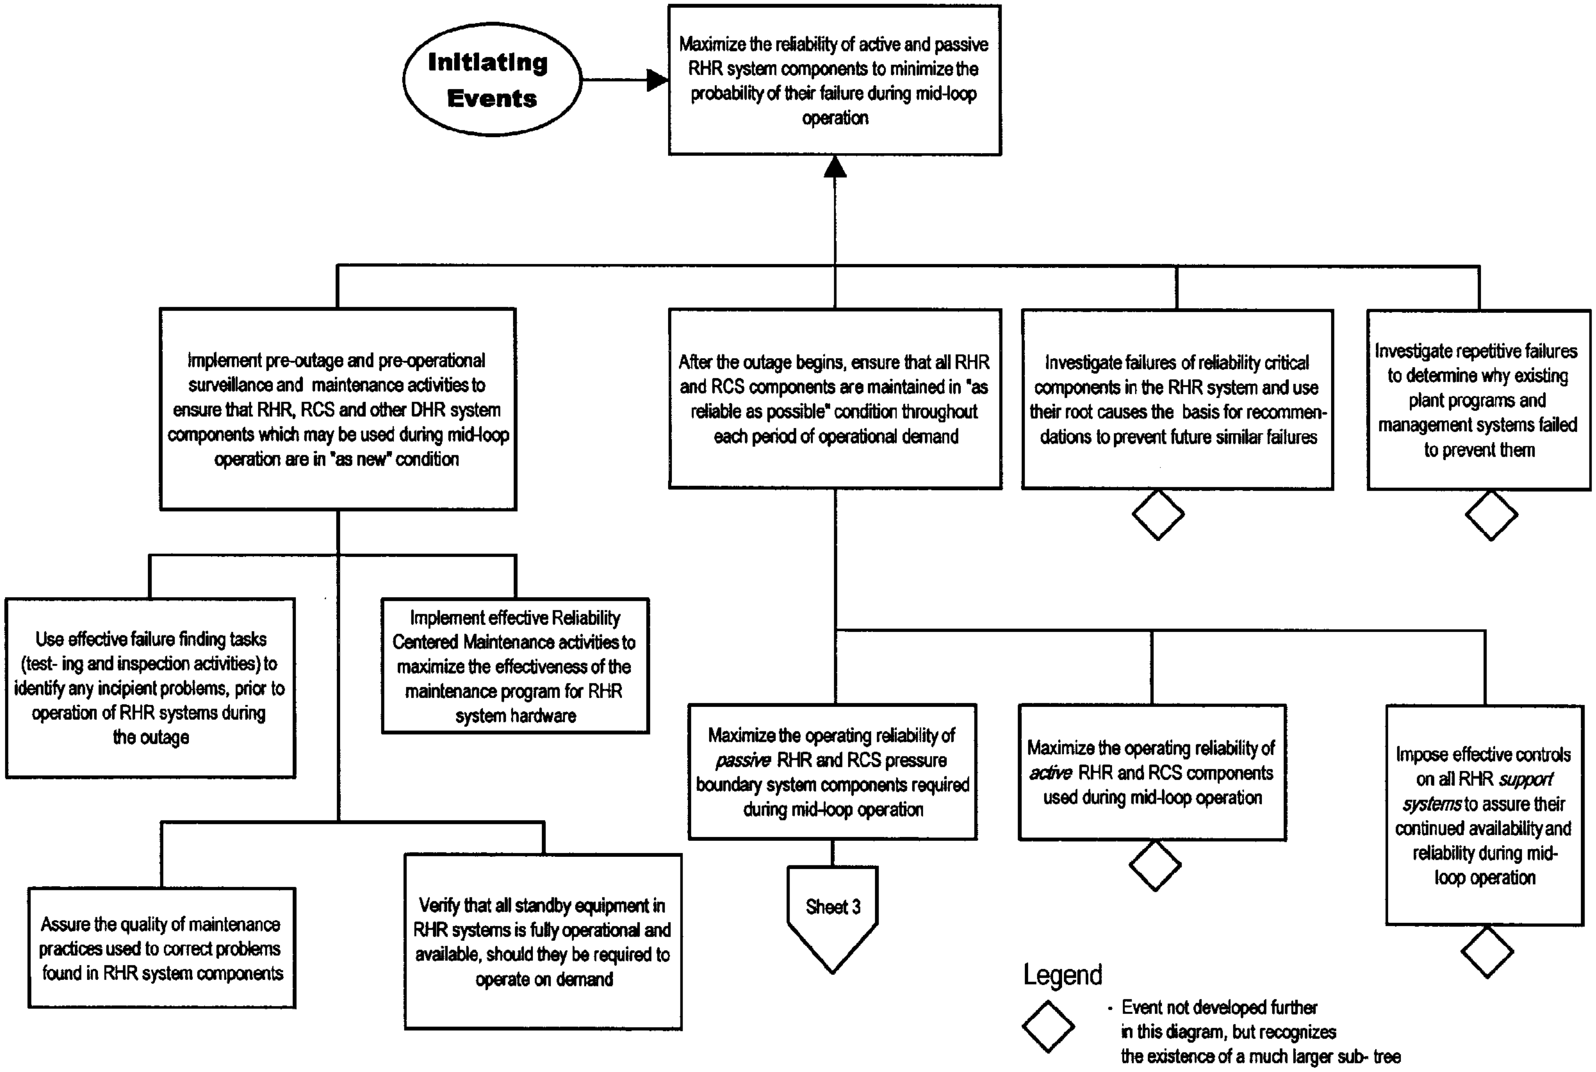
Initiating Events
Maximize the reliability of active and passive RHR system components to minimize the probability of their failure during mid-loop operation
Implement pre-outage and pre-operational surveillance and maintenance activities to ensure that RHR, RCS and other DHR system components which may be used during mid-loop operation are in "as new" condition
After the outage begins, ensure that all RHR and RCS components are maintained in "as reliable as possible" condition throughout each period of operational demand
Investigate failures of reliability critical components in the RHR system and use their root causes the basis for recommen-dations to prevent future similar failures
Investigate repetitive failures to determine why existing plant programs and management systems failed to prevent them
Use effective failure finding tasks (test- ing and inspection activities) to identify any incipient problems, prior to operation of RHR systems during the outage
Implement effective Reliability Centered Maintenance activities to maximize the effectiveness of the maintenance program for RHR system hardware
Assure the quality of maintenance practices used to correct problems found in RHR system components
Verify that all standby equipment in RHR systems is fully operational and available, should they be required to operate on demand
Maximize the operating reliability of *passive* RHR and RCS pressure boundary system components required during mid-loop operation
Sheet 3
Maximize the operating reliability of *active* RHR and RCS components used during mid-loop operation
Impose effective controls on all RHR *support systems* to assure their continued availability and reliability during mid-loop operation
Legend
· Event not developed further in this diagram, but recognizes the existence of a much larger sub- tree

Figure 8. (Sheet 2 of 8) Initiating Events and Their Institutional Influences (Active Components).

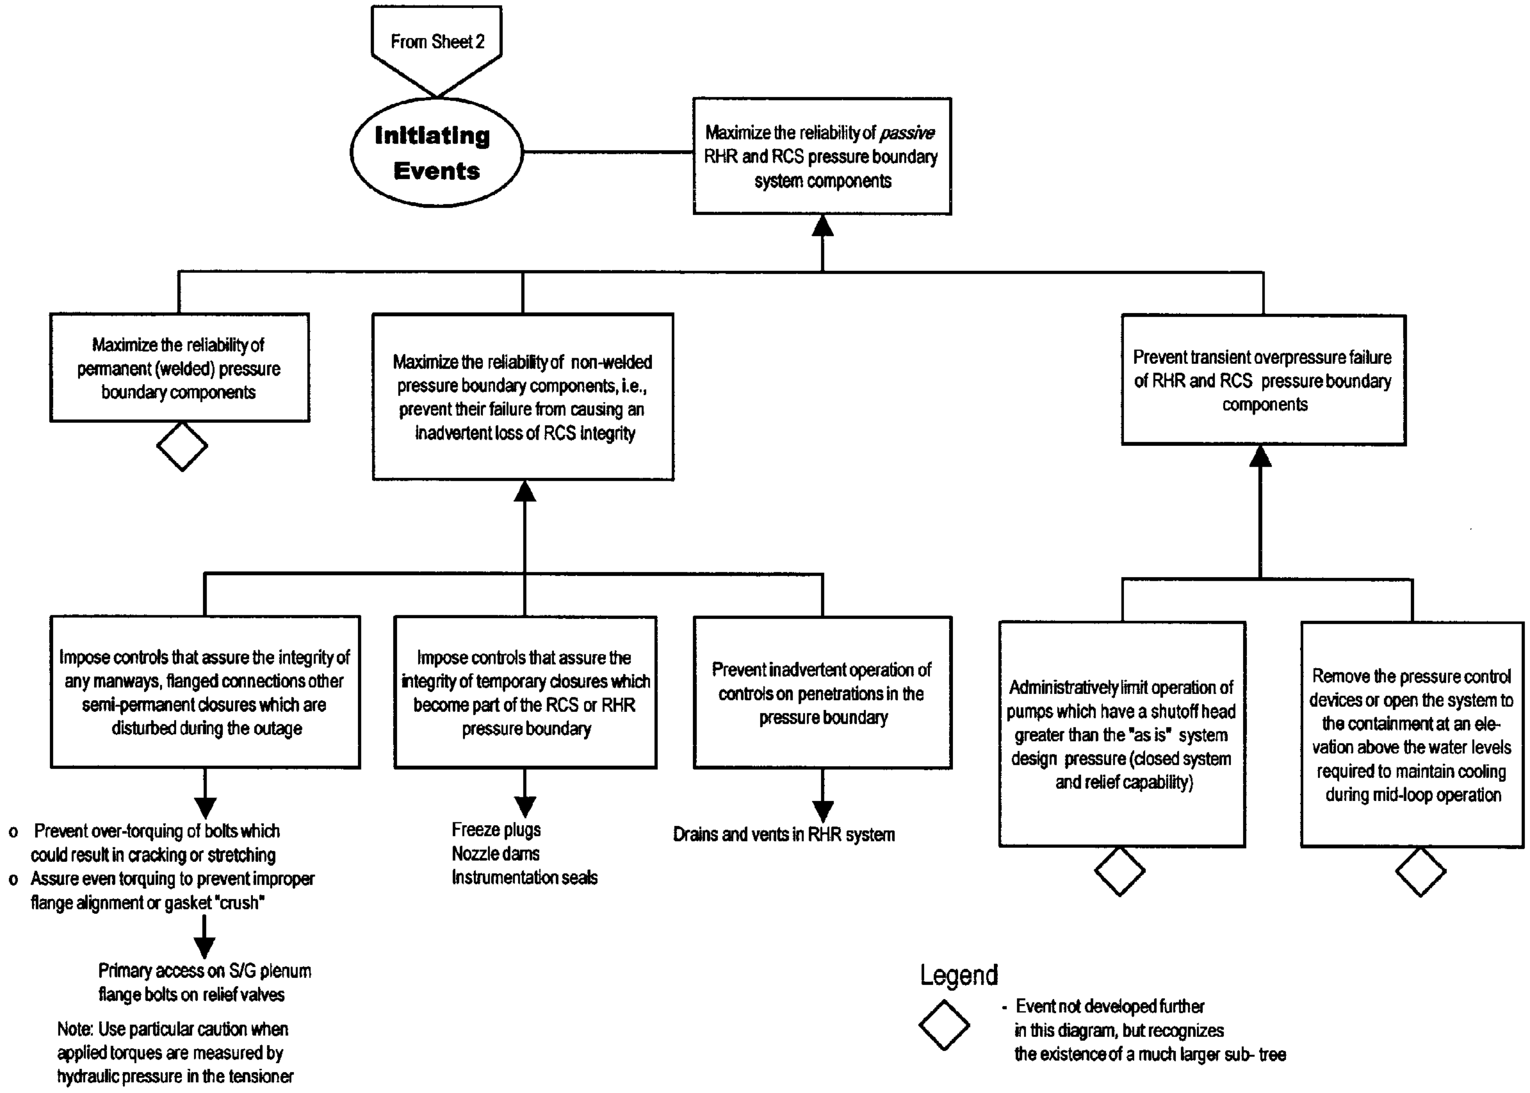
From Sheet 2
Initiating Events
Maximize the reliability of *passive* RHR and RCS pressure boundary system components
Maximize the reliability of permanent (welded) pressure boundary components
Maximize the reliability of non-welded pressure boundary components, i.e., prevent their failure from causing an inadvertent loss of RCS integrity
Prevent transient overpressure failure of RHR and RCS pressure boundary components
Impose controls that assure the integrity of any manways, flanged connections other semi-permanent closures which are disturbed during the outage
Impose controls that assure the integrity of temporary closures which become part of the RCS or RHR pressure boundary
Prevent inadvertent operation of controls on penetrations in the pressure boundary
Administratively limit operation of pumps which have a shutoff head greater than the "as is" system design pressure (closed system and relief capability)
Remove the pressure control devices or open the system to the containment at an elevation above the water levels required to maintain cooling during mid-loop operation
o Prevent over-torquing of bolts which could result in cracking or stretching
o Assure even torquing to prevent improper flange alignment or gasket "crush"
Primary access on S/G plenum flange bolts on relief valves
Note: Use particular caution when applied torques are measured by hydraulic pressure in the tensioner
Freeze plugs
Nozzle dams
Instrumentation seals
Drains and vents in RHR system
Legend
◇ - Event not developed further in this diagram, but recognizes the existence of a much larger sub- tree

Figure 8. (Sheet 3 of 8) Initiating Events and Their Institutional Influences (Passive Components).

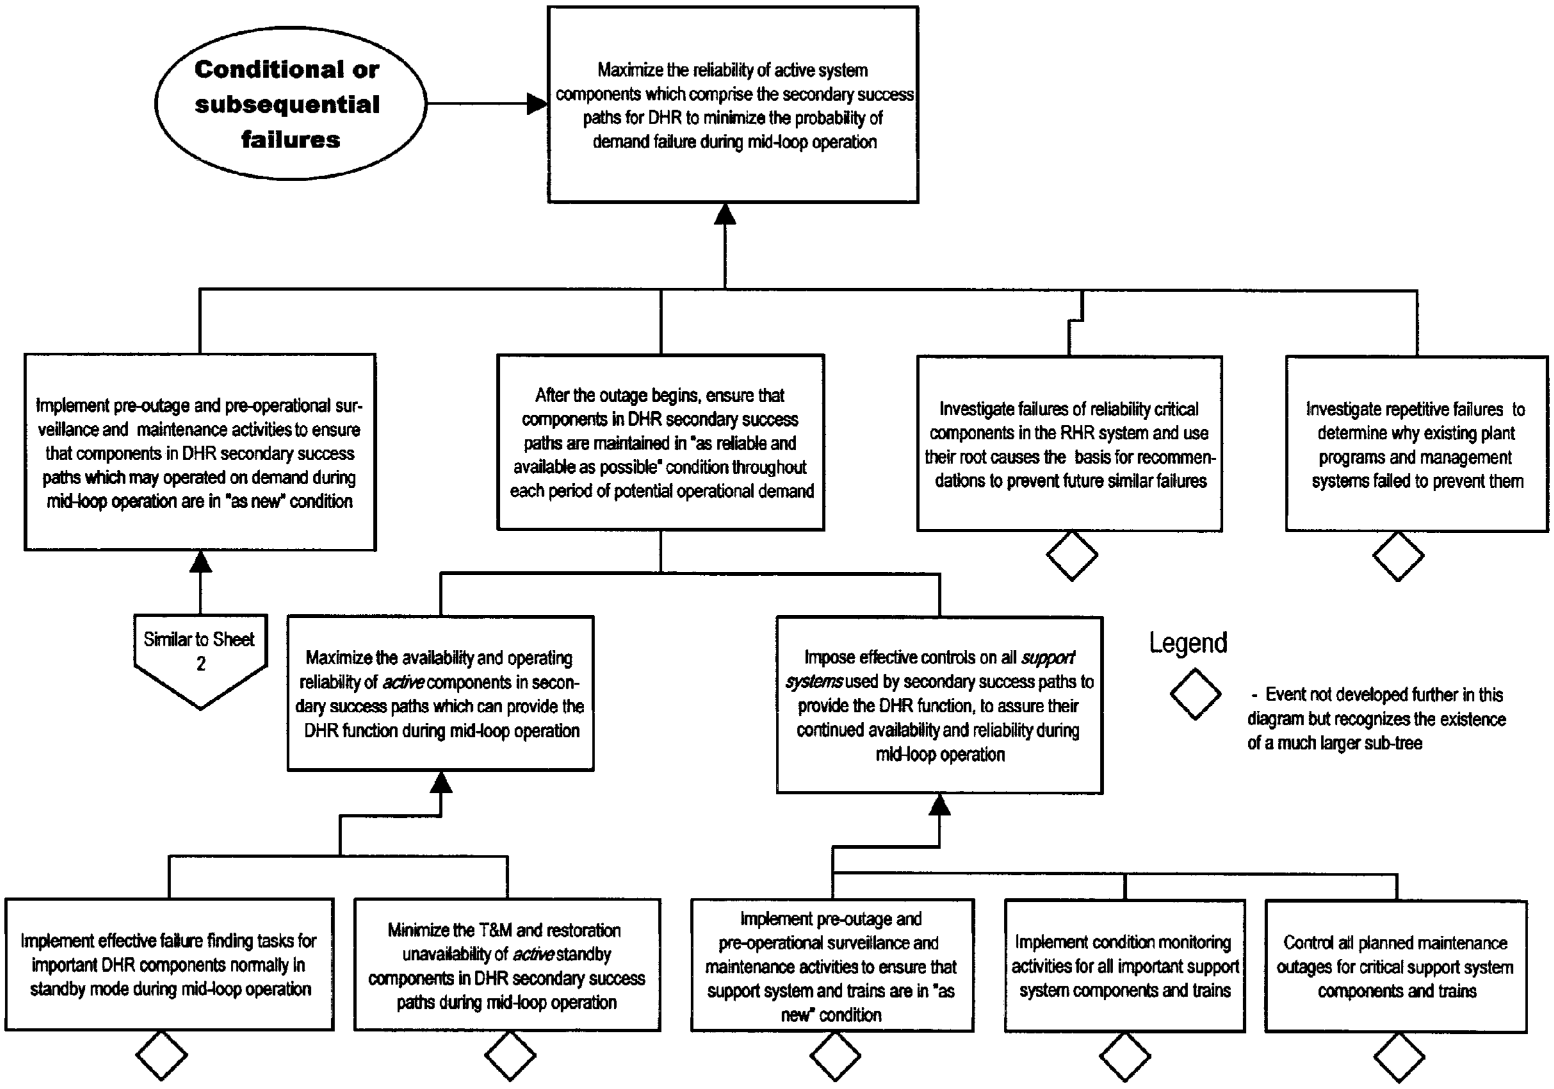
Conditional or subsequential failures
Maximize the reliability of active system components which comprise the secondary success paths for DHR to minimize the probability of demand failure during mid-loop operation
Implement pre-outage and pre-operational surveillance and maintenance activities to ensure that components in DHR secondary success paths which may operated on demand during mid-loop operation are in "as new" condition
After the outage begins, ensure that components in DHR secondary success paths are maintained in "as reliable and available as possible" condition throughout each period of potential operational demand
Investigate failures of reliability critical components in the RHR system and use their root causes the basis for recommendations to prevent future similar failures
Investigate repetitive failures to determine why existing plant programs and management systems failed to prevent them
Similar to Sheet 2
Maximize the availability and operating reliability of *active* components in secondary success paths which can provide the DHR function during mid-loop operation
Impose effective controls on all *support systems* used by secondary success paths to provide the DHR function, to assure their continued availability and reliability during mid-loop operation
Legend
◇ - Event not developed further in this diagram but recognizes the existence of a much larger sub-tree
Implement effective failure finding tasks for important DHR components normally in standby mode during mid-loop operation
Minimize the T&M and restoration unavailability of *active* standby components in DHR secondary success paths during mid-loop operation
Implement pre-outage and pre-operational surveillance and maintenance activities to ensure that support system and trains are in "as new" condition
Implement condition monitoring activities for all important support system components and trains
Control all planned maintenance outages for critical support system components and trains

Figure 8. (Sheet 4 of 8) Conditional and Subsequential Failures and Their Institutional Influences.

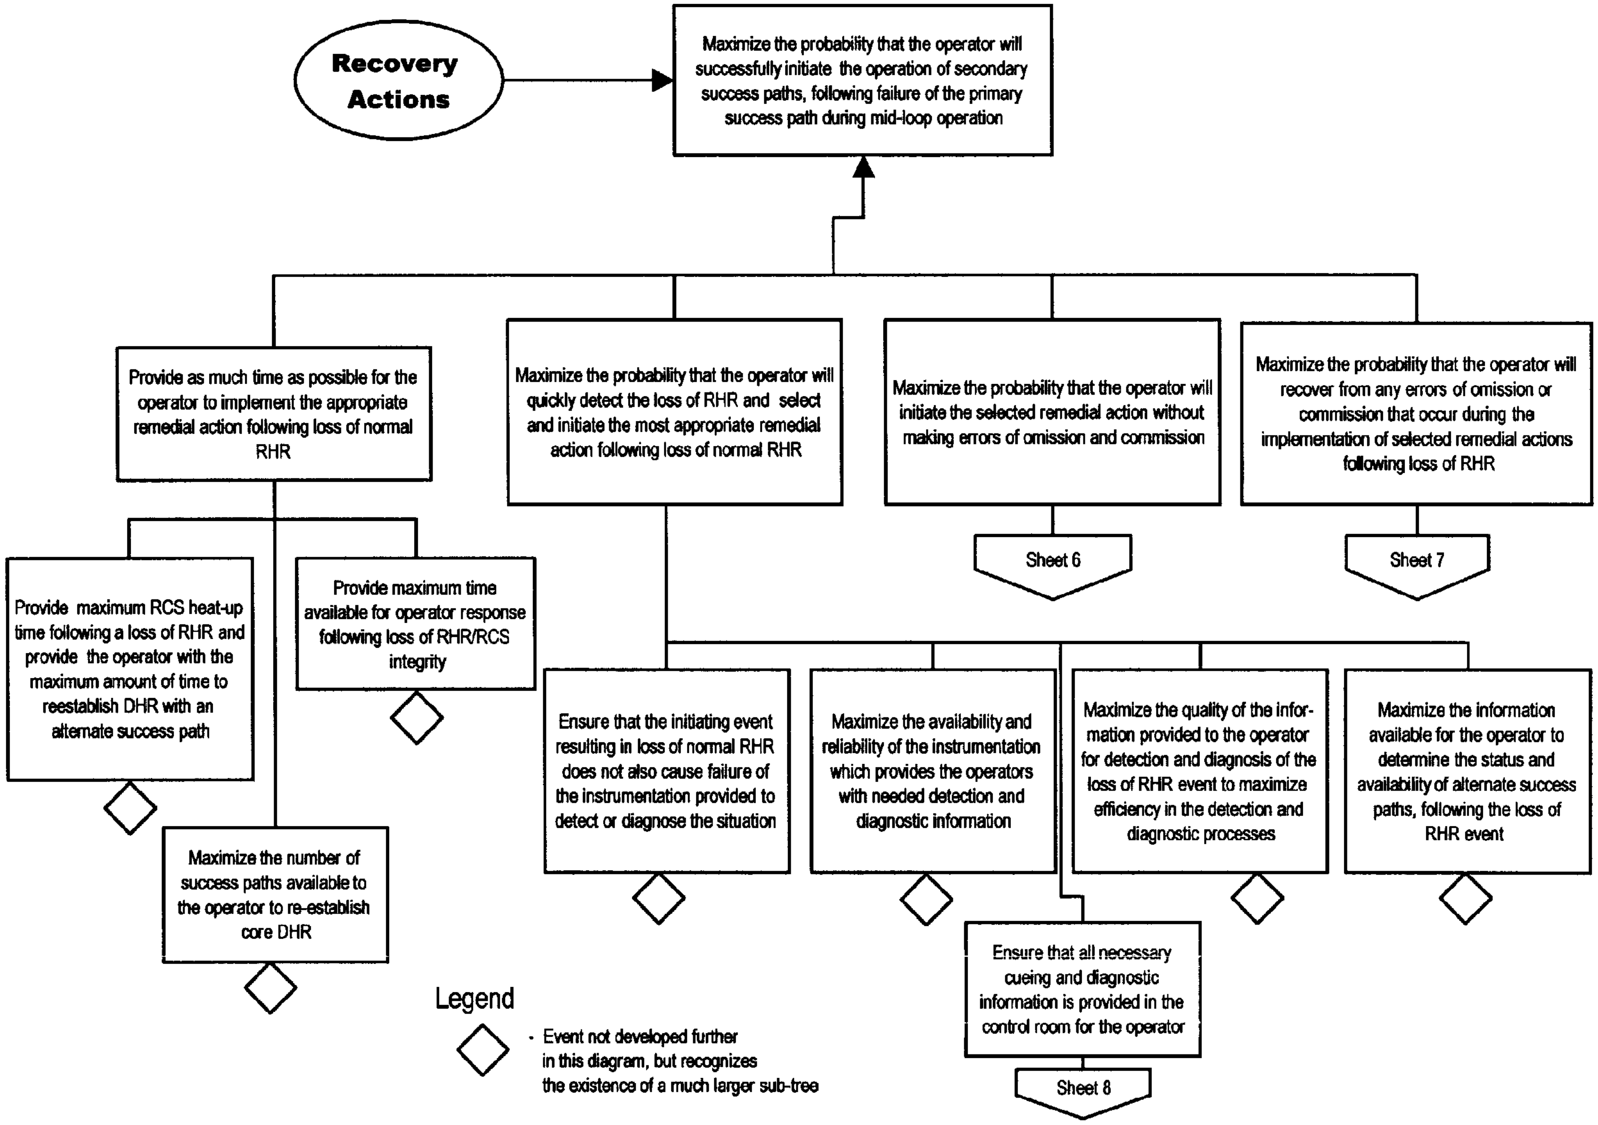
Recovery Actions
Maximize the probability that the operator will successfully initiate the operation of secondary success paths, following failure of the primary success path during mid-loop operation
Provide as much time as possible for the operator to implement the appropriate remedial action following loss of normal RHR
Maximize the probability that the operator will quickly detect the loss of RHR and select and initiate the most appropriate remedial action following loss of normal RHR
Maximize the probability that the operator will initiate the selected remedial action without making errors of omission and commission
Maximize the probability that the operator will recover from any errors of omission or commission that occur during the implementation of selected remedial actions following loss of RHR
Sheet 6
Sheet 7
Provide maximum RCS heat-up time following a loss of RHR and provide the operator with the maximum amount of time to reestablish DHR with an alternate success path
Provide maximum time available for operator response following loss of RHR/RCS integrity
Maximize the number of success paths available to the operator to re-establish core DHR
Ensure that the initiating event resulting in loss of normal RHR does not also cause failure of the instrumentation provided to detect or diagnose the situation
Maximize the availability and reliability of the instrumentation which provides the operators with needed detection and diagnostic information
Maximize the quality of the information provided to the operator for detection and diagnosis of the loss of RHR event to maximize efficiency in the detection and diagnostic processes
Maximize the information available for the operator to determine the status and availability of alternate success paths, following the loss of RHR event
Ensure that all necessary cueing and diagnostic information is provided in the control room for the operator
Sheet 8
Legend
◇ · Event not developed further in this diagram, but recognizes the existence of a much larger sub-tree

Figure 8. (Sheet 5 of 8) Recovery Actions and Institutional Influences.

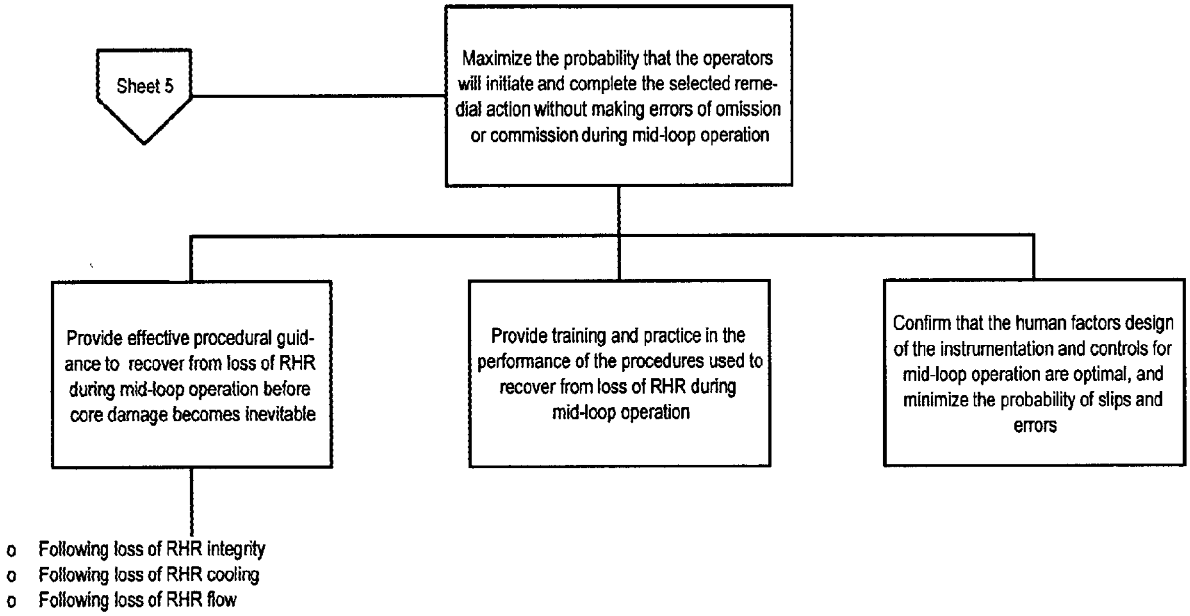
Sheet 5
Maximize the probability that the operators will initiate and complete the selected remedial action without making errors of omission or commission during mid-loop operation
Provide effective procedural guidance to recover from loss of RHR during mid-loop operation before core damage becomes inevitable
Provide training and practice in the performance of the procedures used to recover from loss of RHR during mid-loop operation
Confirm that the human factors design of the instrumentation and controls for mid-loop operation are optimal, and minimize the probability of slips and errors
o Following loss of RHR integrity
o Following loss of RHR cooling
o Following loss of RHR flow

Figure 8. (Sheet 6 of 8) Recovery Actions and Institutional Influences (Errors of Omission and Commission)

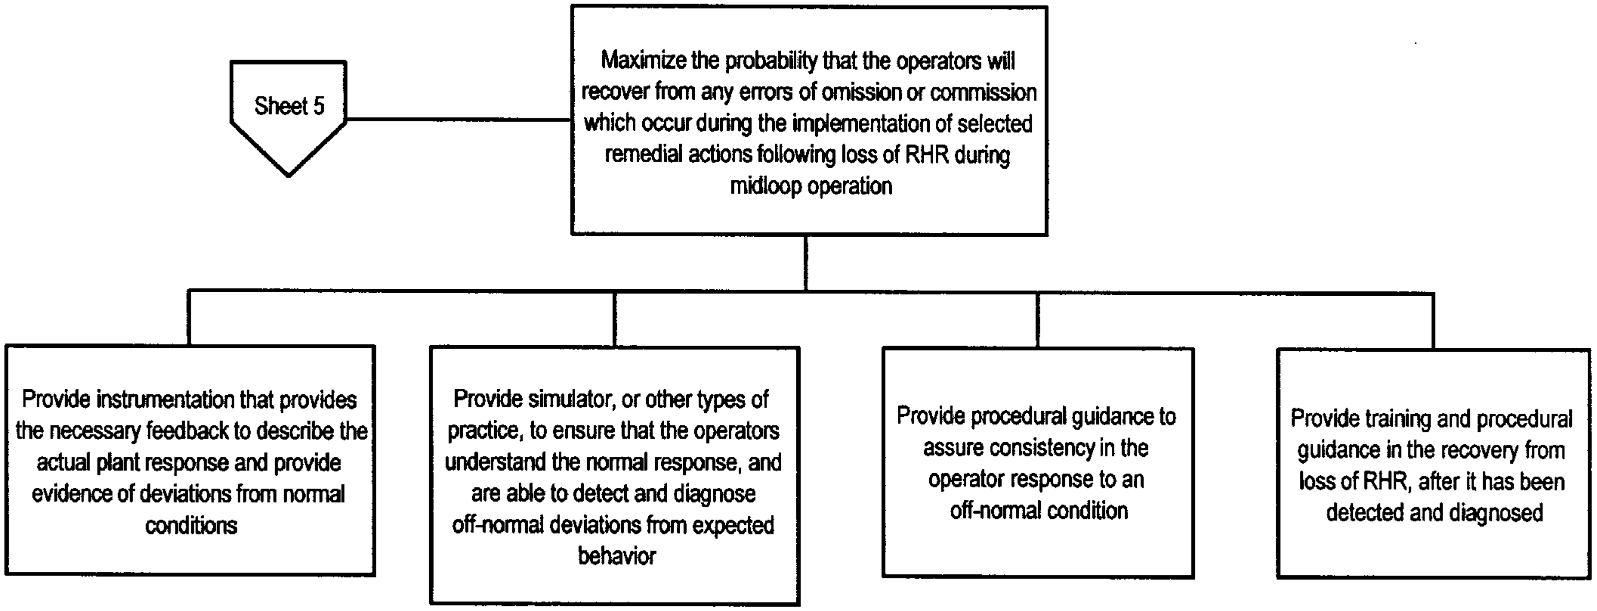
Sheet 5
Maximize the probability that the operators will recover from any errors of omission or commission which occur during the implementation of selected remedial actions following loss of RHR during midloop operation
Provide instrumentation that provides the necessary feedback to describe the actual plant response and provide evidence of deviations from normal conditions
Provide simulator, or other types of practice, to ensure that the operators understand the normal response, and are able to detect and diagnose off-normal deviations from expected behavior
Provide procedural guidance to assure consistency in the operator response to an off-normal condition
Provide training and procedural guidance in the recovery from loss of RHR, after it has been detected and diagnosed

Figure 8. (Sheet 7 of 8) Recovery Actions and Institutional Influences (Recovery from Errors of Omission and Commission).

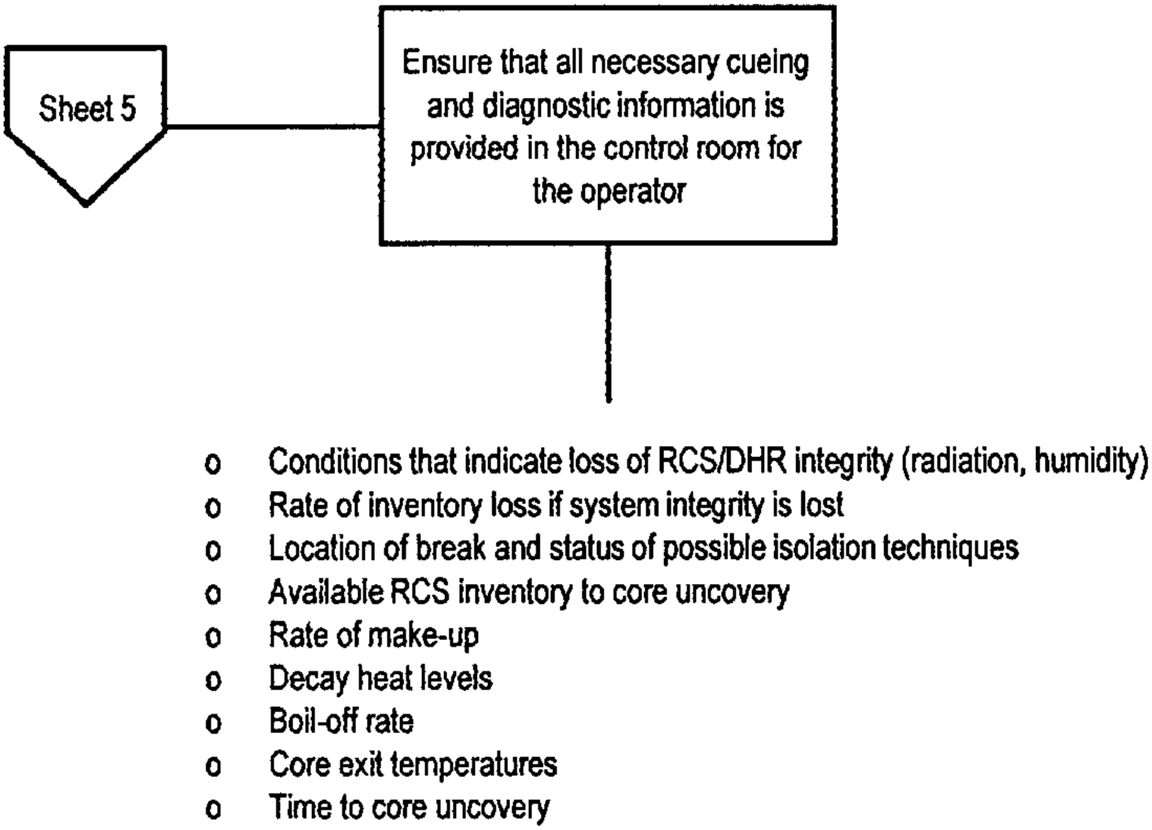
Sheet 5
Ensure that all necessary cueing and diagnostic information is provided in the control room for the operator

- Conditions that indicate loss of RCS/DHR integrity (radiation, humidity)
- Rate of inventory loss if system integrity is lost
- Location of break and status of possible isolation techniques
- Available RCS inventory to core uncovery
- Rate of make-up
- Decay heat levels
- Boil-off rate
- Core exit temperatures
- Time to core uncovery

Figure 8, (Sheet 8 of 8) Information Needed for Detection and Diagnostic Processes.

