Boraflex, RACKLIFE, and BADGER

Description and Uncertainties

Manuscript Completed: June 2012
Date Published: September 2012

Prepared by
Thomas C. Haley, Consultant

April L. Pulvirenti, NRC Project Manager

NRC Job Code, V6073

Office of Nuclear Regulatory Research
Division of Engineering

Foreword

The degradation of neutron absorbing materials in spent fuel pools has prompted the staff to undertake a review of the tools and methods used by the industry to ascertain the condition of panels. Among neutron absorbing materials used in spent fuel pools, Boraflex degraded most severely. Boraflex is a neutron absorber material comprised of silicone polymer and boron carbide powder. When gamma-irradiated by spent nuclear fuel, Boraflex is prone to degradation and dissolution in the aqueous environment of the spent fuel pool. Consequently, the subcriticality margins that existed when Boraflex was first installed have eroded. As margins erode, one must have confidence in the neutron absorber data used in criticality analyses to demonstrate compliance with the subcriticality requirements specified by Title 10 of the Code of Federal Regulations Section 50.68 (10 CFR 50.68) or General Design Criteria (GDC) 62. Therefore, when using surveillance methods to determine the areal density of neutron absorbers in spent fuel pools, it is necessary that associated uncertainties be catalogued in a systematic manner.

Therefore, the staff has focused on the surveillance methods for Boraflex; specifically, the RACKLIFE computational tool and the Boron Areal Density Gauge for Evaluating Racks (BADGER) in-situ measurement technique. Uncertainties associated with RACKLIFE are the primary subject of the current report. Uncertainties associated with the BADGER in-situ measurement techniques are only covered to the extent they impact the validation of RACKLIFE, but are discussed in further detail in a separate report.

The RACKLIFE program was developed by Northeast Technology Company (NETCO). RACKLIFE is based on a silica transport model, and calculates the degradation of Boraflex in terms of percent loss of B-10 areal density. There are three principal conclusions from this review of RACKLIFE. First, there are numerous variables such as bulk pool temperature, bulk pool pH, the volume of panel cavities where the Boraflex panel resides, and the exposure of the panel to gamma radiation, which contribute to uncertainty in the results of the RACKLIFE calculations. A simulation of the RACKLIFE program, developed for this report, found that relative uncertainties in percent loss of B-10 areal density for each source of uncertainty may range from 1% to over 100%. The most substantial contributor to overall uncertainty in RACKLIFE output is uncertainty in bulk pool temperature. Uncertainties are mitigated when the principal degree of freedom in the RACKLIFE silica transport model, the escape coefficient, is varied before the calculation such that the predicted bulk silica concentration matches measured bulk pool silica concentration.

The second result is that RACKLIFE does not recognize local or heterogeneous degradation. Because the RACKLIFE model is based in part on gamma exposure which is calculated only over an entire panel, RACKLIFE can only calculate the value of B-10 loss averaged over an entire panel, and therefore does not recognize common modes of heterogeneous degradation of Boraflex such as gaps or scallops. The third result is that uncertainties in the RACKLIFE results increase as degradation increases. At average degradation levels below approximately 20% loss of B-10 areal density, the degradation model is dominated by near-uniform dissolution, which is well

characterized by the silica transport model and escape coefficient tuning procedure used by RACKLIFE. At B-10 loss levels higher than approximately 40%, degradation modes tend to be dominated by large-scale local degradation such as wide gaps, scallops or flow paths caused by local dissolution, where varying the escape coefficient may not be as effective at mitigating uncertainties.

Executive Summary

Boraflex is a neutron absorber material, composed of silicone polymer and boron carbide powder, that is incorporated into high-density spent fuel storage racks for criticality control in the spent fuel pool. When gamma-irradiated by spent nuclear fuel, Boraflex is prone to degradation and dissolution in the aqueous environment of the spent fuel pool. The resulting loss of neutron absorber, typically over several decades, can erode the spent fuel rack's design margin and bring into question compliance with subcriticality requirements specified by Title 10 of the Code of Federal Regulations Section 50.68 (10 CFR 50.68) or General Design Criteria (GDC) 62.

Initially, Boraflex degradation is primarily manifested as shrinkage of the panel, although some local swelling also occurs. Shrinkage will reduce the dimensions of a panel of Boraflex, increasing the neutron line of sight between fuel assemblies around the edges, but does not result in any overall loss of material. When a panel of Boraflex is mechanically restrained axially, either as a result of spent fuel rack design and fabrication or as a result of local swelling and adherence of the panel in its structural cavity, shrinkage stresses can induce breaks and gaps across the length of the panel in a typically nonuniform distribution. These gaps create a neutron line of sight with a greater effect on reactivity than end shrinkage. As the absorbed dose increases, however, the total amount of shrinkage and gap formation will saturate.

As the accumulated dose increases, dissolution becomes the primary mechanism of degradation. Dissolution causes a loss of panel material as the irradiated polymer itself dissolves in the aqueous spent fuel pool environment, releasing particles of boron carbide from the racks. The rate of dissolution increases very rapidly with dose and, unlike shrinkage, does not saturate. Because of the high bulk losses of the panel matrix material and ^{10}B neutron absorber, the dissolution mechanism poses a greater challenge to criticality than shrinkage degradation. The dissolution rate also is significantly higher at higher temperatures. Some amount of dissolution is nearly uniform across large areas of the panel that contact the spent fuel pool water in the fabricated panel cavity. However, stochastically distributed flow paths through the panel cavity also will create a nonhomogeneous distribution of loss that tends to have a positive feedback: a higher local flow results in a higher local loss, which in turn increases the amount of local flow because of the increased area for flow.

The RACKLIFE computer code was developed in the mid-1990s to track and predict the loss of Boraflex and to manage the storage patterns of spent fuel in the racks to maximize the useful lifetime of Boraflex. The premise behind the RACKLIFE program is that the Boraflex degradation mechanism releases soluble silica into the bulk pool volume. Panel-specific data such as irradiation history and bulk pool data such as temperature, pH, and filtration efficiency are input into the software. Another important variable known as the escape coefficient, which approximates the rate of pool water flowing over the panel in the rack, also is input into the software. RACKLIFE then employs a silica transport kinetics model to calculate a hypothetical value for bulk silica concentration. The model is tuned over several iterations via the escape coefficient until the predicted bulk silica concentration matches the measured bulk silica

concentration. The tuned version of RACKLIFE is then run to calculate as output the average percent of degradation of each panel in the pool. Because the RACKLIFE code is governed by a transport model based on silica dissolved in the bulk pool, RACKLIFE can calculate only an average degradation over each panel; it does not recognize the nonhomogeneous degradation described above. In addition, because the RACKLIFE program is predicated on the specific degradation mechanism of Boraflex, it is generally not applicable to other neutron-absorbing materials.

Numerous uncertainties are associated with results from RACKLIFE calculations. To quantify the effect of these uncertainties on the ultimate output from RACKLIFE, a RACKLIFE simulation program was written for this report and stochastic techniques were applied for the following illustrative cases where input data may vary from nominal: polymerization rate, bulk temperature, bulk pH, rack cavity volume, panel surface area, panel mass, panel dose, and escape coefficient. Assuming reasonable bounding values for variation in these pool and panel input data and no tuning of the RACKLIFE program via the escape coefficient, the RACKLIFE simulation returned uncertainties that ranged from 1 percent to 107 percent. With escape coefficient tuning, uncertainties calculated by the simulation program were substantially reduced and ranged from 1 percent to 31 percent. The input datum that produced the most uncertainty was bulk pool temperature, which returned an uncertainty of 60 percent without tuning and, with tuning, 31 percent if temperature deviates by 20°F (11°C) from nominal for 10 years.

The results of the RACKLIFE uncertainty simulation can guide analysis efforts to target the largest uncertainties in plant-specific models. A major conclusion of this report is that the uncertainties depend strongly on variables specific to each spent fuel rack or pool, such that the simulation results presented in this report necessarily apply to only one hypothetical panel in one hypothetical spent fuel rack. Global bounds on RACKLIFE uncertainties over the entire U.S. spent fuel pool fleet would typically be so large as to be useless. Moreover, the random and nonhomogeneous nature of the dissolution and the negative synergistic effects of the resulting random anomalies on the negative reactivity credit of Boraflex demand stochastic techniques to quantify uncertainties definitively.

The Boron Areal Density Gauge for Evaluating Racks (BADGER) is an in-situ measurement device developed to measure the loss of Boraflex. The BADGER instrument comprises two heads suspended on rigid poles, one containing a neutron source and one containing four laterally positioned BF_3 neutron detectors. The two heads are lowered to the bottom of two adjacent rack cells such that the panel targeted for measurement is located between the two heads. Neutrons pass from the source head through the rack walls and target panel and toward the detector head where they are detected and counted by the four detectors. A low neutron count is an indication of high areal density, while a high neutron count is indicative of low areal density. After sufficient neutron counts are obtained, the BADGER heads are then raised in a 2-inch (5.1 cm) increment and the process is repeated such that a typical panel scan might produce 288 data points arranged in a grid 72 cells high and 4 cells wide. A calibration curve relating neutron count to ^{10}B areal density is constructed by scanning a calibration cell and a "zero-dose" panel of nominal

boron areal density. The grid of raw data points is then processed through the calibration curve, resulting in a separate areal density measurement for each of the 288 cells in the panel. Unlike RACKLIFE, which calculates only a panel average degradation, BADGER measurements provide a two-dimensional representation of spatial degradation of each scanned panel. In addition, because the theory of BADGER is not dependent on any specific mechanism of degradation, it may be used to measure the neutron absorption of any material provided that a suitable calibration curve can be obtained.

As with RACKLIFE, numerous uncertainties are associated with BADGER results. A narrow simulation model developed specifically for this report was used to quantify uncertainties due to variation in operational aspects of BADGER (e.g., head misalignment, detector head effects on calibration, and deviation from nominal ^{10}B areal density of the Boraflex materials measured during calibration). For these sources of uncertainty, the simulation model returned uncertainties ranging from ±4 percent to ±15 percent. Analysis of other uncertainties were not performed because analysis techniques for analyzing BADGER data are generally proprietary, and only a limited amount of BADGER data is available in publicly available references. Even without extensive uncertainty analysis, the presented uncertainty analysis can guide the user in targeting sources of BADGER uncertainty. These uncertainties are presented along with recommendations for scan repeatability testing to guide efforts in analyzing uncertainties.

RACKLIFE is a commercial software program, based on the degradation mechanism of Boraflex, developed to calculate the average degradation of a Boraflex panel and to predict the future degradation of Boraflex. BADGER is a commercial in-situ nondestructive diagnostic tool for measuring the areal density of ^{10}B in a neutron absorber as installed in fuel racks. Overall, RACKLIFE and BADGER approach the loss of ^{10}B areal density from Boraflex from opposite sides—RACKLIFE calculates the amount of ^{10}B that is lost while BADGER measures the amount of ^{10}B that remains. To quantify the loss of B_4C in spent fuel pool racks, the RACKLIFE and BADGER methods are combined.

Acknowledgements

The author gratefully acknowledges the helpful comments provided by NRC staff during the preparation of this report. In particular, the extensive comments, edits, and contributions of April Pulvirenti of the Office of Nuclear Reactor Research are gratefully acknowledged.

Thomas C. Haley
2012

Table of Contents

Foreword ii
Executive Summary iv
Acknowledgements vii
List of Figures xi
List of Tables xiii
1 INTRODUCTION 1-1
2 BACKGROUND 2-1
2.1 Historical Context 2-1
2.2 Problems for Boraflex – A Chronology 2-3
2.3 Boraflex Rack Designs and Spent Fuel Pools 2-6
3 THE PHYSICO-CHEMICAL PROPERTIES OF BORAFLEX 3-1
3.1 The Composition of Boraflex and its Boron Carbide Loading 3-1
3.2 The PDMS Binder and its Response to Radiation 3-3
3.3 Stresses that Lead to Gap Formation 3-8
3.4 Irradiated PDMS in Water 3-10
3.5 Silica Solubility and Polymerization 3-13
3.6 Measuring Aqueous Silica Concentration 3-15
3.7 The Non-PDMS Components 3-16
3.8 Characteristics of Boraflex Panel Dissolution 3-17
4 RACKLIFE 4-1
4.1 Overview of RACKLIFE 4-2
4.2 RACKLIFE Theory 4-3
4.2.1 Radiation Exposure History 4-3
4.2.2 Silica Kinetics 4-10
4.2.3 The Escape Coefficient 4-16
4.2.4 Cleanup Systems 4-17
4.3 Using RACKLIFE 4-17
4.3.1 Estimating Average B_4C Losses and Pool Silica Concentration 4-18
4.3.2 Managing the Movement of Spent Fuel in the Pool 4-19
4.3.3 Planning a BADGER Measurement Campaign 4-19

4.3.4 Real-Time Corroboration of BADGER Measurements 4-22
4.3.5 Predicting Future Boraflex Losses 4-23
4.4 Analysis of Uncertainties in RACKLIFE 4-29
4.4.1 Silica Kinetics Model Uncertainties 4-29
4.4.2 Methodology to Assess Loss Variation due to Variation in Kinetics Model Variables 4-31
4.4.3 Polymerization Rate Uncertainty – Static Analysis 4-32
4.4.4 Identifying Condition Parameters Which Most Influence Uncertainty 4-34
4.4.5 Polymerization Rate Uncertainty – Dynamic Analysis 4-36
4.4.6 Temperature Uncertainty – Static Analysis 4-37
4.4.7 Temperature Uncertainty – Dynamic Analysis 4-40
4.4.8 pH Uncertainty 4-42
4.4.9 Panel Cavity Volume Uncertainty 4-42
4.4.10 Panel Surface Area Uncertainty 4-43
4.4.11 Panel Mass Uncertainty 4-43
4.4.12 Panel Dose Uncertainty 4-44
4.4.13 Average Escape Coefficient Uncertainty 4-47
4.4.14 Panel Specific Uncertainty 4-48
4.4.15 Summary of RACKLIFE Uncertainties 4-49
5 BADGER 5-1
5.1 Using BADGER 5-5
5.2 Analysis of BADGER Data 5-8
5.3 BADGER Calibration 5-13
5.4 Uncertainties in BADGER 5-16
5.4.1 Scan Repeatability 5-16
5.4.2 The Calibration Cell Scan and Calibration Curve Slope 5-20
5.4.3 The Zero Dose Panel Scan 5-24
5.4.4 The Unattenuated Scan 5-25
5.4.5 The Attenuated Scan 5-26
5.4.6 Summary of BADGER Uncertainties 5-27
6 CONCLUSIONS 6-1
7 REFERENCES 7-1