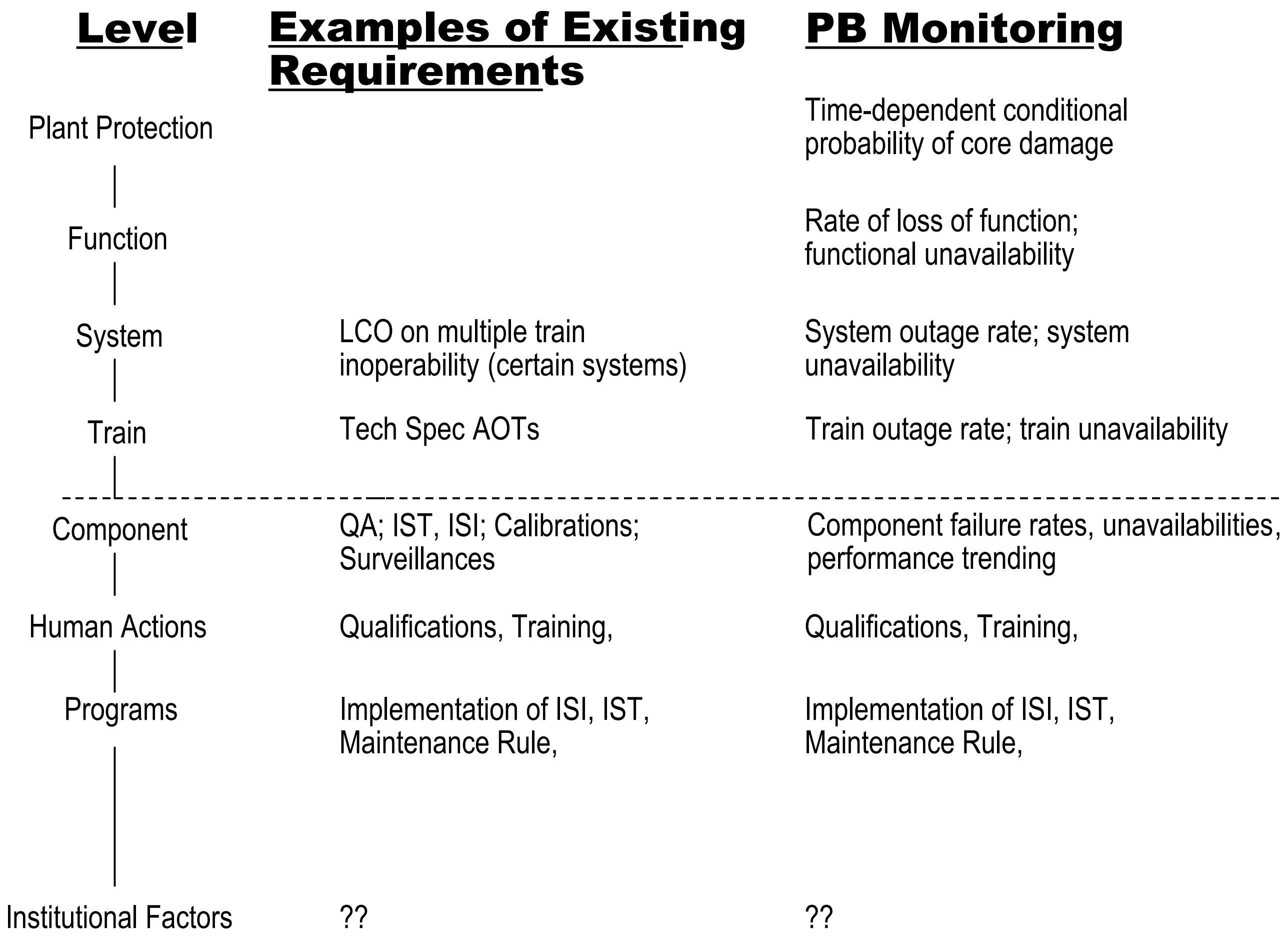

Level	Examples of Existing Requirements	PB Monitoring
Plant Protection		Time-dependent conditional probability of core damage
Function		Rate of loss of function; functional unavailability
System	LCO on multiple train inoperability (certain systems)	System outage rate; system unavailability
Train	Tech Spec AOTs	Train outage rate; train unavailability
Component	QA; IST, ISI; Calibrations; Surveillances	Component failure rates, unavailabilities, performance trending
Human Actions	Qualifications, Training,	Qualifications, Training,
Programs	Implementation of ISI, IST, Maintenance Rule,	Implementation of ISI, IST, Maintenance Rule,
Institutional Factors	??	??

Figure 9. Levels and Kinds of Requirements.

4 A STRUCTURED PROCESS FOR DEVELOPING CANDIDATE SETS OF PERFORMANCE MEASURES FOR APPLICATION IN PERFORMANCE-BASED REGULATION

This section presents a process that aims at choosing a set of monitorable plant characteristics that could provide assurance of satisfaction of top-level safety objectives, with preference given to performance-based measures. The process outlined here is essentially a formalization of points made in the discussion of the shutdown example provided earlier.

The general idea is that since it is impractical to measure fulfillment of top-level measures (e.g., CDF) directly, we decompose them into a hierarchy of lower-level objectives, and eventually identify a set of those lower-level objectives having the properties that (1) satisfaction of the set should correlate strongly with satisfaction of the top-level objective, and (2) some periodic assurance can be derived regarding the satisfaction of these lower-level objectives, either through performance-based oversight or through oversight of fulfillment of prescriptive requirements.

The process is motivated by a need to develop performance-based approaches, and tends to identify performance-based approaches preferentially. By this, we mean that the approach begins by considering the most results-oriented measures possible, and proceeds to less results-oriented measures only if the more results-oriented measures are inappropriate for some reason. Where performance-based approaches are arguably inappropriate, the process is intended to address this also.

4.1 Motivation for This Approach

Consider the portion of the generic diamond tree in Figure 10 Sheet 1. In order to have reasonable performance of the top objective, we have to have reasonable performance at all lower levels. This is not true at every instant of time or with absolute deterministic force between all levels; CDF does not necessarily fluctuate daily as a result of upper-management activities. Nor do we need to have reasonable performance at every node of every level; we simply need to have enough performance across each level, enough of the time, to ensure adequate performance at the top. Bad performance across any level is expected to correlate, sooner or later, with bad performance at the level above; that is part of the point of the diamond tree. Bad performance at any single node of the diamond tree may erode the performance of the nodes above it to which it is coupled, but it should not propagate unattenuated to the top of the tree if other nodes are performing adequately. In a sense, we almost expect something like a conservation law to operate: a high level of performance at one level should correlate with, if not absolutely guarantee, a high level of performance at the next highest level.

Allocation

Consider the level of the diamond tree at which “systems” appear. From the above discussion, it is clear that we need to achieve an overall level of performance at this level in order to satisfy the

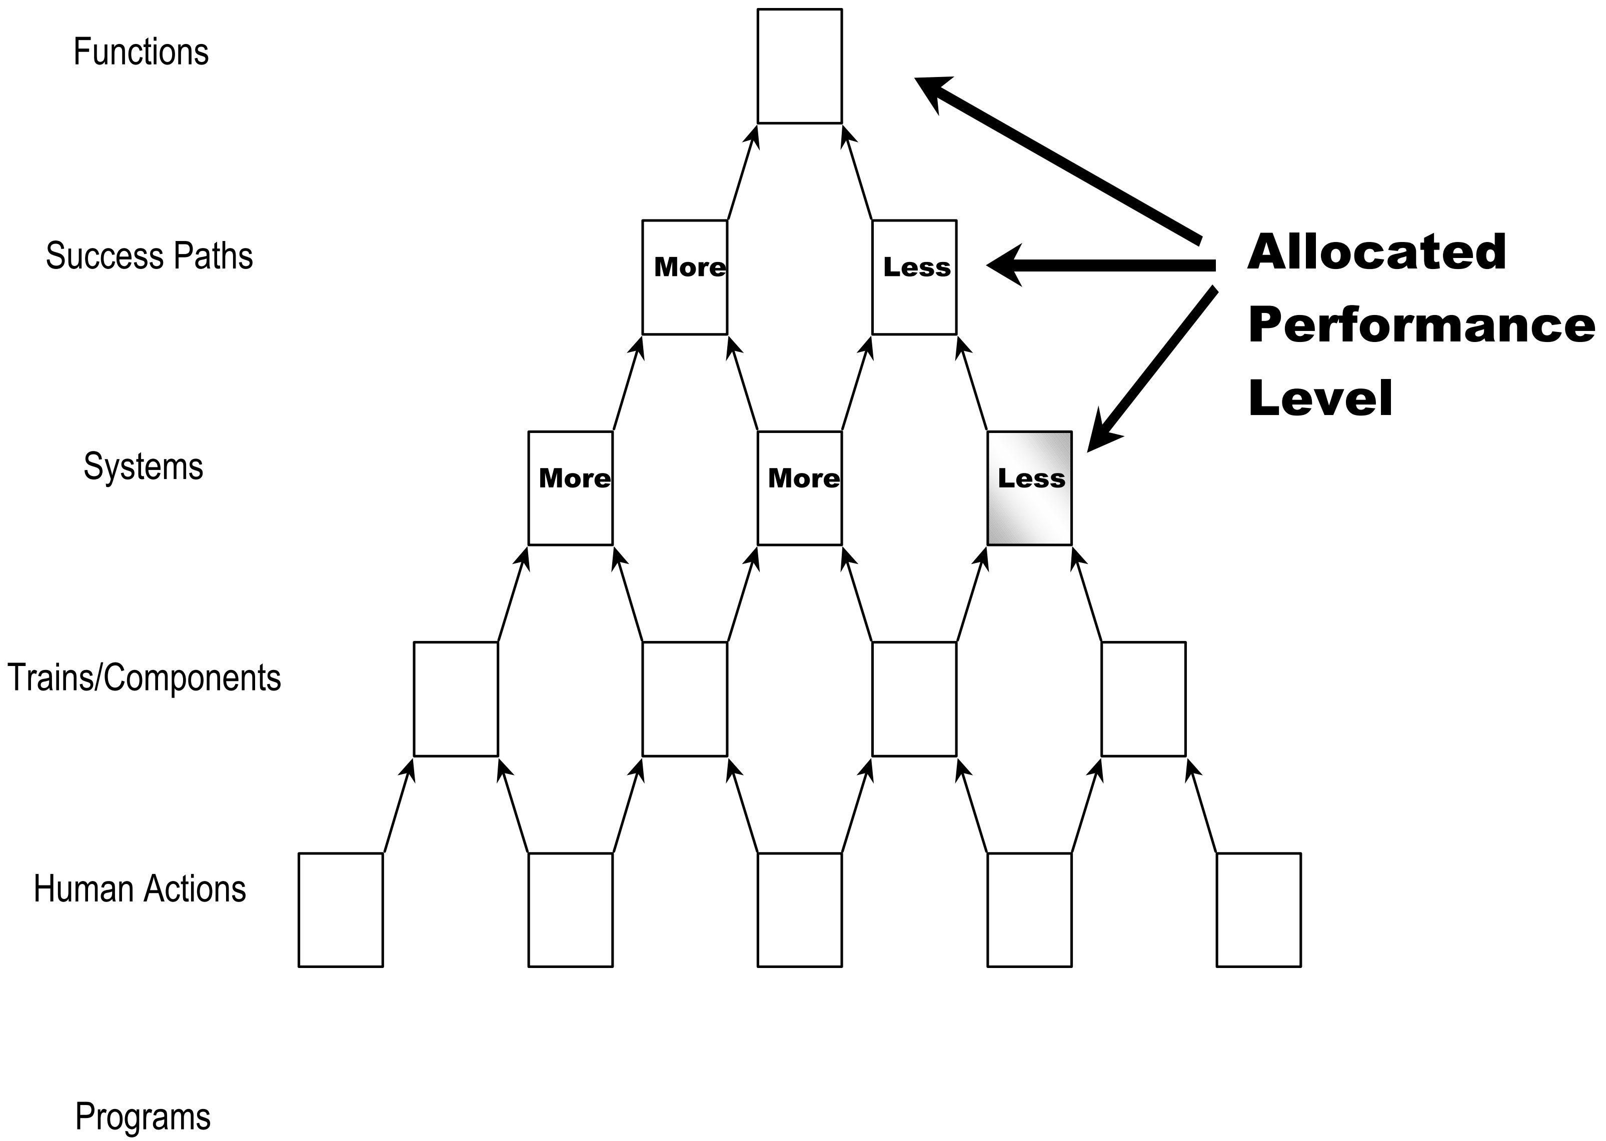

Figure 10. (Sheet 1 of 2) Deriving a Monitoring Scheme Based on Allocated Performance and Practicality of Monitoring.

top-level objective. But we have freedom in how it is achieved—freedom in deciding which systems will carry the burden. If some of these systems are shown to perform very well, then we can afford to be more relaxed about others. The process of deciding how much of the safety burden will be carried by each system is called "allocation." By now, it is widely accepted that allocation should be the licensee's prerogative.

The licensee's allocation process is essentially to optimize the distribution of safety resources, consistent with satisfying the top-level objective together with whatever prescriptive requirements survive to constrain his solution space.

The process discussed below essentially assumes that performance will be allocated by the licensees at least down ("down" in the diamond tree sense) to the system level. The discussion makes the most sense if risk-informed performance goals are articulated for systems, but the process has been formulated in a fairly general way to apply to any meaningful specification of performance.

It is in the process of performance allocation that defense in depth is addressed. A requirement to incorporate defense in depth can take the form of a proscription of allocating excessive performance to combinations of barriers that lack suitable redundancy and diversity.

Monitoring: Deriving Assurance of Performance

Suppose, for purposes of the following discussion, that we have carried out the allocation down to the system level. We therefore know what performance we need. Refer to Figure 10 Sheet 1. As indicated on the figure, we have decided that we need "more" performance from the left success path, and "less" from the right success path. Based on this, we have allocated "more" performance at the system level to the left and middle systems, and "less" performance to the system on the right. In the spirit of the shutdown heat removal example, the left system might have a reliability target or a target frequency of loss of the function; the other two systems might have targets for probability of functional failure, given a demand.

In having allocated at this level, we have specified target levels of performance that translate into a performance level at the top that we can live with, and translate into a resource commitment that we can also live with.

Now: what will we monitor in order to derive whatever assurance we decide to live with? In principle, the choice of monitoring scheme would be iterated with the allocation, but for simplicity, the present discussion will be carried out as if the allocation were fixed, and the task is simply to decide how to monitor performance.

Each small square on Figure 10 Sheet 1 is a performance node with allocated performance indicated in the square (for present purposes, either "more" or "less" performance). Figure 10 Sheet 2 adds large squares, first at the system level, to indicate whether monitoring is a practical option for each node. (As discussed elsewhere in this report, there are many performance nodes for which monitoring is not a practical option; for example, monitoring is not practical if "early warning" cannot be derived.) We see in this example that at the system level, direct monitoring of the performance node is considered practical only for the leftmost system on the diagram. This means that we need to derive assurance for the other two systems by going further down in the tree. At this next level, we have one performance node that is considered practical to monitor, and one that is not; so we need to go one more level down. At that level, we have one performance node that can be monitored, and one that cannot; in this last case, we simply impose prescriptive requirements aimed at ensuring performance, rather than look for something to monitor.

At this point, we can argue that we have some assurance of allocated performance at the system level. The assurance is not ironclad by any means, but is comprehensive in the sense that every performance node at the system level that has performance allocated to it, also has some "assurance" associated with it: the node is either monitored directly, or supported by a node that itself has assurance. The heavy arrows on Fig. 10 Sheet 2 originate at nodes that are monitored or prescribed, and flow upwards through higher-level nodes, ultimately combining at the top-level objective.

In the case of the node at the lower right, we settle for assurance derived from satisfaction of prescriptive requirements. This simply illustrates that a more or less performance-based approach can be realized in each area separately: that one can blend prescriptive and performance-based approaches to assuring performance.

We can argue that for a given performance allocation, this top-down process tends to lead to the most results-oriented combination of monitoring points that we can pick. If this set of monitoring points seems onerous or has some other undesirable feature, it is possible to change the allocation.

Consistent with the rather associative character of the diamond tree, we have not tried to formulate rigorous arguments about deriving minimal monitoring sets, or anything similar. A minimal monitoring set would be a set of performance nodes that spans the diamond tree with no gap, but which no longer spans the diamond tree if any node is removed from the monitoring scheme. This would mean that assurance at and above some level is comprehensive in some sense; the generic example presented above has this property. Certainly the set derived by this process is "minimal" in the sense that removing elements from it open up gaps in the coverage of nodes to which performance has been allocated. Going beyond minimality, we may find that for each monitored node, we may need to monitor several figures of merit. We may also need to supplement the set derived in this way with other monitoring activities. Clearly, more assurance is derived if more things are monitored; this, of course, comes at a price. However, it is argued later that it is very desirable that certain aspects of maintenance be monitored, if not by the regulator then by the licensee, with some oversight by the regulator.

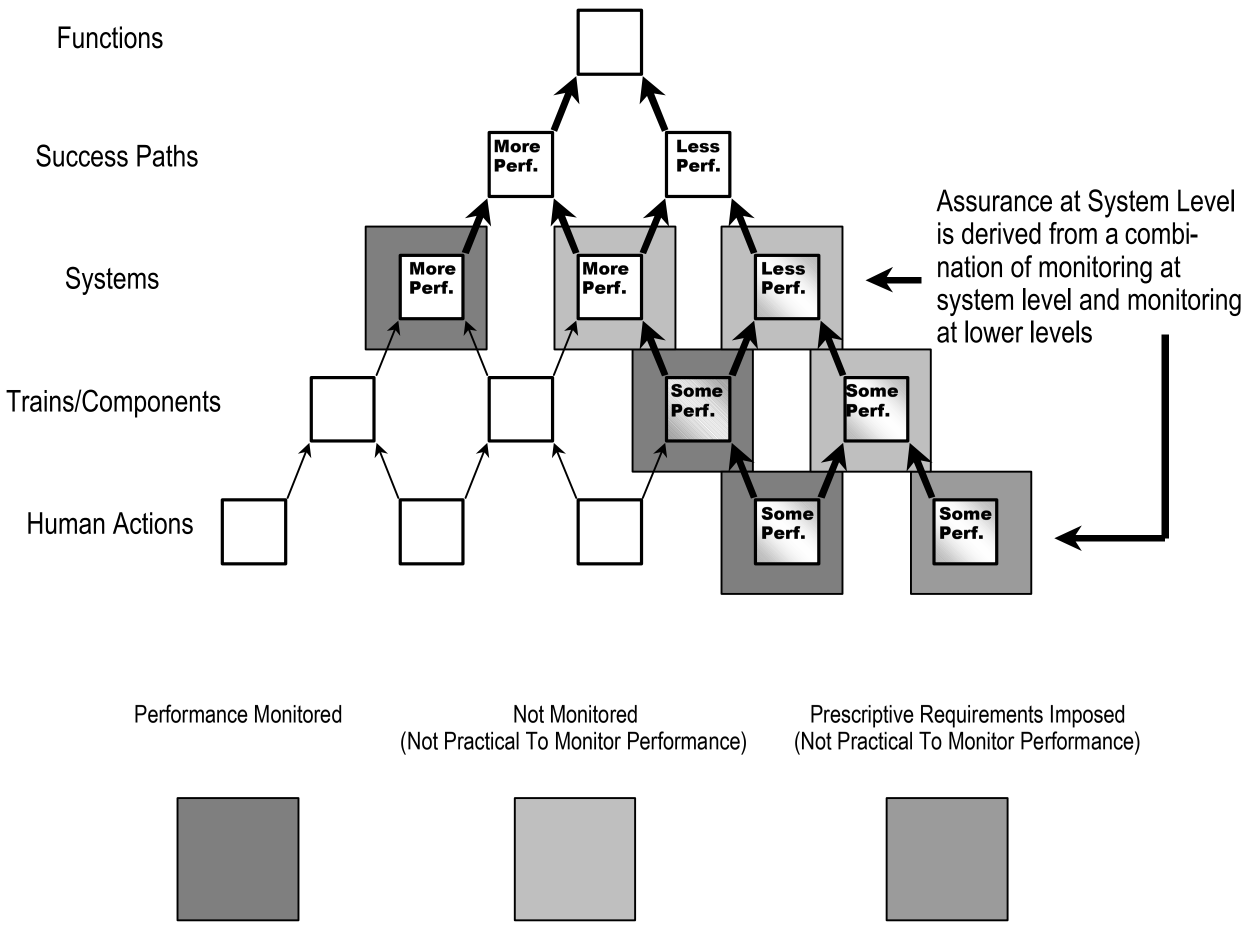
Functions
Success Paths
Systems
Trains/Components
Human Actions
More Perf.
Less Perf.
More Perf.
More Perf.
Less Perf.
Some Perf.
Some Perf.
Some Perf.
Some Perf.
Assurance at System Level is derived from a combination of monitoring at system level and monitoring at lower levels
Performance Monitored
Not Monitored (Not Practical To Monitor Performance)
Prescriptive Requirements Imposed (Not Practical To Monitor Performance)

Figure 10. (Sheet 2 of 2) Deriving a Monitoring Scheme Based on Allocated Performance and Practicality of Monitoring

4.2 Steps

The following steps are intended to implement the above-described approach.

I. Identify Top-Level Safety Objectives

The subsidiary objectives of CDF and LERF are natural candidates as top-level safety objectives. Most notions of performance-based oversight rely at least implicitly on a frequency scale of some kind, and correspondingly imply a frequency of observed events that would be considered to deviate from expectations sufficiently to warrant increased attention. The following steps can be assessed with CDF and LERF in the background as examples, but other examples could be considered.

The top-level safety objectives are ultimately reflected in the upper levels of the diamond tree.

II. Development of the Technical Basis

II.1 Build Integrated Safety Case

In this step, key results are abstracted from risk models (fault tree/event tree, etc.) and other analyses to formulate a “safety case” that explains what the challenges are to plant safety, what the plant capabilities are for responding to these challenges, and how reliably these capabilities are required to respond in order to satisfy top-level safety objectives. The risk-informed (PRA-based) version of this would be seen as a generalization of what the current licensing process does: the risk-informed process considers more initiators than current licensing, credits more success paths than current licensing (i.e., allows credit for non-safety equipment), and considers different end states (core damage, release,…). However, a version of this could be developed that would look like the current licensing process.

A comprehensive discussion of this idea is provided in a forthcoming paper [12].

It is possible to refrain from making this “safety case” a separate step, and to address some of its content as part of the diamond tree development. For now, it is considered useful to discuss formulating the safety case as a separate step. Some of the items summarized in the appendix made the point that one essentially initializes the process of safety management through formulation of the safety case and an implementation plan, and the monitoring phase is distinct. This is not to say that the safety case is immutable, but rather that a comprehensive statement of it is a roadmap to the implementation plan.

Formulation of the safety case corresponds loosely to several of the steps provided in the NEI example [4], which calls for identifying key SSCs supporting each safety function and establishing controls commensurate with safety. The difference between the NEI steps and the present “safety case” is that the present safety case is explicitly articulated in terms of success paths and the associated performance requirements, both reliability/availability and physical capacity, while NEI contents itself with identification of SSCs and stops well short of explicitly identifying needed performance objectives.

Sub-steps:

- Construct a comprehensive risk model.
- Identify key success paths that need to be credited in safety case in order for objectives to be satisfied.
- Allocate. Develop a set of performance requirements (perhaps at the initiating event level and the mitigating system train level) whose collective satisfaction corresponds to satisfaction of the top-level safety objectives, and whose fulfillment corresponds to an appropriate allocation of utility resources. In other words, decide the most efficient way to take credit for performance that satisfies high-level objectives.

  Note that much “risk-informed” work tries to use PRA to “show” that a facility meets the safety goals on either a “best-estimate” or mean-value basis, i.e., taking credit for either mean-value or best-estimate performance by components and operators. Developing a set of performance requirements is NOT the same as writing down estimates of how reliable the plant functions ARE; it is simply writing down how reliable they NEED TO BE for them to satisfy top-level objectives.

  In general, there may be many allocations that satisfy the top objectives; but there are probably relatively few allocations that make sense in terms of the resources that would have to be allocated to achieving them. (Under typical cost/reliability assumptions, gold-plating one system to near-perfection may well be more expensive than having two diverse brass-plated systems in parallel.) Developing this set of performance requirements is the utility’s prerogative and concern. It is a straightforward problem of optimizing an objective function over a large number of variables. (See [12] for a discussion.) This is separate from the regulatory problem. Note that we do not require a rigorous “optimum,” we only require an allocation that the utility is willing to live with, and that both the utility and the regulator see as translating into satisfaction of the top-level targets.

At the conclusion of this step, one will have specified success paths in considerable detail, including passive components. Wherever common cause is an issue, the common cause basic events in the model are logically equivalent to components whose failure must be prevented at some stated level of probability. Wherever human performance is an issue, failure in this area must likewise be prevented at some stated level of probability.

Note that support systems, including instrumentation, are most emphatically a part of this. PRAs generally identify support systems, but may not identify instrumentation; the general idea here is that one uses the PRA to identify key success paths, and then fills in the PRA blanks in each success path by walking it down carefully, noting instrumentation, supports, etc. This includes paths that are success by virtue of post-accident recovery.

At this stage, we have a logical construct, roughly equivalent to a portion of a PRA: equivalent to a more-thorough presentation of some of the success paths of the PRA. This construct represents what we are taking credit for, and how much credit we are taking, including credit for things like prevention of common cause failures. However, identification of all the important things that affect performance is not complete at this stage. Considerations such as institutional factors need to be addressed through the Diamond Tree.

The process of allocation is an interesting subject in itself. Significant work has been devoted to it previously, but the practicalities of optimizing the decomposition are beyond the scope of this report. (For a recent discussion, see Youngblood (1998) [12] and references cited therein.) For present purposes, it suffices to note that there may be several qualitatively different decompositions that should satisfy the regulator, and choosing among them is the licensee's prerogative. This report does not specify in detail the process by which the regulator would come to an agreement that the licensee's allocation is appropriate. The focus of this report is how to formulate a performance-based approach to deriving ongoing assurance of satisfaction of these objectives, whatever they are.

II.2 Build Key Portions Of The Diamond Tree Corresponding To This Safety Case

The diamond tree does not specify the details of all the logical relationships between performance nodes, or quantify the performance at each node; but the tree qualitatively shows many relationships that fall through the cracks of a typical PRA. The top half of the diamond tree would follow straightforwardly from the safety case called out in step II.1. The bottom half of it would address aspects of performance that are not generally addressed satisfactorily in PRAs.

It was seen in a previous section that a full, comprehensive diamond tree is capable of becoming quite large, as many performance aspects are identified. The formulation of II.2 above (Build "key portions") concedes that a full diamond tree may be more than is initially necessary. The character of the search in steps described below is such that for some purposes, branches can be added to the tree incrementally as needed. A comprehensive assessment of institutional factors, however, probably warrants a fuller development; the hybrid approach mentioned in the last section, for example, would appear to call for development of large portions of a diamond tree.

Note that if the tree is formulated with due regard for the hierarchical character that it is supposed to have, then higher levels of the diamond tree are more results-oriented. A regulatory approach that focuses at higher levels is therefore more performance-based.

Note also that depending on the level at which reliability goals have been articulated in the safety case, it allocates frequency objectives to many of the safety functions identified on the diamond tree. The safety case does not specify everything that is needed, but properly formulated, it specifies functional objectives that are arguably necessary in order to satisfy the top objective.

III. Identify Performance Nodes That Are Possible Measuring/Monitoring Points

III.1 Find the highest (most results-oriented) level on the diamond tree at which one can consider using the performance nodes as monitored quantities

We begin at the highest, and therefore most results-oriented, level of the tree. In general, trying to monitor performance at this level would be a mistake, because significant events at this level would have unacceptable adverse consequences. Here, "unacceptable adverse consequences" means that there is an immediate and significant threat to safety (or worse). An event at the level of the top of the tree would be a failure of a top-level objective, corresponding (for example) to core damage.

Therefore, we move down the diamond tree one level at a time, considering each level to see whether the performance nodes identified at this level are candidates for monitoring. Recall that we need comprehensive assurance across each upper level, but we do not require that every node be expected to perform perfectly. We consider progressively lower levels on the diamond tree until we reach a level at which a failure in at least one node would not be considered to have unacceptable consequences.

Example: Assume that the top goal corresponds to various combinations of safety functions, and that having a top-goal failure would be core damage. Obviously we don't want that. Moving down a level, one of the performance nodes is "secondary heat removal." Arguably we don't want to experience failure of that, either. Nor do we want to monitor at the level of complete failures of the AFWS. Presumably, we can tolerate train failures within the AFWS.

Example: If a goal is removal of heat at shutdown, we can at least consider using RCS temperature as a monitored quantity. But see below for more discussion.

III.2 All performance nodes at this level should then be considered individually to see whether they are suitable for use as monitored quantities

III.2.a Below each performance node that is not suitable, continue the search downwards through the tree until suitable performance nodes are identified

In principle, we need some assurance of performance at each level: assurance that the top objective is being satisfied. We do not insist on perfect assurance at each level, but we require that at each level, there be a reasonably substantial basis for believing in a plant's safety performance. At the upper level, this assurance is derived from the lower-level results; at the lower levels, assurance is derived either from still-lower levels, or from direct monitoring of performance, or perhaps from monitoring of compliance with prescriptive requirements, or perhaps from some combination of these.

Each performance node that is a feasible monitoring point is a possible stopping point for that branch; for each performance node that is not a suitable monitoring point, the search continues downward through the diamond tree for a level that would be suitable.

A "suitable performance node" has the following properties:

Its failure does not by itself constitute unacceptable adverse consequences.

The performance node is challenged often enough that a "zero failures" statistic has some meaning.

Some performance nodes are not challenged often enough for their non-failure to be a particularly informative datum. Such an objective needs to be decomposed further.

The shutdown example has a characteristic that is probably generic to situations in which a given performance node is addressed by more than one layer of defense in depth. In the case of the shutdown example, we have multiple layers of defense in depth (performance nodes) appearing at the level just below "prevent failure of heat removal"; but the monitoring activity effectively addresses only one of the layers of defense in depth: the continuity of function of the RHR system. None of the backups (feed and spill, steam generator cooling,…) is addressed, because the criterion is triggered before most of them can be brought into play. Arguably, we want assurance of their performance before they are needed.

III.2.bBefore an apparently suitable node is accepted, examine the node closely to ensure that it meets certain tests. Even if it meets those tests, lower-lying performance nodes should be examined to see whether they would be more advantageous.

Even if a performance node seems suitable, closer examination may reveal a need to go further down, or a significant benefit to doing so. For one thing, there may be lower-level performance nodes that are easier to measure and still provide most of the needed assurance. For another, closer analysis may show that a significant fraction of events that trigger at this level will have unacceptable conditional probabilities of going too far. The separate discussion of the NEI shutdown example illustrates this point.

Monitoring the continuity of the RHR function has certain things to recommend it. Loss of this function has been experienced often enough that maintaining it successfully is actually an informative statistic: it means that certain combinations of human errors are apparently not being committed. However, there are infrequent causes of its loss (support system initiators) that also adversely affect other systems, raising the possibility that for events like this, the conditional probability of going to the unacceptable consequence may be unacceptably high. This has to be decided in light of the expected frequency of those events. That expected frequency, in turn, may be related to a regulatory expectation that we need to sort out on this diamond tree.

This step therefore calls for examination of each performance node to see whether assurance is available that triggering events will not go on to unacceptable consequences. This involves looking at lower levels, especially for things that influence multiple performance nodes.

IV. From all the candidate performance nodes that are candidates for monitoring or oversight, choose a set of monitoring points that best balances the competing needs of providing the needed assurance to regulators while minimizing adverse impact on plant performance (economics)

IV.1 Synergize

Assurance of performance in one area may be partially transferable to another area. Makeup systems, for example, may play a role in LOCA analysis, in bleed and feed, in safety of shutdown operations, and even in normal operations. This is important because comprehensively assuring performance of a given system will enable the assurance of its performance to support numerous objectives, including especially instances in which a given system is practically never challenged in one context but may be challenged in another. (We may trust a makeup system more in the context of LOCA if we have a lot of operating experience with it in a normal makeup function.) In order to make appropriate use of this information in the diamond tree, it is necessary to ascertain precisely which aspects of a given system’s performance are actually being tested by a given challenge (including which trains, which suction sources, which support requirements).

The present step is to look for instances of this.

When this process has gone to completion, a collection of items on the diamond tree will have been flagged as candidate monitored quantities. Some of these may be performance measures of the kind intended by the NEI example, and others not. The task is now to work with this collection to see whether all of it is necessary, or whether some subset of it would suffice.

The above process should have identified sufficient monitored quantities that there is no path from the bottom to the top that does not encounter at least one monitored quantity. If the diamond tree is comprehensive, this could be a very exhaustive list of monitored performance areas. The task is now to examine the necessity and sufficiency of this collection: its suitability as a tool for regulatory oversight. It is likely that some items can be eliminated from the collection, on grounds that they do not add enough information to the rest to be worth the trouble. The next step is therefore:

IV.2 Prioritize

At some point, a set of lower-level performance nodes will have been identified whose satisfaction collectively provides assurance that the top-level objective is satisfied. Some of these are amenable to performance monitoring; others are not amenable to performance monitoring, and assurance of satisfactory performance must be based on programmatic activities. Even here, there is room for approaches that are in a sense “performance-based,” but this is a separate topic.

The present point is that in the comprehensive set of objectives that we have identified, some are more important than others. Some provide meaningful information, but information that would come at too high a price relative to the value of the information provided. The present step is aimed at identifying objectives that either do not provide much information, or objectives that provide information but are difficult or onerous to implement. In the case of an objective that is useful but onerous, it is possible that at some lower level, a sub-goal of the objective retains much of the usefulness and is less onerous.

5 SIGNIFICANT EVENTS, PRA, AND PERFORMANCE-BASED REGULATION

Nowadays, inquiries into “significant” events are considered lacking if they do not proceed to a level of discussion that goes significantly beyond pure hardware failures or pure procedural errors. It is revealing at this point to test the ideas and suggestions expressed in preceding sections by confronting a real event. A fairly recent AIT report [5] addresses an event at shutdown that is not quite within the scope of our Appendix B model but close enough to be a useful comparison.

5.1 Brief Description of Event

For a full description, the reader is referred to the AIT report. Following are key points chosen to support the present discussion.

The event occurred at cold shutdown with one train of RHR operating to remove heat. In this event, nitrogen gas was inadvertently bled continuously into the RCS in such a way as to build up a significant volume of nitrogen in the vessel, pushing out water and causing a reduction of water level above the core. In this event, the loss of level did not proceed to the point of actually causing loss of RHR, but the sense of the AIT report is that the potential was there.

The event was caused when an operator error in valve alignment caused nitrogen to be injected into the RCS through the CVCS. The error was noticed immediately and instructions were given to correct it, but a valve that was required to close completely did not do so, and continued to leak nitrogen into the RCS at a rate that exceeded the capacity of the vent to eliminate it, leading to a net buildup. Owing to the quality of plant status information available to the staff, and some missed opportunities to determine what was happening, the staff allowed this situation to continue for some days. According to the AIT, the state of the RCS was eventually discovered in the course of the staff trying to determine why so much nitrogen was being used up.

Had this particular process gone to the point of causing loss of the operating RHR pump, the probability of recovering the operating pump would have been somewhat reduced. In this event, the other RHR train was already failed, as a result of a “marginal design” and “manufacturing defects.” The AIT indicates that makeup by low pressure injection was feasible, but owing to the particular way in which the event was caused, CVCS makeup might not have been feasible (the pumps being airbound as a result of the nitrogen). The AIT also doubted whether steam generator cooling would have been feasible, citing a lot of nitrogen in the system and the unavailability of the main coolant pumps to assist. In short, some capability remained, but significant capability had been lost, partly because of the initiating event and partly because of other performance issues.

The team concluded that the combination of these events was safety significant. The operation of the RHR decay heat removal system is contingent upon maintaining adequate level in the reactor vessel. The accumulation of nitrogen gas in the reactor vessel head significantly reduced the water level in the reactor vessel. The decrease in reactor vessel level went undetected by plant operators for nearly 4 days. The team determined that systems used to mitigate a loss of the RHR system were also adversely affected by the nitrogen gas intrusion.

5.2 Performance Aspects Noted by AIT

The following paragraphs consider various levels of a diamond-tree hierarchy in turn, and briefly mention performance aspects corresponding to this level that were explicitly mentioned in the AIT report for this event. Ultimately, the point is to note the pluses and minuses of attempts by regulators or licensees to track performance at these various levels, and the discussion is therefore carried out in terms of what significant aspects of this event do or do not show up at these levels.

These levels are called out on Figure 11, which also then associate the performance aspects noted by AIT with an appropriate hierarchical level.

It should be noted that while the AIT responded to a particular event, its coverage of the evolutions leading up to the event led it to address a number of perceived deficiencies that, while interesting and perhaps meaningful, did not significantly or directly influence the conditional probability of damage in this particular event. They do, however, seem significant as "indicators" or symptoms of underlying conditions that did influence the conditional probability of damage in this event.

Functional Level

If a criterion for level control were defined to match NEI's criterion for heatup, this event would have tripped it, and NRC should have been notified. In this particular event, the reportability of this event under prevailing criteria was evidently not clear to the plant, and this was actually pointed out in the AIT report in connection with management performance. Basing a monitoring scheme at this level is arguably unsuitable because events at this level tend to be "safety significant," i.e., they tend to have a higher than desirable conditional probability of becoming serious. Moreover, it seems fairly clear from the discussion that numerous pre-existing conditions were available to signal the potential for this event, and a scheme that picked up pre-existing conditions more quickly would be desirable.

System Level

Several systems come into play here. In this event, the RHR system did not fail. However, the CVCS was effectively failed, and the vent system was degraded. These circumstances arguably would not have been deemed noteworthy but for the occurrence of this event. The reactor vessel level indication system (RVLIS) was not as available as it should have been.

This is really a human performance issue, but it would manifest itself at this level, if that system were monitored at this level. A goal for RVLIS availability might have led to more "timely" reinstallation of it. It is doubtful that having a CVCS goal or an RHR system-level goal would have helped in this event; CVCS was lost because of the character of the initiator, and the RHR train unavailability would not be quickly picked up by a system-level indicator. (But see the next paragraph.)

Train Level

One train of RHR was failed because of pump problems. This train took weeks to repair, for reasons that were criticized in the AIT report. It is not easy to distinguish this level from "component" level in all respects; the point of citing this level is that tracking at this level is more directly meaningful than tracking all components in the train and summing them up. A scheme that tracked train availability would have flagged the "weeks to repair" condition, and perhaps some of the organizational conditions leading up to it.

Component Level

Many valves were leaking, not only allowing the nitrogen leak to go on but also contributing to other diversions of inventory. Parts were unavailable. Post-maintenance failures occurred. These conditions are suggested as also having been signalling underlying conditions.

Human Actions

As is true in many incident reports, this AIT discussion cites numerous instances of things that should have been done differently. These include omissions, delays, and errors. The AIT report also repeatedly points to "lack of a questioning attitude." This could be binned at the "human action" level but there are other levels to consider for it: training, for one. Given that vessel level measurements should have been obtained and were not, it seems fair to bin this at the human action level.

Supervision

This level does not appear in previous sections' discussions of the diamond tree concept. However, the AIT report's characterization of this event seems to warrant inclusion of it. Given that junior operators will occasionally be performing tasks, it seems unfair to bin all of their "errors" at the human action level. Perhaps "supervision" and "procedures" should be at the same level: both essentially furnish guidance to the humans who act on components.

Procedures

This level refers to the development of appropriate procedures per se, and not the implementation of them. The AIT report is critical of some of the procedures, of the review and approval process for them, and finally is critical of the staff for not balking at implementing unreviewed procedures.

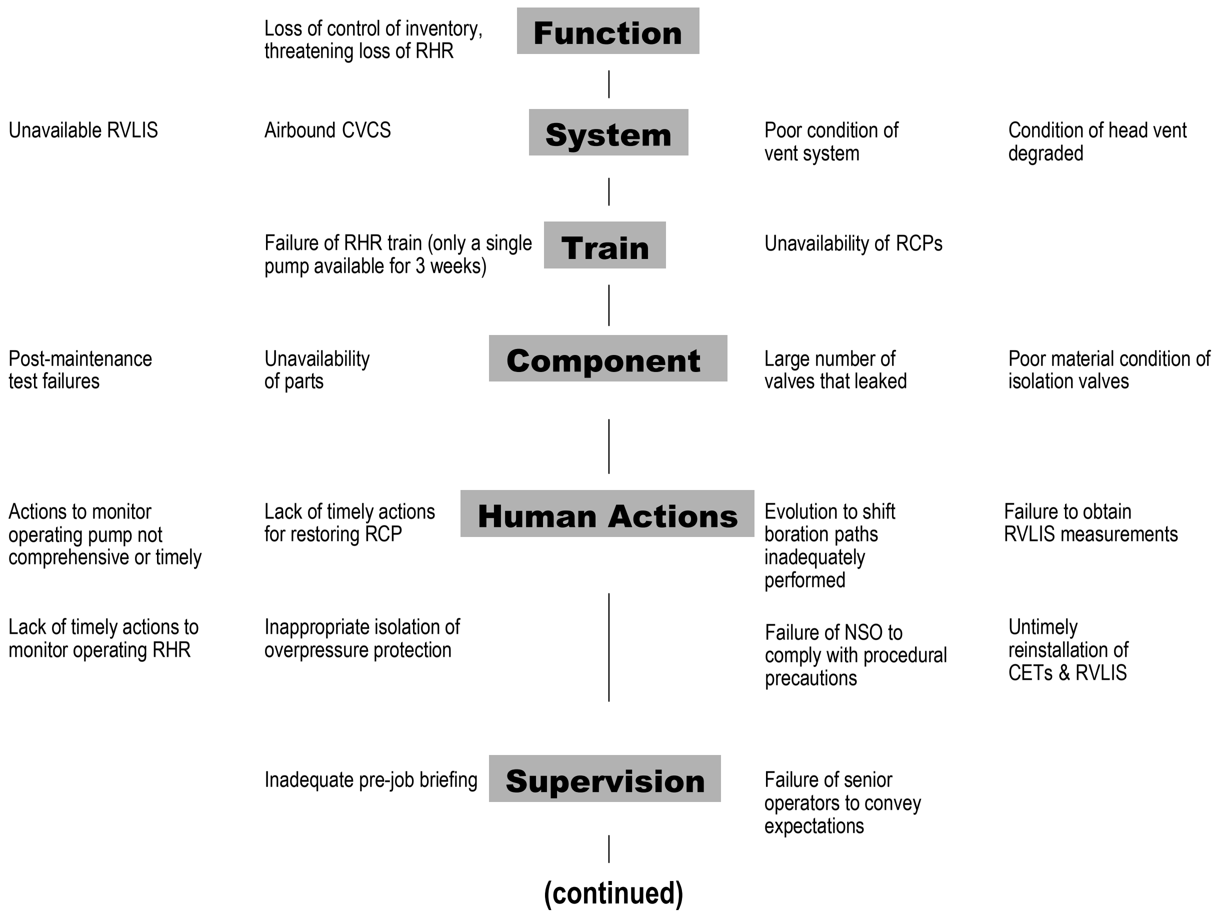
Function
Loss of control of inventory, threatening loss of RHR
System
Unavailable RVLIS
Airbound CVCS
Poor condition of vent system
Condition of head vent degraded
Train
Failure of RHR train (only a single pump available for 3 weeks)
Unavailability of RCPs
Component
Post-maintenance test failures
Unavailability of parts
Large number of valves that leaked
Poor material condition of isolation valves
Human Actions
Actions to monitor operating pump not comprehensive or timely
Lack of timely actions for restoring RCP
Evolution to shift boration paths inadequately performed
Failure to obtain RVLIS measurements
Lack of timely actions to monitor operating RHR
Inappropriate isolation of overpressure protection
Failure of NSO to comply with procedural precautions
Untimely reinstallation of CETs & RVLIS
Supervision
Inadequate pre-job briefing
Failure of senior operators to convey expectations
(continued)

Figure 11. Hierarchical Levels of Performance Aspects Noted in AIT Report.

(continued)

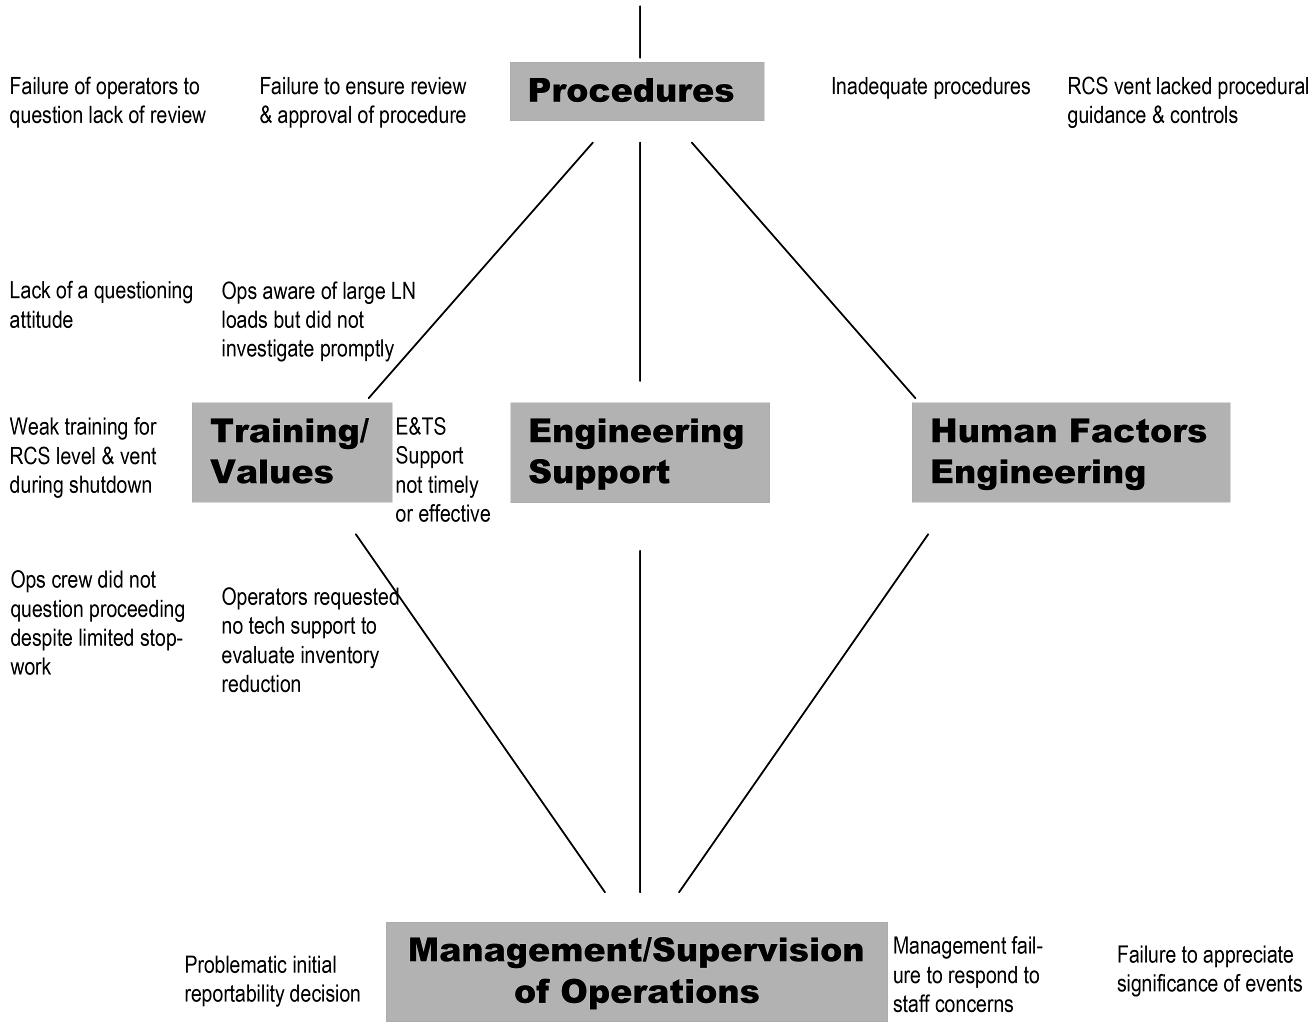
Failure of operators to question lack of review
Failure to ensure review & approval of procedure
Procedures
Inadequate procedures
RCS vent lacked procedural guidance & controls
Lack of a questioning attitude
Ops aware of large LN loads but did not investigate promptly
Weak training for RCS level & vent during shutdown
Training/ Values
E&TS Support not timely or effective
Engineering Support
Human Factors Engineering
Ops crew did not question proceeding despite limited stop-work
Operators requested no tech support to evaluate inventory reduction
Problematic initial reportability decision
Management/Supervision of Operations
Management failure to respond to staff concerns
Failure to appreciate significance of events

Training/Values

This category is intended to capture the organizational influence more general and less explicit than procedures, but more explicit and less general than general management supervision. The AIT report calls out weak training for RCS level management during shutdown. The “lack of a questioning attitude” seems also to belong at this level.

Engineering Support

The AIT report is critical both of the support received by operations staff, and of the failure of operations staff to make proper use of engineering support. One of the significant things about this event is that it went on undiagnosed for so long, and the AIT seems to feel that engineering support being in the loop would have led to earlier diagnosis. Also, the inadequacy of the vent system as a contributing factor is laid partly to inadequate engineering support.

Management/Supervision of Operations

The AIT report actually uses the term “management” in several of its criticisms. Ultimately, high levels of management are implicitly involved, but it seems that the AIT report is addressing a level of management sufficiently technical to be responsible for understanding the significance of events, and particularly their reportability. It is also a level responsible for appropriately evaluating and acting on staff concerns regarding the capabilities of the vent system. This level was cited for inadequate performance in these areas. The confusion regarding reportability is noteworthy in the present context; this event was evidently significant to the AIT, but apparently was almost not reported at all.

5.3 Observations

The essential event was that a recoverable error in conjunction with a leaky valve and a poor venting system led to a reduction in inventory that should have been diagnosed, but was not diagnosed for some days, and then only in a roundabout way.

Essentially all of the significant comments of this AIT can be placed into correspondence with a performance node on a diamond tree. It seems reasonable to expect this comment to generalize: one expects in principle to be able to plot an AIT report on a diamond tree (or at least to be able to adapt a diamond tree structure to permit this.) As we saw above, some items could be considered for more than one level; above, the choice was generally driven by the preference to place the observation as high on the diamond tree as it could reasonably have been detected before the fact.

Intuitively, it seems that the protracted failure to diagnose is a more significant performance aspect than the initiating error. (In fact, steps were taken to recover the error as soon as it was made.) The failure to diagnose the ongoing evolution was due in part to a lack of valid information regarding vessel level. A performance objective on RVLIS or the equivalent level-indication function might have helped here, in the sense that chronic problems maintaining a desired level of RVLIS performance would have flagged a need for attention before the occurrence of an event in which level was lost. In the process presented earlier, a carefully formulated safety case would be the basis for allocating performance over systems, and system or train-level monitoring would result; such a process might well have led to a monitoring criterion on RVLIS, and thereby helped to prevent this event. That is, injection of nitrogen might still have occurred, but it would quickly have been realized that injection was continuing, and the event would have been much less significant.

More generally, fairly low performance in various areas was being tolerated; a lot of valves leaked, a low level of vent system performance was tolerated, RCP function was not restored promptly when a potential need for it could be identified, and repair of the RHR pump (poorly designed and defective) was taking a long time despite the plant being on RHR. Unless these issues had suddenly cropped up in this particular outage, it appears that performance objectives at the train or component level, such as maintenance backlog, train availability, etc., would presumably have identified difficulties. One less leaky valve (the one that didn't seal when the initiating error was corrected) would have stopped the nitrogen injection promptly. It is difficult to argue for component-level monitoring programs that the regulator would routinely examine, but if it is true that there were many leaky valves, then it can be argued that the licensee should have been doing a better job in that area, and the licensee program for doing that would be of interest to the regulator. "Many leaky valves" might show up in a monitored quantity based on a maintenance backlog or a frequency of valve repair, quantities that are farther away from the middle of the diamond-tree than the individual-component level, but in what is presumably the wrong direction on the diamond tree.

Some might question the post-facto placement of RVLIS at such a high level in the hierarchy, and might question the argument that the recommended process would have picked it up. They might argue that RVLIS is really a support system, and that in a top-down approach to allocation of monitoring effort, the front-line elements are typically seen as more important. Indeed, rather than a logic tree covering RVLIS and procedures, a simple PRA model covering this part of the operation might well have had a simple diamond event, "failure to diagnose in X hours," quantified with a screening number. The text in the PRA might have had a simple discussion essentially taking credit for RVLIS and for generally alert operators. This basically illustrates what would be wrong with trying to derive a monitoring scheme directly from an unfiltered set of PRA "importances." This is why the recommended process entails a careful formulation of a safety case, including steps to make sure that all elements important to a success path are identified and receive priority according to the performance that is required from them. In the present context, it is unreasonable to expect to drive the challenge frequency of level control down to zero; level instrumentation therefore has to be seen as important. In the previous discussion of loss of heat removal, we derived a similar conclusion regarding support systems generally, noting from our model application that it is impractical to monitor their performance meaningfully by waiting for them to cause failure of front-line systems. Here, we note that it is similarly impractical to wait for important instrumentation systems to fail key safety functions when they are demanded. Key support systems and key instrumentation systems need to receive priority.

There is a simpler path to the decision to give some priority to level measurement. In the NEI conceptual rule, the shutdown functions to be addressed included heat removal, inventory control, reactivity control, and containment control. Quite apart from the process of top-down allocation from safety functions, as argued above, there is an argument for monitoring the quantity that is called out in the safety function (in this case, inventory). To identify inventory as important is to identify measurement of it as important. Absent a reliable measurement, how would performance of it be reported, within an NEI-type approach? Even without dwelling on the allocation exercise, one would allocate some importance to RVLIS and level measurement generally.

5.4 Conclusions

This event was “safety significant.” This does not mean that the event was a “near miss,” but means that it is believed that quantification of a conditional core damage probability for this event, along the lines of an ASP evaluation, would yield a number high enough to make it worthwhile to study the event.

Some of the systems intended to mitigate loss of RHR were adversely affected by the initiating event.

The AIT noted many shortcomings; it appears that the significance of this event could have been reduced in many different ways. One item of particular interest is more timely response, which in turn would have been promoted by better information to the operators. System-level and train-level performance nodes can be identified whose monitoring at a high level (available to regulators) might well have promoted more timely response, by leading to better performance by systems and staff. Specifically, attention to having access at all times to good level information would have affected this event. It is argued that the process recommended in a previous section has a good chance of identifying these performance aspects ahead of time, based on the general approach and on the specific recommendation to pay special attention to instrumentation and supports in the formulation of the safety case.

Certain performance aspects that were criticized by the AIT seem significant, and tend to confirm suggestions that monitoring certain trends (maintenance backlog, long repair times, spare parts availability,…) would in fact provide early warning. But they seem also to be too far down in the hierarchy to be considered suitable for direct monitoring by the regulators. They might be valuable elements of a licensee performance monitoring scheme.

Some other performance aspects are probably very difficult to measure, even for licensees. “Lack of a questioning attitude” was detected after the fact by the AIT, but monitoring for it would pose a special challenge. This kind of thing is discussed more in the performance-indicator literature.

6 SUMMARY OBSERVATIONS

6.1 General

This report has presented elements of an approach to performance-based regulatory oversight. It is widely believed, and essentially assumed in this report, that replacing some current prescriptive requirements with performance-based oversight—that is, oversight based on a licensee's achievement of results, rather than compliance—would represent a significant gain in efficiency. Accordingly, the approach has emphasized identification of a set of performance measures that are results-oriented, objective, and measurable, and that are collectively comprehensive in the sense of largely spanning the safety case. Areas for which performance-based oversight is not practical are identified in the course of the process, and are to be addressed by other means.

The current approach is required to address two issues. One of these issues relates to the need for performance criteria that provide early warning of safety deficits ("failure to meet a performance criterion, while undesirable, will not in and of itself constitute or result in adverse consequences"). There are several facets of this, and addressing it comprehensively requires care. The other issue is the need to address performance aspects that typical safety modeling (including conventional PRA) does not address adequately, including not only numerous hardware items taken for granted in PRA but also institutional factors, which are clearly important but not generally modeled.

In addressing the second issue, we were led to consider applying a construct called the "diamond tree." The diamond tree displays performance of all kinds at all levels in a unified, hierarchical structure, with the most general, most results-oriented performance nodes at the top, component-specific performance nodes in the middle, and institutional performance nodes below that. This tree serves as a vehicle for qualitatively discussing many kinds of performance (housekeeping, delays in procuring spare parts,…) that ultimately affect safety performance but are difficult to model quantitatively.

To illustrate further the influence of institutional factors on safety, and to show how a diamond tree might deal with these, a recent AIT report covering a significant event at shutdown was discussed with reference to a diamond tree hierarchy. It was argued that essentially all significant findings of this report correspond cleanly to a diamond tree hierarchy, and that many of the performance nodes highlighted in the AIT report are typically neglected in logic model developments.

The work needed to carry out the program described here would be justifiable only if the outcome replaced enough currently-imposed requirements to lead to a real saving. To the degree that existing requirements are somewhat interlocking and mutually reinforcing, it is impractical to change only one or two isolated requirements; therefore, it will be difficult to adopt a piecemeal approach to trying out the ideas presented here. It may be necessary to take a sizable step, in order to make any real progress at all.

6.2 Recommended Process

Performance-based regulation is seen as an ongoing process of monitoring the ongoing implementation of a comprehensive safety case. Accordingly, the recommended process calls initially for careful formulation of a safety case, followed by development of a diamond tree covering at least key portions of the safety case. Careful formulation of the safety case in success space begins the process of picking up important performance elements that are not typically picked up in PRA. Carrying the development further, in the diamond tree, picks up institutional influences on performance that a PRA-based safety case would not naturally address.

In addition to picking up institutional influences, the diamond tree also arranges performance nodes in a hierarchy. This turns out to be useful in more than one way. It supports a systematic approach to choosing results-oriented measures for consideration in a performance-based scheme, and it provides an interesting way to look at significant events.

Making use of the hierarchical character of the diamond tree, the process begins with a trial monitoring scheme pegged at the most results-oriented level of the diamond tree (the top), and systematically works its way down the tree, looking for nodes at which the issues can be adequately addressed. Since the search encounters high (more results-oriented) nodes first, the collection of performance nodes that it tends to identify as a regulatory approach is the most performance-based approach that is possible, consistent with addressing the issues.

The purpose of a performance-based approach is to replace prescriptive regulations, based on our ability to assess the ongoing fulfillment of a significant portion of the safety case by monitoring a carefully-selected set of performance measures. Although this set of measures should be as results-oriented as possible, the set should be very comprehensive, including instrumentation and supports. (Within the diamond tree approach, comprehensiveness means that at a high level on the tree, all nodes are "assured" to an adequate extent.) The recommendation in this report for the use of a comprehensive set is a basic difference between the current approach and previous performance-indicator work, which overlies current regulations, and whose different purpose can be served by more of an audit approach.

Once we have a collection of more or less performance-based nodes, we need to decide what actual measures to associate with each node. This depends on how this node influences the upper levels (through its reliability, availability, quality, efficiency,…). Even a train-level node in a standby system has several figures of merit potentially worth measuring: unavailability averaged over short and long time scales, failure rates, mean time to restore, reliability, and so on. Note that the choice is not simplified just because (for example) a PRA depends only on unavailability of this train. For one thing, other indicators at that node may provide earlier warning of problems; for another, we also need to consider that small changes occurring simultaneously in a wide range of performance nodes may correlate with an emerging problem several levels down in the diamond tree.

For each node whose performance we intend to monitor, we need a model relating lower-level performance aspects to the performance of the monitored node. We need this because we need to understand whether there is a significant class of contributors to failure of this node that have a high conditional probability of going on to an unacceptable consequence. There are multiple possibilities for this kind of thing. Some involve hardware common cause, perhaps coupling failure at this node to failure at another, or functional relationships that mask performance shortcomings at a given node until a significant problem exists. Others involve institutional factors, which can simulate a negative common cause influence across diverse systems. We can model the hardware aspects of this up to a point, but still need a structured way to address the institutional factors.

The diamond tree idea has not been widely applied; it is not a basis for a computational algorithm (comparable to a fault tree calculation) and there is no standard format for it. Filling one out completely would be a time-consuming exercise. However, its thought process seems to drive the analyst straight into the essentials of plant programs, and to elicit identification of useful performance aspects; moreover, its hierarchical format lends itself not only to comparison of regulatory approaches but also to discussion of significant events. Further experimentation leading to standardization of this construct might be worthwhile.

6.3 Comments on Performance Measures

A brief discussion was given above to the effect that it is inappropriate to try to base a monitoring scheme on a rare event, because non-observation of the event does not tell us very much about the underlying situation. If a redundant system is considered to be important, then we need to be monitoring at or below the train level. In the safety system unavailability work mentioned in the appendix (Boccio et al., Azarm et al.), this point was discussed at some length. That work argues for an unavailability indicator monitored at the train level (requantified as appropriate at the system level), reflecting various contributors to downtime. That recommendation is consistent with the discussion given in this report.

Focusing on train availability immediately brings up the topic of AOT restrictions. Note that over a series of train outages, significant downtime can accrue with no violations of an AOT. On the other hand, stringent limiting conditions of operation can cause inconvenience to a plant that needs a longer-than-usual downtime to address a particular problem. It has been noted many times that a fixed AOT is capable in principle of inflicting cost without necessarily promoting high availability. This suggests that AOTs are candidates for replacements by some kind of performance-based regulation.

Unavailability, computed according to accepted ground rules and averaged over a suitable time scale, appears to meet the tests of measurability and objectivity called out in the definition of “performance-based.” One could imagine a set of agreed-upon target train indicators. A comprehensive set of such indicators would tell a great deal about a plant’s safety situation. In a way, it would not tell enough, because joint outages of certain pairs of trains would be unacceptable, so it would be necessary to supplement single-train quotas with guidance limiting the flexibility to have multiple outages. This leads naturally into the notion of general use of an on-line risk model to assign a weight to a given train outage, based on how much increased risk was caused by the outage (relative to the risk that would obtain if the train were in fact operable).

6.4 Institutional Factors

Portions of this report have mentioned approaches that have some potential to identifying important performance nodes that are controlled by institutional factors. It was shown that the diamond tree thought process lends itself to identification of many useful indicators of institutional performance; compared to a fault tree process, for example, the diamond tree development readily identifies many important indicators. But the process described above does not lead to an emphasis on institutional factors as monitoring points, largely because it is so strongly biased towards results-oriented measures. This is deliberate, and is a consequence of the formulation of the concept of performance-based regulation, but it is not clear that a comprehensively satisfactory scheme will result.

Appendix A discusses approaches to regulatory oversight applied by agencies other than NRC, which basically oversee their plants' *processes* of safety management, rather than trying to monitor performance directly in the manner suggested here for results-oriented performance nodes.

In light of the above, one is led to ask whether it is possible to formulate a hybrid performance-based approach, in which a comprehensive set of results-oriented performance nodes (situated in the upper half of the diamond tree) would be monitored along lines discussed throughout this report, while institutional factors (in the lower half of the diamond tree) would be addressed through NRC oversight of a licensee process, rather than through NRC oversight of low-level performance nodes such as "questioning attitude" or "delays in spare parts procurement." This is the subject of the following section.

7 OPTIONAL APPROACHES

7.1 Discussion of Options

For purposes of this report, it has been assumed that there are significant benefits to performance-based regulatory oversight, defined by its key attributes of

- measurable parameters to monitor acceptable plant and licensee performance;
- objective performance criteria;
- licensee flexibility to determine how to meet established performance criteria;
- leading-indicator character, that is, failure to meet a performance criterion must not result in significant adverse consequences.

The present project has been conducted as a search for a collection of outcome-oriented parameters that adequately reflect the current level of plant safety, and concurrently a search for a decision process that accepts these parameters as inputs and furnishes appropriate recommendations to the regulator as outputs.

The fourth attribute listed above has been discussed extensively in this report. It is related to the concept of margin: it says that there should be margin, as measured both by key plant physical parameters and by conditional probability, between failure to meet a performance criterion and the occurrence of significant adverse consequences. If the monitored parameters are pitched at too high a level, then unacceptably adverse consequences may result before performance problems show up in the monitored parameters. A closely related issue—really, one aspect of the leading-indicator issue—is that of the adequacy of conventional PRA information for purposes of performance-based regulation. Conventional PRA does not reflect many details of the interaction between the staff and the physical plant, and neglects institutional factors in general. The approach to performance-based regulation needs to address the potential for institutional factors to increase the conditional probability of adverse consequences, in ways that might not be revealed in all monitoring schemes.

Setting aside the costs of accidents, and the costs associated with regulatory intervention in situations where problems have been allowed to develop for too long, it seems clear that the greatest economies are achieved if regulatory oversight is pitched at the highest (most results-oriented) level possible. In choosing to focus on safety function performance, the NEI proposal provides an example of a very high-level, results-oriented performance node (keeping RCS temperature below a target temperature during a particular phase of shutdown). (Recall that "performance node" refers generically to almost any aspect of performance that appears on a diamond tree.) However, the leading-indicator issue and the institutional factors issue point to a need for oversight at a lower level; indeed, pure cost considerations might support lower level choices as well, if a convincing prospective analysis could be performed of the costs of belated regulatory intervention (extended shutdowns,…). (Recall that "level" in this discussion usually refers to hierarchical level in the diamond tree sense. Examples are given below.) The process outlined in Section 4 is essentially a stepwise approach to balancing these competing needs, assuming that the regulatory needs would drive the monitoring scheme to a lower level than the pure cost needs would.

A possible first-cut result from the process outlined in Section 4 is indicated schematically on Figure 12 Sheet 1. This figure sketches a hierarchy of levels of performance (omitting some of the details presented in earlier sections) and schematically indicates performance-based regulatory oversight pegged above the dotted line (i.e., at train-level performance or above). The process given in Section 4 will drive towards an outcome like this, because the process is biased to stop when some assurance of train-level performance is likely to be forthcoming. The process will also identify performance nodes at which prescriptive requirements are the only realistic option, because it is impractical to monitor outcomes. Therefore, some prescriptive requirements remain, as indicated on the figure.

In light of previous discussion of the importance of institutional factors and human performance at the interface with plant SSCs, one can ask whether even train-level performance monitoring is enough, or whether we need to address lower levels somehow. One way of addressing lower levels would be to adopt a scheme like that shown in Figure 12 Sheet 2. This figure shows the regulator monitoring, to some degree, at all levels. The only reason to consider such an invasive scheme, even for purposes of discussion, is that at present, we do not know how to convincingly model quantitative performance at or above the component level in terms of institutional performance below that level (except for certain human errors), and cannot therefore be sure whether the leading-indicator issue is addressed. This comment is intended to include many aspects of performance, including root cause analysis of failures, and prevention of common cause failures. It can be argued—indeed, is implicitly argued in the NEI proposal—that declining performance in the lower levels will show up at the higher levels before adverse consequences follow. If we could model this claim convincingly, and reconcile the model with events documented in AIT reports, the claim would be easier to accept. Indeed, after some experience with the scheme outlined below, it may become easier to accept.

Although the scheme in Figure 12 Sheet 2 has at least the potential of addressing performance more convincingly, such a scheme is not results-oriented, and would presumably sacrifice many of the efficiencies that motivate inquiry into a performance-based approach. Indeed, if resources permitted, a scheme like this could be invasive enough to be potentially much less efficient than the current approach. This is not to say that performance at the lower levels should not be monitored by the licensee, for licensee purposes; the benefits of such monitoring were the point of the early work on the diamond tree. But regulatory oversight at those levels appears to have drawbacks for both licensees and regulators.

The scheme in Figure 12 Sheet 3 may provide an interim trial solution. The point of the scheme in Figure 12 Sheet 3 is to allow the regulator to address lower-level performance areas indirectly, and not generally in an invasive way (unless, of course, problems were identified). In this scheme, regulatory oversight would be implemented at the lower levels not by regulatory monitoring of the details of licensee manipulations of plant hardware, but by regulatory oversight of licensee management processes aimed at promoting high levels of performance.

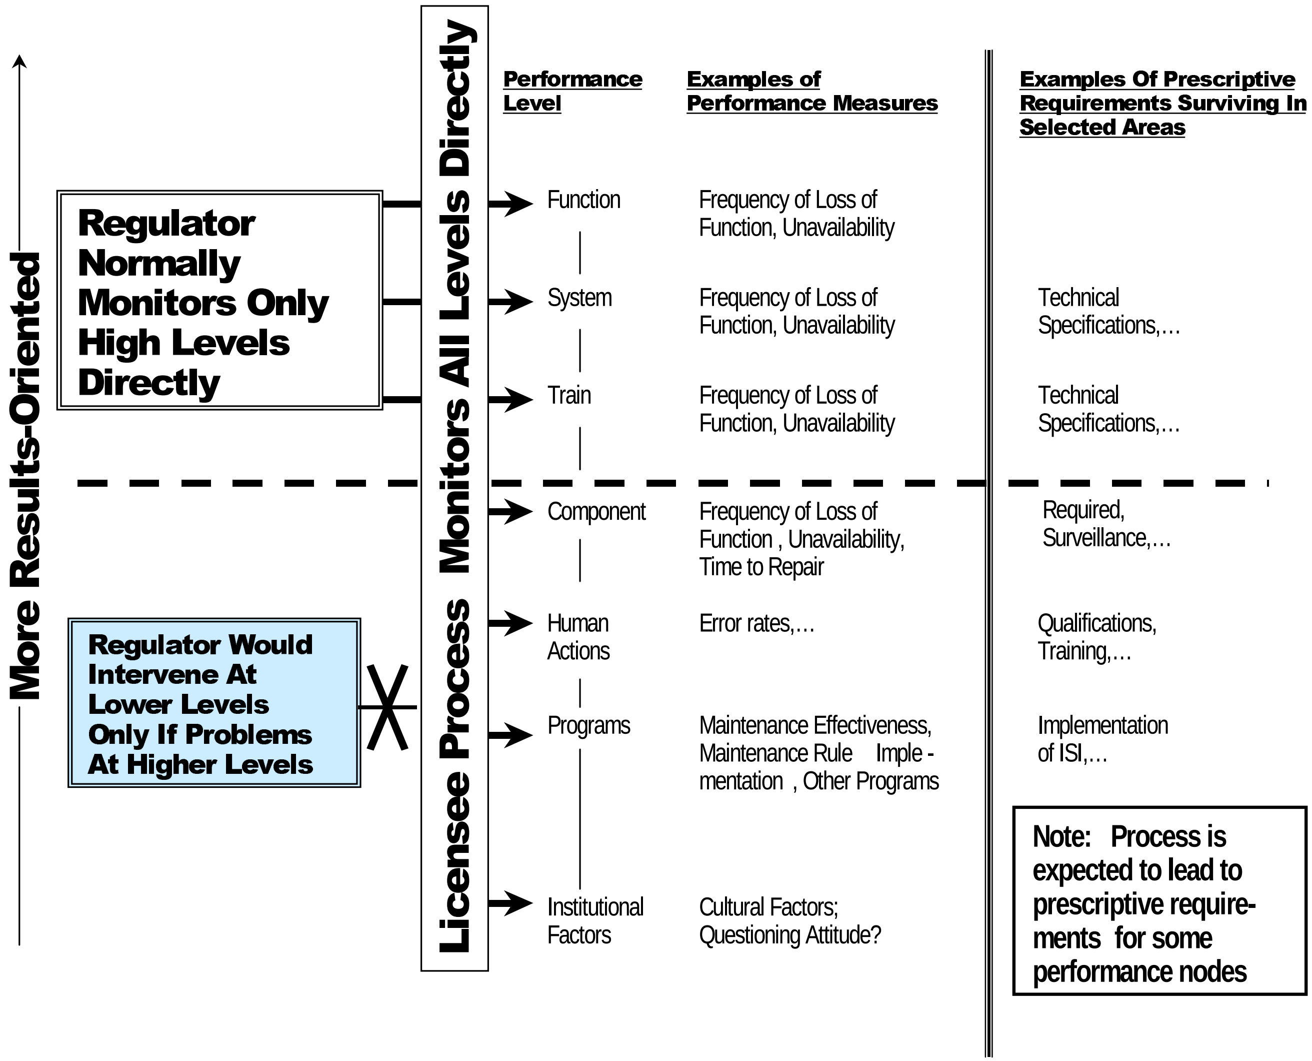

More Results-Oriented

Regulator Normally Monitors Only High Levels Directly

Regulator Would Intervene At Lower Levels Only If Problems At Higher Levels

Licensee Process Monitors All Levels Directly

Performance Level	Examples of Performance Measures	Examples Of Prescriptive Requirements Surviving In Selected Areas
Function	Frequency of Loss of Function, Unavailability	
System	Frequency of Loss of Function, Unavailability	Technical Specifications,…
Train	Frequency of Loss of Function, Unavailability	Technical Specifications,…
Component	Frequency of Loss of Function , Unavailability, Time to Repair	Required, Surveillance,…
Human Actions	Error rates,…	Qualifications, Training,…
Programs	Maintenance Effectiveness, Maintenance Rule Imple - mentation , Other Programs	Implementation of ISI,…
Institutional Factors	Cultural Factors; Questioning Attitude?	

Note: Process is expected to lead to prescriptive require-ments for some performance nodes

Figure 12. (Sheet 1 of 3) Monitoring High Levels and Prescriptive Requirements.

There is partial precedent for this kind of monitoring. One precedent is the maintenance rule. The maintenance rule tells licensees to monitor the performance or condition of SSCs against licensee-established goals, and act according to performance with respect to those goals. Although this is sometimes informally said to be a “performance-based” rule, it is the licensee process that is performance-based; regarding the relationship between regulator and licensee, the rule could almost be said to be meta-performance-based.* That is, it is a rule that prescriptively requires a performance-based licensee process, rather than mandating performance-based regulatory oversight. The regulator audits the licensee process, and not just the safety performance of the SSCs within the scope of the program.

Other precedents were mentioned in Appendix A. Two in particular are EPA’s CAM (Compliance Assurance Monitoring) and OSHA’s PSM (Process Safety Management). The EPA rule is a performance-based way of keeping emissions within allowable limits, and seems not to be a close parallel to the present application. The PSM rule is a bit closer. As discussed in Appendix A, this rule requires licensees to proactively identify and manage safety issues; the standard mandates employee participation in PSM programs, preparation of process safety information, process hazards analysis evaluations, written operating procedures, employee and contractor training, pre-startup safety reviews, maintenance of the mechanical integrity of critical equipment, written procedures for managing change, evaluation of incidents and near misses, emergency action plans, and so on. Licensees have flexibility in how to address these elements, and OSHA may audit their processes for doing so.

Although monitoring of licensee process sounds familiar, it is reiterated that the background of the Figure 12 Sheet 3 scheme is a far more explicit and detailed allocation of performance than has been done in previous applications. The point is, after all, to support a safety case, and to use it to eliminate some existing requirements. For example, one would look at various aspects of maintenance activities in light of explicit expectations of component unavailabilities. In light of the allocation, long repair times for particular components might show up as violating a prior expectation, and licensees would be expected to address such issues within their processes. In order for licensees to address an issue like this before it showed up at the train level, where the regulator would notice it directly, the licensee would need to formulate more ambitious goals at the lower levels than would be needed just to support the safety case.

7.2 Best Option

The best option seems to be to implement the recommended process (Section 4) with a view towards implementing a scheme like that shown in Figure 12 Sheet 3 on a specific plant, preferably with a participating licensee.