7.1 NRC References 7-1
7.2 EPRI References 7-2
7.3 Other References 7-4

List of Figures

Figure 2–1. Timeline of some documented events in the history of Boraflex since 1978. 2-5
Figure 2–2. Representative examples of basic types of spent fuel storage rack designs that used Boraflex..... 2-7
Figure 2–3. Examples of Boraflex retainer methods. 2-8
Figure 3–1. Material constituents of Boraflex. Mass proportions are for illustration and will vary among batches..... 3-1
Figure 3–2. Chemical diagram for the left end of a PDMS polymer chain..... 3-3
Figure 3–3. Helical PDMS molecule..... 3-3
Figure 3–4. An elastomer responds to stress by elastic deformation. 3-4
Figure 3–5. Flexible O-Si-O bonds open and lengthen under stress. 3-4
Figure 3–6. Radiation-induced abstraction of a hydrogen radical allows cross-linking with an adjacent polymer chain that also has an abstracted hydrogen. 3-5
Figure 3–7. Radiation-induced abstraction of a methyl group allows cross-linking with an adjacent polymer chain that also has an abstracted methyl group. Note that the lower pendant arm is likely perpendicular to the flat page..... 3-5
Figure 3–8. Illustration of a panel with top end shrinkage and gaps (white areas). Adapted from EPRI Report TR-107335 (1997), with permission from Electric Power Research Institute..... 3-9
Figure 3–9. Hydroxide surface and conversion of siloxane to amorphous silica..... 3-10
Figure 3–10. Crystalline silica and amorphous silica and siloxane structures. 3-11
Figure 3–11. An amorphous silica and siloxane structure with a hydroxyl layer at the material interface..... 3-12
Figure 3–12. The hydroxide has successfully catalyzed the extraction of the silicon atom into solution as monosilicic acid..... 3-12
Figure 3–13. Silicic acid monomers condense to form a silicic acid dimer plus a water molecule. 3-14
Figure 3–14. Silicic cyclic trimer. 3-14
Figure 3–15. Acute contact angles give small radii of curvature at PDMS-particle interfaces that can enhance local dissolution..... 3-16
Figure 3–16. Illustration of an extensively degraded coupon..... 3-17
Figure 3–17. A to-scale illustration of a local dissolution "scallop" in an axial section of a Boraflex panel. Adapted from EPRI Report TR-1003414 (2002) with permission from Electric Power Research Institute..... 3-19
Figure 4–1. Relative amount of dose absorbed by a Boraflex panel over time from an assembly placed in the racks 7 days after shutdown. 4-5
Figure 4–2. Flux-trap rack showing "near" and "far" rack cells..... 4-6
Figure 4–3. Example gamma dose exposure histories for select individual panel in PWR Region I and Region II racks..... 4-8
Figure 4–4. Example B_4C loss histories for the same PWR Region I and Region II panels as in figure 4.3..... 4-8
Figure 4–5. Silica mass balance kinetics model for a spent fuel pool 4-11

Figure 4–6. Silica release rate per unit area of Boraflex. Adapted from EPRI Report TR-107333 (1997), with permission from Electric Power Research Institute. 4-12
Figure 4–7. Example of a RACKLIFE simulated spent fuel pool showing regions of high (red), medium (yellow), low (green), and near-zero (blue) loss. 4-20
Figure 4–8. Sample cumulative dose (left) and percent loss (right) distribution curves. 4-21
Figure 4–9. Interaction between RACKLIFE and BADGER. 4-24
Figure 4–10. Use of RACKLIFE and BADGER over time. 4-25
Figure 4–11. Matching silica data in RACKLIFE and projecting the future. 4-27
Figure 4–12. Uncertainty in panel loss due to ±50% variation in polymerization rate - static test. 4-33
Figure 4–13. Conditional effect plot for the polymerization rate, k. 4-36
Figure 4–14. Uncertainty in panel loss due to ±20°F (11°C) variation in temperature - static test. 4-39
Figure 4–15. Conditional effects plot for bulk pool temperature, T. 4-39
Figure 4–16. RACKLIFE predicted pool silica (lower plot) and panel loss (upper plot) for static and dynamic variation of temperature and escape coefficient. 4-41
Figure 4–17. Comparison of BWR dose calculation methodologies to verification and validation data. 4-45
Figure 4–18. Comparison of PWR dose calculation methodologies to verification and validation data. 4-45
Figure 4–19. Uncertainty in panel loss due to ±52% variation in dose - static test. 4-46
Figure 4–20. Conditional effect of panel dose. 4-46
Figure 4–21. Uncertainty in panel loss due to 10% variation in escape coefficient. 4-47
Figure 5–1. Axial view of BADGER in egg-crate racks. 5-2
Figure 5–2. Lateral view of BADGER in egg-crate (top) and flux-trap (bottom) rack cells. 5-3
Figure 5–3. BADGER in operation. 5-4
Figure 5–4. A BADGER scan trace. Adapted from EPRI Report TR-107335 (1997), with permission from Electric Power Research Institute. 5-5
Figure 5–5. A hypothetical annotated BADGER scan trace. Adapted from EPRI Report TR-107335 (1997), with permission from Electric Power Research Institute. 5-9
Figure 5–6. Grid view of the degraded panel shown in Figure 5-5. Adapted from [IP2 2001 p56] to correspond with Figure 5-5. 5-12
Figure 5–7. A hypothetical simulated calibration curve for a hypothetical calibration cell, adapted from [EPRI 1998a; Table 1-2]. 5-14
Figure 5–8. BADGER *inner* and *outer* detectors. 5-17
Figure 5–9. BADGER source / detector head misalignments. 5-18
Figure 5–10. PWR Region I - rack versus calibration cell and experimental vs. numerical... 5-21
Figure 5–11. PWR Region II - rack versus calibration cell and experimental versus numerical. 5-23
Figure 5–12. BWR - rack versus calibration cell and experimental versus numerical. 5-23

List of Tables

Table 3-1. Effects of Radiation Dose on PDMS in Boraflex. 3-7
Table 3-2. Equilibrium Si concentrations from irradiated Boraflex in water for select temperatures and pH. 3-13
Table 4-1. Summary of RACKLIFE uncertainties. 4-49
Table 5-1. Summary of BADGER uncertainties. 5-28

1 INTRODUCTION

Boraflex is a criticality controlling neutron absorber used in high-density spent fuel storage racks. When irradiated by spent nuclear fuel, it is prone to dissolving in the aqueous environment of the spent fuel pool. The resulting loss of absorber from the racks can challenge the $k_{eff} \leq 0.95$ subcriticality limit for the spent fuel pool. The RACKLIFE computer code and the BADGER in-situ measurement device were developed to predict and measure, respectively, the loss of Boraflex.

The objectives of this report are to:

- detail the composition of Boraflex, the mechanisms by which it is lost, and the types of loss observed;
- describe the models applied in RACKLIFE, how RACKLIFE is used in practice, and describe and quantify the uncertainties in a RACKLIFE prediction; and
- describe the use of BADGER in measuring Boraflex loss, how BADGER measurement data are translated into Boraflex losses, how BADGER is calibrated to quantify losses, and describe and quantify the uncertainties in BADGER measurements.

As background, Chapter 2.0 summarizes the history of spent fuel management, the chronology of the industry and regulatory responses to the problems with Boraflex, and typical spent fuel rack designs that use Boraflex. Chapter 3.0 then summarizes the physico-chemical characteristics of Boraflex that make it prone to dissolution in the spent fuel pool water. It goes on to detail the somewhat random nature of the dissolution mechanism and the negative synergistic effect the resulting random anomalies have on the reactivity credit of Boraflex.

Chapters 4.0 and 5.0 detail the theory and operation of RACKLIFE and BADGER, respectively. In particular, Sections 4.4 and 5.4 describe numerous sources of uncertainty in RACKLIFE and BADGER and show that some can be difficult to understand and quantify. Simulation models developed specifically for this report are used in each chapter to illustrate the nature of various uncertainties and to quantify their magnitude in a typical application.

2 BACKGROUND

This chapter is a summary intended for those less familiar with how Boraflex came to be selected as the neutron absorber of choice in spent fuel racks (Section 2.1) and how the less desirable consequences of that choice came to be recognized over time (Section 2.2). In addition, some background on the design of spent fuel racks using Boraflex is presented (Section 2.3) to clarify the terminology in later chapters and to give the reader a mental picture of a Boraflex panel in its panel cavity in a spent fuel rack module in a spent fuel pool.

2.1 Historical Context

Every 12 to 24 months, between one third and one half of a light-water reactor (LWR) nuclear core is discharged as spent fuel. Plans for what to do with this spent nuclear fuel date back to at least 1957 when scientific consensus became focused on burial in a repository inside deep underground salt deposits. In 1970, the Atomic Energy Commission announced plans for such a repository in Kansas, only to withdraw them 2 years later based on the site's unsuitability. Reprocessing was also an expected option for spent fuel but, the U.S. domestic commercial reprocessing option was suspended in 1976 and banned in 1977 (although the ban was subsequently lifted in 1981). The Nuclear Waste Policy Act of 1982 scheduled a national repository to open by 1998 and was amended in 1987 to focus specifically on the Yucca Mountain, Nevada, site. This approach was reaffirmed in 2002. In 2006, the U.S. Department of Energy (DOE) estimated that, conditional on licensing, it could begin to receive fuel in 2017. However, the Yucca Mountain project was defunded starting in 2009.

Nuclear utilities looked for other options as spent fuel pools were approaching capacity. This was mandated by the Nuclear Waste Policy Act of 1982 that, in addition to specifying DOE as responsible for spent fuel disposal, also specified that the nuclear industry had responsibility for spent fuel storage until the Federal Government could receive it. The industry solution was to "rerack"—to replace existing spent fuel storage racks with higher density racks (where "density" is quantified by the number of spent fuel assemblies stored per unit lateral area of a spent fuel rack module). Alternative processes and technologies were all considered (e.g., away-from-reactor storage facilities, spent fuel consolidation, or a second layer of horizontal storage cells above the racks in the pool), but reracking quickly became a proven effective solution.

The spent fuel racks that "came with the plant" when built were typically open-lattice frames rigidly attached to the pool floor and walls. These racks relied on a large spacing, about 50 cm (20 inches) between the spent fuel assemblies to maintain subcriticality. Because pools were originally designed for short-term storage prior to shipping to a reprocessing facility, many of these designs allowed for only a full core offload plus one cycle of discharged fuel—around 400 assemblies for a large pressurized-water reactor (PWR) or 1,100 assemblies for a large boiling-

water reactor (BWR). Low-density open-lattice designs were common even into the 1980s despite the prospect of future reracking; with the construction delays and high interest rates of the 1970s, plants being constructed chose to retain the original low-density racks just to avoid changing any terms of the contract. Plants would then rerack immediately upon taking possession of the reactor. Some early plants even went through multiple reracks, going to higher and higher density racks each time.

However, reracking was not an inexpensive solution. A typical high-density rack module could be 3 meters (9 feet) square, 5 meters (17 feet) high, and weigh 14 metric tonnes (30,000 pounds), mostly stainless steel, with between 12 and 20 such rack modules in a typical spent fuel pool. These rack modules sit unanchored on the spent fuel pool floor and are not attached to the pool walls or each other. External steel bars are typically used to maintain spacing between racks during seismic events. The cost of a typical full-pool rerack was measured in tens of millions of dollars.

To achieve higher densities of spent fuel assemblies, the spacing between fuel assemblies was reduced. To maintain subcriticality with smaller spacing, neutron absorbers were introduced into the racks. As neutron absorbers and rack fabrication techniques improved, and criticality analysis methods became more precise, the assembly spacing became nearly as close as that in the reactor while still maintaining the required margin of subcriticality. Thus, the highest density racks could increase the number of storage cells available by a factor of five or more, while the number of refueling cycles accommodated by the increased capacity could increase by a factor of 15 or more. For example, for a BWR on a 1-year / 1/3rd-core refueling cycle, a low-density pool of 1100 cells could accommodate 800 assemblies of a full core offload and 267 assemblies of one cycle of spent fuel. A high-density pool of 5500 cells could accommodate 800 assemblies of a full core offload and 4005 assemblies of 15 fuel cycles of fuel. The pool could then accept fuel for at least 15 years instead of 1 year before losing full-core offload capability. Designing racks to maximize the use of the available spent fuel pool area further increased the time the pool could accept fuel, as did extending the operating cycle length by increasing fuel feed enrichments and final burnup.

The neutron absorbers originally used included borated stainless steel, then sintered boron carbide powder/carbon matrix composites (e.g., Carborundum™), and later boron carbide aluminum alloy composites (e.g., Boral™). In the late 1970s, a new neutron absorbing material, Boraflex, was applied to spent fuel racks. Boraflex is a light, flexible, easy-to-install combination of neutron-absorbing boron carbide and silicone rubber developed by Brand Industrial Services Company (BISCO). Because reracking was so expensive, the use of an inexpensive polymer matrix to hold the boron carbide powder proved attractive, with bids for reracking using Boraflex coming in with substantially lower costs. Boraflex quickly became the neutron absorber of choice for reracks. At its peak, Boraflex was being used in about 75 spent fuel pools worldwide. Despite the escalating number of problems with Boraflex identified in the mid- to late-1980s (discussed in the next section), lingering ambiguity about the nature and severity of the problem and the low price of Boraflex reracks meant that spent fuel rack modules using Boraflex continued to be installed until 1991.

2.2 Problems for Boraflex – A Chronology

It is now well established that Boraflex degrades by a combination of total gamma dose and the subsequent dissolution of silica into the water of the pool. However, the elucidation of the degradation mechanism (described in detail in Chapter 3) was a decade-long process. In 1978, the U.S. Nuclear Regulatory Commission (NRC) issued Generic Letter 78-11, "Guidance for Spent Fuel Pool Modifications," and the enclosure, "OT Position for Review and Acceptance of Spent Fuel Storage and Handling Applications" (dated 14 April 1978). The Generic Letter states that

> "Methods for verification of long-term material stability and mechanical integrity of special poison material utilized for neutron absorption should include actual tests… to assure long-term safety and integrity of the pool and fuel rack system."

Moreover, Section III Paragraph 1.5 (1) of the enclosure states that

> "For those facilities which employ a strong neutron absorbing material to reduce the neutron multiplication factor for the storage pool, the licensee shall provide the description of onsite tests which will be performed to confirm the presence and retention of the strong absorber in the racks. ...Coupon or other type of surveillance testing shall be performed on a statistically acceptable sample size on a periodic basis throughout the life of the racks to verify the continued presence of a sufficient amount of neutron absorber in the racks to maintain the neutron multiplication factor at or below 0.95."

Boraflex coupon surveillance programs vary from plant to plant. Examples of coupon dimensions include 50 mm (2 inch) squares, 200 × 300 mm (8 × 12 inch) rectangles, and 115 mm (4.5 inch) long dogbone-shaped tensile test specimens with 6 mm (0.25 inch) reduced sections. A few plants also had full-size removable panels that could be inspected. A coupon is generally enclosed in a stainless steel cladding, which can be two sheets attached by bendable tabs or welds, allowing varying degrees of water ingress into the clad cavity. Typically, coupons were scheduled to be permanently removed and tested yearly and then every few years for the first 10 years, and then every 5 years thereafter. Testing programs typically called for measuring the coupon's dimensions, weight, hardness, and neutron attenuation.