It might be possible to test this approach working with part of a safety case. As remarked previously, it is difficult to do a meaningful performance allocation on too narrowly-defined a scope, because many requirements in seemingly different areas are mutually supporting; too

* Definition 3 of “meta”: More comprehensive: transcending <*meta*psychology>—used with the name of a discipline to designate a new but related discipline designed to deal critically with the original one <*meta*mathematics> [Webster’s Ninth New Collegiate Dictionary]

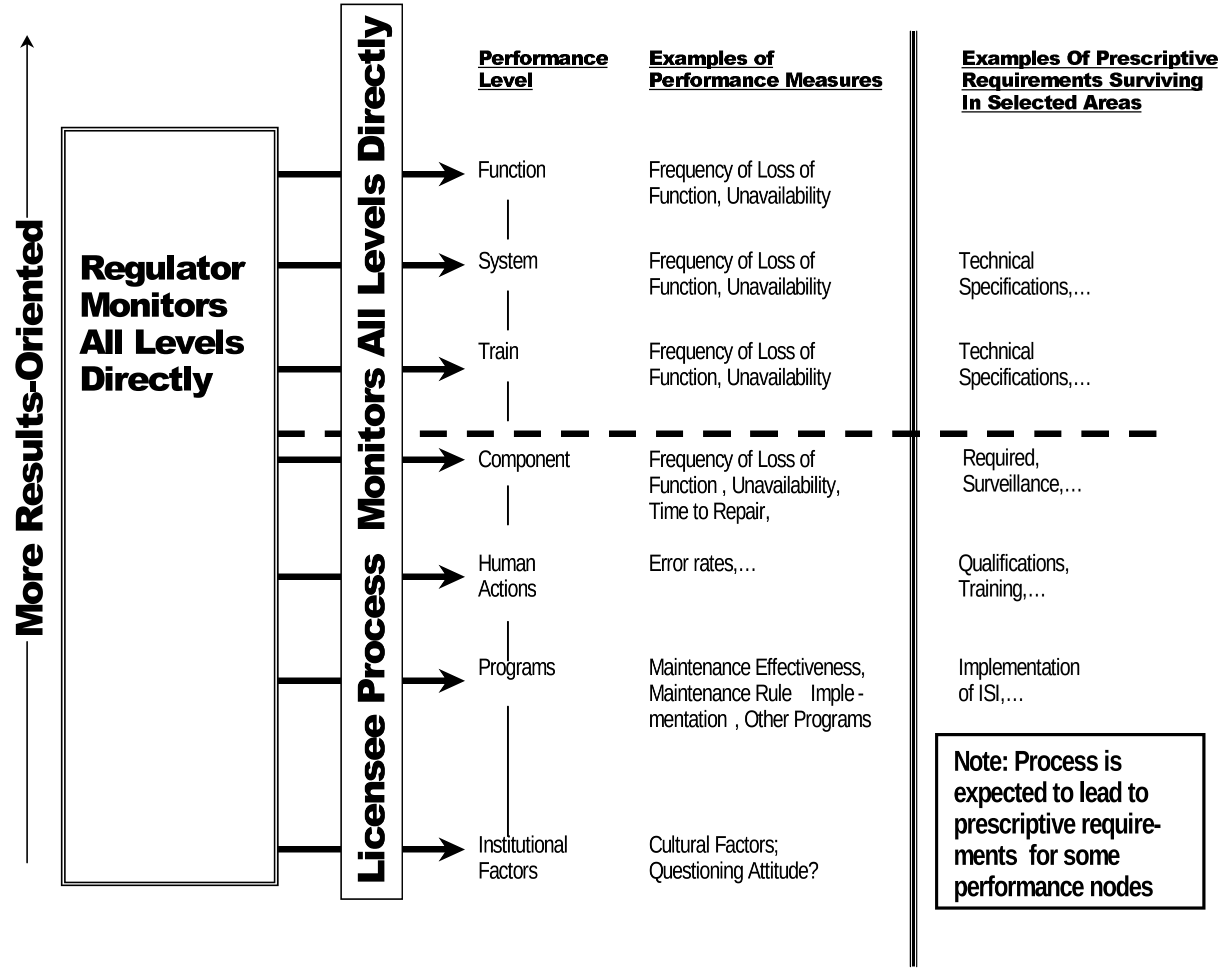
More Results-Oriented
Regulator Monitors All Levels Directly
Licensee Process Monitors All Levels Directly
Performance Level
Examples of Performance Measures
Examples Of Prescriptive Requirements Surviving In Selected Areas
Function
Frequency of Loss of Function, Unavailability
System
Frequency of Loss of Function, Unavailability
Technical Specifications,…
Train
Frequency of Loss of Function, Unavailability
Technical Specifications,…
Component
Frequency of Loss of Function , Unavailability, Time to Repair,
Required, Surveillance,…
Human Actions
Error rates,…
Qualifications, Training,…
Programs
Maintenance Effectiveness, Maintenance Rule Imple - mentation , Other Programs
Implementation of ISI,…
Note: Process is expected to lead to prescriptive require- ments for some performance nodes
Institutional Factors
Cultural Factors; Questioning Attitude?

Figure 12. (Sheet 2 of 3) Monitoring of All Levels and Prescriptive Requirements.

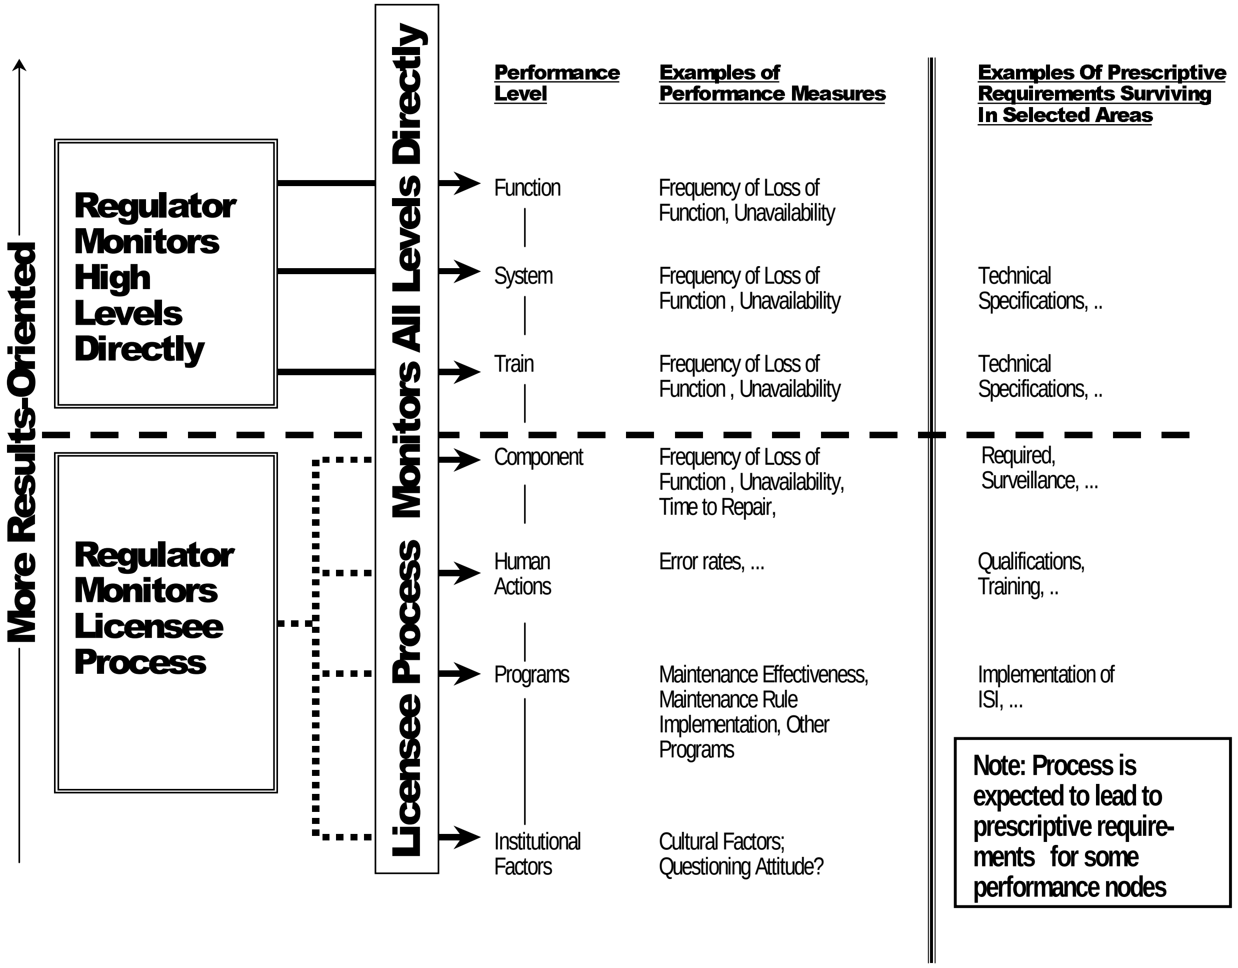
More Results-Oriented
Regulator Monitors High Levels Directly
Regulator Monitors Licensee Process
Licensee Process Monitors All Levels Directly
Performance Level
Examples of Performance Measures
Examples Of Prescriptive Requirements Surviving In Selected Areas
Function
Frequency of Loss of Function, Unavailability
System
Frequency of Loss of Function , Unavailability
Technical Specifications, ..
Train
Frequency of Loss of Function , Unavailability
Technical Specifications, ..
Component
Frequency of Loss of Function , Unavailability, Time to Repair,
Required, Surveillance, ...
Human Actions
Error rates, ...
Qualifications, Training, ..
Programs
Maintenance Effectiveness, Maintenance Rule Implementation, Other Programs
Implementation of ISI, ...
Institutional Factors
Cultural Factors; Questioning Attitude?
Note: Process is expected to lead to prescriptive require-ments for some performance nodes

Figure 12. (Sheet 3 of 3)Monitoring of High-Level Nodes, Licensee Process, and Prescriptive Requirements

narrow a scope has the additional disadvantage of eliminating fewer requirements for a given amount of work. While it is asking too much to overhaul all of regulation in a single sweep, it appears to be hoping for too little to make only a few changes to isolated requirements. For purposes of a trial, it might work to perform a reasonably broadly scoped performance allocation, and then consider changing the regulatory approach within a more narrowly defined scope, based on performance allocated in a specific functional area.

Initially, one would follow a process like that discussed in Section 4, subject to the above caveat regarding scope, and with perhaps more iteration and looping back than was made explicit in the discussion provided in Section 4. This would lead to a set of goals that would mostly appear above the dotted line in Figure 12 Sheet 3, along with prescriptive requirements imposed at performance nodes to which performance has been allocated, but for which outcome-based schemes cannot be expected to work. Such prescriptive requirements will tend to be pitched at the train/component level.

Ideally, within the scope of the trial, many other prescriptive requirements presently appearing above the dotted line would be eliminated. In order to justify this, we need assurance that overall satisfactory performance can be monitored effectively, and there is some doubt that we can do that based only on nodes appearing above the dotted line. Therefore, as suggested above, we would address lower-level requirements in terms of oversight of licensee process. For purposes of the trial, discussions would need to be held with the participating licensee regarding his process for managing performance at levels below the dotted line. All licensees have some management activities at those levels, and some appear to have very significant management activities. The question would be how to portray these activities to the regulator in a way that provides convincing evidence of high performance, without essentially involving the regulator in plant decision-making.

In the near term, such a trial would be expected to provide insight into:

- the workings of the allocation process;
- what levels of allocated performance actually make sense for the subject plant;
- what performance nodes appear to constitute a reasonable and monitorable safety case for the functional performance areas chosen;
- what degree of lower-level performance appears to be needed to provide adequate assurance to support train/component level performance measures;
- how the regulator derives assurance of performance at the lower levels.

The last item, of course, is potentially specific to the formulations emerging from the application of the process in the specific example.

In the longer term, one would hope to verify whether the regulatory oversight of licensee process at the lower levels is in fact necessary, or whether an adequate diagnostic tool can be formulated at the level of the train/component level indicators.

8 REFERENCES

1. "Strategic Assessment Issue: 12. Risk-informed, Performance-based Regulation," USNRC, Release Date: September 16, 1996 (http://www.nrc.gov/NRC/STRATEGY/ISSUES/dsi12isp.htm).
2. "Risk-Informed, Performance-Based Regulation," Attachment to Memorandum, Jackson to Seale, "Discussion on Risk-Informed, Performance-Based Regulation," February 20, 1998 (USNRC).
3. NRC Information Notice 95-36: Potential Problems With Post-Fire Emergency Lighting (USNRC, August 29, 1995).
4. "Improving the Regulatory Process Through Risk-Based and Performance-Based Regulation," Nuclear Energy Institute, Draft 10/30/95, attachment to letter from Rasin (NEI) to Milhoan (USNRC), November 14, 1995.
5. "NRC AUGMENTED INSPECTION TEAM REVIEW OF THE UNDETECTED INTRODUCTION OF NITROGEN GAS INTO THE REACTOR VESSEL DURING PLANT SHUTDOWN REPORT NO. 50-213/96-80," Enclosure to Letter, Hubert J. Miller (NRC) to Mr. Ted C. Feigenbaum, dated October 30, 1996.
6. See, for example, "Precursors to Potential Severe Core Damage Accidents: 1994/A Status Report," NUREG/CR-4674 Vol. 21, R. J. Belles et al. (ORNL/NOAC-232, ORNL, 1995). On page 3-4, under the heading "Important Precursors," it is stated that "Events with such conditional probabilities [greater than or equal to 10^{-4}] have traditionally been considered important in the ASP Program."
7. "A framework for assessing influence of organization on plant safety," M. Modarres et al., Reliability Engineering and System Safety 38, 157 (1992).
8. *What Every Engineer Should Know About Reliability and Risk Analysis*, M. Modarres (Marcel Dekker, 1993).
9. "Use of the 'Goal Tree Methodology' to Evaluate Institutional Practices and Their Effect on Power Plant Hardware Performance," R. N. M. Hunt, M. Modarres, and M. L. Roush, in *12th Inter-RAM Conference for the Electric Power Industry*, p. 304 (1985).
10. "Application of Goal Trees in Reliability Allocation For Systems and Components of Nuclear Power Plants," M. Modarres, M. L. Roush, and R. N. M. Hunt, in *12th Inter-RAM Conference for the Electric Power Industry*, p. 337 (1985).
11. See also Draft NUREG/CR, "A Framework and Method for the Amalgamation of Performance Indicators at Nuclear Power Plants," Vols. 1 and 2, Wreathall et al. (draft dated April 15, 1992). This report was prepared under NRC Contract No. NRC-04-87-070, and is docketed under Accession Number 9805210039.
12. "Applying Risk Models To Formulation Of Safety Cases," R. W. Youngblood, Risk Analysis, 18, No. 4, p.433 (1998).

Appendix A

Issues in Performance-Based Regulation:

A Review and Discussion of Some Pertinent Literature

A.1 Introduction

In a recent paper[1], the staff identified certain characteristics of a performance-based regulatory approach:

- There are measurable parameters to monitor acceptable plant and licensee performance.
- Objective performance criteria are established to assess performance.
- There is licensee flexibility to determine how to meet established performance criteria.
- Failure to meet a performance criterion must not result in unacceptable consequences.

That is, “performance” is assessed in terms of monitored parameters. Based on this, the present project is seen as a search for such parameters that actually bear on the current level of plant safety, and concurrently a search for a decision process that accepts these parameters as inputs and furnishes appropriate recommendations to the regulator as outputs.

This appendix includes a brief literature survey and a discussion of certain key concepts that came out of the search. The present emphasis is not on detailed criticism and assessment of each idea, but rather on characterizing the spectrum of ideas and tools that are available and that might be applicable to performance-based regulation.

Before discussing the papers and books that have been examined, it is useful to introduce certain essential ideas by example. This is done with reference to an NEI treatment [2] that included a sample rule as an example of a performance-based regulation. The NEI example appears to have some features to recommend it, and other features that would be seen as shortcomings, unless they were addressed by program elements rather different from those mentioned in the example. After this example is presented, some key issues are noted, and variations on the basic ideas are mentioned briefly.

Based on concepts and concerns emerging from discussion of the NEI example, the organizing principle of the literature search is that it looks for items that bear on the basic ideas or key issues. Highlights of that search will be summarized. First, the discussion covers some of the literature aimed at figures of merit to be considered. Next, a very brief discussion is provided of some of the ideas that might inform the formulation of the integrating tool. Finally, preliminary insights are summarized for purposes of supporting a discussion of near-term work.

A.2 Key Concepts and Immediate Issues

The industry has explained its views in many places. In one of the fullest expositions [2], a draft example “rule” is provided to illustrate the industry’s idea of what is desirable. In general, industry wants measurable parameters to monitor plant and licensee performance, objective criteria to assess performance based on risk insights, deterministic analyses, or performance history, and licensee flexibility to determine how to meet established performance criteria.

Note that these are the first three of the four characteristics recognized by the staff in the DSI paper quoted above. The fourth element, added by the staff, is intended to assure that performance criteria provide early warning of impending problems.

The industry presentation provides an illustration in the form of a “conceptual rule” entitled “Monitoring the Performance of Shutdown Safety Functions.” The “rule” requires the licensee to identify SSCs performing specified safety functions, establish contingency plans to address situations when the minimum planned equipment is not available to perform those functions, monitor the performance of these systems against licensee-established performance criteria in a manner that provides reasonable assurance that safe shutdown conditions are being maintained, and take corrective actions. The core of the example is the idea that the performance criterion for decay heat removal would be maintaining 80% of the thermal margin to boiling. Regulatory oversight would depend on performance in the following way.

> If the licensee consistently maintains RCS temperature below 154 degrees F, i.e., meets its performance criterion, it is a results-oriented manner of demonstrating that its controls and contingency plans associated with the decay heat removal function are adequate and that a high margin of safety is being maintained. Regulatory oversight for this level of performance should be minimal…. If the licensee has an event where the decay heat removal function is impacted temporarily, but RCS temperature remains below 154 degrees F, the licensee would perform a cause determination and take appropriate corrective action to address the cause. This cause determination and corrective action should be documented.…If the licensee has a loss of decay heat removal event and fails to meet the criterion,…[regulatory oversight may expand at this point to include an assessment of the cause and corrective action taken by the licensee.

This example very usefully focuses a number of ideas and issues that seem significant. A key phrase is

> “a results-oriented manner of demonstrating that…controls and contingency plans associated with [each safety function] are adequate and that a high margin of safety is being maintained.”

Does maintaining RCS temperature over a period of weeks really demonstrate that controls and contingency plans are adequate? What does “adequate” mean? How might a criterion like this one be linked to CD and LR objectives? Presumably, in this example, vessel inventory would be regulated under the “inventory” safety function, rather than the DHR safety function. But if not—if regulation of inventory even in non-LOCA situations is within the purview of this function—then one would need to ask whether bulk temperature spans the space of relevant considerations.

In this example, many events triggering NRC attention would essentially correspond to “precursors,” in the sense of NRC’s Accident Sequence Precursor program. Setting aside momentarily the question of whether the threshold of regulatory significance is set too high in this particular example, it is an interesting point that PB regulation has some relationship to precursor analysis, especially if we generalize the notion a bit. This is taken up in the following subsection.

Performance-based Regulation as Generalized Precursor Analysis

For many years, NRC has sponsored a project [3] to analyze events that occur at operating plants to see whether they are “precursors.” The sense of this term is that for some classes of events, it is reasonable to expect severe events to have had some partial precedents in previous operating experience. Suppose that failure events A, B, and C each have a probability of 0.1 conditional on initiating event I, and that $I*A*B*C$ is a core damage sequence. Assuming that A, B, C are independent (other than being conditional on I), we will see $I*A*B$ ten times as often as we will see $I*A*B*C$; thus, there is a very high likelihood that we will experience the precursor event $I*A*B$ before we will ever see the whole sequence. Indeed, there were at least two events that were precursors to the TMI accident, in that they involved a stuck-open PORV and perhaps other failures, and the sense of a precursor program is that understanding events like that may help us to prevent the complete scenario (e.g., by fixing the relative vulnerability that drives the scenario’s likelihood).

In order for a precursor event to seem significant, it must in some sense violate an expectation. Suppose that our prior assessment for events A, B, and C is that their probabilities conditional on I are 0.01, rather than 0.1, and that I occurs on the order of once per year. Then we do not expect to see $I*A$ or $I*A*B$ in any one plant (we might see $I*A$ in the fleet), and observation of such events might be a signal that our model of the facility is wrong. Of course, we may suppose initially that the observed instance of $I*A*B$ was just a fluke, but at least our attention will be drawn to it.

How does common cause affect this picture? If we postulate a CCF that fails all three elements with a probability not much lower than failing two of them, then the likelihood of the full sequence occurring before any precursor is changed very significantly. CCF is a significant issue for performance-based regulation, and we will return to it later.

For present purposes, the point is that the industry proposal can be read at a high level as a proposal to regulate by a kind of precursor analysis. That is, “performance” is the prevention of precursors, where a “precursor” is an event or condition that causes the probability of the undesired consequence to be higher than some threshold; if precursors are being prevented, the NRC can conclude that plant programs are in good shape.

One apparent difference between the way in which events are characterized in the NEI proposal and the way in which events are characterized in the NRC precursor program is that in the latter, most of the thresholds are articulated in terms of number and kind of equipment failures, or challenges to a function, rather than in terms of things like RCS temperature. It would appear that some events that NRC would call precursors (e.g., challenges to safety systems, loss of safety systems,…) might not trip the RCS temperature criterion, if recovery of function was accomplished promptly. There is the potential, in other words, that the conditional probability of core damage could be appreciable in some events that did not actually cause much heatup.

The above discussion focuses on the physical state of the plant as measured by failure-counting or RCS condition; the condition of plant programs is not explicitly mentioned. It could be argued that some sort of breakdown in plant culture could affect many basic events simultaneously, causing a significant increase in effective CDF. This would eventually manifest itself in observed failures; but would the manifestation occur in time to do any good? Or would a performance-based program need to try to measure institutional factors? If we were convinced that a significant programmatic deficiency would manifest itself in the more hardware-oriented figures of merit, then we would be assuming that operational events would signal the problem before a CD sequence ever went to completion. Some of the extant performance-indicator work appears to implicitly assume this; scram rate, for example, is believed by some to provide an indication that things are not as they should be. A single scram event with nothing else wrong is not necessarily a precursor event in any accepted sense of that term, but (according to some) a sequence of scrams may indicate that multiple basic-event probabilities are higher than one had thought. In this sense, a sequence of scrams signals an underlying condition that is arguably a kind of precursor.

It appears that institutional factors and scram rate fall outside the envelope of “results” within which the industry had apparently hoped to focus staff attention. The assumption in this project is that we should look at a broad set of measures; we may find that a subset of these can be chosen that has all the properties that the regulator needs and also meets industry preferences, but it presently seems undesirable to limit our review to end-result indicators, because there are reasons to believe that real end-result indicators are not sufficiently leading.

Implicit in some of the above discussion is the idea that we are driving towards a risk-informed process. As mentioned previously, the present scope emphasizes “how do we regulate,” and not how to decide what is important. It is taken for granted here that, consistent with Commission policy, risk insights will be used to the extent practicable, and deriving insights from PRA is the scope of other projects and not this one. But as noted in recent draft regulatory guidance, even in the context of risk-informed applications, we are obliged to go beyond PRA, to what the staff presently calls “integrated decision-making.” Existing PRAs are not considered suitable as a sole basis for decision-making, because of incompleteness, some aspects of which will persist for fundamental reasons. If we accept that the need for integrated decision-making is real, then we must expect that a performance-based scheme will need to address performance elements that are not modeled explicitly.

However, it appears that we lose nothing, and perhaps gain some focus, from thinking of performance figures of merit as measuring the precursor potential, where “precursor” is understood to mean “an event or condition that implies a higher-than-previously-assumed conditional probability of an undesired consequence (core damage or large release),” and where it is understood that the threshold of significance is set to a more conservative level than in NRC’s ASP program. This definition should be understood to include institutional factors among conditions implying a higher-than-previously-assumed conditional probability of undesired consequence. At this stage, it is not assumed that measurement of institutional factors will required; this question is left open, and will be decided based on whether this is necessary in order to provide an adequate assessment of the state of the institution.

Phases in Performance-Based Regulation

Some papers in this area do not distinguish between regulatory considerations at licensing and regulatory considerations during a plant’s operational lifetime. One that does is [4]. This paper is a good read in general; in the interest of brevity, only two excerpts are given:

> The Robens Committee concluded [20 years ago??], first, that the single most important cause of accidents was apathy—on the part of all concerned in industry; and second, that a major cause of this was that there was simply too much law.

Confirmatory instances come readily to mind; one wonders, for example, whether some of the perceived problems at Millstone are a reflection of this syndrome.

> He or she [i.e., licensees] may…be required to prepare a safety case demonstrating how the risks will be prevented or otherwise controlled, and to set out a safety management system showing how the safety case will be implemented and maintained.

Evidently, it is the “safety management system” that must “perform” in the present discussion. Thus, one envisions a process in which a plant’s safety case is first accepted by the regulators; then, in a related but distinct finding, the safety management system derived from it is found to be adequate; finally, a safety-case and safety-management-system-specific set of performance goals can be derived and applied. The present project begins at the point where the safety case has been formulated, and we assume that we are concerned primarily with regulatory assurance of satisfactory safety performance of an already-licensed entity.

Competing Notions of “Performance-Based”

It seems clear from the NEI paper that their notion of “performance-based” is aimed at the relationship between the regulator and the licensee—at the timing and nature of regulatory visits—and not just at basing regulatory requirements on a calculated performance requirement. Not everyone uses the terms in this way. Staff work on Appendix J, for example [5], explicitly considered that the then-current performance requirement on containment leakage could be substantially relaxed without much risk impact. On this basis, one could argue for an increased interval between tests. If all that is changed is the interval of a test that is dictated in a fundamentally prescriptive regulation, then the new regulation is “performance-based” only in the sense of having taken explicit account of a safety-case performance requirement. Even an algorithm for varying the test interval as a function of experienced performance would not qualify as a performance-based regulation on NEI’s definition, if the staff audited the licensee in a manner qualitatively similar to the old way. It would appear that according to NEI, a truly performance-based regulation in this area would replace existing regulations with an analog of the DHR example given above; in this analog, NRC would not inspect in today’s manner, and would stay away unless the licensee consistently had difficulty meeting a leak rate criterion. In fact, the staff could have kept the old conservative leak rate criterion, and changed regulatory practice to audit if, and only if, the licensee had trouble meeting it, and this would have been a performance-based rule on the NEI definition. It is noted that the paper cited above does not say “performance-based,” it says “performance-oriented,” and has therefore not strictly confused this terminology; however, it appears that confusion on this terminology is widespread, and people who do not distinguish “based” from “oriented” may be misled. Indeed, a staff SECY [6] explicitly refers to this performance-oriented work as performance-based, even though the descriptions of the work reviewed so far focus on the relative generosity of the new requirement, rather than on the qualitative nature of the relationship between regulator and licensee.

In the present report, “performance-based” is construed a bit more widely than NEI would suggest, but we adhere to the NEI notion (as we understand it) that the defining characteristic of performance-based regulation is that the interactions between staff and licensees are determined by how well things are going, and not by audits of compliance with prescriptive how-to guidance. In this appendix, we deviate from the NEI definition only because, at this stage, we wish to consider a broader range of schemes that do this; this is probably necessary because it remains to be established whether end-results indicators are leading enough. But we join with NEI in choosing to focus on the NRC/licensee relationship. Quantifying more sensible performance requirements is already a facet of a great deal of ongoing work in risk-informed regulation, and we have other priorities in this project.

Summary Statement of the Problem

Performance-based regulation is formulated here as a regulatory approach in which the regulators satisfy themselves for each plant on an ongoing basis that the likelihood of undesirable events is sufficiently low, not by enforcing prescriptions on how certain plant programs must be carried out, but by accepting certain indications of performance as evidence that plant programs are being carried out satisfactorily. “Performance-based regulation” is here understood to apply during the operational life of the facility, after the design has been reviewed and a license has been granted. For our purposes, an essentially valid safety case is supposed to exist, and the staff and licensees have a thorough mutual understanding of what is required in order to make the safety case come true. Performance-based regulation is oversight of the implementation of the safety case. The major questions to be answered here are (1) what performance indications should the regulator work with? and (2) what form should the integrating tool assume? (i.e., how are these indications to be used to make a decision?) As a subset of (2), how or where should threshold values be set (exceedance of which would signal a need for increased regulatory attention)?

The industry example discussed above argues implicitly that direct measurement of safety function performance (no complete safety function failures) provides adequate assurance to the regulators, and that the integrating tool is really just the conjunction of a set of safety-function indicators. With reference to insights from the precursor program, it is asked in this project whether this kind of indicator is sufficiently “leading” for regulatory purposes. A “yes” answer would mean that we would expect indications at the safety function level to highlight developing problems before public safety was threatened. A “yes” answer would require statistics on failure events at an individual plant to be fairly unambiguous indicators, and would require things like common cause failures to be identified with high probability before a complete accident sequence occurred. It is not within the scope of the present report to pronounce definitively on the shortcomings of such an end-result-oriented approach, but there appear to be some shortcomings, at least in the example proposed by industry, and it is unreasonable at this stage to focus on industry's approach alone.

A major question is whether we need to try to measure institutional factors directly (rather than indirectly, through safety function performance). If we consider that institutional factors can affect many basic events at once, then institutional factors can be seen as a kind of CCF, which is to say among other things that institutional factors are extremely important. Again, a question to be answered about CCF in general, and institutional factors in particular, is whether they show up adequately in end-results-oriented indicators.

A.3 Performance Figures of Merit

We turn now to papers whose main emphasis tries to shed light on the formulation of performance goals that might serve present purposes.

A.3.1 Performance-Indicator and Precursor Literature

Example: Goals from IPE Parameters

We begin with excerpts from a paper by Burns and Turcotte [7][*]. This paper presents a derivation of performance criteria that is about as straightforward as one can get, and runs into basic issues immediately. Key points from the paper:

> "Acceptable" is defined as performance that is consistent with IPE component failure rates.
>
> Recommendation for goal-setting: given a point value for [failure] rate, assuming this is the "true" rate, calculate the probability that n or more failures will occur, and choose n having a reasonably low (10–20 %) chance of occurring. "Thus, if we see n or more failures, then the component failure rate(s) are significantly higher than expected, and additional measures (goal setting, etc.) are warranted."
>
> However: "Since isolated failures are expected to occur, we believe that n=1 is not an appropriate performance criteria [sic] for any system or subsystem. Thus, we set n=2 as the performance criteria for systems or subsystems that included only standby components. In our review…we found that every system and subsystem included some standby components, hence, 2 is the lowest performance criteria [sic] applied to any system or subsystem. […apparently over a three-year period…]"

They conclude that the criterion should be

PC=S+N+I,

[*] A paper by Hunt and Modarres, cited as reference 1 in chapter 1 in Modarres's book [8], is mentioned as the origin of the diamond tree, but the citation appears to be in error; this paper does not exist in the location given. It is possible that the citation was intended to refer to [9] or [10], but they do not address the diamond tree by name. [9] all but does so, describing what we today call the bottom half of the diamond tree as "a structure somewhat like an inverted pyramid with multiple top events, each of which represents the performance goal for a specific piece of hardware…." The first public mention of the diamond tree by that name that we can identify was in a short course taught by Hunt (1987).

where PC is performance criterion, S=standby failure allowance (always 2), N=normally operating equipment train allowance (1 for each train of the system/subsystem {i.e., operating systems are allowed more failures than standby systems}), I is instrumentation allowance (0, 1, or 2; higher number for more instrumentation stuff, like RPS).

This work is interesting because it sets out to establish goals directly from IPE-credited numbers, but quickly encounters what we might call the “fluke occurrence” issue and, in order to avoid a high false alarm rate, reflects an arguably very high random-failure allowance in the criterion. For some programmatic purposes, this is appropriate. On the other hand, most of the rest of the work mentioned in this section is aimed at addressing issues in performance assessment that this approach does not address.

Performance Indicators

A very large amount of NRC-sponsored work—too much to be comprehensively mentioned here—has been done in this area under the rubric of “performance indicators.” The point of most of this work has been to learn how to gain insight into whether plants are maintaining safety performance, and not particularly to try to redefine the relationship between regulators and licensees. Nevertheless, insight into what indicators actually correlate with safety is of fundamental interest to this project. It is interesting and significant that the work mentioned below proceeds immediately beyond the kind of failure-counting described in the previous paper, and does not seriously consider functional failures such as RCS overheating.

One multi-year NRC-sponsored effort, led mostly by Azarm and Vesely [8–12], assessed various system and train-level indicators. They considered different unavailability indicators at the system level and at the train level. They also addressed common cause to some extent. They favor a “Safety System Function Trend” indicator. A major aspect of this is that it does not suffice to count failures: downtime is important, among other things. The work, documented both in published NUREG/CR reports and in internal BNL reports, is expected to be quite useful in the process of formulating performance measures for actual use. (This is not to say that their conclusions are automatically accepted here, but rather that they addressed some of today’s important questions, and furnished a considerable basis for the answer that they obtained.) Their most recent summary report also provides a useful statement of the regulatory basis for the work, including excerpts from NRC memoranda that articulate a need to identify important symptoms, trends, or precursors before they lead to precursor events. Their work was not explicitly motivated by performance-based regulation as we currently use that phrase, but simply by a perceived need to monitor plant safety performance and compare it against expectations to help identify potential problems early on. A particularly interesting quote observes that “…a problem illustrated by the Davis-Besse event [of 1985] is the degradation of system and component reliabilities below those we have been typically calculating in estimating core-melt frequencies….”

Another performance-indicator effort, led by Wreathall [13], appears to have raised and addressed an even broader spectrum of issues. This effort was documented in draft unpublished NUREG/CR reports. An indication of the breadth of issues addressed in these reports is provided by the types of indicators considered:

1. direct indicators, which measure plant safety performance (such as the kind of indicator addressed by Azarm et al.).
2. programmatic indicators, which measure the effectiveness of plant programs (maintenance, training,…) that "provide leading indications of trends in safety that would be confirmed by changes in direct indicators"
3. organizational indicators, which represent "factors of overall behavior, such as organizational learning, allocation of resources, and corporate experience; these have been found to be associated with safety performance…"

The use of the term "leading" in bullet 2, together with the diversity of indicators considered, signals clearly that the report at least tried to consider many of the issues raised in the introductory sections of the present report.

A more narrowly-focused look at certain issues was provided by Wreathall and Appignani [14], who looked at maintenance effectiveness. Some excerpts from their paper follow:

> [Previously-proposed] measures of maintenance…were subjected to a screening evaluation…two measures [were identified] as consistently correlating with other NRC safety performance measures such as the frequencies of scrams and the frequencies of significant events. These two measures were the frequencies of inadvertent emergency safeguards features (ESF) actuations from test and maintenance errors, and gross heat rate (a measure associated with plant thermal-hydraulic performance).
>
> Inadvertent ESF
>
> [plants with a high rate of this were also those whose reported corrective actions most frequently mentioned "discipline or counseling."] This kind of reliance on punishing the final human link in the chain is a characteristic of what Westrum has called "pathogenic" organizations—those that deny they have problems, and avoid the opportunity to improve the standards of safety performance…. In engineering terms, this is the equivalent to not searching for root causes…. Pathogenic organizations…are those that present the greatest hazards because of their denial of problems, and therefore root cause mechanisms are commonly left in place…

<u>Daily Power Level (DPL)</u>

> DPL is simply the average net power generated in each 24-hour period, plotted through time…[the paper does not explicitly argue the significance of this measure, but offers a discussion pointing out that short-term wild fluctuations versus occasional long outages would have different implications, and that a DPL plot would show things that an integrated availability would not…the measure is still being evaluated…].

This work is interesting not only because it correlates quantitative and measurable things with each other, but also because it correlates them with institutional factors.

A.3.2 Performance-Based Approaches to Evaluation of Non-Reactor Facilities

A.3.2.1 DOE

DOE currently uses 23 ES&H performance indicators organized into four major categories, as follows:

Accidents/Events that have already happened:

1. Lost Workday Case Rate
2. Occupational Safety and Health Cost Index
3. Electrical Safety
4. Industrial Operations Safety
5. Transportation Safety
6. Chemical Hazard Events
7. Reportable Occurrences of Releases to the Environment
8. Cited Environmental Violations
9. Environmental Permit Accedences
10. Radiation Dose to the Public
11. Worker Radiation Dose
12. Radiological Events

Precursors to accidents and near misses

13. Near Misses and Safety Concerns
14. Inadequate Procedures/Procedures Not Followed
15. Safety System Actuations
16. Safety Equipment Degradation

ES&H Management

17. Environmental Compliance Milestones Met
18. Open DNFSB Recommendations

Hazards level of material at risk

19. Spent Nuclear Fuel and Plutonium Vulnerabilities Resolved
20. Plutonium Stabilization
21. Toxic Chemical Releases
22. Pollution Prevention
23. HEU Vulnerabilities

Comments

The DOE Performance Indicators are focused on measuring DOE's overall ES&H performance and do not focus on the ES&H performance of individual facilities [15,16]. Also these indicators include occupational safety and health concerns and compliance issues, and do not strictly focus on radiological release accident prevention. DOE has plans to improve their performance indicators by developing risk-based approaches to performance indicator selection, prioritization, and analysis, and by developing performance indicators that indicate how well DOE is doing in reducing hazards or vulnerabilities and in safety management. DOE Order 210.1 establishes DOE's requirements regarding performance indicators, and mandates the use of ES&H performance indicators to improve the performance of DOE facilities, programs, and organizations.

DOE Operational Readiness Reviews

DOE requires that operational readiness reviews (ORRs) or readiness assessments be performed prior to the restart of an existing nuclear facility or the startup of a new nuclear facility. DOE Order 425.1 [17] and DOE Standard 3006 [28] provide the requirements and guidance related to ORRs. DOE Order 425.1 establishes a number of core requirements that must be met to assess whether a facility can be operated safely. DOE Standard 3006 subdivides these core requirements into core objectives and discusses how to develop Criteria Review and Approach Documents (CRADs) to provide a basis for evaluating the core requirements. DOE Standard 3006 also provides examples of management oversight and risk trees (MORT) to define when a facility is ready for operation. The primary example divides the facility into plant and hardware, administrative programs, and personnel and training.

DOE Performance Objectives and Criteria

This document [19] contains ES&H Performance Objectives and Criteria (POCs) designed to assist in the evaluation of DOE Safety Management Systems. This document is intended to be used as a tool in the planning and conduct of ES&H oversight activities. The POCs are organized around the DOE "Safety Management Template" and the Department's three guiding principles for safety:

- Line managers are responsible and accountable for safety
- Comprehensive requirements exist, are appropriate, and are effectively implemented
- Competence exists commensurate with responsibility

POCs derive from a distillation of the full range of requirements and good management practices applicable to a particular Criterion and encompass the key programmatic concerns. As they are not overly prescriptive, line management has reasonable flexibility in adapting programs and processes to the facility's unique mission, hazards, and life stage. For each Criterion, a discrete number of concisely stated performance objectives are listed. Listed directly below the performance objective is a set of criteria by which that objective is to be evaluated. Preceding each set of POCs is a short description of the scope of the technical program or system and a listing of requirements governing its development and implementation, including regulations, DOE Orders, and industry standards.

The handbook on measuring performance [20] provides guidance on developing performance measurement systems. The guidance is written generically and is thus applicable to a wide variety of processes. The handbook offers three disciplined, systematic approaches for measuring performance:

- The first approach, the Performance Measurement Process, was developed by the DOE Nevada Family Quality Forum. This approach is quite detailed and outlines an 11-step process for measuring performance.

- The second approach, Developing Performance Indicators . . . A Systematic Approach, was used at Sandia National Laboratories. It is less detail-oriented than the first, and uses a fictitious company, the Hackenstack Firewood Company, for anecdotal purposes.
- The third approach, Developing Performance Metrics-the University of California Approach, was developed by the University of California. This method is broadest in scope.

Comments on DOE Performance Assessment

A point frequently made about the DOE is that it is both manager and regulator of its facilities. Performance assessment therefore means more to the DOE than just satisfaction of a public-safety objective or even a worker-safety objective; performance in DOE space also bears on mission and resource issues, in fact all of the issues that a company management faces in general. Comparisons between DOE "performance-based" approaches and those taken by other agencies must be weighed in light of this characteristic of the DOE mission. In light of this, it appears that while DOE may be making agency-wide use of indicators 1-16 in the list furnished above, its implementation of performance-based safety regulation of individual DOE facilities does not appear to have advanced to where it can meaningfully inform the present effort.

A.3.2.2 EPA

The EPA is about to promulgate its long-awaited Compliance Assurance Monitoring (CAM) regulation [21, 22, 23]. This rule will require industry to identify plant performance indicators that will provide a reasonable assurance of compliance with existing limitations and standards; and then to monitor plant emission controls based on these performance based indicators. Essentially the performance indicators are *assigned by the facility owner* such that he can reasonably assure that as long as he is within the performance monitor range (meeting the indicator), he is meeting the air emission standards. In the event the performance standard is not being met, the operator takes self-directed corrective action.

The basic monitoring approach of the CAM rule is to monitor the performance of pollution control technology instead of directly monitoring actual emissions. The CAM rule would require owners or operators of emission sources to increase performance awareness of the operational status of pollution control technology and to act on discrepancies in that operation with the net result of reducing emissions.

The history of the EPA Compliance Assurance Monitoring (CAM) rule follows President Clinton's public commitment to simplifying, streamlining, and reducing the costs of the Federal Government. EPA followed the administration's lead with a commitment to "simplify, streamline, and reduce the costs of Federal EPA programs to better protect the environment and public health by applying a common-sense approach to smarter environmental and public health protection," especially to regulated entities. This rule has been in development for over 6 years, and has seen considerable public review.

EPA is focusing on new “Results-Based Tools” that they believe, in the long term, will shift their focus, and the focus of regulated industries, away from meeting narrowly defined regulatory requirements and toward the achievement of environmental performance-based results. As the EPA puts it “…. thinking about results unleashes innovation and helps the public and private sector find new solutions to old problems.”

This proposed rule will introduce additional flexibility into the EPA’s compliance-assurance monitoring program through user self-defined programs and focuses on preventing pollution rather than imposing additional command-and-control regulations. No doubt, this is a significant departure in Agency direction for implementation of the monitoring and compliance certification requirements in titles V and VII of the 1990 Clean Air Act Amendments. The EPA states that “…The goal of the action is to provide reasonable assurance of compliance rather than a direct connection between monitoring and certification and will reduce the emphasis on assuring compliance through the threat of enforcement. Instead, this approach emphasizes assuring compliance by placing the burden on regulated sources to monitor their performance and take proactive steps to minimize emission accedences.”

Corrective actions in response to accedence of monitored performance-indicator “trigger limits” require problem investigation and perhaps a Quality Improvement Plan (QIP). The analog with the NRC Maintenance Rule (10 CFR Section 50.65) is clear: CAM trigger values are specified to indirectly monitor the performance of pollution control equipment, corresponding to maintenance rule determination of risk and performance criteria and goals; CAM monitoring against the trigger performance-indicators is like maintenance rule performance monitoring against the performance criteria; e.g., MPFF, unplanned reactor scrams, availability/reliability goals, etc.; if CAM triggers are exceeded, a corrective action program may result, while unacceptable performance with respect to maintenance rule criteria occasions a cause determination followed by goal setting and disposition from performance category (a)(2) to (a)(1)).

A.3.2.3 OSHA

Major industrial companies around the world have redoubled their commitment to assure the safe and economical operation of their facilities. In fact, a proactive approach to accident prevention, process safety management, and emergency management is fast becoming required “good business practice.” The concern that industrial facilities be operated safely and effectively is also reflected in Federal guidelines such as the Occupational Safety and Health Administration’s (OSHA) management of highly hazardous chemicals (OSHA 20 CFR 1910.119), and in the Environmental Protection Agency’s (EPA) Risk Management Programs (RMP) under the Clean Air Act (EPA 40 CFR 68).

On May 26, 1992 the OSHA standard [24], covering the Process Safety Management of Highly Hazardous Chemicals (29 CFR 1910.119), became law. The primary objective of the PSM OSHA standard was to prevent unwanted releases of hazardous chemicals, especially to locations that would expose employees and others to serious hazards. To ensure that the safety of both plant and contractor employees is considered, the PSM Standard clarifies the responsibilities of employees and contractors involved in work affecting processes covered by the regulation. The standard mandates 14 separate PSM management control elements that specify employee participation in PSM programs, preparation of process safety information, process hazards analysis evaluations, written operating procedures, employee and contractor training, pre-startup safety reviews, maintenance of the mechanical integrity of critical equipment, written procedures for managing change, evaluation of incidents and near misses, emergency action plans, compliance audits, and trade secret protection.

Some would say that the OSHA PSM regulations are performance-based standards for the industrial community. (We would tend not to.) The PSM regulation establishes minimum requirements for each covered facility that define "what" is to be done for the various elements of a Process Safety Management program for process safety control, without specifying in too much detail the implementing plans, programs, and procedures to get there. It is perhaps OSHA's intention to allow industry considerable flexibility in developing "how to" procedures to achieve the requirements stated in the PSM Standard 29 CFR 1910.119. Indeed, given the enormous variety of facilities that come under this rule, it would be remarkable if a completely prescriptive requirement were considered feasible. Even as it is, typical or suggested features of the 14 OSHA elements and implementation methodology are provided.

Perhaps the center of the OSHA standard is the element covering Process Hazards Analysis (PHA). The process hazards analysis involves a systematic review of what can go wrong and what safeguards are in place or need to be in place to prevent or mitigate a hazardous chemical release. The intent here is to put in the hands of the chemical industry (actually require the use of PHA techniques) risk assessment tools to help them make risk informed decisions on plant process safety and achieve performance-based goals of safely operating processes. The OSHA defines acceptable PHA methodologies to include one or more combinations of 6 analysis techniques: What-if, Checklist, What-if/Checklist, Hazard and Operability (HAZOP) studies, Failure Modes and Effects Analysis (FMEA), Fault Tree Analysis, or an appropriate equivalent methodology. OSHA's audit guidance for this area is actually quite explicit [25].

The HAZOP and What-If process are widely practiced in the chemical, petrochemical, and oil & gas and manufacturing industries and are acclaimed for the risk-informed process safety insights that are possible with these PHA methods. Most frequently, PHA insights are based on a team review and represent non-numerical qualitative assessments of hazards and their risk, allowing a user defined calibration of relative hazards and their importance at the plant. The PHA reviews provide a structured review of the plant considering: process design, process technology, maintenance activities and procedures, conduct of operations, and non-routine activities and procedures.

It is also interesting to note that the PSM standard requires the evaluation of accident and near miss scenarios involving hazardous materials. The post analysis of accidents can serve as an assessment displaced in time of process safety flaws, while the analysis of near miss scenarios can be interpreted as a process safety leading indicator of precursors to more serious events.

Based on the above discussion, this regulatory approach is not a pure form of what NEI would mean by "performance-based regulation." In this case, it is performance of PSM that is monitored by the regulator, not plant safety functions. However, it is still of considerable interest, first because it gets into institutional factors, and second because the regulatory position appears to have a great deal in common with guidance to the "expert panels" that play an inevitably significant role in risk-informed applications. This includes assessment of near-misses. In fact, this rule basically tells the industry to do what we call integrated decision-making, and audits the process. It corresponds in this sense to the Maintenance Rule.

A.4 Tools for Integrated Consideration of Performance Figures of Merit

The theme of this section is the process of making findings to support regulatory decisions, based on the potentially large number of performance parameters to be considered.

Some years ago, it was assumed in most discussions that a really good risk analysis is all the integrating tool that one needs. Today, risk analysis is seen as a major input, but not the only input, to something called “integrated decision-making.” Expert panels are seen as playing an important role in this process. Why is this necessary? For one thing, because of the very way in which risk analysis models are constructed, it is difficult to build into them all of the necessary considerations. The very term “analysis” connotes a logical decomposition of the problem into smaller elements, and depending on how this is done, it is easier or harder to reflect particular considerations in the modeling. Fault tree analysis is an example of this. As usually employed, it provides a picture of logical relationships between events (in particular, how functional failures follow from combinations of more “basic” events). This has proven to be an extremely important tool for understanding not only the functional properties of plant designs, but also possible scenarios; but the psychological processes involved in development of fault trees does not directly elicit consideration of broad influences such as institutional factors in the same way that it elicits information about the topology and physical connectivity of plant systems. Given a good fault tree, one can graft onto it, in a qualitative way, influences like “plant culture,” or “housekeeping (combustibles),” provided that one can codify this in logical rules. But some of this is not naturally described in logic expressions, and some of it must be thought of ad hoc. Indeed, some of the work mentioned in the last section and in the present section is aimed not only at working with important figures of merit but also at identifying them in the first place.

This section, then, briefly discusses a spectrum of tools that in one way or another link basic performance figures of merit to more conclusionary top-level outputs. We will begin with tools that are equivalent to today’s risk models, and work up from there.

“Simple” integrating tools are things like on-line plant models, precursor-like analyses of plant risk, and train or component-level models that relate trended parameters to component status. These tools are really just plant models that indicate a current estimate of CDF or of some plant functional status parameter; they only indirectly support decisions, in the sense that any model can support some kind of decision. A step up from these models is “integrated decision-making,” whose very name suggests application of decision analysis tools. Such tools can directly suggest findings, such as weaknesses in particular areas as revealed by a conjunction of indicator inputs, or by a conjunction of trends, etc. Finally, three less generic tools are discussed briefly: fuzzy logic, abductive inference, and “integrated approach methodology.” Each of these is aimed at supporting conclusions based on integrating information in the form of a collection of parameters, but in a way that is to some extent non-standard (compared to decision theory). Some of the work discussed previously under “performance indicators” also has an integrating quality to it, and it could have been placed in this section as well.

A.4.1 Simple Integrating Tools

A.4.1.1 Precursor Analysis

It is not the purpose of this subsection either to repeat or to alter the discussion provided earlier of precursor analysis, or to add to remarks regarding the drawbacks of trailing indicators, etc. Rather, it is the purpose here to note a body of more sophisticated statistical analysis of precursor events aimed at getting insights on inter-system dependences that are not modeled. Conventional precursor analysis analyzes a particular event by taking a conventional risk model, setting the appropriate failed events to "true," and calculating the modeled conditional probability of the top event. This provides a measure of how "near" the "near miss" was. This is one way to learn from precursor events, but not the only way. The observed rate at which precursors occur also has implications for the plant model and for performance-based insights, and relatively little work is done along these lines, as compared to the usual conditional-probability exercise. An exception is a body of work by Prof. Bier; a recent instance is "Models for the Analysis of Precursor Data that Preserve Inequality Relationships," by Yi and Bier [26].

The practical application of this work to performance-based regulation is not firmly established in the present author's mind, but it seems worthwhile to note in a literature survey that there is literature trying to extract potentially important insights from performance data in this particular way. Some of the performance-indicator work mentioned previously has some elements in common with this work, but most of the PI work is essentially applying models, while this work has aspects of rising above them.

A.4.1.2 Plant Diagnostic Tools

There is also a significant body of work that seeks to draw conclusions about the status of plant components based on monitoring of a few key physical parameters. That is, one monitors component performance in particular ways. Given the hard line taken here on the meaning of "performance-based" regulation, the reason to cite this work is not to advocate monitoring of component status as part of a performance-based regulatory scheme. Rather, it is simply to note another in the seemingly endless variety of tools that can be used to establish relationships between one set of variables and another set of variables, even if the actual functional relationship is very complicated.

An example that turned up in this search is by Grenzebach and Marx [27]. Excerpts follow:

> Can the objective of assessing the viability of the system at any given point be met by using a small number of variables (four indicators of Component Cooling Water (CCW) temperature) to predict the behavior of a far larger number of variables (forty-four measures of the Reactor Coolant Pump RCP) [sic][system)? The statistical procedure of canonical correlation analysis can help answer this question.

This may be germane. "Canonical correlation may be viewed as an extension of regression analysis." The attempt was to predict RCP variables in terms of CCW variables, and they claim success. The possible relevance of it is the formalism relating one set of performance variables to another. There is a plethora of work doing somewhat similar things with neural networks. A possible point of this kind of thing for present purposes is that there seems to be an abundance of literature around suggesting ways to achieve a mapping between a collection of input performance variables and one or more integrated performance variables.

This is also the right place to note that ASME and the staff have been working together to propose an improved test regimen for IST components, one that registers key aspects of physical performance instead of just verifying operability on a pass/fail basis. If careful analysis supports a claim that this improved regimen can compensate for extended test intervals, then this kind of monitoring is an essential element of the performance feedback process that is considered essential to risk-informed applications. As mentioned previously, however, from the present point of view, this is not an aspect of "performance-based regulation."

A.4.2 Applicable Ideas from Decision Theory

A likely outcome of the analysis of different performance measures is that each of several distinct combinations of performance indicators will appear more or less to span the necessary set of technical issues, and would each give about the same picture of institutional performance. At some point, it would be necessary to pick one scheme or another: to make a decision how to proceed. By "scheme," we mean a set of indicators and an integrating tool that accepts them as input and furnishes some kind of recommendation as an output.

The schemes would differ in:

- scope of plant covered in the scheme
- cost to licensee (collecting, certifying, submitting information)
- cost to regulator (processing, reviewing, integrating the information)
- degree of invasiveness (degree to which regulator crosses the line into the plant management function, for example in assessing programmatic efficacy)
- level of uncertainty associated with indications of satisfactory performance (false negative rate)
- level of uncertainty associated with indications of unsatisfactory performance (false positive rate)
- timeliness of status updates
- degree of qualitative insight expressed in diagnostic output (statement that the plant state is unsatisfactory vs. statement of how unsatisfactory the plant state is, vs. in what way the plant state is unsatisfactory, etc.)

One use of decision theory, then, would be to choose among the possible schemes, weighing the very different kinds of pluses and minuses that appear to operate.

At another level, some tools of decision theory have been used in situations that might appear to map onto decision problems, but are in part just modeling problems that need to be solved in order to support a decision. Influence diagrams are an example of this. They have been used to chart the influences of physical variables on other physical variables in various phenomenological modeling areas. Influence diagrams, or the moral equivalent, suggest themselves for use in formulating the integrating tool.

A.4.3 Less-Familiar Modes of Inference

If we stick entirely to hardware performance indicators, we can imagine propagating them through a risk model to arrive at an assessment of the plant state. However, if we choose to work with institutional factors, we are dealing with items that we do not know how to couple directly into a risk quantification. We arguably have enough knowledge and experience to reason in defensible ways based in part on these indicators; the problem is simply that we do not have a simple functional form that maps them into CDF. The following two subsections are examples of areas of inquiry that are aimed partly at helping us reason defensibly in situations like this.

A.4.3.1 Neural Networks, Fuzzy Reasoning

There is a vast literature on neural networks. Time has not been spent here to review it.

Previous work has established convincingly that neural networks can capture complex and subtle functional relationships between variables, such that given a suitable information base of inputs and corresponding outputs, a neural network can be trained to respond to particular combinations of inputs with adequate estimates of combinations of outputs. An appeal of this kind of thing is that one can imagine providing a neural network with a data base comprising combinations of inputs including assessments of institutional factors, and training the network to respond with an assessment of the plant condition. One can thus end up with a functionally useful sort of “correlation,” as T/H people use that term, even lacking an explicit underlying model of how the various factors combine causally to yield the output. The catch, of course, is that one needs a number of valid training examples: examples in which one knows the right answer for each of many combinations of inputs. It is therefore possible, but challenging, to develop a tool like this. One would need to proceed by taking a plant, working up something like the “integrated approach methodology” discussed below (integrating that plant’s risk model with institutional factors), and training a network using plant history as input and plant SALP rating as output (note: this is intended more as a metaphorical description of a project, not literally a proposal at this stage.)

Another kind of development is also possible. Work has been done based on the idea that one should try to capture some of the broader influences on risk in a model that, in some ways, resembles a neural net, but rather than being trained from scratch to reproduce data, is more of a systematic organization of the user’s own thoughts, fuzzy though they may be. The following discussion is based on *Fuzzy Sets, Natural Language Computations, and Risk Analysis*, by Kurt J. Schmucker [28].

This book would be expected to bear on the present work because “risk analysis” appears prominently in its title: it is more directly about risk analysis than are many works on fuzzy logic. The present section will summarize a few of its thoughts.

The general flow of the book is: review of set theory, presentation of fuzzy set theory, presentation of natural language computation, and presentation of the “fuzzy risk analyzer.” The following excerpts point to the book’s basic ideas:

> The traditional approaches to risk analysis are based on the premise that probability theory provides the necessary and sufficient tools for dealing with the uncertainty and imprecision which underlie the concept of risk in decision analysis.
>
> The theory of fuzzy sets calls into question the validity of this premise. More specifically, it suggests that much of the uncertainty which is intrinsic in risk analysis is rooted in the fuzziness of the information which is resident in the database, and, more particularly, in the fuzziness of the underlying probabilities. Viewed in this perspective, then, it is the failure of classical probability theory to come to grips with the issue of fuzziness of data that limits its effectiveness in dealing with a wide variety of problem areas—including risk analysis—in which some of the principal sources of uncertainty are nonstatistical in nature.
>
> (The above is from the foreword by L. A. Zadeh.)
>
> The intellectual task of analyzing the risk present in any large undertaking is an endeavor that abounds both with inherent imprecision and with a scarcity of historical data. Traditional mathematical and computational methods offer little to aid the analyst in work beset with either of these two difficulties, let alone work that is plagued by both of them. This is because the basic philosophical system upon which our mathematics and computer science is based is discrete and adheres strictly to the principle of the excluded middle: a statement must either be true or false. Unfortunately, this is rarely the case in risk analysis.
>
> Fortunately, there is an alternative to this philosophy. This alternative, fuzzy set theory, is aimed at the development of tools for the solution of problems too complex or too ill-defined to be susceptible to analysis by conventional methods.

The views expressed in this book regarding statistical quantification of likelihoods have some overlap with those expressed in the book on abductive inference (see below). In Schmucker’s book, there is the added implication that we somehow need to redress the conceptual errors made in the name of the excluded middle. This is aimed in part at problems like the data-pooling problem alluded to in the book on abductive inference; the implication is that particular historical events in the data base bear *to some degree* on the event whose likelihood we are trying to quantify today, rather than being either applicable or inapplicable.

What the book means by “fuzzy risk analyzer” turns out to be a relatively simple thing and one that does not deal convincingly with the kind of engineering model to which we are accustomed in risk analysis. The job of breaking down a top-level index (e.g., core damage) into lower-level indices (e.g., system unavailabilities or train unavailabilities or…) is not much discussed. The tool discussed in this book is natural language quantification: it works with phrases like “some, very little, a lot” instead of numbers, and propagates these expressions through a hierarchy loosely analogous to a logic tree to get a high-level estimate of overall risk. Much of the book discusses a particular software implementation of this called the “fuzzy risk analyzer.”

Comment

This work does not claim or appear to have a magic bullet for completeness uncertainty. (It would be remarkable if it did.) The problem that this work sees itself as addressing is that the data that we propagate through the usual models are inapplicable or overly precise, and perhaps that we combine these things using arithmetic that may not adequately capture our thought processes. Rebuttal papers have been written by Bayesians who say about fuzzy logic essentially that those elements of it that are new are wrong, and that the elements of it that are not wrong are not new. In fact, Bayesians would take exception to the claims that there are no extant tools for dealing with “a scarcity of historical data.” Bayesians would argue strongly that Bayesian analysis does precisely this. The fuzzy proponents, though, appear to mean in part that the applicability of the data to the problem at hand is fuzzy at best, and the Bayesian coping mechanism for this probably depends on something like expert elicitation.

As an integrating tool, the fuzzy risk analyzer presented in the book seemed to be a disappointment. However, the problems in data analysis that were mentioned above are real, and interestingly, this book’s take on them is not inconsistent with that of the abductive reasoning group. While the Bayesians appear to control the mainstream of current thought in this area, it has to be admitted that the best-known examples of Bayesian treatments seem to depend on the pertinence of particular items of evidence, and are carried out for situations in which the pertinent variables are adequately controlled. In short, where applicability is murky, the Bayesians appear to have left some problems unsolved. It is an interesting and open question whether fuzzy logic might help us to think more clearly about how to use marginally applicable data. (For that matter, it would be of interest to learn that hitherto-unnoticed Bayesian treatments could do the same thing, less controversially.) This issue would affect performance-based regulation in at least two ways. One way is to show how to use marginally relevant evidence in characterizing the likelihood of events in a calculation, and another way (perhaps) is to show how one might use institutional factors in assessing the current level of an institution’s safety.

A.4.3.2 Abductive Inference: Inference to the Best Explanation

The discussion in this subsection is based on *Abductive Inference,* edited by Josephson and Josephson [29].

A lot of work has been done in recent years in “intelligent” diagnostic systems and in medical informatics. The idea is to help medical professionals make diagnoses, based on information about the symptoms presented by a patient. It is beyond the scope of the present work to try to judge the success of these developments, but an outsider’s sense of the area is that progress is being made. It is worth comparing the ideas involved in this area to the ideas that come into play in performance-based regulation.

Simply put, the problem in medical diagnosis is to figure out what, if anything, is wrong with a patient, based on the symptoms presented. These symptoms are continuous-valued, perhaps judgmental-assessed or qualitatively characterized parameters (some, a little, lots)$*$(fever, spots, pain, hormone levels,…). It is an option to perform various more or less invasive, more or less expensive medical tests, if the readily-available information appears to warrant this. Serious work has been done on intelligent diagnostic systems, whose purpose is to formulate a best diagnosis, or perhaps to formulate several possible diagnoses and rank them according to likelihood of being correct.

According to some workers, the logical processing involved in diagnosis is “abductive” reasoning, which is contrasted with deductive and inductive. Deductive reasoning is exemplified by syllogisms, such as “all men are mortal, Socrates is a man, therefore Socrates is mortal.” Inductive reasoning draws general conclusions from an examination of special cases; in formal mathematical induction, one formulates a proposition that depends on n, and shows that (a) it is true for n=1, (b) its truth for n=m implies its truth for n=m+1; then one can conclude that it is true for all n>=1. Some would say that examination of a large number of individual life histories leads to the conclusion that all men are mortal, and that this, too, is induction. Abduction is more like “people who have measles generally have fever and always have spots; Joe has fever and spots; it’s likely that Joe has the measles.” Two normal reactions to this cartoonish example are that real diagnosis is much more complex, and that this kind of reasoning is not ironclad. It appears that progress has been made on much less cartoonish examples, and that the uncertainty is an essential feature of the problem and something that the reasoning algorithm must deal with.

The following excerpt is from Josephson’s characterization.

> Abduction, or inference to the best explanation, is a form of inference that goes from data describing something to a hypothesis that best explains or accounts for the data. Thus abduction is a kind of theory-forming or interpretive inference….
>
> We take abduction to be a distinctive kind of inference that follows this pattern pretty nearly:
>
> > D is a collection of data (facts, observations, givens).
> >
> > H explains D (would, if true, explain D).
> >
> > <u>No other hypothesis can explain D as well as H does.</u>
> >
> > Therefore, H is probably true.

Let us compare the problem of understanding the condition of a plant [in the institutional sense, including its staff] with the problem of medical diagnosis. In the medical case, we have history, and we have results from things like blood tests, a plethora of imaging techniques, some knowledge of various risk factors, etc. We also have a lot of semi-phenomenological models of how certain ailments cause certain problems. In the NPP case, we have information in the form of recent operational history, plus some body of knowledge based on interviews with staff, perhaps an assessment of organizational factors, financial information, knowledge of recent employee histories, etc. And of course we have a plant-level risk model, many engineering assessments, and so on. Based on an assessment of all this information, one could imagine concluding that the plant is essentially normal, or alternatively that there are weaknesses in the management of maintenance that explain certain unavailability trends and certain maintenance outage durations, or alternatively that because of the significance of a recent precursor event (analogous to a fainting spell), more detailed examination of the institution’s state is warranted. Described in this way, assessment of a plant somewhat resembles medical diagnosis. In both the medical context and the plant context, we wish to formulate the best possible H to explain the observations.

It is implicit in many discussions of performance-based regulation that we are trying to synthesize an overall figure of merit (call it “safety”) and, if this is found wanting for some particular plant, the regulators will bear down harder in an effort to bring about improvement in “safety.” If we treated the medical diagnosis in a completely analogous way, the idea would be to quantify the “health” of the patient, and try to make the patient “healthier” if necessary. But in medical practice, the idea seems closer to deciding not just whether the patient is healthy, but also what ailment(s) the patient currently has. If we consider that interpretive inference, and not just a Delphic assessment of “degree of safety,” is an objective of the performance assessment, then it seems that work in abductive inference is potentially germane to performance-based regulation. Symptom-based medical intervention is arguably analogous to performance-based regulatory intervention.

Based on the rather simplistic discussion furnished above, one can ask whether inferential problems of this kind can be mapped onto multidimensional regression analysis, and if so, whether there is anything really new about this. Based on the present author’s impressions, the answer appears to be that the justification for work in the field is not to provide new verbiage to describe traditional data-fitting, but rather to develop an equivalent to the way in which real people think when they are trying to figure out what’s going on in real problems, and to develop tools for working on problems that are sufficiently complex that iterative synthesis and evaluation and comparison of complex hypotheses (as a human would do) is essential to the solution of the problem. In other words, in abductive inference, there is some evaluation of the explanatory power of competing hypotheses, and an implicit rejection of a potentially large number of candidate explanatory hypotheses that are formulable within the space of considerations being addressed. Some narrowly-defined class of problems of this kind may be reducible to parameter-fitting, but that class of problems is not the reason for the existence of the field.

As an interesting aside, Josephson’s comment on Bayesian analysis seems germane to the present inquiry.

> It has been suggested that we should use mathematical probabilities to help us choose among explanatory hypotheses. (Bayes’s Theorem itself can be viewed as a way of describing how simple alternative causal hypotheses can be weighed.) If suitable knowledge of probabilities is available, the mathematical theory of probabilities can, in principle, guide our abductive evaluation of explanatory hypotheses to determine which is best. However, in practice it seems that rough qualitative confidence levels on the hypotheses are enough to support abductions, which then produce rough qualitative confidence levels for their conclusions. It is certainly possible to model these confidences as numbers from a continuum, and on rare occasions one can actually get knowledge of numerical confidences (e.g., for playing black-jack). However, for the most part numerical confidence estimates are unavailable and unnecessary for reasoning. People are good abductive reasoners without close estimates of confidence. In fact it can be argued that, if confidences need to be estimated closely, then it must be that the best hypothesis is not much better than the next best, in which case no conclusion can be confidently drawn because the confidence of an abductive conclusion depends on how decisively the best

> explanation surpasses the alternatives. Thus it seems that confident abductions are possible only if confidences for hypotheses do not need to be estimated closely.

The discussion goes on to note a problem that is very significant for many applications: that if data from a broad cross-section of instances are pooled, the applicability of the pool to an individual case at hand is compromised, while if one focuses only on cases that are closely comparable to the individual case at hand, the statistics suffer. Thus,

> There is a Heisenberg-like uncertainty about the whole thing; the closer you try to measure the likelihoods, the more approximate the numbers become. In the complex natural world the long-run statistics are often overwhelmed by the short-term trends, which render the notion of a precise prior probability of an event inapplicable to most common affairs.

A.4.4 "Integrated Approach Methodology"

This discussion is based on a report, "Integrated Approach Methodology," prepared on DOE funding for Sandia [30]. The report discusses development of a very comprehensive representation of a plant together with its staff and management. The abstract, in its entirety, is this:

> A top-down Integrated Approach model is used to rigorously and completely describe a nuclear power plant and its operation. This model provides a framework within which all decisions concerning plant activities can be made. Plant Information System Improvement activity is used as an example of an application of this approach.

The report clearly was not prepared with performance-based regulation in mind, but, being aimed at plant information systems, creates a hierarchical framework (goal, function, system, train, component,…) that, in principle, establishes the relationship between performance at any given level to performance at all other levels. This explicitly includes performance in the area that we have here called "institutional factors"; indeed, the emphasis in the present report on performance of the institution, as opposed to just the physical plant, is due to the influence of this Sandia report. The report addresses the influence of maintenance activities on plant performance and so on. The report does not call its development an influence diagram, but it may turn out to be possible to establish a clean correspondence.

For present purposes, the method offers an integrated approach to identification and perhaps selection of portfolios of performance indicators, ranging from the high-level safety-function indicators favored by NEI down through the performance indicators analyzed in earlier PI work, and on to the institutional factors discussed intermittently in this report. It appears that this approach would also sort out whether there is any logic to working with indicators such as scram rate, according to its particular reasoning process. Its disciplined approach to formulation of goals would tend to militate against observation of indicators simply on grounds that they might correlate with an interesting characteristic; if we are doing performance-based regulation, then the approach wants to see indicators linked through a causal chain to "safety."

A.5 Summary Observations

1. It is important to be clear about the definition of performance-based regulation. The definition used here is consistent with the Commission’s recent guidance and with the NEI white paper. Approaches that simply admit realistic component performance considerations into the details of regulatory requirements that are fundamentally prescriptive in nature are essentially risk-informed adjustments to the current regulatory approach.

2. It appears that OSHA implements the Process Safety Management rule in a manner that is meaningfully comparable to the manner in which the Maintenance Rule is implemented. Both involve oversight of the process used by the regulated entities to manage risk at their plants: how they set priorities, how they monitor, how they respond to failures, how they develop and apply insights from operational experience.

3. More than one source suggests that high-level safety-function indicators, such as those proposed by NEI, are insufficiently leading. This is seen as a major issue.

4. It appears that we can address this issue without settling on an integrating tool, by establishing the degree to which an unsatisfactory situation can be consistent with transiently satisfactory high-level safety function indicators. This is seen as a near-term priority.

5. There is precedent in the literature for measuring “performance” of programmatic activities (process safety management,…). This kind of thing appears to be very different from the NEI proposal. One can bend the idea of “results” to encompass institutional results other than plant performance, but it must be questioned whether this is consistent with the Commission’s intent for performance-based regulation. However, it is interesting that the OSHA PSM rule in particular seems to be an instance of “integrated decision-making” by expert panels convened by the regulated entities. In fact, the PSM rule appears to bear a significant resemblance to the Maintenance Rule.

6. It is useful to view performance-based regulation as regulation essentially by precursor analysis, provided that “precursor” is interpreted in a sufficiently general way. The classical definition of “precursor” focuses on particular events that raise concerns; the more general definition would regard a breakdown of a key plant program as a “precursor,” whether or not a classical precursor event occurred as a consequence of the breakdown, and the threshold of significance would be adjusted to attract regulatory attention to events that are considerably short of significant. This usage of “precursor” is slightly different from that in the ASP program, but is considered useful because it focuses our thinking simultaneously on “results” and on the need for these results to be “leading indicators.” Thinking about performance-based regulation in this way helps to explain what we mean by performance-based, as opposed to the less strict definitions mentioned above, which count as “performance-based” the PSM/Maintenance Rule approach involving performance-oriented audits of licensee process. This presentation carries the risk of creating the impression that the idea is to confine regulatory oversight to reaction to severe events. Reaction to severe events is not the idea; the idea would be to set the threshold of regulatory interest well short of the occurrence of near misses.

7. There is relatively little mention of CCF in the work surveyed. This is arguably a serious omission. If we believe that CCF is risk-significant, then a scheme that does not address it somehow is defective. It deserves a larger role than it appears presently to enjoy in the literature of performance-based regulation, even if only as a surrogate for the kind of institutional factors discussed in some of the work cited. Measures to prevent CCF may be different from those needed to prevent the average train failure, and omission of CCF is therefore potentially serious.

8. The process of abductive inference (the logic of diagnosis) appears to have a lot in common with the objective of performance-based regulation. Implementing it would require establishing correlations between symptoms and underlying conditions, and because it is different from what has been going on, it would require work to formulate the tools. It is noteworthy because it suggests that one can do better than assessing “safety performance” in shades of gray. This kind of tool is seen by many as having tremendous potential in medicine, and would appear to have analogous potential here, although it would be difficult to realize that potential here in the near term.

9. It seems that revealed-preferences literature explicitly addresses only a small part of the present effort. One slice that it addresses would be, for example, deciding whether land contamination is something that we really care about in itself. Another might be helping us to choose among different possible schemes on the basis of attributes such as invasiveness, “leading” character, objective quantifiability, etc. It might help us to decide what subset of all possible performance indicators does the best possible job of supporting the regulatory decision that we wish to make.

References

1 “Strategic Assessment Issue: 12. Risk-informed, Performance-based Regulation,” USNRC, Release Date: September 16, 1996 (http://www.nrc.gov/NRC/STRATEGY/ISSUES/dsi12isp.htm).

2 “Improving the Regulatory Process Through Risk-Based and Performance-Based Regulation,” Nuclear Energy Institute, Draft 10/30/95, attachment to letter from Rasin (NEI) to Milhoan (USNRC), November 14, 1995.

3 See, for example, “Precursors to Potential Severe Core Damage Accidents: 1994/A Status Report,” NUREG/CR-4674 Vol. 21, R. J. Belles et al. (ORNL/NOAC-232, ORNL, 1995).

4 “Risk-Based Regulation Setting Goals for Health and Safety,” by Jenny Bacon, in *Proceedings of PSAM-II*, Volume 1, section 21, pp. 15–20 (1994).

5 Dey, M., “Performance-oriented and risk-based regulation for containment testing,” Nucl. Eng. and Design 166, 305–309 (1996).

6 SECY 96-218, “Quarterly Status Update for the Probabilistic Risk Assessment (PRA) Implementation Plan, Including a Discussion of Four Emerging Policy Issues Associated with Risk-Informed Performance-Based Regulation” (October 11, 1996).

7 “A Method for Setting Reliability Performance Criteria,” Burns, Kevin J., and Turcotte, Richard T., *Proceedings of the ASME-JSME 4th International Conference on Nuclear Engineering*, Volume 4, 123–126 (1996).

8 “Validation of Risk-Based Performance Indicators: Safety System Function Trends,” Boccio, J. L., et al., NUREG/CR-5323 (October 1989).

9 “Methods for Dependency Estimation and System Unavailability Evaluation Based on Failure Data Statistics, Summary Report,” Azarm, M. A., et al., NUREG/CR-5993 Vol. 1 (July 1993).

10 “Methods for Dependency Estimation and System Unavailability Evaluation Based on Failure Data Statistics, Detailed Description and Applications,” Azarm, M. A., et al., NUREG/CR-5993 Vol. 2 (July 1993).

11 “System Unavailability Indicators, Volume 1. Summary Report,” Azarm, M. A., and Vesely, W. E., BNL Report A-3295 (April 1994).

12 “System Unavailability Indicators, Volume 2. Details Analysis & Appendices,” Azarm, M. A., and Vesely, W. E., also BNL Report A-3295 (April 1994).

13 Draft NUREG/CR, “A Framework and Method for the Amalgamation of Performance Indicators at Nuclear Power Plants,” Vols. 1 and 2, Wreathall et al. (draft dated April 15, 1992). This report was prepared under NRC Contract No. NRC-04-87-070, but appears not to have been published. See also “A Framework for assessing influence of organization on plant safety,” M. Modarres et al., Reliability Engineering and System Safety 38, 157 (1992).

14 “One Search for Measures of Maintenance Effectiveness in Safety,” Wreathall, J., and Appignani, P., *Probabilistic Safety Assessment and Management*, Volume 1, Elsevier, pp. 31–35 (1991).

15 DOE Order 210.1, Change 2, Performance Indicators and Analysis of Operations Information, May 1, 1996.

16 DOE/EH-0531 (4Q96), Performance Indicators for ES&H, Report Period Ending December 1996, June 1997.

17 DOE Order 425.1, Change 1, Startup and Restart of Nuclear Facilities, October 26, 1995.

18 DOE-STD-3006-95, Planning and Conduct of Operational Readiness Reviews (ORR), November 1995.

19 *Environment, Safety and Health Performance Objectives and Criteria*, DOE Office of Environment, Safety and Health Evaluations, July 1996.

20 *How to Measure Performance: A Handbook of Techniques and Tools*, prepared by the Training Resources and Data Exchange (TRADE) Performance-Based Management Special Interest Group for DOE, October 1995.

21 “Compliance Assurance Monitoring (CAM) Rule Discussion and Rulemaking, Summary and Discussion of the Draft CAM Rulemaking (40 CFR Parts 64, 70, and 71),” 8/2/96, EPA Website file designation CAM-PKG.WPF, site address: http://ttnwww.rtpnc.epa.gov/html/emtic/cam.htm.

22 “Prepare to Implement Compliance Monitoring,” Singer, James M. and Golla, Scott W., Chemical Engineering Progress, June 1997.

23 “Clean Air Act Credible Evidence Rule; Final Rule” (http://es.epa.gov/oeca/ore/aed/cepreamb.html) .

24 “Process Safety Management of Highly Hazardous Chemicals; Explosives and Blasting Agents; final rule,” Occupational Safety and Health Administration (OSHA), 29 CFR 1910.119, February 24, 1992.