When Boraflex coupon surveillance programs were devised, the primary concern was material reaction to radiation dose. Material interaction between the Boraflex, the pool water, and the stainless steel was typically a secondary and independent concern. Accordingly, some plants always placed their coupon trains next to freshly discharged assemblies at every outage to accelerate their dose testing. This sometimes led to doses more than an order of magnitude greater than dose to the Boraflex in the rest of the pool, and even beyond doses that Boraflex in service would actually encounter in its license lifetime. Moreover, while the coupons had the

same material characteristics as the racks they were simulating, they did not necessarily have the same pool water flow characteristics around the Boraflex samples. For example, some coupons had large inspection holes in the cladding on each side to allow for visual examination of the coupons. The portion of the coupon positioned under the inspection hole was exposed to a higher flow rate than the cladded portion of the coupon. In this case, the visible portion of a Boraflex coupon in a high-dose/high-flow environment could dissolve entirely, while in-rack Boraflex, in a lower-dose/lower-flow environment, could have remained in very good condition. In contrast, the inspection of normal-dose/low-flow coupons might have no negative findings, and yet the higher-flow service racks could have experienced significant degradation. This discrepancy between coupon conditions and service conditions often made coupons an unreliable predictor of Boraflex performance in the same spent fuel pool.

"Blackness testing" was another method used to verify the continuous presence of a neutron absorber. See [Tsao et al. 2006] for a publicly available review of a blackness testing system and related theory. This in-situ nondestructive test used a single source/detector head that would travel the length of a rack cell while continuously counting neutrons. The source emitted fast neutrons that could easily pass through an absorber panel into empty adjacent cells, thermalize in the water there, and then backscatter. A relatively low neutron count rate indicated that the absorber was intact, while a relatively high count rate indicated that at least some portion of the absorber was missing. This technique provided a "go/no-go" test for the local axial presence of a strong absorber: it could detect missing absorber panels, gaps, and axial end shrinkage. Later refinements to scan procedures and analysis techniques allowed rough quantitative estimates of the amount of absorber present although this mode of testing was not widely used.

In 1987, both coupons and blackness testing revealed problems with Boraflex that were reported to NRC. At one plant surveillance, the coupons showed "considerable degradation" although, for reasons described above, it was concluded that the coupons did not accurately represent the conditions of the spent fuel racks. At another plant, a BWR, blackness testing revealed numerous axial gaps in the panels, up to 100 mm (4 inches) in size, with the average gap 40 mm (1.5 inches). The axial distribution of gaps was bimodal and truncated: gaps were more likely to occur near the midplane and toward the top, and no gaps occurred in the bottom-third of the panels. (At the time it was conjectured by some that this unexpected pattern indicated a rack manufacturing problem more than a Boraflex material problem. Subsequent tests at other plants showed that this distribution, while sharing common features with similar rack designs, was specific to these racks.) As a result of these reports, NRC Information Notice 87-43 was issued in September 1987.

These initiating events are shown in Figure 2-1 along with other events in the history of Boraflex. Directly above the time line, in green, are NRC communications with respect to Boraflex, culminating in Generic Letter 96-04 in June 1996 that requested each licensee "that uses Boraflex as a neutron absorber in its spent fuel storage racks" to

"(1) assess the capability of the Boraflex to maintain a 5-percent subcriticality margin and (2) submit to the NRC a plan describing its proposed actions if this subcriticality margin cannot be maintained by Boraflex material because of current or projected future Boraflex degradation."

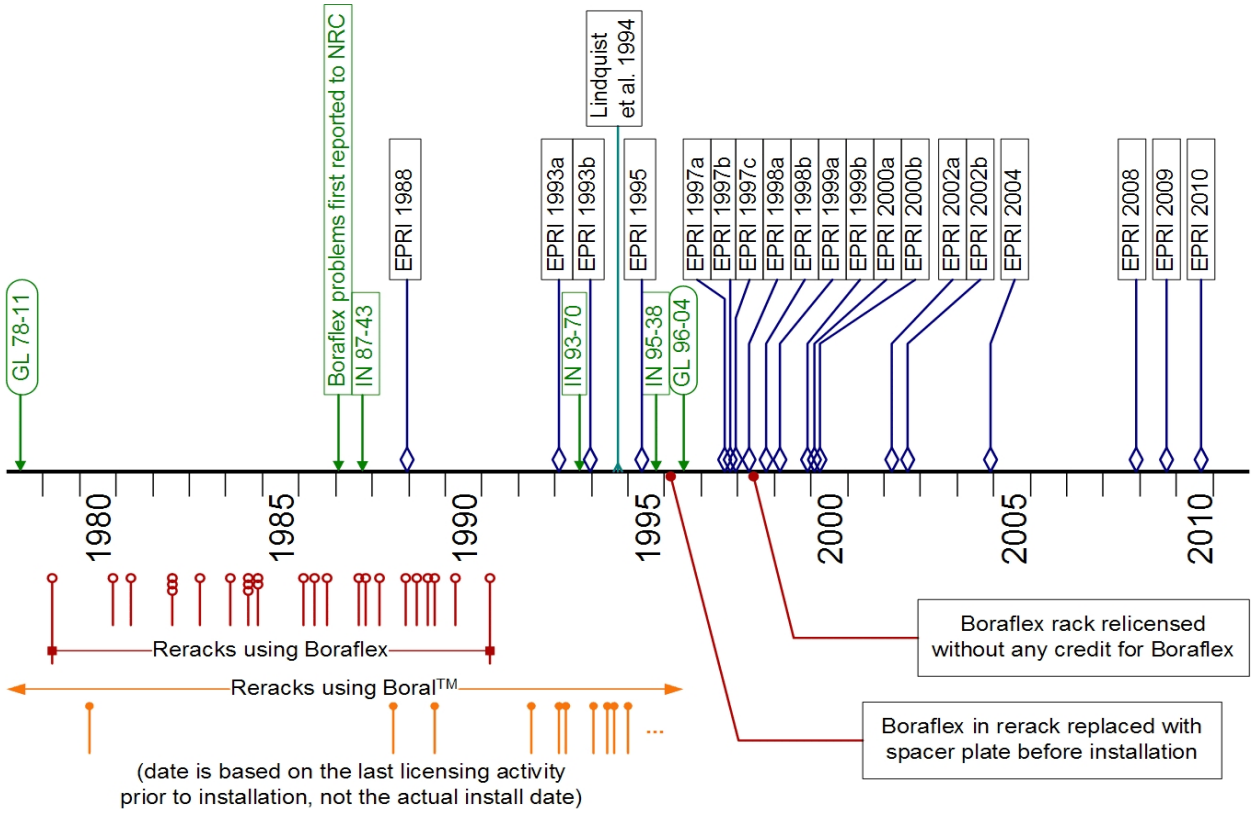

Figure 2–1. Timeline of some documented events in the history of Boraflex since 1978.

Industry action is represented in the timeline by Electric Power Research Institute (EPRI) publications directly related to Boraflex and its problems, indicated by the blue lines. In addition, a peer-reviewed article in *The Journal of Nuclear Materials* by Lindquist et al. is shown at the top. All of these publications are cited in full in the References (Section 7).

Below the time line are shown the approximate dates of various final licensing actions that led to the installation of a Boraflex rack. (Note that these are not the actual installation dates.) These dates occurred between 1979 and 1991. For comparison, licensing actions for Boral™ racks are also shown (only through 1994). It is noted that in 1996 a rerack designed to use Boraflex was installed after replacing the Boraflex with spacer plates. Beginning in 1998, some utilities began to relicense their racks with no subcriticality credit for whatever amount of Boraflex remained in

them. Subcriticality credit was assigned to storage patterns, burnup credit, and, in the case of PWRs, soluble boron in the pool water.

As noted at the end of Section 2.1, there was considerable uncertainty about the nature and extent of the Boraflex problem when it first emerged. Coupon findings could be shocking but bear no relation to actual rack conditions. "Large" (100 mm; 4 inch) gaps were found, but their implications for criticality safety were generally small and well within a rack's design margin to the $k_{eff} \leq 0.95$ limit. The fact that the first gaps observed were so nonuniform in their distribution, as noted above, pointed to a rack fabrication issue more than to a Boraflex material issue. NRC Information Notice 87-43 in 1987 emphasized the state of uncertainty but also highlighted the differences between racks with different manufacturers as had been done in some industry analyses.

BISCO moved forward to conduct some further testing of Boraflex. BISCO's test results appeared to confirm earlier suspicions that the observed gaps in Boraflex were due to shrinkage of the elastomeric component, followed by breakage of the elastomer concentrated near points of attachment of the Boraflex to the rack, where tensile stress was greatest. However, that shrinkage appeared to be limited to 2-2.5 percent (now known to be up to 4.1 percent), after which the material transformed to a "stable ceramic" (now known to dissolve). Other analysts expressed skepticism of BISCO's laboratory data and the conclusions drawn from them. Because of the lack of conclusive data, industry groups were reluctant to discontinue the use of Boraflex. Boraflex thus continued to be used in reracks. However, as evidence against the material mounted, reracks then shifted to using Boral.

In closing this chronology, it should be noted that Boraflex problems were actually observed as early as 1984. Dissolved silica levels in some spent fuel pools had been increasing, but at the time no one knew to attribute it to the dissolution of Boraflex (as described in Section 3.4). While the presence of silica incurred maintenance costs of tens of thousands of dollars per month for filters and demineralizer resins, utilities were generally able to maintain their reactor water silica content below the 1 ppm limit to avoid voiding the warranty on the reactor fuel. Thus, it was not until the 1987 NRC Information Notice and the 1988 EPRI report that the connection between the degradation of Boraflex and the appearance of pool silica was made in print and subsequently confirmed by laboratory testing (as discussed in Chapter 3.0).

2.3 Boraflex Rack Designs and Spent Fuel Pools

Boraflex is used in three basic types of rack designs—egg-crate, flux-trap, and hybrid (the latter used in only one spent fuel pool, subsequently reracked). Representative examples of these types are shown in Figure 2-2. Egg-crate racks, shown in Figure 2-2a, have one sheet of Boraflex between assembly storage cells and achieve the highest assembly storage density in the spent fuel

pool with an assembly pitch similar to that in the reactor. Note the density of the four assemblies, denoted by the diagonal lines, shown in each design. Egg-crate racks are used for BWR fuel and for PWR fuel with burnup credit. In PWRs, they are commonly referred to as "Region II" racks. The reactivity of lower burnup and fresh PWR fuel is such that additional neutron absorption is necessary. Flux-trap racks, shown in Figure 2-2b, achieve this by providing a water gap to

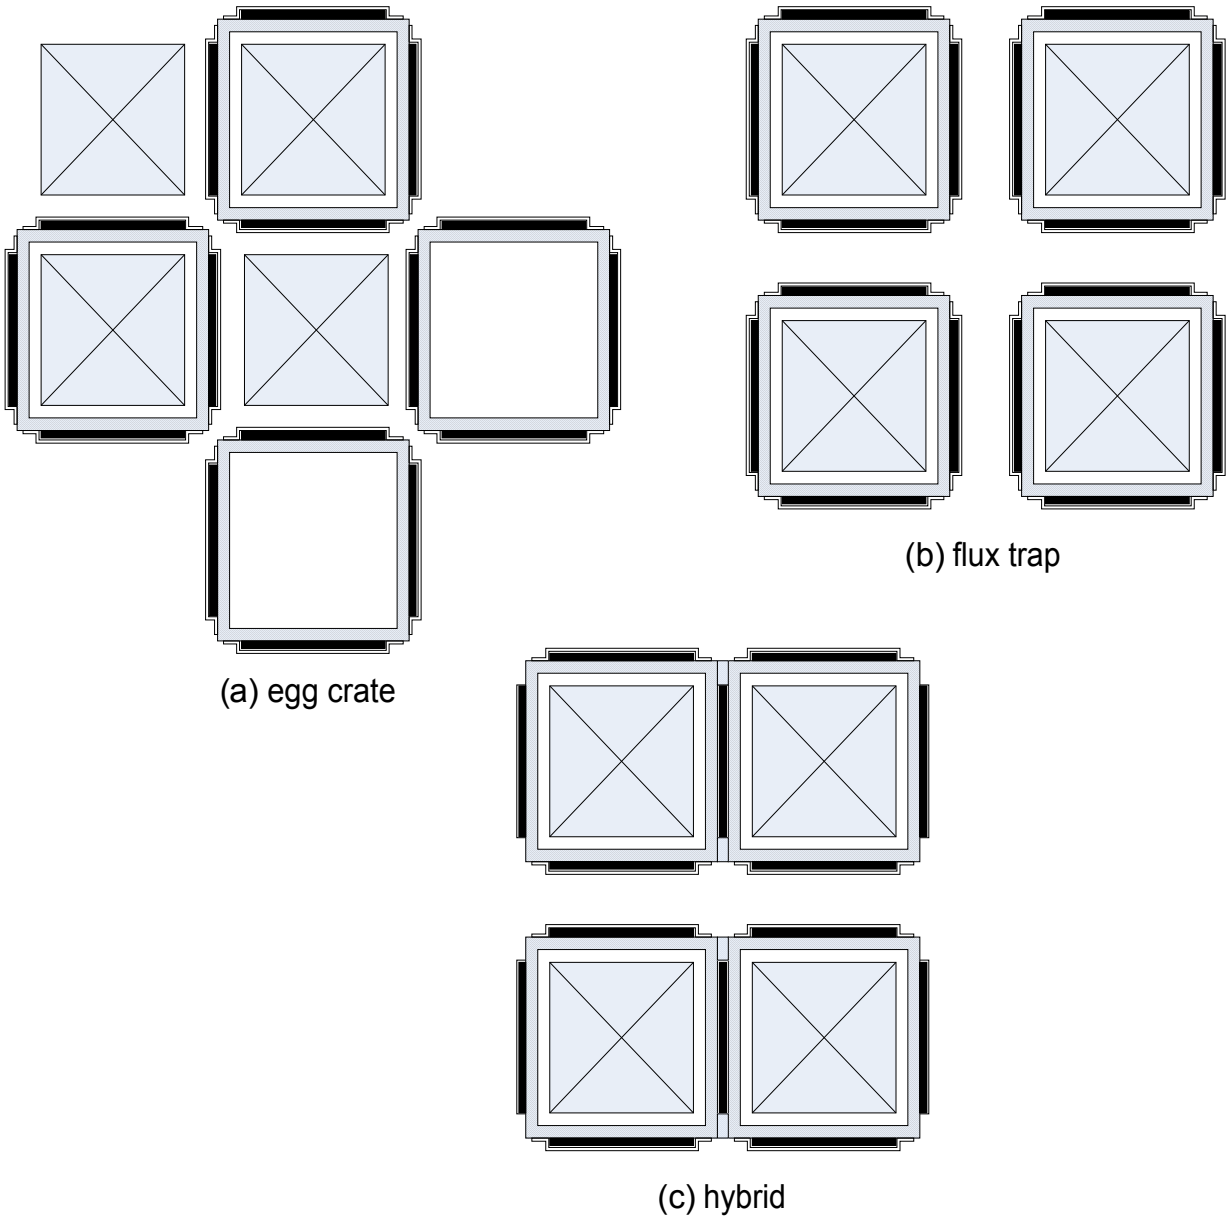

Figure 2–2. Representative examples of basic types of spent fuel storage rack designs that used Boraflex.

thermalize neutrons such that the neutron importance of the two Boraflex panels between assemblies—one on each side of the water gap—increases. These are referred to as "Region I" racks. Fresh fuel and low-burnup fuel to be returned to the PWR can safely be accommodated in Region I racks, and so this region of racks is typically near the fuel transfer canal for easy reloading. An intermediate solution, the hybrid rack, is shown in Figure 2-2c.