25 APPENDIX A, PSM AUDIT GUIDELINES, OSHA Instruction CPL 2-2.45A, Directorate of Compliance Programs.

26 “Models for the Analysis of Precursor Data that Preserve Inequality Relationships,” Yi, W., and Bier, V. M., Proceedings of PSA ’96, International Topical Meeting on Probabilistic Safety Assessment, Park City, Utah, p. 1158 (ANS, 1996).

27 “Canonical Correlation Analysis of the Reactor Coolant Pump (RCP) and Component Cooling Water (CCW) Systems,” Grenzebach, William S., and Thomas J. Marx, *Probabilistic Safety Assessment and Management ’96: ESREL ’96—PSAM III*, Volume 2, Springer-Verlag, pp. 1018–1023 (1996).

28 *Fuzzy Sets, Natural Language Computations, and Risk Analysis,* by Kurt J. Schmucker (Computer Science Press, 1984).

29 *Abductive Inference,* edited by J. R. Josephson and S. G. Josephson (Cambridge University Press, 1996).

30 “Integrated Approach Methodology,” M. Roush, M. Modarres, R. N. M. Hunt, D. Kreps, and R. Pearce, prepared for Sandia National Laboratories under Contract 64-7956 (January 1987).

APPENDIX B

SHUTDOWN RISK MODEL

B.1 Introduction

In order to furnish perspective on certain key issues, a risk model was developed for the operating condition for which the performance measure is being evaluated. The risk model covered shutdown “mid-loop” operations for three initiators, loss of the operating residual heat removal (RHR) system, loss of the nuclear component cooling water (CCW) system which provides cooling to the RHR as well as other systems, and a station blackout (SBO). “Mid-loop” operation is with the primary system inventory reduced to approximately the middle of the hot leg nozzle for access to the steam generators for testing or other operations.

The risk model is for a hypothetical 4 loop PWR. The event trees were developed based on current understanding of shutdown mid-loop thermal hydraulic success criteria for typical plants. System fault trees and failure data used for the initiators as well as the mitigating functions were obtained from models of actual plants. Operator error and non-recovery probabilities for AC power and other failures were added to the model with typical time dependent values used in the quantification. Because of the expected importance of recovery of failures that lead to the initiating events, the model was developed to allow the application of individual non-recovery probabilities to sequence cut sets that include the initiator cut sets.

The purpose of this development is to support discussion of a general process, and in particular, to illustrate the point that the significance of a given plant state is not necessarily determined by an observation of a physical parameter deviation (e.g., slightly elevated temperature), but also depends sensitively on what failures occurred in order to cause the deviation. For these purposes, a fault tree for various initiating events has been linked with a more typical accident sequence model structure. However, it has not been the purpose of this development to generate actual results for any particular plant or class of plants. Therefore, the presentation of the model is carried out at a high level. It is our opinion that this model is suitable for purposes of illustration; however, validation and refinement of this model for any plant or class of plants is beyond the scope of the current project.

B.2 Event Tree Development

The event trees for mid-loop operation represent the success or failure of the various functions that are needed to prevent core damage due to overheating for situations where normal residual heat removal is interrupted for some reason. Because the performance measure being evaluated is specifically for failure of the heat removal function and another measure would be applicable to the loss of inventory (and other functional failures), it is assumed that there is no loss of inventory except that due to the consequences of the loss of heat removal (that is, boiling in an open vessel).

The event trees for the three initiators (Figures B-1, B-2, and B-3) are structurally identical. The difference in initiating events and in availability of mitigating systems is represented by appropriate differences in the functions that describe the initiators and each branch point and the linking of these functions. The Boolean reduction of these linked functions properly accounts for mitigator failures caused by the initiator.

The first question on the event tree is whether the RCS is open or not. An open system allows water to be lost once boiling begins and prevents the build up of pressure in the RCS. This in turn prevents the effective use of the steam generators to remove decay heat.

For a closed system, if secondary heat removal works, then nothing else needs to function. For secondary heat removal to work, at least one steam generator in an unblocked loop must contain water, be vented to the atmosphere by opening secondary relief valves, and, in the long term, have feedwater makeup. Without heat removal, it is estimated that for early mid-loop operation (2 days after shutdown) the primary system heats up at a rate of approximately 9 deg. F per minute. Therefore, for an initial temperature of 140 deg. F boiling occurs in about 8 minutes. Further, for a closed system, it is estimated that the pressure will reach approximately 200 psig in ½ hour. This model assumes that at this pressure, the low pressure RHR system would be isolated or would fail. Therefore, in a closed system, preventing core damage without using an alternative cooling path requires RHR recovery in less than ½ hour.

If the RHR is not recovered in ½ hour (and assuming the low pressure RHR is isolated), then for either an initially open system or one in which the operator opens a PORV, in approximately one hour the level in the primary system drops about 6", after which time it is assumed that the RHR would not be able to take suction from the hot leg. Therefore, one hour is that latest time the RHR can be recovered and core damage prevented without coolant injection. If the RHR is not recovered in one hour, it takes approximately 4 hours to boil enough water (through the initially open system or vented system, i.e. success in bleed) for the water level to drop to the top of the core. Four hours is therefore the latest time injection can occur to prevent core damage.

With successful injection, the time to recover closed cycle heat removal via the RHR is considerably extended. Because borated water is being injected and water is leaving as steam without a significant boron concentration, the boron concentration in the primary system is increasing. It is estimated that at about 17 hours after the loss of heat removal, the boron concentration in the core would correspond to saturation at the initial temperature of 140 deg. F. At this point, any temperature decrease would have to be managed carefully to avoid boron precipitation.

Recovery of cooling is modeled at each time period either as failure of the operator to initiate the cooling or non-recovery of the hardware, which for the station blackout is the non-recovery of AC power (either offsite or onsite). Failure of the bleed function is either operator or hardware failure. For the closed vessel with success in bleed, injection success is injection from one of the two high pressure safety injection pumps, again considering both operator and hardware failures. Bleed is required for successful injection. For an open vessel, injection is from either high pressure, low pressure or gravity injection. A common mode operator error to fail to initiate injection is modeled along with individual hardware failures. Failure to recover injection is also included for either high pressure or low pressure injection.

The initiating event for each tree is modeled with a fault tree that includes failure of the running pump (or other piece of equipment) and various combinations of failure of the standby equipment. The tree includes all relevant support systems except those that are covered by a separate initiator (otherwise there would be double counting of failures). As indicated above, the linking of an initiator fault tree with trees for each of the subsequent functions produces core damage sequences that include the specific cause of the initial loss of cooling. Each cut set was reviewed to determine those that would be more likely to be recovered than the others and the cut sets edited to include a lower non-recovery probability as a function of time. For example, a mode of RHR failure is failure of an automatic temperature controller. This is considered highly recoverable for times beyond ½ hour; hence, the 1 and 17 hour non-recovery probabilities for any cut set involving this initiator were reduced.

B.3 Functional Fault Trees

The failure of the various events on the event tree is modeled by a functional fault tree that includes the various hardware, operator and non-recovery failure events. The trees for the three initiators are provided in Figures B-4, B-5, and B-6. These trees transfer to the appropriate detailed system models. For example, in Figure B-4 for the loss of RHR initiator (TTJ is the plant system acronym), failure to recover cooling in ½ hour (function H-1/2), is due to either operator error or non-recovery of hardware failure. The hardware failure is the same as that giving the initiating event modeled by tree IE-TTJ. For subsequent times, conditional operator error or non-recovery probabilities are “ANDed” with the earlier failures. The injection failures are considered on sheet 2 of the tree. System tree D-TOP includes both high pressure injection (at gate GD-DO-130) and low pressure injection (at gate GD-DL-121). The latter includes failures of some of the same pumps, valves, and other hardware that caused the initiating event. The duplicate events in a given cut set are eliminated in the Boolean algebra producing the event tree sequence cut sets. It can be seen in the tree that the same event for the operator failing to inject in 4 hours is used for each injection source. This is in accordance with the assumed common mode failure of injection.

B.4 Overall Results

The event trees in Figures B-1, B-2, and B-3 include the CDF of each sequence that leads to core damage. The branch points are annotated with the corresponding failure probabilities resulting from the solution of the functional fault trees. It should be noted that these function unavailabilities may include either disallowed maintenance events or events that are included in other functions. These are eliminated in the solution of the event tree and do not appear in the final cut sets. The sequence frequencies cannot necessarily be obtained by simply multiplying the branch point values.

Table B-1 provides the total CDF for the tree initiators evaluated as well as the top 15 cut sets. A few of the cut sets will be discussed to provide an understanding of the results. The cut sets are made up of basic events, each of which is a 16 (or less) string of characters. In cut set #1 (and in others) the events with TJ022 and TJ024 are the two RHR pumps. The term -FR indicates a failure to run. The term -INIT is included to allow the inclusion of a one year mission for the operating system running failure and a shorter mission for the standby pump failure. The events beginning with NONRECOV- are the non-recovery probabilities, which for cut set #1 combine to give non-recovery in 17 hours. Cut set #1 is therefore a loss of RHR sequence initiated by running failures of both RHR pumps followed by failure to recover cooling in 17 hours. It is a contributor to sequence S13 in Figure B-1.

Cut set #3 is similar to #1 except that an AC bus failure causes failure of pump TJ024. Cut set #4 is the same except the pump failure results from two ventilation fan failures. Cut set #6 is due to the failure of temperature controllers on the two RHR heat exchangers (TJ031 and TJ032). Two of the non-recovery probabilities have a suffix “A” indicating that the values were reduced during the cut set editing process discussed above.

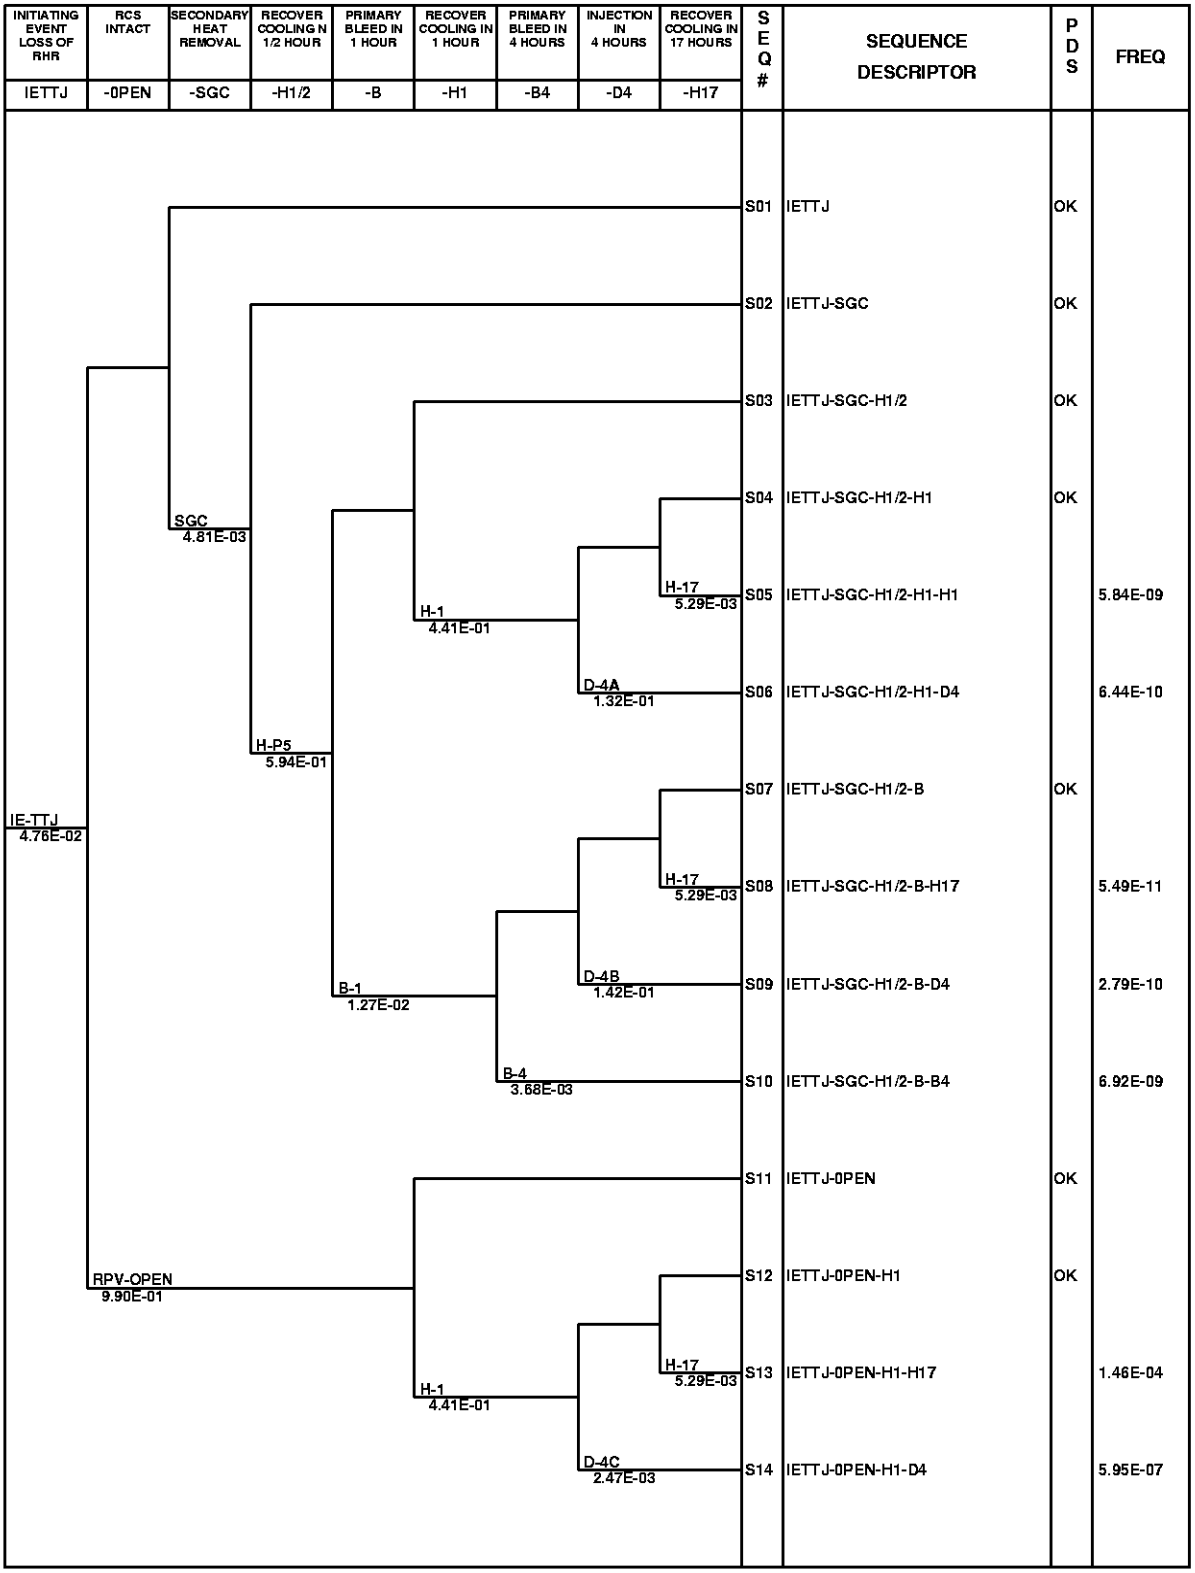

SEQ #	SEQUENCE DESCRIPTOR	PDS	FREQ
S01	IETTJ	OK	
S02	IETTJ-SGC	OK	
S03	IETTJ-SGC-H1/2	OK	
S04	IETTJ-SGC-H1/2-H1	OK	
S05	IETTJ-SGC-H1/2-H1-H1		5.84E-09
S06	IETTJ-SGC-H1/2-H1-D4		6.44E-10
S07	IETTJ-SGC-H1/2-B	OK	
S08	IETTJ-SGC-H1/2-B-H17		5.49E-11
S09	IETTJ-SGC-H1/2-B-D4		2.79E-10
S10	IETTJ-SGC-H1/2-B-B4		6.92E-09
S11	IETTJ-0PEN	OK	
S12	IETTJ-0PEN-H1	OK	
S13	IETTJ-0PEN-H1-H17		1.46E-04
S14	IETTJ-0PEN-H1-D4		5.95E-07

IETTJ.EVT 21 DEC 98

FIGURE B-1 Loss of RHR Initiator

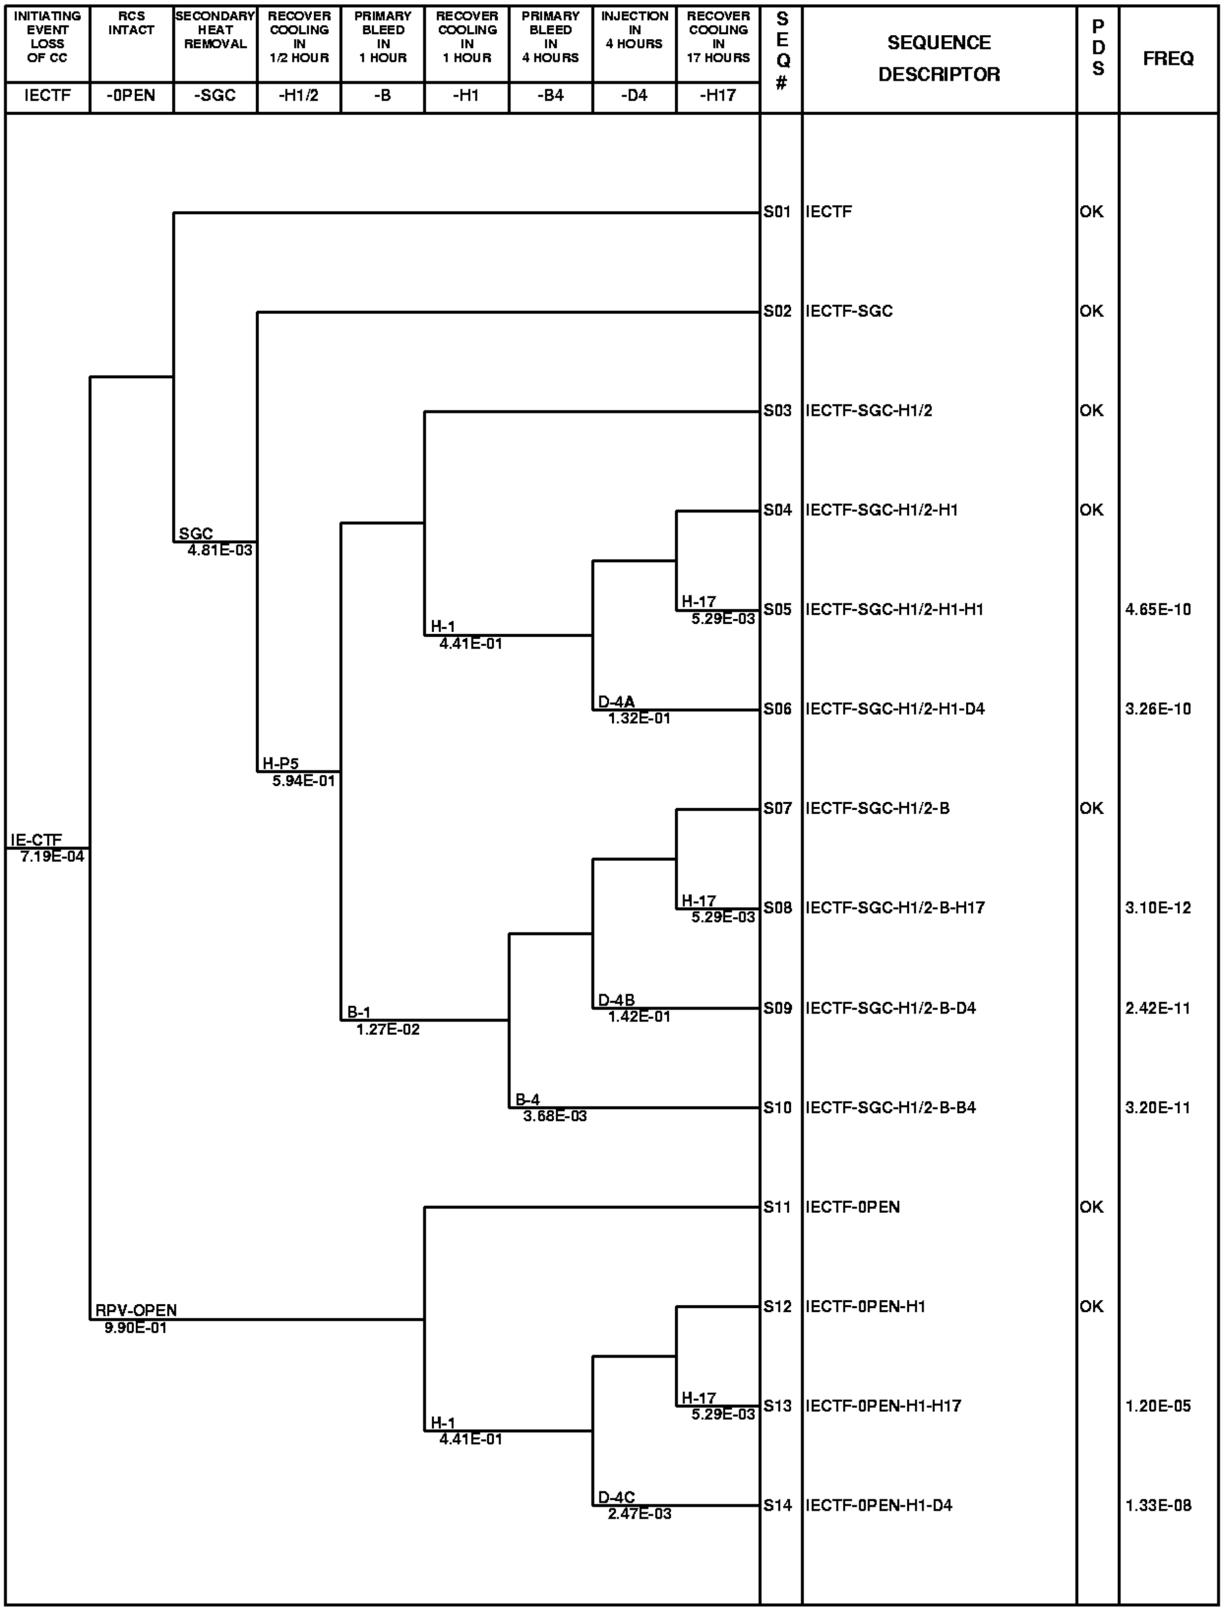

SEQ #	SEQUENCE DESCRIPTOR	PDS	FREQ
S01	IECTF	OK	
S02	IECTF-SGC	OK	
S03	IECTF-SGC-H1/2	OK	
S04	IECTF-SGC-H1/2-H1	OK	
S05	IECTF-SGC-H1/2-H1-H1		4.65E-10
S06	IECTF-SGC-H1/2-H1-D4		3.26E-10
S07	IECTF-SGC-H1/2-B	OK	
S08	IECTF-SGC-H1/2-B-H17		3.10E-12
S09	IECTF-SGC-H1/2-B-D4		2.42E-11
S10	IECTF-SGC-H1/2-B-B4		3.20E-11
S11	IECTF-0PEN	OK	
S12	IECTF-0PEN-H1	OK	
S13	IECTF-0PEN-H1-H17		1.20E-05
S14	IECTF-0PEN-H1-D4		1.33E-08

IECTF.EVT 21 DEC 98

Figure B-2. Loss of CCW Initiator

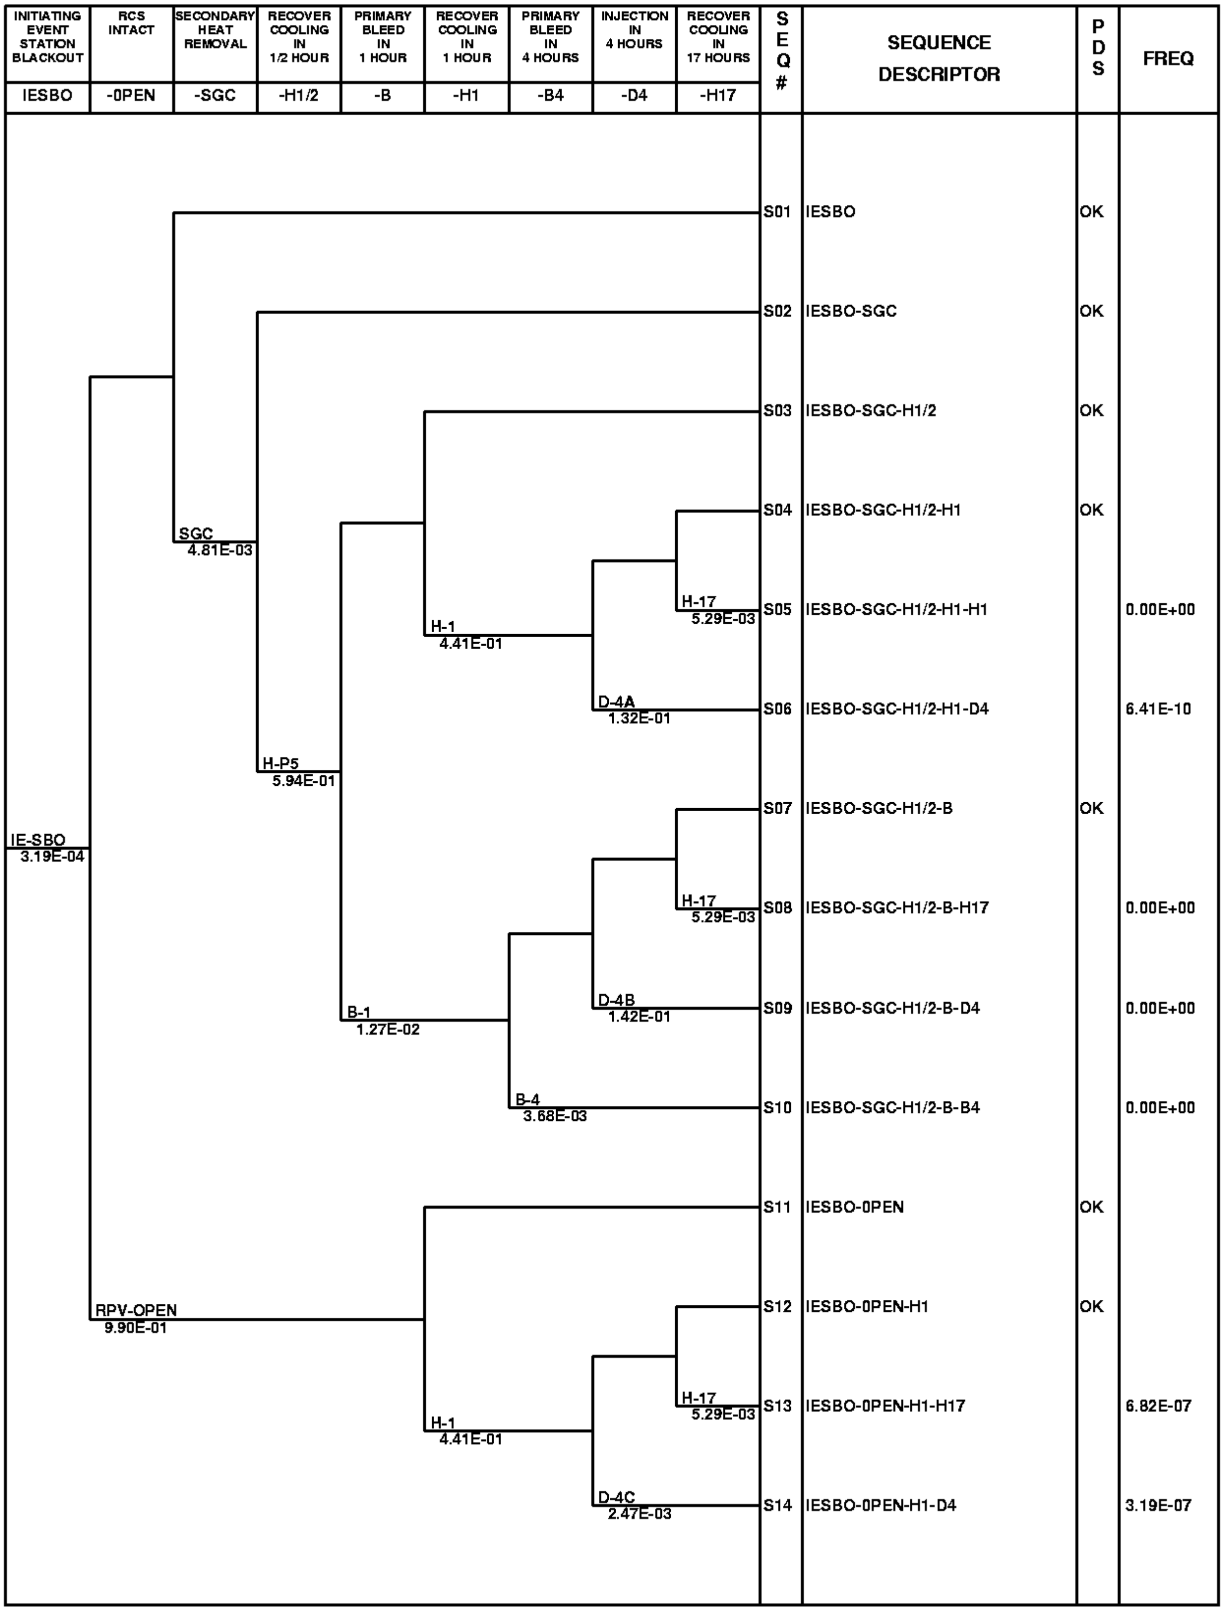

SEQ #	SEQUENCE DESCRIPTOR	PDS	FREQ
S01	IESBO	OK	
S02	IESBO-SGC	OK	
S03	IESBO-SGC-H1/2	OK	
S04	IESBO-SGC-H1/2-H1	OK	
S05	IESBO-SGC-H1/2-H1-H1		0.00E+00
S06	IESBO-SGC-H1/2-H1-D4		6.41E-10
S07	IESBO-SGC-H1/2-B	OK	
S08	IESBO-SGC-H1/2-B-H17		0.00E+00
S09	IESBO-SGC-H1/2-B-D4		0.00E+00
S10	IESBO-SGC-H1/2-B-B4		0.00E+00
S11	IESBO-0PEN	OK	
S12	IESBO-0PEN-H1	OK	
S13	IESBO-0PEN-H1-H17		6.82E-07
S14	IESBO-0PEN-H1-D4		3.19E-07

IESBO.EVT 21 DEC 98

Figure B-3. Station Blackout Initiator

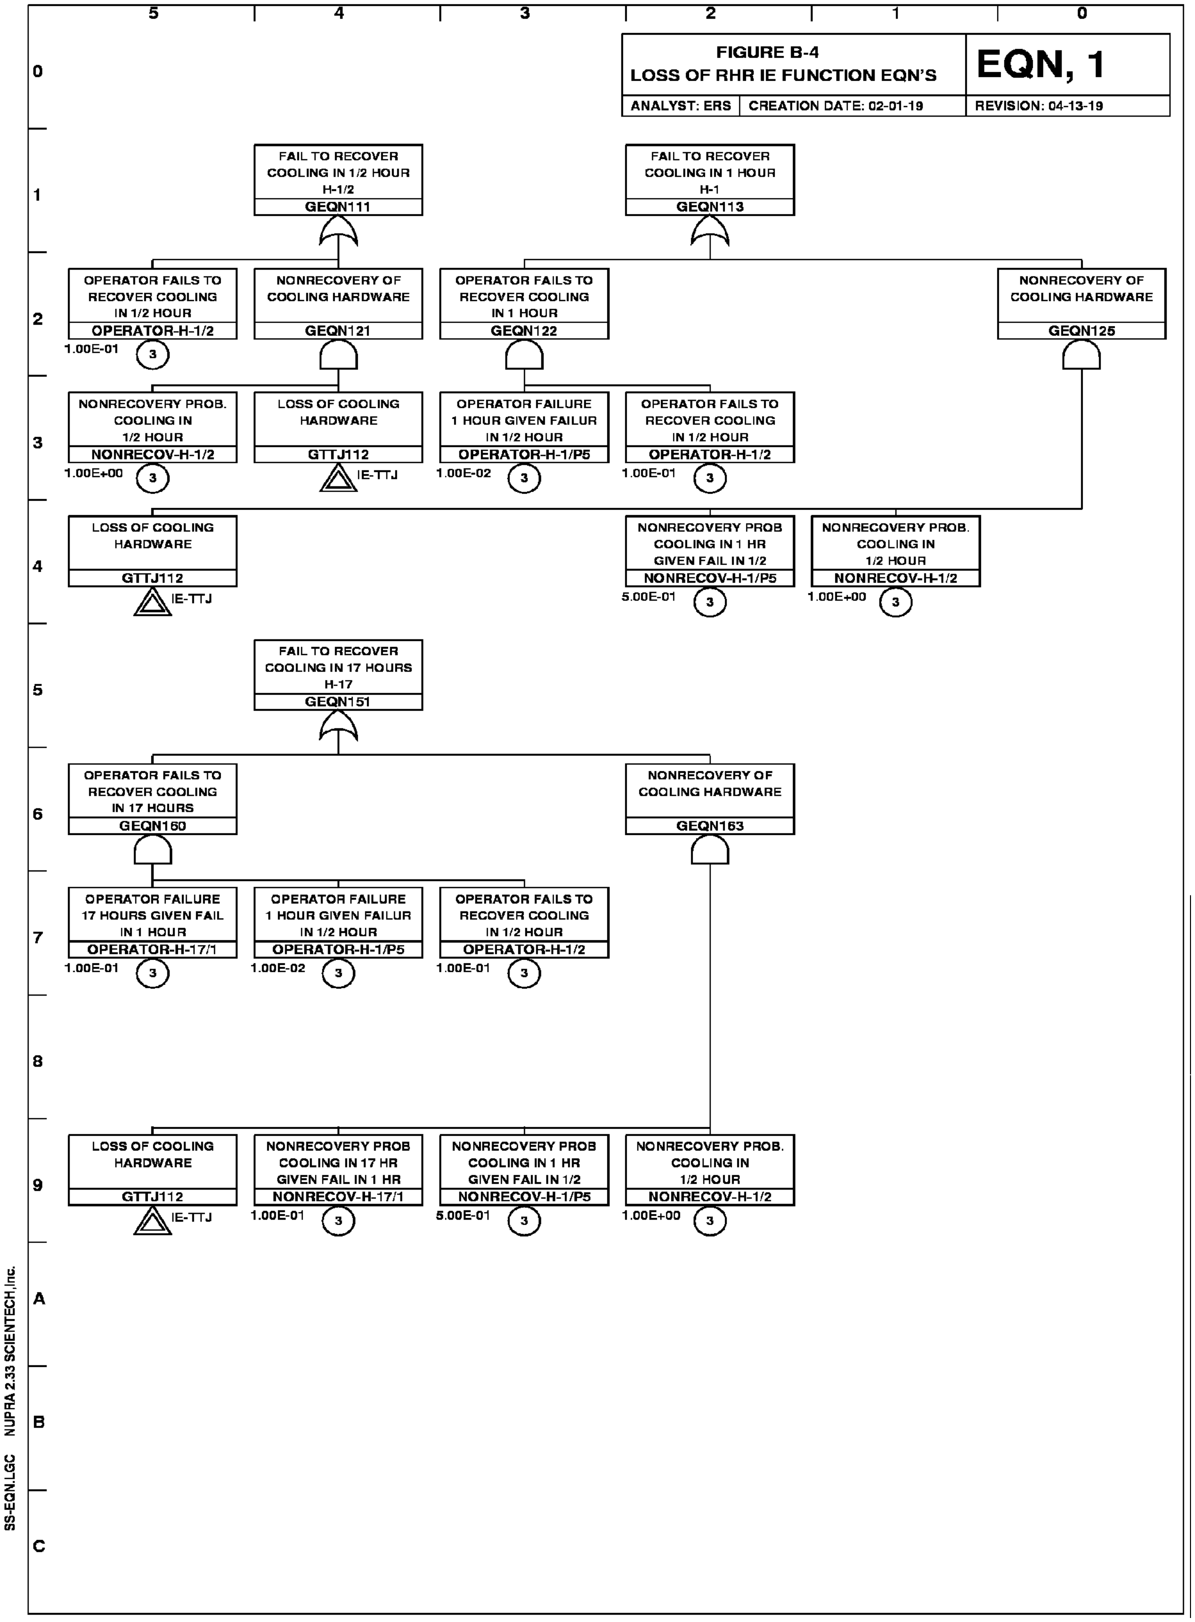
5
4
3
2
1
0
FIGURE B-4
LOSS OF RHR IE FUNCTION EQN'S
EQN, 1
ANALYST: ERS
CREATION DATE: 02-01-19
REVISION: 04-13-19
FAIL TO RECOVER COOLING IN 1/2 HOUR H-1/2
GEQN111
FAIL TO RECOVER COOLING IN 1 HOUR H-1
GEQN113
OPERATOR FAILS TO RECOVER COOLING IN 1/2 HOUR
OPERATOR-H-1/2
1.00E-01
3
NONRECOVERY OF COOLING HARDWARE
GEQN121
OPERATOR FAILS TO RECOVER COOLING IN 1 HOUR
GEQN122
NONRECOVERY OF COOLING HARDWARE
GEQN125
NONRECOVERY PROB. COOLING IN 1/2 HOUR
NONRECOV-H-1/2
1.00E+00
3
LOSS OF COOLING HARDWARE
GTTJ112
IE-TTJ
OPERATOR FAILURE 1 HOUR GIVEN FAILUR IN 1/2 HOUR
OPERATOR-H-1/P5
1.00E-02
3
OPERATOR FAILS TO RECOVER COOLING IN 1/2 HOUR
OPERATOR-H-1/2
1.00E-01
3
LOSS OF COOLING HARDWARE
GTTJ112
IE-TTJ
NONRECOVERY PROB COOLING IN 1 HR GIVEN FAIL IN 1/2
NONRECOV-H-1/P5
5.00E-01
3
NONRECOVERY PROB. COOLING IN 1/2 HOUR
NONRECOV-H-1/2
1.00E+00
3
FAIL TO RECOVER COOLING IN 17 HOURS H-17
GEQN151
OPERATOR FAILS TO RECOVER COOLING IN 17 HOURS
GEQN160
NONRECOVERY OF COOLING HARDWARE
GEQN163
OPERATOR FAILURE 17 HOURS GIVEN FAIL IN 1 HOUR
OPERATOR-H-17/1
1.00E-01
3
OPERATOR FAILURE 1 HOUR GIVEN FAILUR IN 1/2 HOUR
OPERATOR-H-1/P5
1.00E-02
3
OPERATOR FAILS TO RECOVER COOLING IN 1/2 HOUR
OPERATOR-H-1/2
1.00E-01
3
LOSS OF COOLING HARDWARE
GTTJ112
IE-TTJ
NONRECOVERY PROB COOLING IN 17 HR GIVEN FAIL IN 1 HR
NONRECOV-H-17/1
1.00E-01
3
NONRECOVERY PROB COOLING IN 1 HR GIVEN FAIL IN 1/2
NONRECOV-H-1/P5
5.00E-01
3
NONRECOVERY PROB. COOLING IN 1/2 HOUR
NONRECOV-H-1/2
1.00E+00
3
0
1
2
3
4
5
6
7
8
9
A
B
C
SS-EQN.LGC NUPRA 2.33 SCIENTECH,Inc.