While these basic types are common to all spent fuel racks using Boraflex, the methods by which a Boraflex panel is held to a structural cell wall vary considerably between manufacturers. Some of the more common methods are illustrated in Figure 2-3. Because of these variations in designs, the dimensions and volumes cited below are for illustration only; while typical values are given, they can vary significantly.

Before proceeding, a note on terminology is in order. The fluid volume that an assembly sits in is referred to as a *rack cell*. This term is also often used inclusive of the rack walls around the cell. The fluid volume between absorber panels in a flux-trap design is referred to as a *flux trap water gap*. The fluid volume that surrounds a panel of Boraflex, to be detailed subsequently, is referred to as a *panel cavity*. In most designs, a panel of Boraflex is enclosed in its panel cavity by using a *cover plate*—a thin sheet of stainless steel—in conjunction with the structural stainless steel that forms the rack cell.

Figure 2-3a shows a wrapper-plate design where each panel of Boraflex is enclosed in its own panel cavity. A panel cavity is formed by using the structural stainless steel cell wall of the rack (e.g., 1.9 mm/0.075 inch thick) as a back wall, and welding on the thin (e.g., 0.5 mm/0.02 inch thick) stainless steel wrapper plate as the remaining sides. The Boraflex rests in the cavity in a clearance fit—the panel cavity is typically around 0.25 mm (0.01 inch) thicker than the Boraflex. For a typical wrapper plate design, this gives a cavity fluid volume (with a fully intact panel of Boraflex inside) of around 0.15 liters—about 2/3rds of a cup of water. As the Boraflex degrades, this volume will increase to around 1.5 liters (~6 cups) of water if the Boraflex were completely gone.

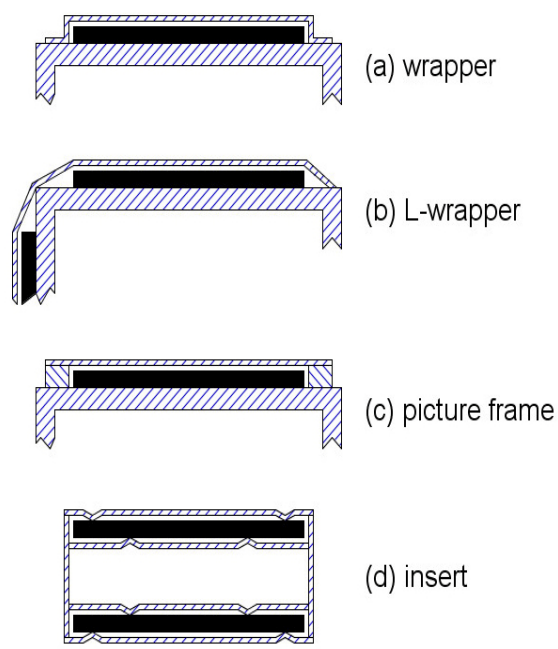

Figure 2–3. Examples of Boraflex retainer methods.

Figure 2-3b shows a variation on the wrapper design where a single L-shaped wrapper plate covers two panels on two

sides of each rack cell. Depending on the geometry in the as-built condition, this could mean that two panels effectively share a common panel cavity.

The method of fixing the Boraflex panel in place can influence the mode of degradation in that a different location of attachment point may cause different modes of stress on the panel or different flow paths in the rack cavity. Therefore, some discussion of attachment methods is warranted. One egg-crate design is to "sandwich" a sheet of Boraflex between two adjacent rack cell walls. Another sandwich retainer method is to use a picture frame as shown in the hybrid rack in Figure 2-2c. Figure 2-3c shows the use of stainless steel bar stock to create a picture frame to which a cover plate is attached. The flux-trap equivalent insert design shown in Figure 2-3d uses contact from V-notches rolled in cover plates to hold the Boraflex. Another method is to press protruding dimples into the walls to provide clearance for the Boraflex as well as attachment points to spot weld the walls together; the otherwise rectangular Boraflex panel has semicircular notches cut out of the vertical sides at various elevations to accommodate these dimples. Other rack designs are typically variations of the retainer methods described here.

In each design, gaps in the weld seams and/or vent holes in the cover plates are always present to allow an escape for off-gassing from irradiated Boraflex (discussed in Section 3.2). This permits a flow of spent fuel pool water into and out of the panel cavity, possibly from all sides. The lengths and spacing of the welds along the seams are relatively consistent within a given rack module and usually for all modules of a given type in a spent fuel pool, but they can vary significantly between designs in a pool (e.g., between Region I flux-trap racks and Region II egg-crate racks), as well as between pools at different plants. On occasion, a seam will be identified by the manufacturer as needing rework, which can change the nature of the seam welds. Because the cover plates are so thin, the manufacturing process can result in some small deformations of the plates. This can significantly increase or decrease the area for flow between the welds, significantly change the small volume of the panel cavity, and possibly even pinch the underlying Boraflex panels.

The Boraflex panels used vary in size from, typically, around

- 0.6 to 3.2 mm thick (0.025 to 0.125 inches, depending on the areal density of ^{10}B needed).
- 120 to 220 mm wide (5 to 8.5 inches for BWR and PWR fuel, respectively).
- 3560 to 3860 mm long (140 to 150 inches).

Measurements have shown that the thickness at various locations of a given panel can vary by up to 10 percent, although 5-percent variation or less is more typical. The "minimum certified areal density" of a panel accounts for this variation as well as for variations in the boron loading (as discussed in Section 3.1). Some rack designs somewhat accounted for the fact that Boraflex shrinks under irradiation by installing sheets longer than the active fuel length by 25 mm (1 inch) on each end. Other designs took credit for axial leakage effects without accommodating shrinkage and used Boraflex sheets shorter than the active fuel length. To hold the long flexible

sheets of Boraflex in place during manufacturing, sometimes a thin layer of silicone rubber adhesive was used to hold it to the rack wall.

While the space between fuel assemblies in higher and higher density racks has decreased, the reactivity of the fuel assemblies has increased. Increased fresh fuel charge enrichments, more reactive fuel assembly geometries (e.g., increased fuel pellet diameters and numbers of fuel rods), and more reactive materials (e.g., increased fuel pellet densities and burnable poison loadings) lead to both increased gamma exposure and increased Boraflex degradation, challenging the $k_{eff} \leq 0.95$ subcriticality limit. Rather than allowing Boraflex to degrade evenly among all rack modules thus rendering all rack modules nearly unusable for all fuel, some plants chose to deliberately control the degradation rates of the racks. A plant could "preserve" a rack module (i.e., minimize its gamma exposure and therefore slow its Boraflex degradation) by storing only older discharged fuel in the rack cells. A plant could "sacrifice" a rack module by storing recently discharged fuel in a checkerboard pattern (e.g., a bundle is stored every other rack cell, providing distance between bundles; other patterns using different proportions of vacant rack cells also have been used). Because subcriticality credit is provided by distance between bundles, the condition of the Boraflex becomes moot and is therefore allowed to degrade more quickly. By managing degradation in this manner, the plant could preserve a significant number of the usable cells and extend the overall life of the rerack.

In addition, a plant may design a spent fuel configuration that would enhance post-accident cooling and damage recovery. An example would be mixing recently discharged high-activity fuel (e.g., less than 2 years cooling time) with older lower-activity discharged fuel. Fuel configurations designed to address accident conditions could conflict with fuel configurations designed to preserve or sacrifice select rack modules.

In addition to criticality safety, the spent fuel pool design must provide shielding against the intense fission product decay radiation and a heat sink for the heat generated by the decay. The large volume of spent fuel pool water—typically over 1 million liters (300,000 gallons)—provides both the shielding and the heat sink. To decrease dissolved mineral contamination reaction rates with Zircaloy clad, the water is purified using a cleanup system that generally involves filters and demineralizers; these are typically in-line with heat exchangers for heat removal. This large pool volume also is the ultimate sink for water exchange with each of the panel cavities. The cleanup systems may remove some of the products of Boraflex degradation from this sink, as discussed in Section 3.6.

3 THE PHYSICO-CHEMICAL PROPERTIES OF BORAFLEX

This chapter reviews the physical and chemical properties of Boraflex. Much of the theoretical material in subsequent chapters relies on an understanding of the vocabulary and concepts presented in this chapter. While the active ingredient of Boraflex is the strong thermal neutron absorber Boron-10 (^{10}B) in the form of a boron carbide powder, the degradation of Boraflex is related to the elastomeric polymer polydimethyl siloxane (PDMS) used to bind the particles of boron carbide. Thus, PDMS is the dominant focus of this chapter.

The following references contain additional details on the physico-chemical properties of PDMS and silica in general and Boraflex in particular: [Iler 1979], [EPRI 1988, Appendices A and B], [O'Donnell 1989], [Lindquist et al. 1994], [EPRI 1997c, Section 2], [EPRI 1998b, Section 2.4], [EPRI 1999b, Section 3], [Graubner et al. 2004], and [Dove et al. 2008]. While some of the chemical phenomena described in this section are still not well understood (see, for example, [Dove et al. 2008]), field observations and experiments have generally agreed with expectations deduced from the literature.

3.1 The Composition of Boraflex and its Boron Carbide Loading

Boraflex is a lightweight and flexible neutron absorber. Its material composition is illustrated by the inner circle of Figure 3-1. By mass, Boraflex is typically a combination of

- ~50% boron carbide (B_4C, with ^{10}B as the strong 1/v thermal neutron absorber),
- ~25% the elastomeric polymer PDMS, and
- ~25% crystalline silica (SiO_2) filler. (Crystalline silica is inert and relatively insoluble, and should not be confused with soluble amorphous silica.)

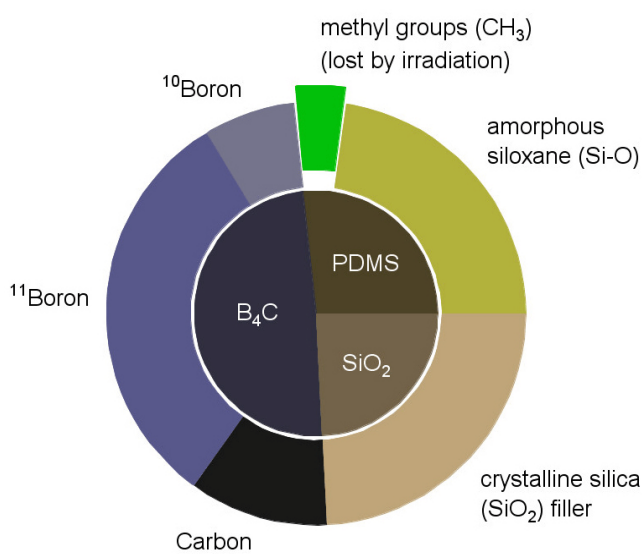

Figure 3–1. Material constituents of Boraflex. Mass proportions are for illustration and will vary among batches.

To manufacture Boraflex, B_4C is added as a finely divided powder to each part of a two-part liquid silicone elastomer "kit." These two liquid parts are then combined with a curing agent and thermally cured to form a flexible solid elastomer. The first liquid elastomer part is today referred to as Dow Corning Sylgard™ 170 Silicon Elastomer (Part A). It is a black liquid and is primarily a combination of dimethylvinyl-terminated PDMS, quartz (crystalline silica), and less than 1 percent carbon black. The second part is today referred to as Dow Corning Sylgard™ 170 Silicon Elastomer (Part B). It is a white liquid and is primarily a combination of PDMS, quartz, and 5 percent–10 percent dimethyl, methylhydrogen siloxane. These chemical species are discussed in detail in Section 3.2. As discussed in Section 3.7, the quartz is added as a filler and to increase the tensile strength of Boraflex. Material certification sheets that accompany Boraflex provide constituent specific gravities and mass fractions so that the Boraflex density, mass fraction of boron carbide, and mass fraction of amorphous silica (PDMS silica, exclusive of the quartz crystalline silica for reasons discussed subsequently) can be computed. Independent chemical testing has confirmed the accuracy of these sheets (see, for example, [EPRI 1993b, Section 2]).

Particle diameters of the finely divided B_4C powder can range from ~2 μm to ~200 μm although the middle 50 percent of the diameter distribution is typically between 30 μm and 55 μm. Typical particle diameters can be characterized as

~75 μm – volumetric mean; i.e., average of the particle-volume-weighted diameters
~45 μm – median; i.e., 50th percentile of diameter distribution.

In Boraflex, the distribution of B_4C particles—and thus the strongly neutron absorbing ^{10}B—is that of a microscopically heterogeneous system. Macroscopic testing of Boraflex samples has generally shown the distribution of ^{10}B to be very uniform and significantly above any stated minimum certified areal density.

In spent fuel rack applications, the natural B_4C loading varies between 40 percent and 50 percent, but most racks used proportions close to 50 percent. This value varies even during the manufacture of a single production run. Throughout the run, samples are tested to determine the margin above minimum certified ^{10}B content. Early in the run, more B_4C may be used to ensure that the contracted minimum certified areal density is achieved. Later in the run, especially as B_4C begins to run out, less B_4C may be incorporated than was used earlier in the run but enough to ensure that the B_4C content is still within the known margin above the minimum certified content. Each delivery of Boraflex would have come with batch specific material certification sheets that, in addition to stating the minimum certified areal density, also identified specific gravities and weight percent loadings of various materials such that the material constituents can be fully characterized.

When the formed surfaces of unirradiated Boraflex are viewed using scanning electron microscopy (see, for example, [EPRI 1999, pp. 3-28, 29]), they appear smooth with small but extensive surface cracks in the PDMS and some small bumps of ~2 μm diameter silica filler

particles evident. When cut edges are viewed [EPRI 1999, pp. 3-30, 31], the B_4C particles clearly protrude out of the surface of the PDMS matrix along with many crystalline silica filler particles.

3.2 The PDMS Binder and its Response to Radiation

The PDMS binds the B_4C particles in its polymer matrix, consisting of siloxane (Si-O) backbones and methyl (CH_3) side groups. The chemical representation of a strand of the PDMS matrix is shown in Figure 3-2. The actual helical spiral structure of the molecule is better shown in Figure 3-3.

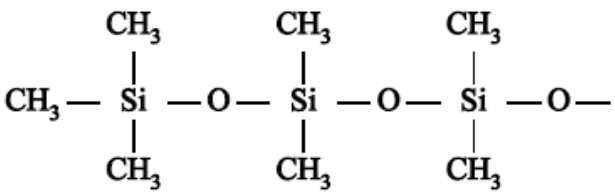

Figure 3–2. Chemical diagram for the left end of a PDMS polymer chain.

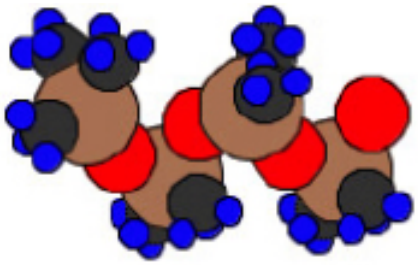

Figure 3–3. Helical PDMS molecule.