Figure B-4 (Sheet 1 of 3) Loss of RHR IE Function EQNs

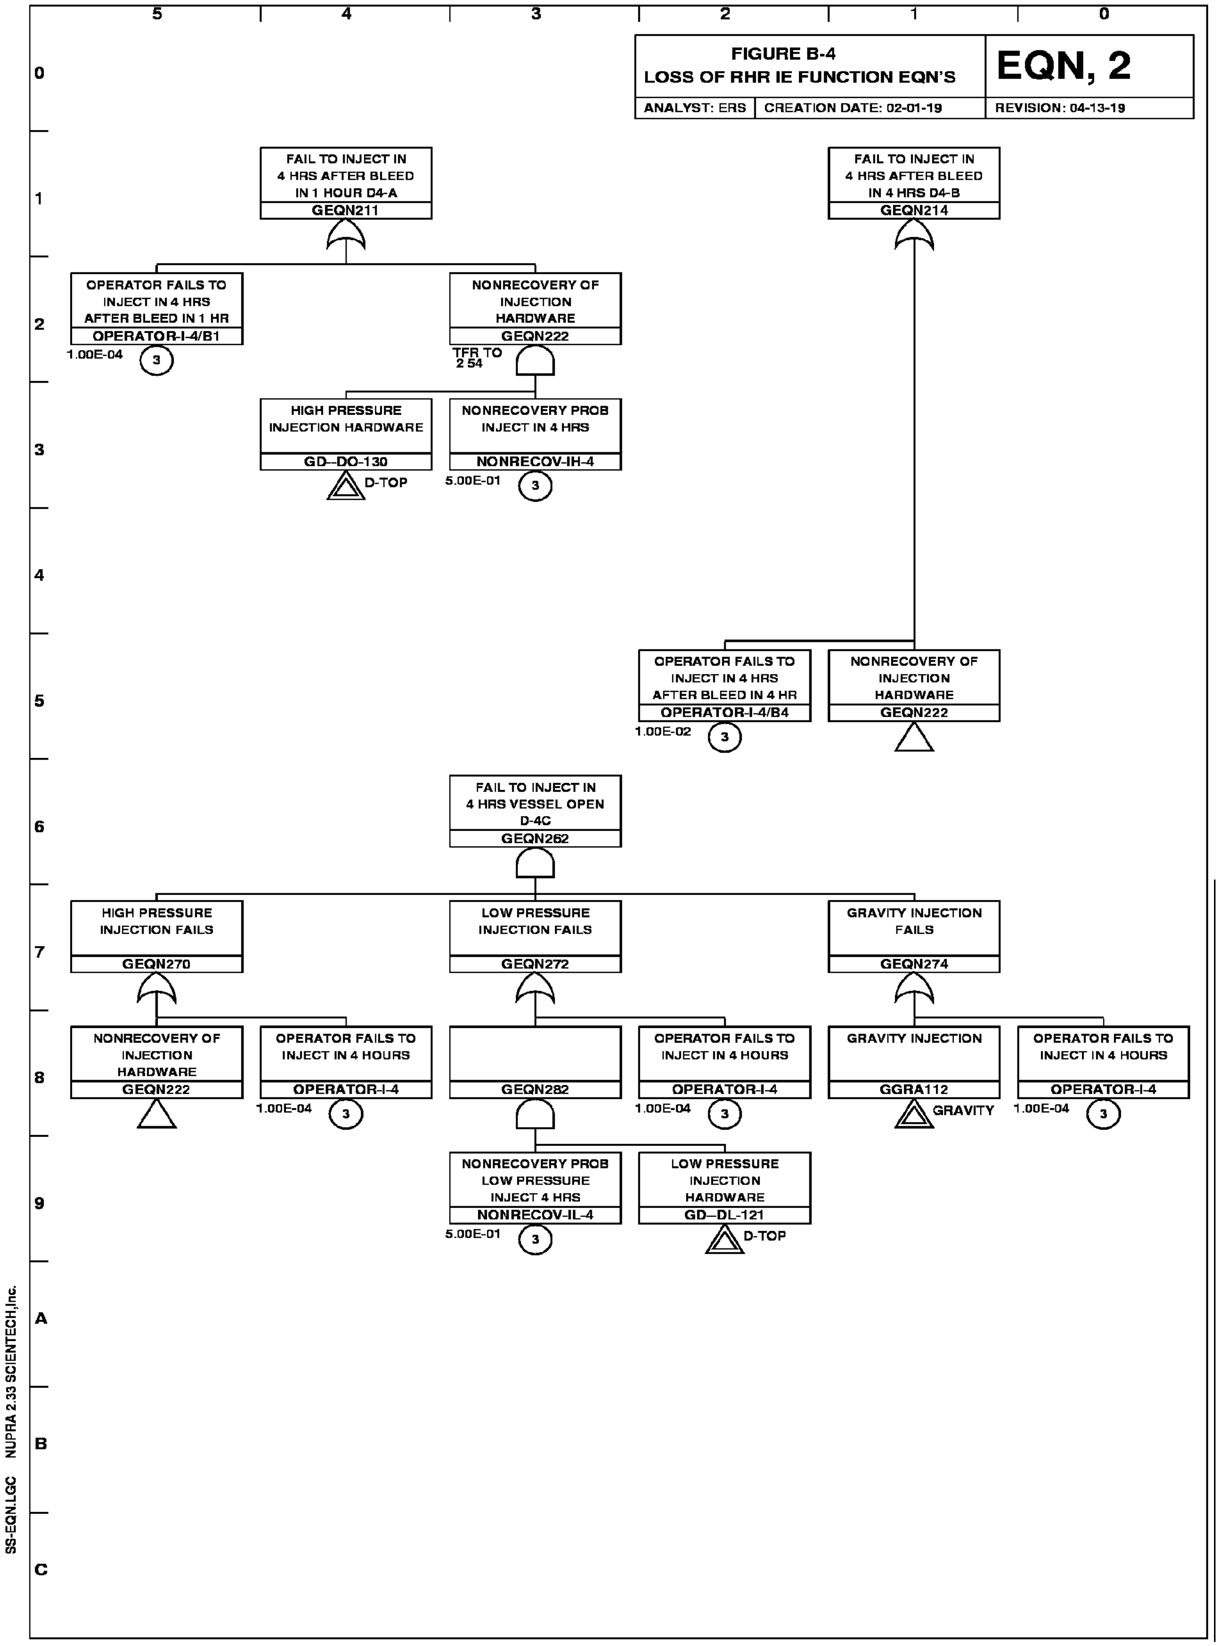
5
4
3
2
1
0
FIGURE B-4
LOSS OF RHR IE FUNCTION EQN'S
EQN, 2
ANALYST: ERS
CREATION DATE: 02-01-19
REVISION: 04-13-19
0
1
2
3
4
5
6
7
8
9
A
B
C
FAIL TO INJECT IN 4 HRS AFTER BLEED IN 1 HOUR D4-A
GEQN211
FAIL TO INJECT IN 4 HRS AFTER BLEED IN 4 HRS D4-B
GEQN214
OPERATOR FAILS TO INJECT IN 4 HRS AFTER BLEED IN 1 HR
OPERATOR-I-4/B1
1.00E-04
3
NONRECOVERY OF INJECTION HARDWARE
GEQN222
TFR TO 2 54
HIGH PRESSURE INJECTION HARDWARE
GD--DO-130
D-TOP
NONRECOVERY PROB INJECT IN 4 HRS
NONRECOV-IH-4
5.00E-01
3
OPERATOR FAILS TO INJECT IN 4 HRS AFTER BLEED IN 4 HR
OPERATOR-I-4/B4
1.00E-02
3
NONRECOVERY OF INJECTION HARDWARE
GEQN222
FAIL TO INJECT IN 4 HRS VESSEL OPEN D-4C
GEQN262
HIGH PRESSURE INJECTION FAILS
GEQN270
LOW PRESSURE INJECTION FAILS
GEQN272
GRAVITY INJECTION FAILS
GEQN274
NONRECOVERY OF INJECTION HARDWARE
GEQN222
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
3
GEQN282
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
3
GRAVITY INJECTION
GGRA112
GRAVITY
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
3
NONRECOVERY PROB LOW PRESSURE INJECT 4 HRS
NONRECOV-IL-4
5.00E-01
3
LOW PRESSURE INJECTION HARDWARE
GD--DL-121
D-TOP
SS-EQN.LGC NUPRA 2.33 SCIENTECH,Inc.

Figure B-4 (Sheet 2 of 3) Loss of RHR IE Function EQNs

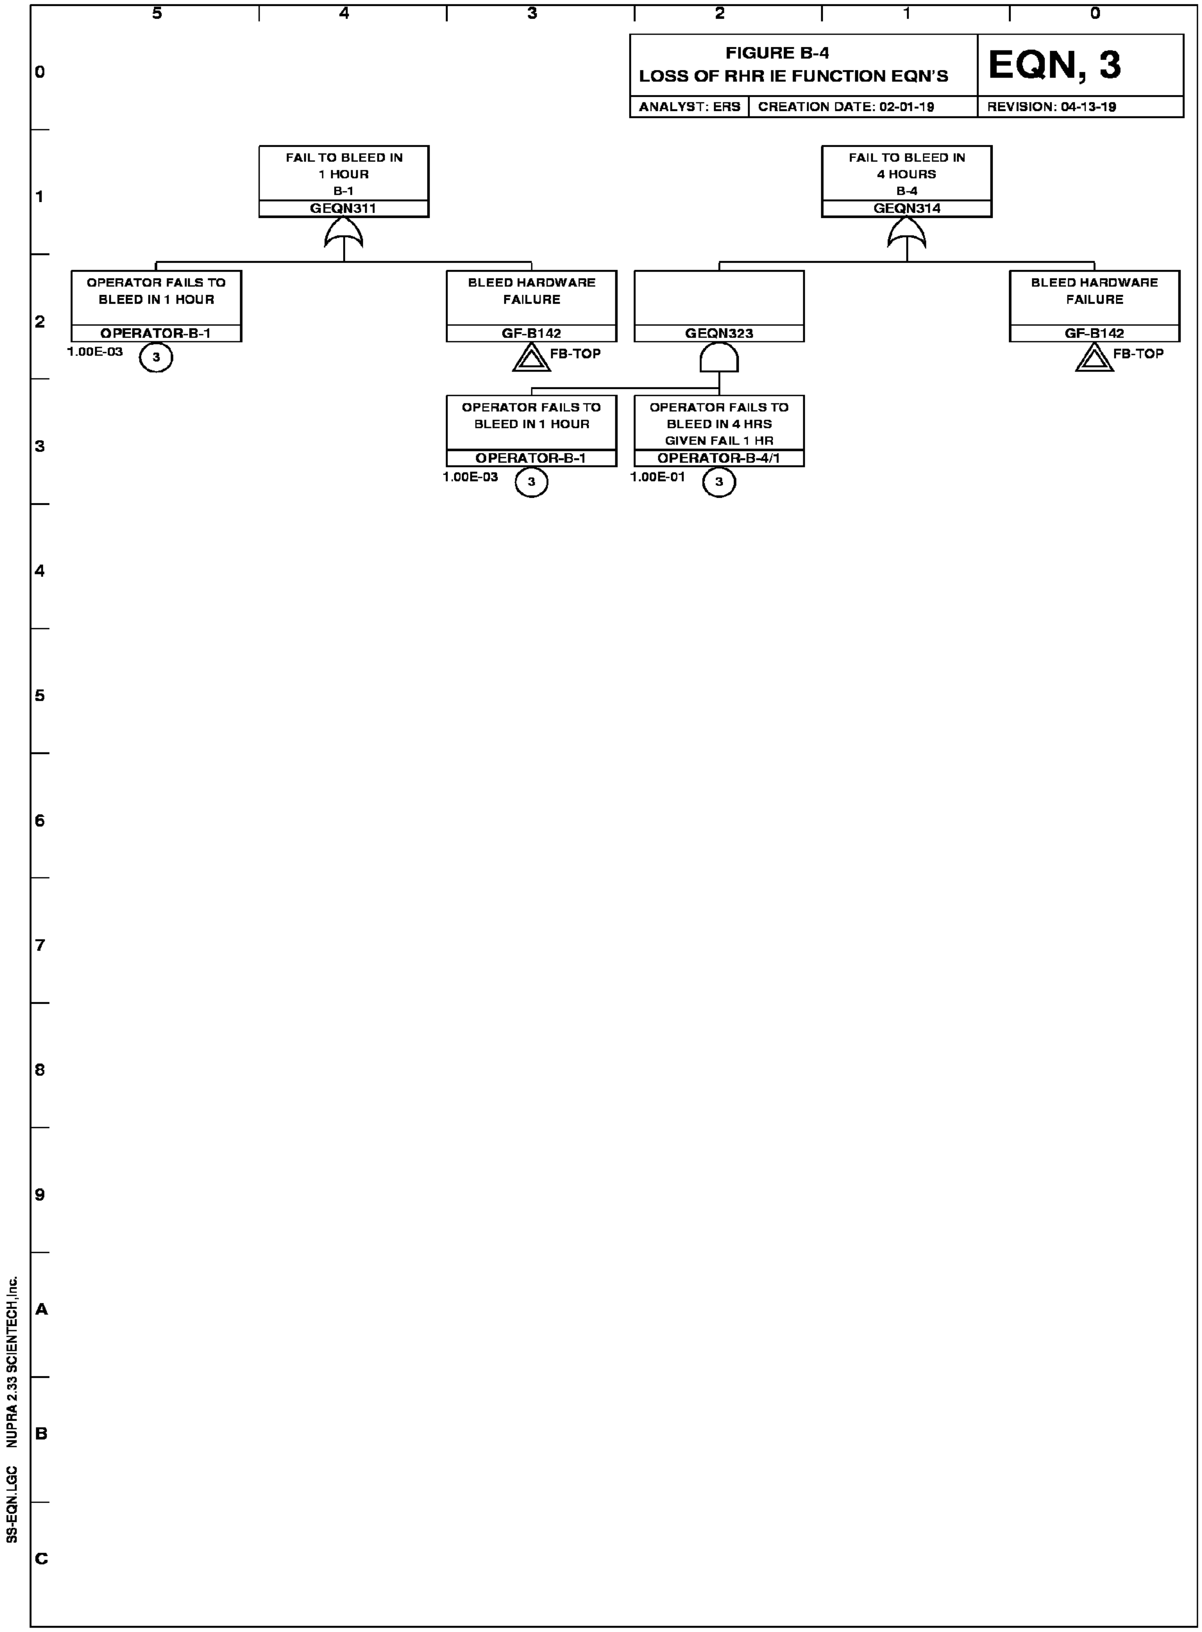

Figure B-4 (Sheet 3 of 3) Loss of RHR IE Function EQNs

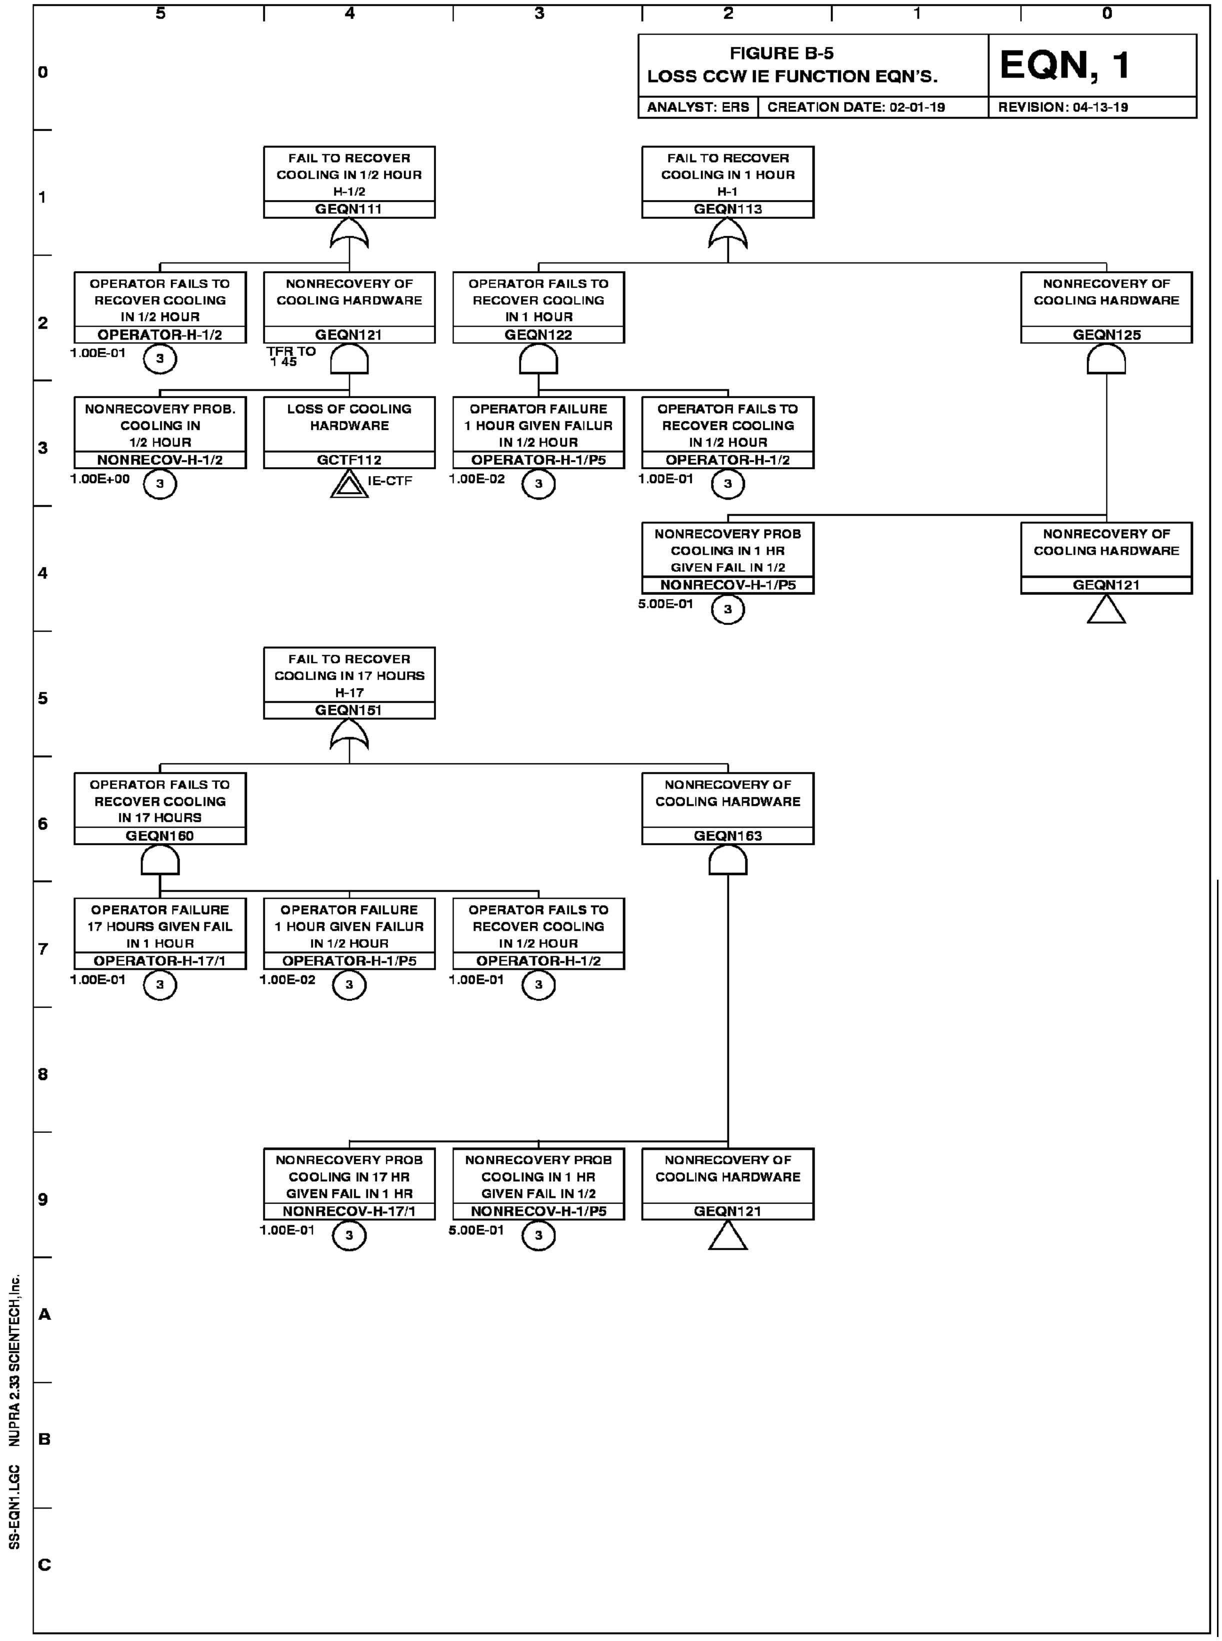
5
4
3
2
1
0
FIGURE B-5
LOSS CCW IE FUNCTION EQN'S.
EQN, 1
ANALYST: ERS
CREATION DATE: 02-01-19
REVISION: 04-13-19
FAIL TO RECOVER COOLING IN 1/2 HOUR H-1/2
GEQN111
FAIL TO RECOVER COOLING IN 1 HOUR H-1
GEQN113
OPERATOR FAILS TO RECOVER COOLING IN 1/2 HOUR
OPERATOR-H-1/2
1.00E-01
NONRECOVERY OF COOLING HARDWARE
GEQN121
TFR TO 1 45
OPERATOR FAILS TO RECOVER COOLING IN 1 HOUR
GEQN122
NONRECOVERY OF COOLING HARDWARE
GEQN125
NONRECOVERY PROB. COOLING IN 1/2 HOUR
NONRECOV-H-1/2
1.00E+00
LOSS OF COOLING HARDWARE
GCTF112
IE-CTF
OPERATOR FAILURE 1 HOUR GIVEN FAILUR IN 1/2 HOUR
OPERATOR-H-1/P5
1.00E-02
OPERATOR FAILS TO RECOVER COOLING IN 1/2 HOUR
OPERATOR-H-1/2
1.00E-01
NONRECOVERY PROB COOLING IN 1 HR GIVEN FAIL IN 1/2
NONRECOV-H-1/P5
5.00E-01
NONRECOVERY OF COOLING HARDWARE
GEQN121
FAIL TO RECOVER COOLING IN 17 HOURS H-17
GEQN151
OPERATOR FAILS TO RECOVER COOLING IN 17 HOURS
GEQN160
NONRECOVERY OF COOLING HARDWARE
GEQN163
OPERATOR FAILURE 17 HOURS GIVEN FAIL IN 1 HOUR
OPERATOR-H-17/1
1.00E-01
OPERATOR FAILURE 1 HOUR GIVEN FAILUR IN 1/2 HOUR
OPERATOR-H-1/P5
1.00E-02
OPERATOR FAILS TO RECOVER COOLING IN 1/2 HOUR
OPERATOR-H-1/2
1.00E-01
NONRECOVERY PROB COOLING IN 17 HR GIVEN FAIL IN 1 HR
NONRECOV-H-17/1
1.00E-01
NONRECOVERY PROB COOLING IN 1 HR GIVEN FAIL IN 1/2
NONRECOV-H-1/P5
5.00E-01
NONRECOVERY OF COOLING HARDWARE
GEQN121
3
0
1
2
3
4
5
6
7
8
9
A
B
C
SS-EQN1.LGC NUPRA 2.33 SCIENTECH,Inc.

Figure B-5 (Sheet 1 of 3) Loss of CCW IE Function EQNs

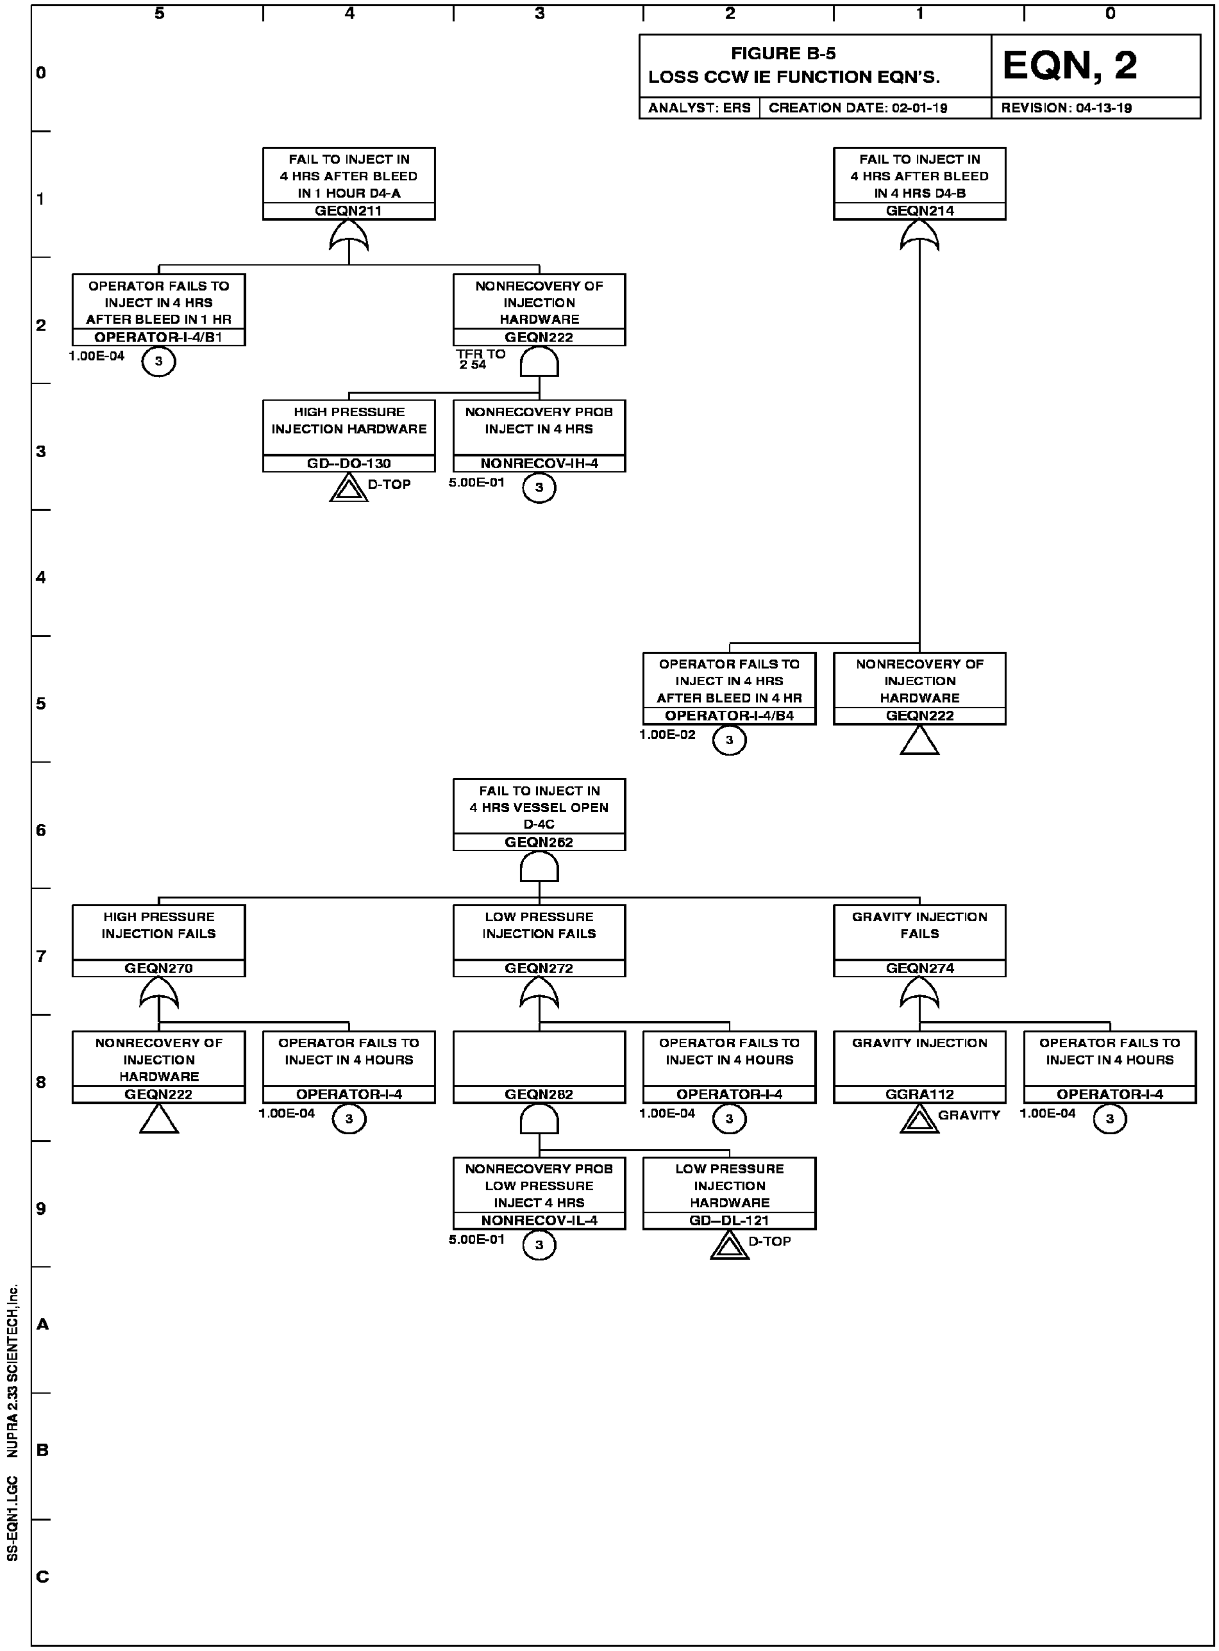
5
4
3
2
1
0
FIGURE B-5
LOSS CCW IE FUNCTION EQN'S.
EQN, 2
ANALYST: ERS
CREATION DATE: 02-01-19
REVISION: 04-13-19
0
1
2
3
4
5
6
7
8
9
A
B
C
FAIL TO INJECT IN 4 HRS AFTER BLEED IN 1 HOUR D4-A
GEQN211
FAIL TO INJECT IN 4 HRS AFTER BLEED IN 4 HRS D4-B
GEQN214
OPERATOR FAILS TO INJECT IN 4 HRS AFTER BLEED IN 1 HR
OPERATOR-I-4/B1
1.00E-04
3
NONRECOVERY OF INJECTION HARDWARE
GEQN222
TFR TO
2 54
HIGH PRESSURE INJECTION HARDWARE
GD--DO-130
D-TOP
NONRECOVERY PROB INJECT IN 4 HRS
NONRECOV-IH-4
5.00E-01
3
OPERATOR FAILS TO INJECT IN 4 HRS AFTER BLEED IN 4 HR
OPERATOR-I-4/B4
1.00E-02
3
NONRECOVERY OF INJECTION HARDWARE
GEQN222
FAIL TO INJECT IN 4 HRS VESSEL OPEN D-4C
GEQN262
HIGH PRESSURE INJECTION FAILS
GEQN270
LOW PRESSURE INJECTION FAILS
GEQN272
GRAVITY INJECTION FAILS
GEQN274
NONRECOVERY OF INJECTION HARDWARE
GEQN222
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
3
GEQN282
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
3
GRAVITY INJECTION
GGRA112
GRAVITY
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
3
NONRECOVERY PROB LOW PRESSURE INJECT 4 HRS
NONRECOV-IL-4
5.00E-01
3
LOW PRESSURE INJECTION HARDWARE
GD--DL-121
D-TOP
SS-EQN1.LGC NUPRA 2.33 SCIENTECH,Inc.