As noted in Section 3.1, the PDMS chains in Boraflex are dimethylvinyl-terminated. This means that each end of the chain is actually terminated with a vinyl group, $H_2C{=}CH{-}$, which cross-links during thermal curing to help convert the liquids into an elastomeric solid. Cross-linking, which is a process that connects neighboring polymer chains, is discussed in more detail below. Dimethyl methylhydrogen siloxane, which is a small but significant constituent of Elastomer 170 Part B, is identical to PDMS except that a small proportion of the methyl CH_3 side groups are replaced by hydrogen to also promote a small amount of cross-linking. The PDMS chains are typically many hundreds of monomer units long, and these variations in end groups and side groups are negligible to the overall chemical composition. When the two-part liquid elastomers containing the boron carbide powder are combined, a volatile curing agent is also added, typically in a 10:1 ratio of base to curing agent. The curing agent is almost completely driven off during the thermal curing process and so is also negligible to the overall chemical composition.

When PDMS is placed in water, it is hydrophobic—like most elastomers, its generally very low level of intermolecular interaction at the surface repels water. When PDMS is placed in a radiation field, three distinct chemical processes occur:

- Abstraction within or of a methyl group (breaking off H and CH_3 radicals).
- Cross-linking between polymer chains (connecting neighboring chains).
- Scissioning of polymer chains (breaking chains).

See [EPRI 1988, Appendix B] for a more thorough review of the literature on the effects of radiation on polymers in general and PDMS in particular.

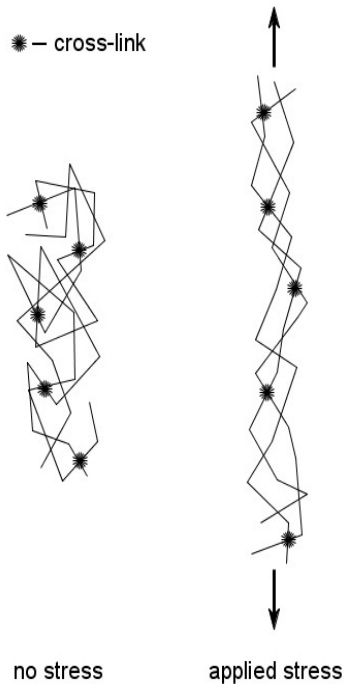

Figure 3–4. An elastomer responds to stress by elastic deformation.

Many polymer chains can slide past each other relatively easily, making the material like a fluid when a stress is applied. This is evident in the more fluid forms of PDMS such as spreadable caulks. Cross-links—atomic bonds between polymer chains—prevent local slipping and help the material hold its shape. A small amount of intentional cross-linking during manufacture gives elastomers such as Boraflex the characteristics of a solid, including the combination of high elasticity and high-yield strain. As shown in Figure 3-4, applying an axial mechanical stress to an elastomer will cause its amorphous tangle of polymer chains to slide so as to distribute the stress. This allows the material to lengthen in the direction of the stress, while it thins somewhat in the directions perpendicular to the stress. It lengthens easily because the silicon-oxygen (Si–O) bonds along each chain are very flexible and can easily change their angles, as shown in Figure 3-5, and even change their lengths. The cross-links prevent the chains from completely slipping apart as they might in a more fluid state. When the stress is removed, the cross-links serve as a "memory", causing the material to elastically return to its original shape.

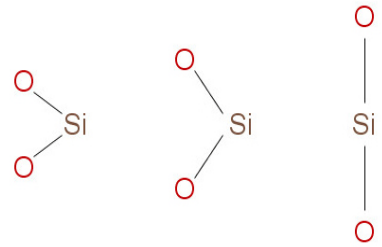

Figure 3–5. Flexible O-Si-O bonds open and lengthen under stress.

Gamma radiation will induce cross-linking in PDMS. In the more likely cross-linking mechanism, shown in Figure 3-6, the radiation abstracts a hydrogen radical (H·) from a methyl group, leaving a methylene radical (CH_2·) on the polymer chain. The unpaired electron in the methylene radical can migrate along the polymer chain. At some point, it may come in close proximity to another

radiation-induced methylene radical on a neighboring chain and form a covalent bond with it, thereby cross-linking the polymer chains at that point. This is shown on the right side of Figure 3-6.

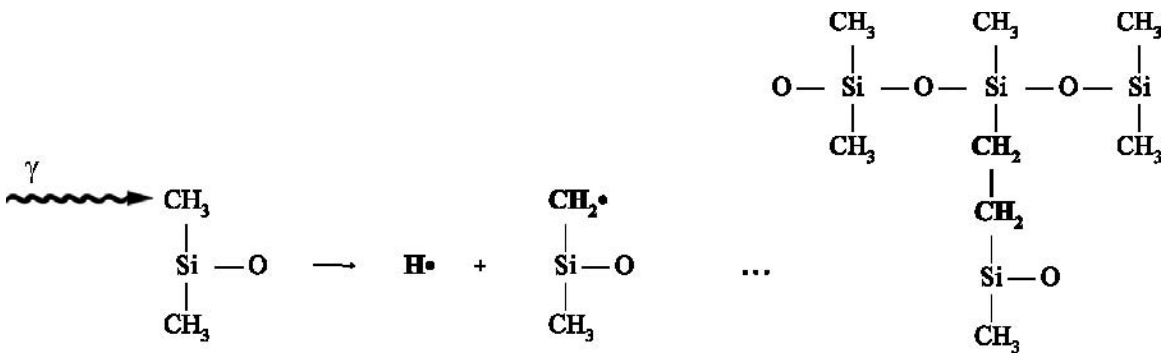

Figure 3–6. Radiation-induced abstraction of a hydrogen radical allows cross-linking with an adjacent polymer chain that also has an abstracted hydrogen.

Figure 3-7 shows an alternative mechanism in which the radiation abstracts the entire methyl group (CH_3·), leaving behind a silicon radical (Si·). If this encounters another silicon radical on a neighboring chain, an Si−Si cross-link will form as shown on the right side of Figure 3-7.

The abstracted H and CH_3 radicals from both cross-link mechanisms combine with each other to form hydrogen (H_2), methane (CH_4), and ethane (C_2H_6)—all colorless, odorless gasses. In a vacuum, these gasses evolve roughly in the proportions of 41 percent H_2, 47 percent CH_4, and 12 percent C_2H_6 [EPRI 1988, p. B-5]. Gasses that are trapped in the polymer matrix can produce local swelling stresses that result in crazing and cracking of the Boraflex. When an irradiated Boraflex spent fuel rack is disturbed (e.g., during BADGER testing, when the BADGER probe is being inserted into a cell with a tight fit), small bubble trains will sometimes be seen venting from the panel cavities. Sometimes bubble trains can be observed in quiescent pools from racks with freshly discharged fuel, which may indicate a panel cavity relatively open to pool water ingress.

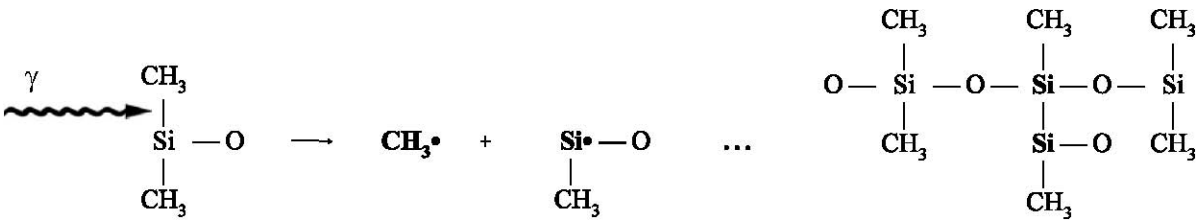

Figure 3–7. Radiation-induced abstraction of a methyl group allows cross-linking with an adjacent polymer chain that also has an abstracted methyl group. Note that the lower pendant arm is likely perpendicular to the flat page.

With larger absorbed radiation dose comes more cross-linking, which will increasingly change the material properties of the polymer (e.g., elasticity and yield strain as discussed above, and density and hardness as discussed below). However, at some dose level the cross-linking will saturate. Saturation can occur with as low as 14 percent cross-linking. This is because cross-links form barriers to free electron mobility, which prevents subsequent radicals formed by abstraction from migrating along the polymer chain to close proximity with another abstraction site on another chain.

Another effect of gamma radiation on PDMS is chain scissioning—breaking a Si–O bond in the siloxane backbone. This produces chain fragments of lower molecular weight that lower the strength of the material. While cross-linking tends to enhance material strength by turning chains into networks, scissioning will tend to break that network into smaller molecules. Scissioning can result in two highly reactive backbone radicals (e.g., Si•, O•) in close proximity. If the chain fragments have limited mobility (e.g., due to a high degree of nearby cross-linking), then recombination is likely.

For PDMS in particular, at low to moderate gamma doses, the rate of cross-linking is much higher than the rate of scissioning and especially nonrecombinant scissioning. However, nonrecombinant scissioning is greatly enhanced by the presence of oxygen, including dissolved oxygen from an aqueous environment. Free oxygen (O_2) from air or in solution can dissociate to combine with silica radicals (produced by scissioned siloxane bonds), forming amorphous silica (SiO_2). (This process will be detailed further in Section 3.5 where the effects of water at the surface are included.) Passivating the siloxane silica radical with oxygen swells the material and makes the scission site nonrecombinant. Thus, scissioning in water can significantly increase PDMS deterioration. Moreover, unlike cross-linking, scissioning does not saturate. Thus, some rate of material degradation will continue with increasing gamma radiation dose even after cross-linking effects have saturated.

As polymers undergo structural changes due to irradiation, it is common for them to change color. It has been observed that as Boraflex is irradiated, it will change from shiny black to dark dull gray to lighter gray. However, this color change is likely due more to the interaction of irradiated Boraflex with the surrounding water of the spent fuel pool. Based on the above, the material effects of radiation on Boraflex can be divided into five dose ranges as shown in Table 3-1.

Table 3-1. Effects of Radiation Dose on PDMS in Boraflex.

Dose [Gy]	Dose [rads]	Cross-linking	Effects on material properties
$<1\cdot10^4$	$<1\cdot10^6$	elastomeric levels	negligible effects
$1\cdot10^4 – 5\cdot10^6$	$1\cdot10^6 – 5\cdot10^8$	increasing	increasing Shore A hardness
$5\cdot10^6 – 1\cdot10^8$	$5\cdot10^8 – 1\cdot10^{10}$	increasing	increasing Shore D hardness, density, and shrinkage
$9\cdot10^7 – 1.5\cdot10^8$	$9\cdot10^9 – 1.5\cdot10^{10}$	becomes saturated	fully hard and brittle, ceramic-like; no further densification or shrinkage
$>1\cdot10^8$	$>1\cdot10^{10}$	saturated	continued scissioning degradation, including increased porosity

As discussed earlier, some level of cross linking is necessary for elastomers; Table 3-1, adapted from data found in [EPRI 1988], shows that at a low absorbed dose PDMS retains its elastomeric properties. As dose increases, the cross-linking impedes the motion of the polymer chains to the extent that the material loses its elasticity. This can be observed as an increase in a Shore A durometer. Above about $5\cdot10^6$ Gy ($5\cdot10^8$ rad), more significant material changes occur. As large numbers of hydrogen and methyl group abstractions occur, the off-gassing reduces the overall mass of a Boraflex panel. However, shrinkage due to increased cross-linking reduces the volume of the panel even more than the mass, such that the material density increases. Finally, at a threshold of about $1\cdot10^8$ Gy ($1\cdot10^{10}$ rad), the number of cross-links saturates; the material is hard and brittle, like a ceramic. Further dose does not significantly increase the amount of cross-linking but does continue to cause scissions, which contribute to continued degradation of the Boraflex material throughout its service life.

These effects are all temperature dependent. At higher temperatures, unpaired electrons are more mobile, which increases migration rates and thus increases cross-link rates. The recombination or coordination reactions of scission radicals also will occur more rapidly at higher temperatures.

3.3 Stresses that Lead to Gap Formation

The rate of PDMS deterioration under irradiation can be accelerated by application of external stresses. External stresses are introduced when the shrinkage of PDMS due to cross-linking is combined with a mechanical point of constraint on the panel. Examples of constraints on the panel include:

- Cover/wrapper plate pinching during manufacturing (recall Section 2.3).
- Residual adhesion from adhesives used in manufacturing (recall Section 2.3).
- Swelling-induced hang-up – a constraint on the panel resulting from local swelling of portions of the Boraflex panel.

While in general the total material volume decreases as gamma-promoted cross-linking shrinks the panel, the nonuniform migration of unpaired electrons that cause cross-linking, and the evolution of gas, can lead to local swelling. As discussed subsequently in Section 3.4, contact with free oxygen in water will cause additional swelling. The small clearance fits for the Boraflex in its panel cavity, combined with small deformations of the wrapper plates, can lead to the swelled Boraflex wedging inside its cavity and causing a constraint and external stress on the panel material.

These stresses are exacerbated by the practice of delivering new Boraflex in rolls. The inside of the roll, with its smaller radius of curvature, will have larger bending stresses than the outside. Unless the Boraflex is unrolled and allowed to relax for some time prior to mounting with adhesive, an axial gradient of stresses will be built-in along the length of the panel.

An unrestrained panel of Boraflex will shrink axially, possibly enough that no absorber is present to absorb neutrons emitted from the ends of the spent fuel rods. If there is sufficient margin below the criticality limit, then the reactivity penalty for this "end shrinkage" is relatively small, one estimate being $\Delta k \approx 0.004$ for 4 percent end shrinkage concentrated at one end [EPRI 1993a, p. 4-19]. As noted in Section 2.3, some rack designs take advantage of axial leakage and use a Boraflex panel that is shorter than the length of the fuel; for these designs, the effects of end shrinkage will be larger. The panel will also shrink in the lateral "width" direction. Because of the synergistic reactivity effects of width shrinkage and gaps, width shrinkage reactivity effects are typically included in analyses of the effects of gaps. The result presented below for gaps includes the effect of width shrinkage. Shrinkage in the thickness direction has a negligible effect on reactivity.

As the PDMS hardens and shrinks with cross-linking and loses strength with scissioning, it becomes prone to brittle fracture. The presence of built-in stresses, combined with the mechanical constraint mechanisms (pinching, adhesion, and hang-up) and shrinkage described above, can cause gaps to form (instead of or in combination with some end shrinkage). Figure 3-8 illustrates a typical 3,660 mm (144 inch) panel with end shrinkage at the top and gaps as indicated by the bands of white across the panel. The end shrinkage lowers the top by about 25 mm (1

inch). About one-quarter down the panel is a 32 mm (1.25 inch) gap, two approximately 9 mm (1/3 inch) gaps are near the midplane, and a 57 mm (2.25 inch) gap is near the bottom. (The other features of this figure will be discussed in detail in Section 3.8; this figure is adapted from BADGER data found in [EPRI 1997b], as also described in Figures 5-4 and 5-5 in Section 5.2.) The shrinkage and gaps are seen to be lateral, as would be expected in unrestrained shrinkage and restrained tensile fracture. Slightly angled gaps, indicative of shear stresses, have been observed, but are rare. The widening of the top gap on the right side, as well as the patch of missing Boraflex above it, are dissolution effects as discussed in Section 3.8.

Note that neither the formation of gaps nor end shrinkage indicates a loss of material (other than the inconsequential off-gassing from scissioning)—it is simply a redistribution of material in a densified form. BADGER testing has shown that Boraflex with only shrinkage-induced gaps has a higher ^{10}B areal density than its as-built certification. However, the ^{10}B in the degraded panel is found only in the densified Boraflex, while the gaps are filled with moderating water and provide an unattenuated "line-of-sight" between fuel assemblies. From a criticality safety standpoint, the reactivity increase due to gaps exceeds the reactivity decrease due to densification.

As gamma dose increases so does shrinkage, and thus gaps, once formed, can grow. However, because cross-linking saturates, so does gap initiation and growth. Thus, the maximum total gamma radiation-induced gapping expected in a Boraflex panel is about 4.1 percent of its axial length. This is independent of the forming direction of the Boraflex sheet when fabricated. (See [EPRI 1997c, Section 2.1] for further details on this and subsequent points.) Larger "gaps" have been observed but, as discussed in Section 3.8, these are dissolution effects and not radiation-induced shrinkage effects.

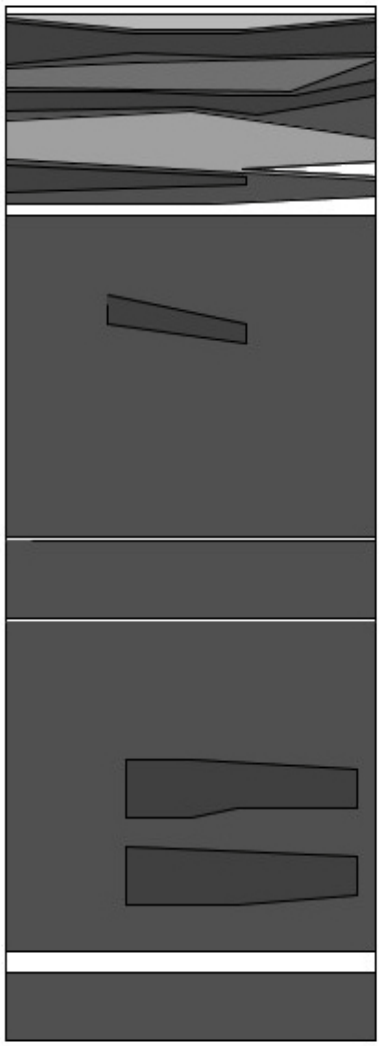

Figure 3–8. Illustration of a panel with top end shrinkage and gaps (white areas). Adapted from EPRI Report TR-107335 (1997), with permission from Electric Power Research Institute.

From a few panels of Boraflex examined by destructive visual inspection and thousands by nondestructive blackness testing and BADGER scanning, the following conclusions are made with respect to radiation induced gaps [EPRI 1997c, pp. 2-6, 2-11].

- The largest gamma radiation-induced cumulative gap plus end shrinkage will be less than 4.1 percent at saturation.
- The largest single gap will be less than 115 mm (4.5 inch); this is 2.8 percent shrinkage in a typical 3,660 mm (144 inch) panel.
- Many panels have multiple gaps, and most of these gaps are 25 to 50 mm (1 to 2 inches).
- The distribution of gaps is chaotic, but not completely random; small differences in manufacturing and dose history can change the distribution, but the distribution is largely deterministic.

The reactivity penalty for gaps can be substantial. A bounding scenario with one 100 mm (4 inch) gap at the midplane of every Boraflex panel gives $\Delta k \approx 0.037$ [EPRI 1993a, Figure 4-9 p. 4-22]. In a scenario specific to an observed distribution of gaps from blackness testing where every panel was still assumed to form a 100 mm (4 inch) gap but now distributed axially such that the gaps are neutronically decoupled, the penalty is $\Delta k \approx 0.011$ [EPRI 1993a, Figure 4-8 p. 4-18]. However, if testing confirms that less than 100 percent of the panels have gaps, and/or if the gaps observed are smaller than the bounding value, then the reactivity penalty can often be accommodated in available design margin to the regulatory limit.

None of the mechanical constraint mechanisms (pinching, adhesion, and hang-up) noted at the beginning of this section has been observed to be necessary or sufficient for gaps to develop with gamma dose. For example, regarding adhesion, a low dose rate exposure may "burn off" the adhesive before panel gaps would begin to form, or the adhesive could yield before the panel does. Panels fixed with adhesive that exhibit only end shrinkage (no gaps) have been observed as have panels with gaps where adhesive was not used.

3.4 Irradiated PDMS in Water

Abstraction and scission ends of PDMS which do not recombine will scavenge nearby species. In particular, at the surface of the Boraflex where it contacts the spent fuel pool water, vacant PDMS coordination sites will attach to hydroxide ions (OH^-) in the water, forming silanol (SiOH) groups. As shown in Figure 3-9, the outside surface will develop a layer of hydroxyl groups (OH) at the Boraflex-water interface. The gray line represents the surface of the Boraflex interfaced to the spent fuel pool water.

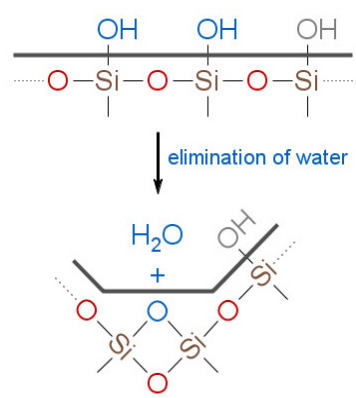

Figure 3–9. Hydroxide surface and conversion of siloxane to amorphous silica.

In a reversible reaction, a water molecule can be shed from the surface to create what is effectively a local silica molecule (i.e., the oxygen to silicon ratio O/Si, which is initially 1 for pure

siloxane [SiO], increases toward 2 [pure amorphous silica, SiO_2]). This is illustrated at the bottom of Figure 3-9. As shown, this process will locally deform the surface, leading to local swelling.

The presence of the silanol groups on the surface transforms the Boraflex from hydrophobic to hydrophilic, thus allowing water to wet and ingress through pores. This greatly enhances the migration of free oxygen (dissolved in the water) into the interior of the Boraflex. Because free oxygen increases cross-link and scission reaction rates and greatly decreases the probability of scission recombination, the progression of this surface change results in an increased overall rate of material degradation.

As the methyl groups in the PDMS are abstracted, what remains are cross-linked and scissioned amorphous siloxane backbones, which is effectively a fragmented network of Si–O bonds. Thermodynamically, in the presence of the excess oxygen, the siloxane would prefer to reorient itself into a network of three dimensional silica tetrahedrons – the crystalline structure of quartz – as shown in Figure 3-10a. However, because of the initially amorphous structure of the PDMS polymer strands and the irregular cross-linking and scissioning process (particularly in blocking migration of unpaired electrons), reorientation is too energy-intensive. Instead, the material becomes an amorphous mixture of siloxane and silica, as illustrated in Figure 3-10b, with O/Si ratios between 1 and 2, but tending toward 2. As can be seen, even amorphous structures can show local structural order over lengths measured in nanometers, but there is no long-range order.

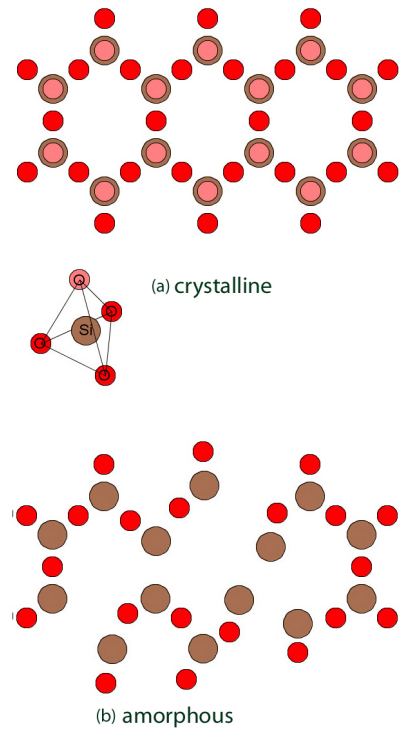

Figure 3–10. Crystalline silica and amorphous silica and siloxane structures.

As Boraflex transforms into amorphous silica, it becomes increasingly prone to dissolution in water. Figure 3-11 shows an irregular surface of Boraflex (the gray line) with silanol groups at the interface between the water and the mixed amorphous silica/siloxane network. Also shown is an additional ambient hydroxide ion that has adsorbed to the PDMS surface. To do so, it has increased the coordination of a silicon atom above four, thereby weakening its other bonds. In a reaction still not fully understood in the current literature (e.g., see [Dove et al. 2008]), three local water molecules (H_2O) participate to extract the silicon atom (Si) into solution as a silicate ion ($Si(OH)_5$). The reaction rate is relatively slow but persistent since the silicate ion will hydrolize in the water to become monosilicic acid and another hydroxide ion,

$$Si(OH)_5^- \rightarrow Si(OH)_4 + OH^-$$

As shown in Figure 3-12, the released hydroxide ion can effectively catalyze another abstraction of Si from the Boraflex surface.

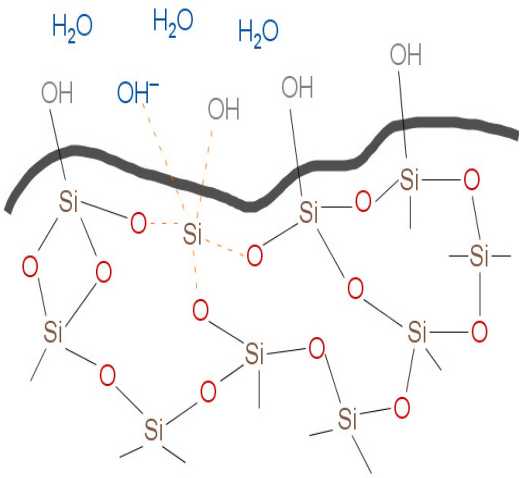

Figure 3–11. An amorphous silica and siloxane structure with a hydroxyl layer at the material interface.

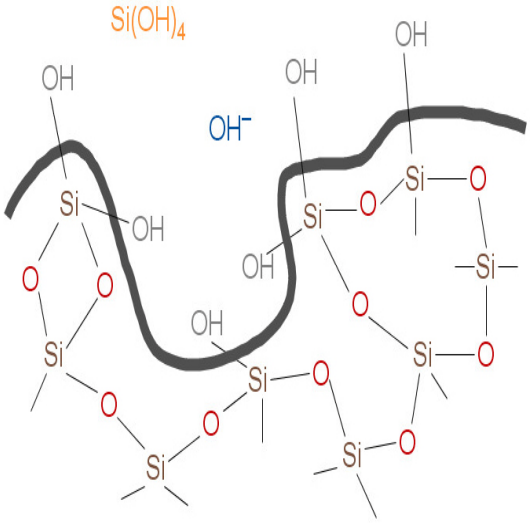

Figure 3–12. The hydroxide has successfully catalyzed the extraction of the silicon atom into solution as monosilicic acid.

The hydroxyl surface layer also is reformed in the process. As illustrated in Figure 3-12, the new surface has smaller radii of curvature, and nearby silicon atoms can now coordinate to two surface hydroxyl groups—effects that increase the dissolution reaction rate. These silicon atoms are removed more quickly, causing the dissolution cavity to spread. This mechanism of more rapid surface degradation continues until (locally) the whole monolayer is removed and a smoother surface is established with singly coordinated hydroxyl groups. This somewhat more resistant surface then dissolves at the slower reaction rate of the original adsorption and coordination of a hydroxyl ion shown in Figure 3-11, which, upon reaction, then proceeds to "peel away" the next monolayer. This is sometimes referred to as the "onion skin" model of dissolution (see, for example, [EPRI 1993b, p. 2-2]). Because cross-linking makes the siloxane structure stronger, higher levels of cross-linking slow the rate of hydroxide ion attack, making the material somewhat less soluble. By breaking up the siloxane network structure, scissioning makes the material more soluble.

One method to mitigate PDMS dissolution is to passivate the surface against hydroxide ion attack. This was explored using zinc and zirconium ions in solution to adsorb onto the surface and repel the hydroxide ions, as described in [EPRI 1999b]. However, field tests were disappointing. Subsequent analysis showed that, for the metal ions to fully adsorb, the pool pH must be carefully controlled at conditions contrary to typical operations (e.g., pH well below neutral in BWR pools

and pH above neutral in PWR pools with dissolved boron as boric acid). Moreover, the presence of monosilicic acid can cause metal-silicates to form, precipitating out the metal ions before they can adsorb. Although cleanup systems can reduce the monosilicic acid concentration in the bulk pool water, the concentration can remain higher in the panel cavities as dissolution continues.

3.5 Silica Solubility and Polymerization

The solubility of a silica/siloxane network in water can vary widely depending in particular on the following factors:

- particle size – solubility increases exponentially as particle diameter decreases
- temperature – solubility increases almost linearly as temperature increases
- pH – solubility is nearly constant between pH 7 and pH 8
  – solubility increases almost linearly as pH decreases below 7
  – solubility increases almost exponentially as pH increases above 8
- bulk impurities – solubility is likely affected, but little data are available on how much and in what direction

For large-scale siloxane systems like irradiated Boraflex, "particle size" is better measured as the amount of porosity and the number of surface silanol (SiOH) groups at the siloxane/water interface, which are indicators of local radii of curvature at the surface. Laboratory measurements of irradiated Boraflex ([EPRI 1997a]) and extrapolations based on fits from the literature ([Iler 1979]) give solubilities such as those shown in Table 3-2.

Table 3-2. Equilibrium Si concentrations from irradiated Boraflex in water for select temperatures and pH.

	25°C (77°F)	38°C (100°F)	60°C (140°F)
pH 5	80 ppm	100 ppm	140 ppm
pH 7 - 8	65 ppm	81 ppm	115 ppm
pH 9	71 ppm	88 ppm	125 ppm

As the concentration of monosilicic acid ($Si(OH)_4$) increases due to dissolution, the "proximity to equilibrium" will reduce the dissolution rate. However, dissolution is not a reversible equilibrium. As monosilicic acid molecules encounter each other in solution, they can combine to form a polymer that removes the silicic acid product from the aqueous system. This lowers the proximity to equilibrium, driving the reaction to the right and allowing additional silicon atoms to dissolve. Therefore, as monosilicic acid concentration increases, so will the amount of

polymerization; a stable equilibrium will never be achieved to limit the dissolution reaction. There is, however, a quasi-equilibrium where the rates of dissolution and polymerization are in equilibrium. The polymerization reaction is reversible, so if the concentration of monosilicic acid decreases, the polymers will depolymerize until quasi-equilibrium is restored. Figure 3-13 illustrates how monosilicic acid molecules in solution—effectively silicic acid monomers—can combine to form a polymer:

$$2\ Si(OH)_4 \rightleftarrows (HO)_3SiOSi(OH)_3 + H_2O$$

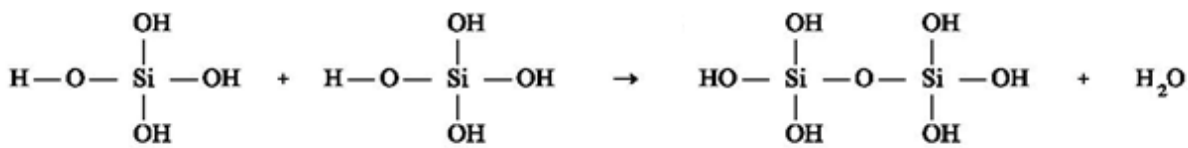

Figure 3–13. Silicic acid monomers condense to form a silicic acid dimer plus a water molecule.