Figure B-5 (Sheet 2 of 3) Loss of CCW IE Function EQNs

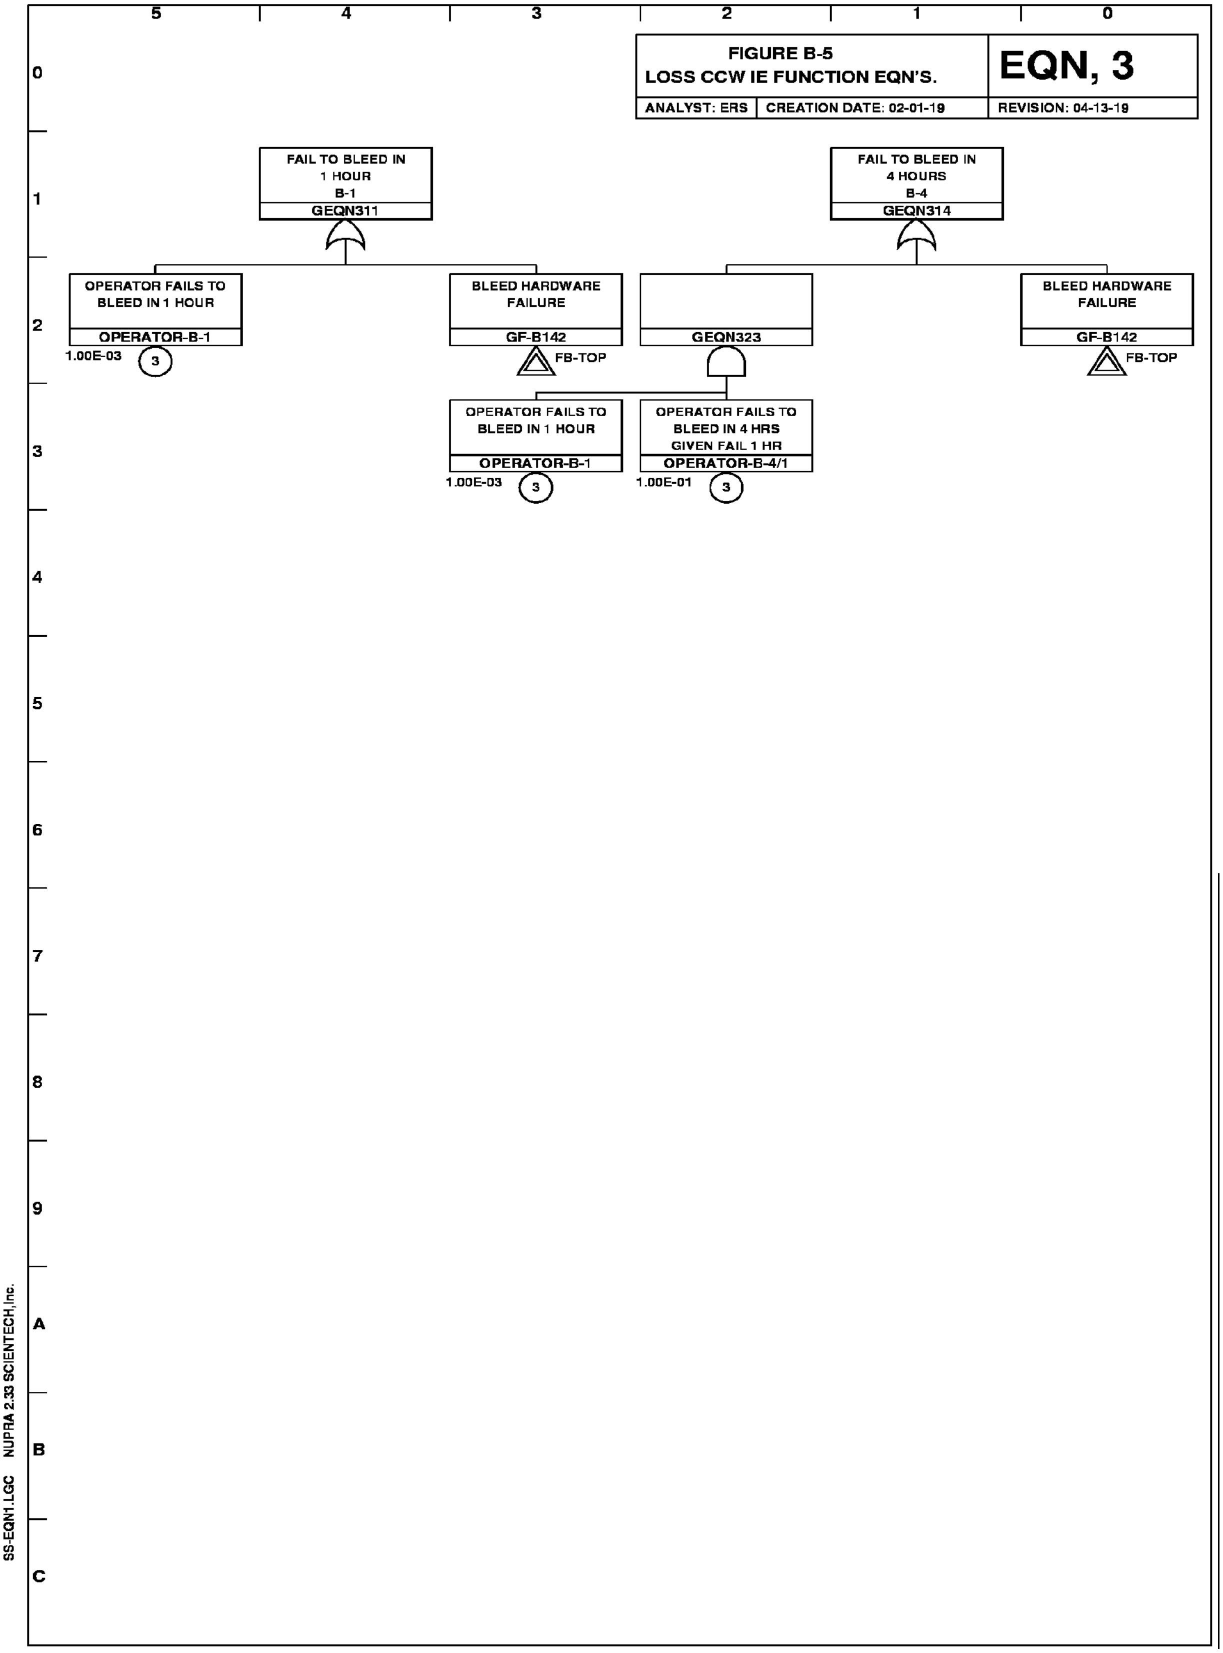

Figure B-5 (Sheet 3 of 3) Loss of CCW IE Function EQNs

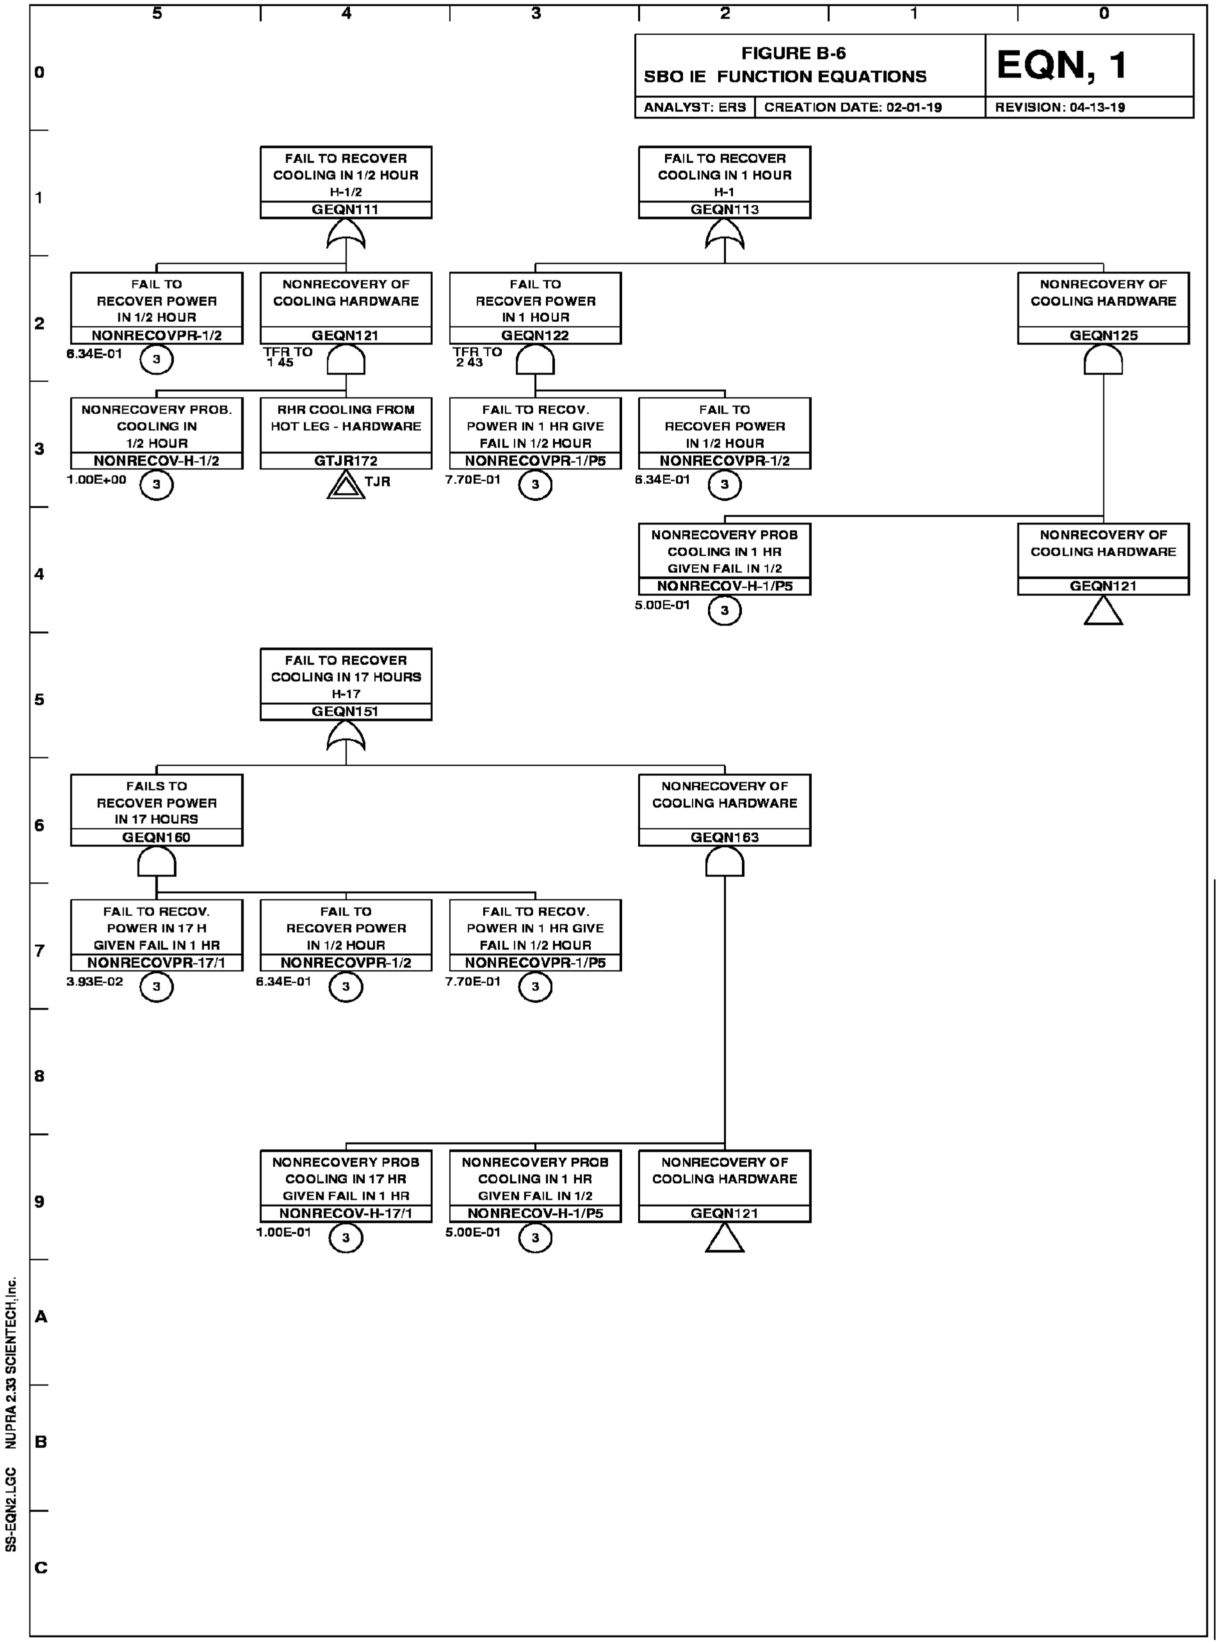
5
4
3
2
1
0
FIGURE B-6
SBO IE FUNCTION EQUATIONS
EQN, 1
ANALYST: ERS
CREATION DATE: 02-01-19
REVISION: 04-13-19
0
1
2
3
4
5
6
7
8
9
A
B
C
FAIL TO RECOVER COOLING IN 1/2 HOUR H-1/2
GEQN111
FAIL TO RECOVER COOLING IN 1 HOUR H-1
GEQN113
FAIL TO RECOVER POWER IN 1/2 HOUR
NONRECOVPR-1/2
6.34E-01
3
NONRECOVERY OF COOLING HARDWARE
GEQN121
TFR TO 1 45
FAIL TO RECOVER POWER IN 1 HOUR
GEQN122
TFR TO 2 43
NONRECOVERY OF COOLING HARDWARE
GEQN125
NONRECOVERY PROB. COOLING IN 1/2 HOUR
NONRECOV-H-1/2
1.00E+00
3
RHR COOLING FROM HOT LEG - HARDWARE
GTJR172
TJR
FAIL TO RECOV. POWER IN 1 HR GIVE FAIL IN 1/2 HOUR
NONRECOVPR-1/P5
7.70E-01
3
FAIL TO RECOVER POWER IN 1/2 HOUR
NONRECOVPR-1/2
6.34E-01
3
NONRECOVERY PROB COOLING IN 1 HR GIVEN FAIL IN 1/2
NONRECOV-H-1/P5
5.00E-01
3
NONRECOVERY OF COOLING HARDWARE
GEQN121
FAIL TO RECOVER COOLING IN 17 HOURS H-17
GEQN151
FAILS TO RECOVER POWER IN 17 HOURS
GEQN160
NONRECOVERY OF COOLING HARDWARE
GEQN163
FAIL TO RECOV. POWER IN 17 H GIVEN FAIL IN 1 HR
NONRECOVPR-17/1
3.93E-02
3
FAIL TO RECOVER POWER IN 1/2 HOUR
NONRECOVPR-1/2
6.34E-01
3
FAIL TO RECOV. POWER IN 1 HR GIVE FAIL IN 1/2 HOUR
NONRECOVPR-1/P5
7.70E-01
3
NONRECOVERY PROB COOLING IN 17 HR GIVEN FAIL IN 1 HR
NONRECOV-H-17/1
1.00E-01
3
NONRECOVERY PROB COOLING IN 1 HR GIVEN FAIL IN 1/2
NONRECOV-H-1/P5
5.00E-01
3
NONRECOVERY OF COOLING HARDWARE
GEQN121
SS-EQN2.LGC NUPRA 2.33 SCIENTECH, Inc.

Figure B-6 (Sheet 1 of 3) SBO IE Function EQNs

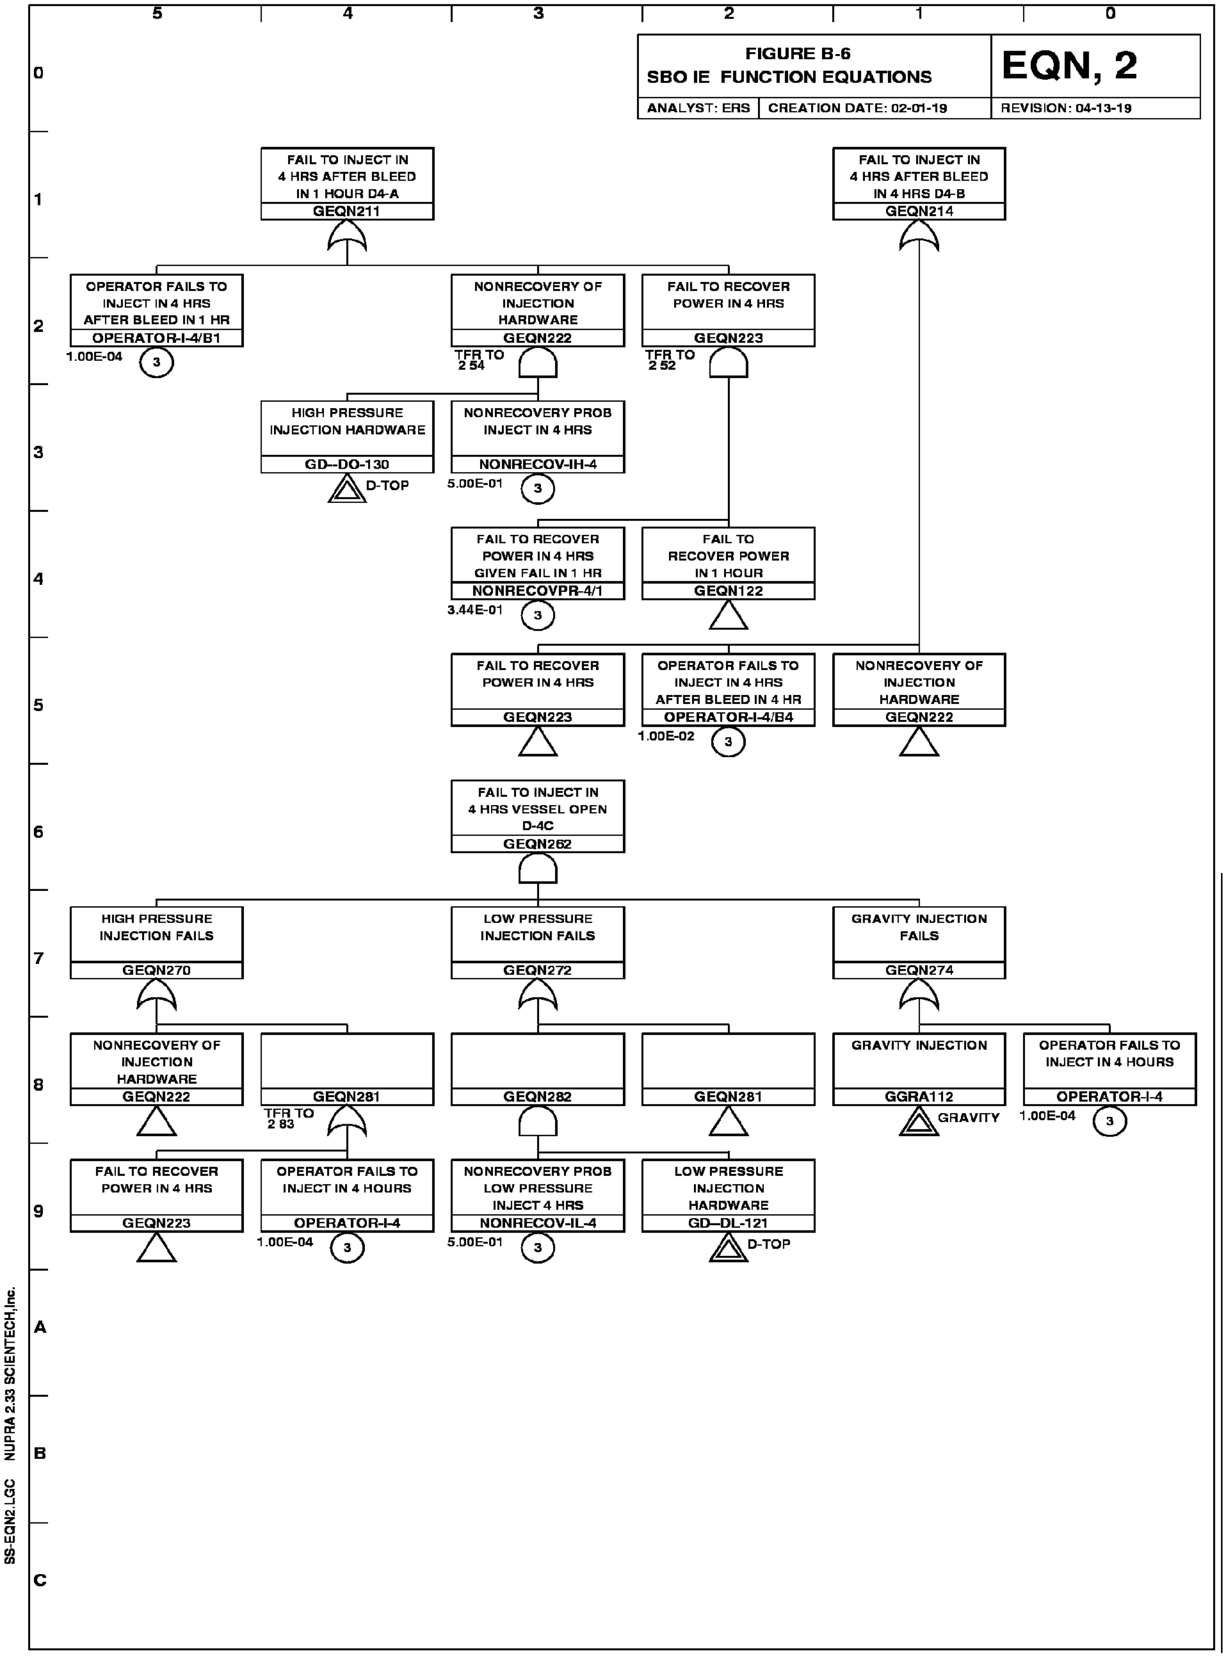

5
4
3
2
1
0
FIGURE B-6
SBO IE FUNCTION EQUATIONS
EQN, 2
ANALYST: ERS
CREATION DATE: 02-01-19
REVISION: 04-13-19
FAIL TO INJECT IN 4 HRS AFTER BLEED IN 1 HOUR D4-A
GEQN211
FAIL TO INJECT IN 4 HRS AFTER BLEED IN 4 HRS D4-B
GEQN214
OPERATOR FAILS TO INJECT IN 4 HRS AFTER BLEED IN 1 HR
OPERATOR-I-4/B1
1.00E-04
NONRECOVERY OF INJECTION HARDWARE
GEQN222
TFR TO 2 54
FAIL TO RECOVER POWER IN 4 HRS
GEQN223
TFR TO 2 52
HIGH PRESSURE INJECTION HARDWARE
GD--DO-130
D-TOP
NONRECOVERY PROB INJECT IN 4 HRS
NONRECOV-IH-4
5.00E-01
FAIL TO RECOVER POWER IN 4 HRS GIVEN FAIL IN 1 HR
NONRECOVPR-4/1
3.44E-01
FAIL TO RECOVER POWER IN 1 HOUR
GEQN122
FAIL TO RECOVER POWER IN 4 HRS
GEQN223
OPERATOR FAILS TO INJECT IN 4 HRS AFTER BLEED IN 4 HR
OPERATOR-I-4/B4
1.00E-02
NONRECOVERY OF INJECTION HARDWARE
GEQN222
FAIL TO INJECT IN 4 HRS VESSEL OPEN D-4C
GEQN262
HIGH PRESSURE INJECTION FAILS
GEQN270
LOW PRESSURE INJECTION FAILS
GEQN272
GRAVITY INJECTION FAILS
GEQN274
NONRECOVERY OF INJECTION HARDWARE
GEQN222
GEQN281
TFR TO 2 83
GEQN282
GEQN281
GRAVITY INJECTION
GGRA112
GRAVITY
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
FAIL TO RECOVER POWER IN 4 HRS
GEQN223
OPERATOR FAILS TO INJECT IN 4 HOURS
OPERATOR-I-4
1.00E-04
NONRECOVERY PROB LOW PRESSURE INJECT 4 HRS
NONRECOV-IL-4
5.00E-01
LOW PRESSURE INJECTION HARDWARE
GD--DL-121
D-TOP
3
A
B
C
SS-EQN2.LGC NUPRA 2.33 SCIENTECH,Inc.

Figure B-6 (Sheet 2 of 3) SBO IE Function EQNs

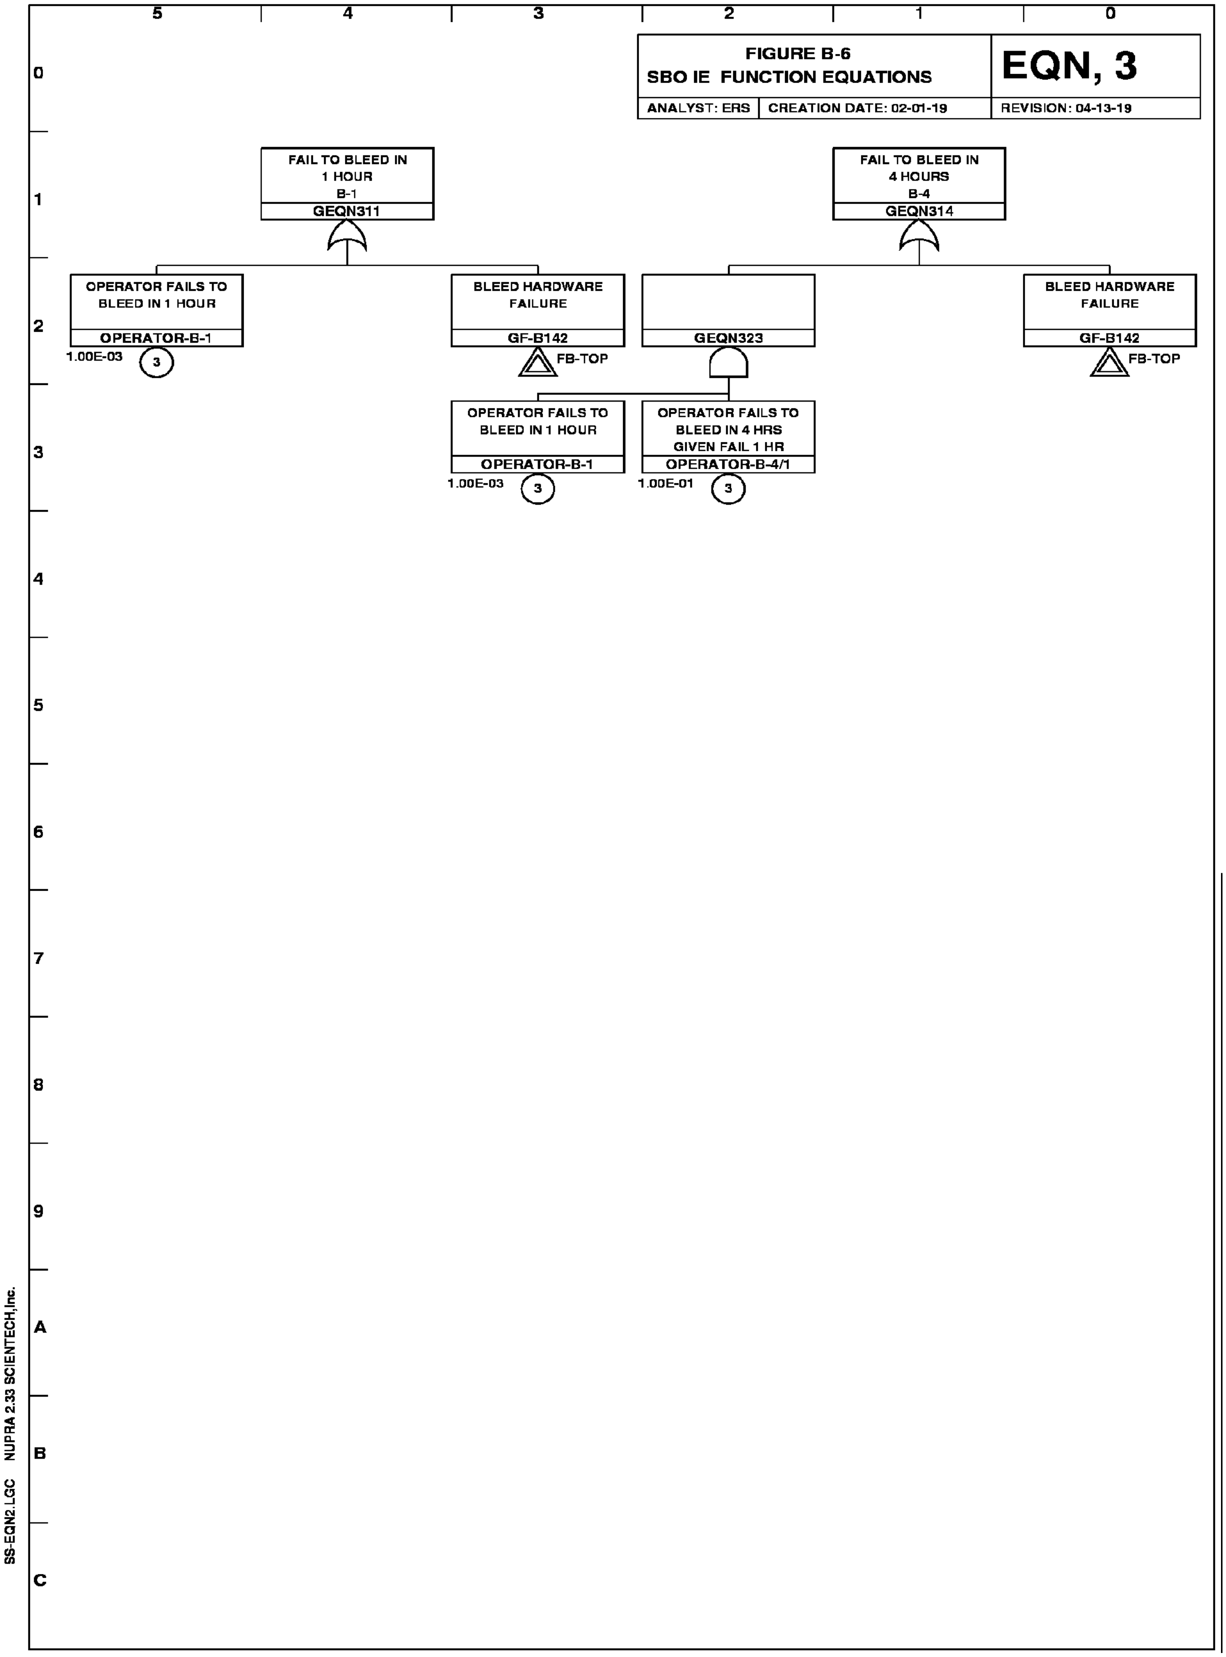

Figure B-6 (Sheet 3 of 3) SBO IE Function EQNs

Table B-1. Top CDF Cut Sets

```
 Top event unavailability (rare event)  = 1.594E-004
 Top event unavailability (third order) = 1.594E-004
   1  8.9006E-005   RPV-OPEN         NONRECOV-H-1/2
                    NONRECOV-H-1/P5  TJ022-MP-FR-INIT TJ022-D001-MP-FR
                    TJ024-D001-SP-FR NONRECOV-H-17/1  TJ024-SP-FR-INIT
   2  1.5171E-005   RPV-OPEN         NONRECOV-H-1/2
                    NONRECOV-H-1/P5  TJ022-MP-FR-INIT TJ022-D001-MP-FR
                    TJ024-D001-SP-FS NONRECOV-H-17/1
   3  1.1745E-005   RPV-OPEN         NONRECOV-H-17/1
                    AC-DD-BUS—BS-LF AC-DD-BS-LF-INIT TJ022-D001-MP-FR
                    TJ022-MP-FR-INIT NONRECOV-H-1/P5  NONRECOV-H-1/2
   4  4.9496E-006   RPV-OPEN         NONRECOV-H-17/1
                    TL058-FN-FR-INIT TL057-D001-FN-FR TJ022-D001-MP-FR
                    TJ022-MP-FR-INIT NONRECOV-H-1/P5  NONRECOV-H-1/2
                    TL058-D001-FN-FR TL057-FN-FR-INIT
   5  4.8505E-006   RPV-OPEN         NONRECOV-H-17/1
                    TJ022-S002-CV-FC TJ022-D001-MP-FR TJ022-MP-FR-INIT
                    NONRECOV-H-1/P5  NONRECOV-H-1/2
   6  4.2589E-006   RPV-OPEN         TJ032-T002-TC-LF
                    NONRECOV-H-1/2   TJ032-TC-LF-INIT TJ031-TC-LF-INIT
                    NONRECOV-H-1/P5A TJ031-T002-TC-LF NONRECOV-H-17/1A
   7  4.2589E-006   RPV-OPEN         OPERATOR-H-17/1
                    OPERATOR-H-1/P5  OPERATOR-H-1/2   TJ031-T002-TC-LF
                    TJ031-TC-LF-INIT TJ032-T002-TC-LF TJ032-TC-LF-INIT
   8  4.2115E-006   RPV-OPEN         NONRECOV-H-1/P5
                    NONRECOV-H-1/2   FLOOD-TF-SEAL    NONRECOV-H-17/1
                    FL-TF-SEAL—INIT
   9  3.9045E-006   RPV-OPEN         NONRECOV-H-1/2
                    NONRECOV-H-1/P5  RECOVERY-MAINTHE NONRECOV-H-17/1
                    MANT-TJ-UPSET-HE MAINT-TJ-UP-INIT
  10  3.7363E-006   RPV-OPEN         NONRECOV-H-17/1
                    RECOVERY-MAINTHE NONRECOV-H-1/2   NONRECOV-H-1/P5
                    MAINTNC-UPSET-HE
  11  2.3166E-006   RPV-OPEN         NONRECOV-H-1/2
                    NONRECOV-H-1/P5  RECOVERY-DRAINHE NONRECOV-H-17/1
                    CCF-OVRDRN-SY-HE
  12  2.2275E-006   RPV-OPEN         NONRECOV-H-17/1
                    YZ-TL04SCC-PT-HE NONRECOV-H-1/P5  NONRECOV-H-1/2
                    RECOVERY-ACTU-HE CCF-SPURYZ-SY-HE
  13  1.8357E-006   RPV-OPEN         NONRECOV-H-1/P5
                    NONRECOV-H-1/2   FLOOD-TJ-TATB    NONRECOV-H-17/1
                    FL-TJ-TATB—INIT
```

```
14  1.0791E-006   RPV-OPEN         NONRECOV-H-17/1
              TJ224-D001-MP-CC NONRECOV-H-1/P5  NONRECOV-H-1/2
15  6.1889E-007   RPV-OPEN         NONRECOV-H-17/1
              TF003-D001-MP-FC TF003-MP-FC-INIT TF002-D001-MP-TM
              TF001-D001-MP-FR TF001-MP-FR-INIT NONRECOV-H-1/2
              NONRECOV-H-1/P5
```

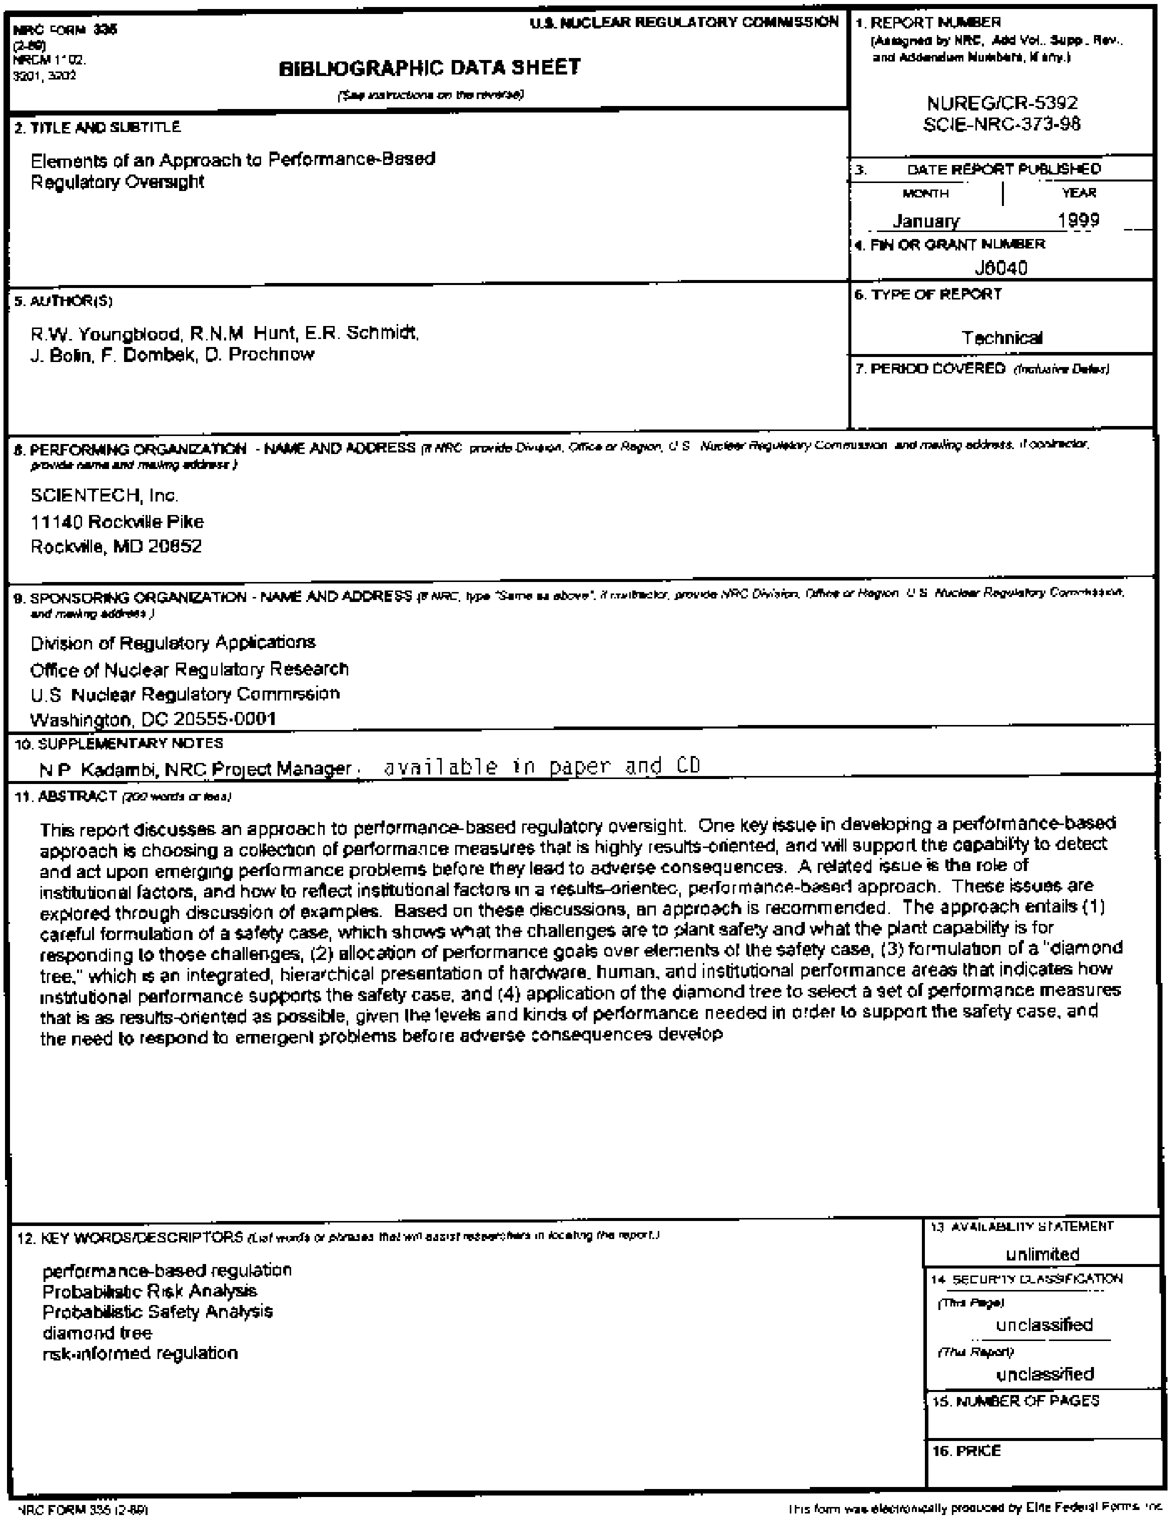

NRC FORM 335
(2-89)
NRCM 1102,
3201, 3202

U.S. NUCLEAR REGULATORY COMMISSION

BIBLIOGRAPHIC DATA SHEET

(See instructions on the reverse)

1. REPORT NUMBER
(Assigned by NRC, Add Vol., Supp., Rev., and Addendum Numbers, if any.)

NUREG/CR-5392
SCIE-NRC-373-98

2. TITLE AND SUBTITLE

Elements of an Approach to Performance-Based
Regulatory Oversight

3. DATE REPORT PUBLISHED

MONTH	YEAR
January	1999

4. FIN OR GRANT NUMBER

J6040

5. AUTHOR(S)

R.W. Youngblood, R.N.M Hunt, E.R. Schmidt,
J. Bolin, F. Dombek, D. Prochnow

6. TYPE OF REPORT

Technical

7. PERIOD COVERED (Inclusive Dates)

8. PERFORMING ORGANIZATION - NAME AND ADDRESS (If NRC, provide Division, Office or Region, U.S. Nuclear Regulatory Commission, and mailing address; if contractor, provide name and mailing address.)

SCIENTECH, Inc.
11140 Rockville Pike
Rockville, MD 20852

9. SPONSORING ORGANIZATION - NAME AND ADDRESS (If NRC, type "Same as above"; if contractor, provide NRC Division, Office or Region, U.S. Nuclear Regulatory Commission, and mailing address.)

Division of Regulatory Applications
Office of Nuclear Regulatory Research
U.S Nuclear Regulatory Commission
Washington, DC 20555-0001

10. SUPPLEMENTARY NOTES

N.P. Kadambi, NRC Project Manager; available in paper and CD

11. ABSTRACT (200 words or less)

This report discusses an approach to performance-based regulatory oversight. One key issue in developing a performance-based approach is choosing a collection of performance measures that is highly results-oriented, and will support the capability to detect and act upon emerging performance problems before they lead to adverse consequences. A related issue is the role of institutional factors, and how to reflect institutional factors in a results-orientec, performance-based approach. These issues are explored through discussion of examples. Based on these discussions, an approach is recommended. The approach entails (1) careful formulation of a safety case, which shows what the challenges are to plant safety and what the plant capability is for responding to those challenges, (2) allocation of performance goals over elements of the safety case, (3) formulation of a "diamond tree," which is an integrated, hierarchical presentation of hardware, human, and institutional performance areas that indicates how institutional performance supports the safety case, and (4) application of the diamond tree to select a set of performance measures that is as results-oriented as possible, given the levels and kinds of performance needed in order to support the safety case, and the need to respond to emergent problems before adverse consequences develop

12. KEY WORDS/DESCRIPTORS (List words or phrases that will assist researchers in locating the report.)

performance-based regulation
Probabilistic Risk Analysis
Probabilistic Safety Analysis
diamond tree
risk-informed regulation

13. AVAILABILITY STATEMENT

unlimited

14. SECURITY CLASSIFICATION

(This Page)
unclassified

(This Report)
unclassified

15. NUMBER OF PAGES

16. PRICE

NRC FORM 335 (2-89)

This form was electronically produced by Elite Federal Forms, Inc.