Similarly, a monomer and a dimer can combine to form a trimer,

$$(HO)_3SiOSi(OH)_3 + Si(OH)_4 \rightleftarrows [(OH)_2SiO]_3 + H_2O$$

then a tetramer $[(OH)_2SiO]_4$, and so on. Note the two hydroxyl groups, $(OH)_2$, in the formulae for the trimer and tetramer, as opposed to the three groups, $(OH)_3$, for the dimer. This is because the polymers form by trying to maximize the number of siloxane (Si-O) bonds and minimizing the number of silanol groups (SiOH). This is best achieved by forming ring structures, as shown in Figure 3-14, and larger 3D structures for higher molecular weight polymers.

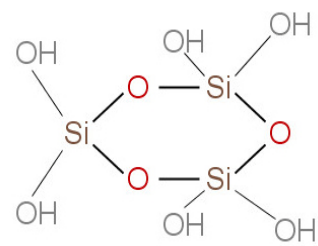

Figure 3–14. Silicic cyclic trimer.

As the polymers in water grow and combine, they effectively form particles and are called silica colloids. Because colloidal silica particles retain the same ratio of water to silica at all scales, the refractive index and the density of the colloidal system is effectively uniform no matter how large the polymer particles become. The presence of these polymer particles is therefore virtually impossible to visually observe in water.

As the colloids form and grow, they can become quite stable against depolymerization by complexing with other molecules in the water. For example, even small silicic polymers have an affinity for dissolved aluminum hydroxide at least a million times larger than the affinity of a silicic monomer (monosilicic acid), and the complex is highly resistant to depolymerization. The decrease in the depolymerization reaction rate prevents a return to quasi-equilibrium of the

dissolution reaction, as discussed above, further promoting the dissolution of silica from irradiated PDMS in the Boraflex panels.

3.6 Measuring Aqueous Silica Concentration

The $Si(OH)_4$ monosilicic acid molecule is often referred to simply as soluble silica. It is also called "reactive silica" because it quickly reacts with ammonium molybdate at low pH, forming a yellow silicomolybdic acid; the intensity of the yellow color is proportional to the concentration of silica in solution. The sensitivity of the colorimetric measurement can be increased by further reducing the solution with ascorbic acid to produce "molybdenum blue." This analysis technique is codified by "ASTM D859 - 10: Standard Test Method for Silica in Water."

The molybdate colorimetric test measures only reactive silica—the monosilicic acid molecule. The remaining polymerized silica is loosely referred to as either "colloidal silica" or "unreactive silica" and is not detected in the colorimetric test. The sum of the reactive silica and colloidal silica is called "total silica." By treating the silica solution to a hot sulfuric acid digestion, the colloidal silica can be depolymerized to reactive silica, and the molybdate blue procedure can then be used to measure total silica. The difference between the molybdate tests for total silica and for reactive silica then calculates the colloidal silica in the original solution.

The molybdate blue procedure is by far the most common method for testing silica concentrations in spent fuel pools as it can easily be performed in an onsite chemistry lab. When performed carefully, the uncertainty in a reactive silica measurement can be as low as 6 percent, although this is difficult to achieve. The low pH needed for the ammonium molybdate reaction means that more and more of the small polymer particles will depolymerize over time, increasing the apparent amount of reactive silica the longer the test takes. For more accurate measurements, and for easier measurements of total silica, inductively coupled plasma spectroscopy (ICPS) is used, with an accuracy as low as below 1 percent. However, the expensive and complex test equipment is usually only found in specialty labs, to which a spent fuel pool chemistry group would send samples. ICPS measures total silica; to measure only the reactive silica, the sample is first ultra-filtered to remove the colloidal components. The distribution of colloidal silica particle sizes can be determined by testing filtered samples from a distribution of filter efficiencies. ICPS was used in experimental analyses of Boraflex dissolution (see, e.g., EPRI 1997a, Figure A-5 p. A-9]), but was rarely used to measure spent fuel pool silica. A review of "The LWR Plant Water Silica Database" (Revision 4) shows that, of 51 spent fuel pools examined (43 with Boraflex), only two reported total silica data (see [EPRI 2000a, p. 2-1]).

Spent fuel pool cleanup systems respond differently to different forms of silica in solution. Colloidal silica acts more like a solid than a dissolved ion and so is filtered easily but will pass through ion-exchange demineralizers. Reverse osmosis (RO) systems are typically at least 99 percent efficient at removing colloidal silica, although it can cause colloidal fouling of the RO membrane. Reactive silica will pass through a filter, and RO systems are only around 90 percent

efficient in removing it. Reactive silica has slight ionic characteristics, but it is not counted as an anion in water analyses—it is part of the total dissolved solids (TDS). Thus, ion-exchange demineralizers can remove it to some extent: with about 45 percent efficiency in BWR pools and only 0.5 percent efficiency in PWR pools with resins optimized for not removing boric acid. The low efficiencies result from the selectivity of monosilicic acid being less than one. This means that not all molecules of monosilicic acid will be exchanged for hydroxide ions on the ion-exchange resins. In addition, almost any other anion will exchange out the monosilicic acid molecule. As the ion exchange resins become full, silica breakthrough is the first to occur, and complete and proper regeneration is necessary to flush out all of the reactive silica removed from the bulk pool water. Finally, combination powder resin filter demineralizers generally will not remove any reactive silica. However, they are typically around 50 percent efficient at removing colloidal silica.

3.7 The Non-PDMS Components

Figure 3-1 shows that a crystalline silica (SiO_2) filler is added to Boraflex. When fine crystalline silica particles are hydroxylated—that is, when one or more OH groups are bonded to the crystalline silica particle's surface—the hydrogen tip of the OH group can form a hydrogen bond to an oxygen atom along the amorphous silica PDMS backbone. This causes an order of magnitude reduction in slipping between local PDMS chains that can increase the tensile strength of PDMS by a factor of 30.

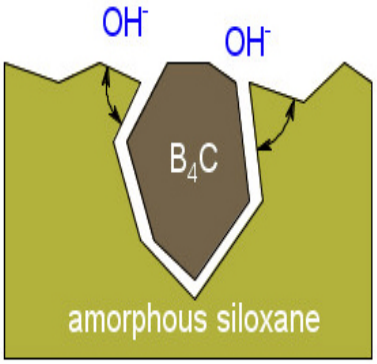

Figure 3–15. Acute contact angles give small radii of curvature at PDMS-particle interfaces that can enhance local dissolution.

As noted in Section 3.5, crystalline silica, like ordinary sand, has a very low solubility. Thus, spent fuel pool measurements of dissolved silica concentration correlate only to the dissolution of the amorphous silica/siloxane network that was the PDMS matrix. While not a contributor to Boraflex degradation, the crystalline silica introduces its own unique problems. The crystalline silica particle diameters in Boraflex are on the order of 2 μm. This makes them easy to entrain in a flow of water, causing turbidity problems. During refueling outages, when assemblies are being moved and natural and forced convection flows are at their maximum, the entrained high reflectivity crystalline silica can turn the water a milky white, obscuring the visual identification of fuel assemblies and cell locations. In a few cases, this has led to temporarily suspending refueling operations (for further details, see [EPRI 2004]). Crystalline silica has no impact on the removal of reactive or colloidal silica except to the extent that cleanup systems become loaded or fouled.

When a boron carbide particle is embedded in PDMS and both surfaces are exposed to the pool water (e.g., at a cut edge), then, as shown in Figure 3-15, the PDMS-particle interface will form a small radius of curvature. As the amorphous silica and siloxane network is irradiated, the *local* dissolution rate will be significantly higher at these cut edges as discussed in Section 3.4. The dissolution rate is also increased at other local defects which produce an acute angle including the natural cracking of the Boraflex surface observed in microscopy scans (as noted at the end of Section 3.1), cracking and crazing from swelling (as described in Section 3.2), and nicks and scratches from fabrication. The global rate of release of the boron carbide particles should be proportional to the dissolution of the amorphous silica and siloxane matrix. After dissolution of the degraded PDMS support matrix, the insoluble boron carbide particles are then free to migrate, generally by gravitationally settling to the bottom of the spent fuel pool. Video observations have shown small piles of black particles at the bottom of spent fuel racks and on the spent fuel pool floor below them. BADGER testing has also observed zones of unusually high attenuation, suggesting that boron carbide has settled in that zone (e.g., due to local Boraflex swelling or cover plate deformation that traps the particles above there). This is shown in Figure 5-4, as discussed in Section 5.2.

3.8 Characteristics of Boraflex Panel Dissolution

The BISCO analyses of Boraflex described in Section 2.2 focused on the mechanical and material effects of gamma radiation dose: stresses, shrinkage, swelling, fracture, etc., which only redistributed the neutron absorbing ^{10}B. As discussed in Section 3.3, subsequent analyses of the nondissolution degradation effects—end shrinkage, width shrinkage, and gaps as detailed in [EPRI 1988] and [EPRI 1993a]—showed that these effects were to some extent predictable (e.g., the 4.1 percent limit on total end shrinkage plus gap heights was well understood and corroborated by extensive data) and often manageable within existing margins to the $k_{eff} \leq 0.95$ limit. However, BISCO had not continued the experiments for sufficient time in an aqueous environment to recognize the subsequent dissolution process.

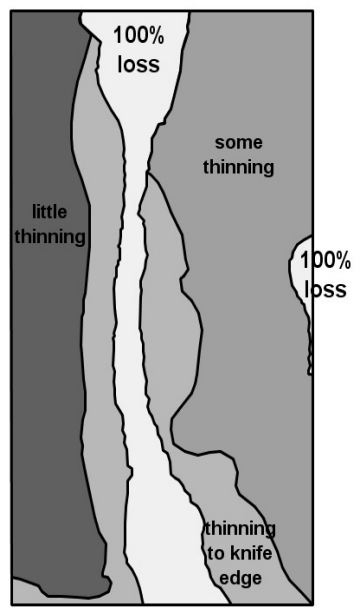

Figure 3–16. Illustration of an extensively degraded coupon.

Sections 3.4 and 3.7 above show that the cause of lost ^{10}B areal density is the *dissolution* of Boraflex: dissolution in the sense of a process of dissolving (the degraded PDMS matrix) and in the sense of disintegration (by the dissolving matrix releasing the embedded boron carbide particles). As discussed in Section 2.2, the degradation of Boraflex coupons was the first recognized evidence that a loss of material (and thus ^{10}B areal density) was taking place. (Recall from the end of Section 2.2 that rising spent fuel pool silica levels were not recognized as coming from Boraflex until later.) Figure 3-16 illustrates a representative coupon

with significant degradation. A 100 percent loss flow path from bottom to top is flanked by regions of extensive loss where the Boraflex has thinned to a knife edge. On the right is a region with some thinning and a "scallop" of 100 percent loss penetrating into the right edge. On the left is a region with very little thinning. Although other coupons from the same batch with the same design and service history might show similar features (e.g., flow paths, scallops, regions with various degrees of thinning), the distribution of those features is likely different: the dissolution effect has a large degree of randomness. It is entirely possible that another coupon, for example one with a tighter fitting clad, might show significantly less dissolution.

Boraflex dissolution occurs because of its contact with water during irradiation. As discussed in previous sections, the primary effects of water are as follows.

- Water provides dissolved oxygen to the panel surface, as well as surfaces available due to porosity. The oxygen combines with scissioned siloxane chain ends to prevent recombination, leading to microstructure deterioration, conversion of siloxane to silica, macroscopic swelling, and increased porosity.

- Water provides its constituent hydroxide ions to attack the surface silanol groups (again, including pore surfaces)—the dissolution mechanism.

Both of these chemical mechanisms are relatively slow. At lower ambient temperatures (e.g., 30°°C or 85 °F) Boraflex irradiated in water to cross-link saturation (>10^8 Gy or >10^{10} rads, per Table 3-1), significantly dissolves over years, not minutes (see [EPRI 1997a, Appendix A]). Thus, the availability of water to propel dissolution (by constituent hydroxides) and enhance dissolution (by dissolved oxygen) is important to the dissolution reaction rate.

As water passes over the Boraflex surface, monosilicic acid ($Si(OH)_4$) is forced into solution by hydroxide ion attack as shown in Figures 3-11 and 3-12. As more silica is released, the local concentration of monosilicic acid will approach its solubility limit and the reaction rate will decrease, though not to zero. Recall from Section 3.5 that the solubility limit is never practically reached because the monosilicic acid can polymerize. In theory, a Boraflex panel fully encapsulated in a sealed clad that is filled with water would release silica into the cavity water until the entire panel cavity was filled with a mixture of hydrated polymerized silica gel, particles of boron carbide and crystalline silica, and a thinner panel of residual Boraflex. However, as discussed in Section 2.3, the panel cavity is always open to the bulk pool water to some extent. Thus, there is always a rate of water exchange between the higher silica concentration of the panel cavity and the lower silica concentration of the bulk pool and, thus, a net outflow of silica.

If there were a relatively slow, uniform, natural convection flow of water from the bottom to the top of a panel cavity, one would expect a higher amount of dissolution at the bottom where the entering concentration is lower, and less dissolution at higher elevations as the monosilicic acid concentration increased. If the flow were relatively fast, this dissolution gradient would not be noticeable. A very low volumetric flow rate (e.g., if an inlet at the bottom of the panel cavity is

"pinched") might achieve quasi-equilibrium quickly and so may result in a higher loss bottom edge followed by low uniform dissolution without a gradient.

However, the flow through a panel cavity is not necessarily uniform from bottom to top. A typical panel cavity will have a chaotic distribution of inlets and outlets (and flow rates) due to:

- inherent random manufacturing variations, particularly in how the cover plate is attached to the spent fuel rack;
- the direction of convective water flow through the rack cell;
- gaps and other shrinkage effects; and
- random installation forces, fuel assembly impacts, and local Boraflex swelling effects (e.g., that can bend, twist, and buckle the cover plate).

As the Boraflex preferentially dissolves along a flow path, the flow path will become even larger, reducing the hydrodynamic resistance to flow and enhancing convective flow patterns, thus further increasing the dissolution rate along that flow path.

The result is thus a mixture of "uniform" dissolution simply due to bulk water being in contact with the surface of the Boraflex in its panel cavity and "local" dissolution due to the chaotic distribution of flow paths. A particularly common local dissolution effect is the "scallop." This is where, for example, the cover plate may be slightly buckled on one side of a short weld line or spot weld attaching the cover plate to the rack structure. This allows an enhanced flow path next to the weld, which tends to carve out a semi-ellipse of Boraflex. This is illustrated in Figure 3-17, a rough but to-scale sketch of a 330 mm (13 inch) axial section of a 216 mm (8.5 inch) wide Boraflex panel. The figure, which is based on an inspection photograph from an in-pool Boraflex surveillance assembly described in [EPRI 2002b], illustrates the largest local dissolution effect observed in that assembly.

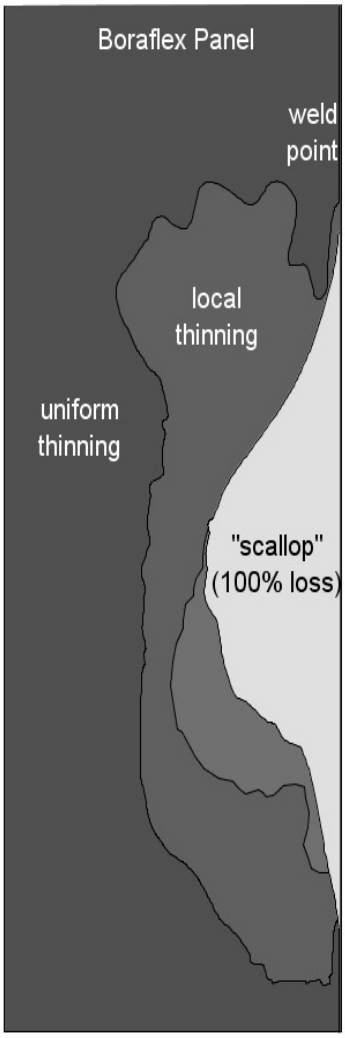

Figure 3–17. A to-scale illustration of a local dissolution "scallop" in an axial section of a Boraflex panel. Adapted from EPRI Report TR-1003414 (2002) with permission from Electric Power Research Institute.

Another example of local dissolution is the dissolved gap. This was illustrated in Figure 3-8, where the upper gap (just below a scallop on the right side) is larger on the right side than

on the left. As noted in Section 3.3, the largest single gap due to shrinkage alone is less than 115 mm (4.5 inches). However, much larger gaps have been observed in significantly degraded Boraflex panels—the result of dissolution, enhanced due to the additional flow area created by the gap. The dissolution is also greatly enhanced because a gap edge formed by fracture will have a porosity and PDMS-particle interface like that of a cut edge (see Figure 3-15 in Section 3.7) rather than a formed surface.

A final consideration is the effect of temperature on dissolution. Table 3-2 shows that silica solubility is a nearly linear function of temperature. If significant gamma radiation heating convects to the flow of panel cavity water, the water temperature may increase as the flow progresses through its flow path. This will increase the dissolution rate as the flow progresses, which will tend to cancel the reduced dissolution rate due to increased monosilicic acid concentration in the water.

In performing a criticality safety analysis it is important to recognize the unpredictable nature of local dissolution effects and the potentially significant reactivity effects associated with them. Like the gaps described in Section 3.3, areas of 100 percent local dissolution allow a greater unattenuated neutron transport path between assemblies, significantly decreasing the reactivity credit of the Boraflex. In one analysis of the effects of local dissolution (see [EPRI 2002a, Section 4.1.2]), it was found that (for an infinite array of spent fuel storage rack cells) the reactivity penalty was the same where:

- each Boraflex panel had lost 44 percent of its boron carbide due to uniform thinning alone, and

- each panel had lost 20.8 ± 2.5 percent (mean ± standard deviation of panel loss) of its boron carbide due to uniform thinning, gaps, and local dissolution.

This illustrates the large reactivity penalty associated with the synergy between the various modes of Boraflex degradation and dissolution. For example, if a RACKLIFE simulation were to predict a 30 percent panel average loss, a criticality safety analysis which interpreted that loss as uniform thinning would greatly underestimate the actual reactivity penalty. Because of the random locations and geometries of local dissolution anomalies, criticality safety analyses often follow a stochastic approach to include this synergistic effect in quantifying the reactivity penalty.

4 RACKLIFE

The timeline in Figure 2-1 shows that problems with Boraflex first came to prominence in 1987. As discussed in Section 2.2, coupons were an indicator of Boraflex degradation, albeit an often unreliable indicator. Based on the review of the literature in [EPRI 1988], degradation was expected to correlate strongly with absorbed gamma dose from spent fuel. However, there were no generally available user-friendly tools for calculating the gamma dose to a coupon. In 1991, as described in [EPRI 1993a], a spreadsheet application was developed to calculate the gamma dose for both plant coupons and for EPRI-sponsored test programs utilizing both coupons and full length panels. As the dissolution mechanism of boron carbide loss became evident, by 1993 this spreadsheet application had evolved into a huge and unwieldy spreadsheet that could simulate the dissolution of a panel of Boraflex, as detailed in [EPRI 1993b].

With utility concerns and NRC interest in Boraflex increasing (leading to a Generic Letter in 1996 [NRC GL 96-04]), it was decided to develop a stand-alone computer code that engineers at individual plants could use to assess the Boraflex in the spent fuel pools they supported. The code was named RACKLIFE, for Boraflex rack life extension, in that it could be used to manage fuel moves in the spent fuel pool to extend the service life of at least some spent fuel pool capacity, thereby delaying or avoiding the capital costs of a rerack.

Prior to the release, substantial beta testing was performed by select users—typically those who felt they had a need to quickly understand the state of their Boraflex racks. Subsequent releases in 1996 (version 1.01) and 1997 (version 1.02) fixed a few minor bugs and introduced enhancements based on user feedback. Because of the large volume of input data required (described subsequently), it was originally envisioned that clerical staff would perform the data input, and so the original code used a familiar database-style user interface. When it was found that the engineers themselves generally had to perform the data input, version 1.10 was released with a spreadsheet-style user interface, as well as further enhancements requested by users. This version was also released with a verification and validation report [EPRI 1999a].

Versions up through 1.10 were developed for PCs running the DOS operating system. However, a number of users were encountering mandates from IT departments that sought to migrate all desktop computer systems to the Microsoft Windows™ platform. Thus, a Windows-based version was later developed as Version 2.0. (The author's direct experience with RACKLIFE ends with the initial planning for this Windows-based version.) Subsequent versions added further enhancements, including a thermal model to track bulk pool and panel cavity temperatures. The current version is 2.1.1, as described in [EPRI 2009].

4.1 Overview of RACKLIFE

RACKLIFE takes advantage of the known dissolution degradation mechanism of Boraflex described in Chapter 3. The RACKLIFE software can track thousands of fuel moves and compute the total dose exposure of each of the Boraflex panels in the pool. The gamma dose deposition models and silica dissolution kinetics models are combined to calculate Boraflex degradation, initially as a loss of silica, and then as a loss of boron carbide. RACKLIFE does not take shrinkage or gaps into account because those degradation mechanisms do not release silica to the bulk pool. Unknown parameters in the model, such as the rate of flow through the Boraflex panel cavities (as discussed in Section 3.8) are adjusted by the user such that RACKLIFE predictions of spent fuel pool silica concentrations match measured values. Adjustments can also be made based on the results of BADGER test measurements of boron carbide loss.

Overall, RACKLIFE is designed to:

- characterize the gamma exposure history of a spent fuel rack absorber panel;
- compute the panel average loss of silica, thus computing the average loss of neutron absorbing boron carbide; and
- manage the movement of spent fuel assemblies in the spent fuel racks to optimize a rack management and preservation strategy as described in Section 2.3.

To compute exposure histories, RACKLIFE uses:

- reactor core power histories in the month prior to fuel discharge,
- a database that tracks the movement of discharged spent fuel assemblies into and out of a rack cell,
- a database that characterizes each assembly moved, and
- geometry data for the spent fuel rack cells and their locations in the spent fuel pool.

When computing the gamma dose to a particular panel, the original intent was for RACKLIFE to compute gamma doses from spent fuel assemblies in racks containing Boraflex where the dominant gamma shields were the stainless steel structures of the racks and the ambient water. It was found that the spent fuel gamma absorption characteristics of Boraflex were sufficiently close to those of water that Boraflex could be modeled as water. Because some pools had Boraflex racks adjacent to non-Boraflex racks (e.g., from multiple reracks), non-Boraflex stainless steel rack module designs were also permitted in RACKLIFE. The movement of fuel in and the gamma dose from non-Boraflex cells to nearby Boraflex panels was tracked and computed, respectively, but no absorber panel doses were calculated for the non-Boraflex racks. By developing and implementing appropriate attenuation correlations for other rack materials, RACKLIFE could in principle preserve its current database and user interface design to extend its application to compute exposure histories for these other materials.

To compute the loss of silica—and ultimately the loss of boron carbide—from Boraflex, RACKLIFE also uses:

- spent fuel pool chemistry data, including silica concentration, as discussed in Section 3.6,
- spent fuel pool cleanup system characteristics and utilization, as discussed in Section 3.6, and
- a tunable parameter (the *escape coefficient*) that characterizes the flow of spent fuel pool water into and out of each absorber panel cavity. It is expressed as the number of panel cavity volume exchanges per day, with units [day^{-1}].

This part of the code is specific to the dissolution of Boraflex as described in Section 3.0 and is not applicable to the analysis of other neutron absorbers.

In its calculations, RACKLIFE makes no attempt to determine the spatial distribution of dose or Boraflex dissolution over the surface of a panel. There is an endless variety of end-of-cycle spent fuel axial assembly power profiles, which theoretically could be estimated by a 3D reactor model at considerable expense. There is also an endless variety of spent fuel pool water flow patterns through a Boraflex panel cavity, which could likely never be determined a priori. Therefore, it is not possible to simulate the dose and dissolution processes with more precision than as a panel average.

The remainder of this chapter further describes the theory, application, and inherent uncertainties in the RACKLIFE computer code. It begins in Section 4.2 by describing the theory behind the simulations of exposure history and Boraflex dissolution. In Section 4.3, this theory is shown in application for tracking and predicting Boraflex dissolution. Section 4.4 then details the inherent computational uncertainties in RACKLIFE as it generates its predictions of further Boraflex dissolution.

4.2 RACKLIFE Theory

4.2.1 Radiation Exposure History

Details on the radiation exposure history algorithms can be found in [EPRI 1997a] as well as in [EPRI 1993a; Section 5 and Appendix A]. The calculation process for determining the dose to a panel of Boraflex proceeds in three steps, as follows.

1. For each reactor that discharged spent fuel to the spent fuel pool, compute the average assembly end-of-cycle (EOC) thermal power.

The spent fuel gamma source is most sensitive to the reactor power during the last thirty days of operation. In computing the EOC reactor thermal power, three different power histories are available in RACKLIFE:

- instantaneous shutdown from a fraction (including 100 percent) of full power;
- operation at one fraction of full power, then a step change to operation at another fraction for some time interval, and then instantaneous shutdown; or
- coast down linearly from one fraction of full power to another fraction for some time interval, and then shut down instantaneously.

The user must choose which of these three scenarios most closely matches the actual EOC reactor operation. RACKLIFE provides feedback on the equivalent reactor EOC power so that a user with a detailed knowledge of the actual EOC operating history might adjust the parameters for a given profile to match the actual equivalent EOC power.

2. For each assembly in the spent fuel pool, compute the spent fuel assembly gamma source term as a function of the assembly:

   - fuel loading [MTU],
   - fuel initial enrichment [w/o ^{235}U],
   - End of Cycle power sharing (1.0 is core average at end of cycle),
   - burnup [MWD/MTU], and
   - type (describing dimensional and fuel rod array specifications)

   The assembly "type" is selected from a database distributed with the RACKLIFE code that includes dimensional and fuel rod array data from the Energy Information Agency (EIA) Fuel Assembly Database (FADB) of the Characteristics Database (CDB). Users can update this database themselves as new fuel types become available.

   Note that this computation gives an assembly average source term; no spatial effects are included. Again, RACKLIFE makes no attempt to determine the spatial distribution of panel dose or dissolution.

3. For each absorber panel, the gamma sources from assemblies in rack cells adjacent to the panel are integrated over time and transported to the panel using correlations of dose rate density (in [rads/hour/assembly]) versus gamma source density (in [gammas/second/MTU]) and intervening stainless steel and water thicknesses (from user input rack dimensional data).

Only spent fuel assemblies in the two rack cells directly adjacent to a Boraflex panel are included in calculating the dose to a panel. If the Boraflex panel is on the outside of a rack module, then the nearest neighbor in an adjacent rack module (if there is one) is identified for the calculations. During the initial development of RACKLIFE, this nearest-neighbors-only approach was necessary to achieve an acceptable computational speed. Subsequent testing during the verification and validation of RACKLIFE, though limited, indicated that the effects of this simplification were well within the uncertainty of the dose calculation methodology.

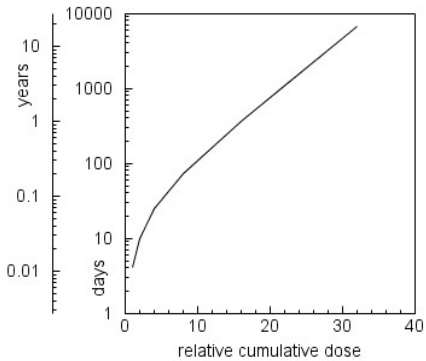

Figure 4–1. Relative amount of dose absorbed by a Boraflex panel over time from an assembly placed in the racks 7 days after shutdown.

The decay of the spent fuel gamma source term with time is tracked as the integration proceeds. The integration takes into account all assemblies which are moved into and out of the adjacent cells, so that the exposure history for each panel is quite unique.

The source term decay makes freshly discharged fuel more significant to Boraflex dissolution than "old" fuel. Figure 4-1 shows a relative cumulative dose curve for an assembly discharged to a rack cell after cooling in the reactor for seven days after shutdown. After 4 days in the rack cell, "1 unit" of relative dose has been absorbed by the adjacent Boraflex panel. After 6 more days—10 days total—"2 units" of relative dose have been absorbed; that is, it took 6 days for the dose to double. The dose doubles again ("4 units") after 25 days, then 72 days, then 1 year, and then 18 years. Thus, it took 6 days for a freshly discharged assembly to double the panel dose, but 17 years for a year-old assembly to double the panel dose. After 18 years (over 6,500 days), the panel has absorbed "32 units" of dose—only 32 times what it absorbed in the first four days after discharge.

Example exposure histories are illustrated in Figure 4-3 (shown on a subsequent page) for a PWR with Region I flux-trap and Region II egg-crate racks. (The small scale jaggedness is an artifact of the data sampling scheme used in the simulations developed for this report; RACKLIFE actually computes the dose rate every day of its simulations.) The blue lines in each plot represent the exposure histories of zero-dose panels— panels between spent fuel rack cells that have never had any spent fuel assemblies discharged to them. These are panels that should not exhibit any dose related dissolution.

Because the Region I racks are flux-trap designs in which there are two Boraflex panels separated by the flux-trap water gap, each panel between two rack cells will have a "near" cell and a "far" cell, as illustrated in Figure 4-2. The blue arrows from source assembly to the red Boraflex panel of interest highlight the neutron transport path length and material differences between near and far. A spent fuel assembly in the far cell will contribute less gamma dose to the panel of interest because of the larger geometric and material gamma attenuation. Thus, compared to a Region II egg-crate rack, a panel in the Region I racks will often accumulate less dose from adjacent fuel assemblies.

However, as discussed in Section 2.3, the Region I racks are designed to accommodate fresh fuel assemblies (having no burnup reactivity credit) and so are generally closer to the fuel transfer canal to expedite refueling operations. To further quicken refueling operations, permanently discharged fuel and any discharged fuel intended for return to the reactor are also placed in the Region I racks. The placements of the fuel assemblies are often optimized by computer programs with the objective of minimizing the fuel transfer time to reduce the outage time. Thus, freshly discharged fuel is more often placed first in the Region I racks, resulting in much higher dose rates to those Boraflex panels. When the refueling outage is over, assemblies may then be transferred from the Region I racks to the Region II racks to make room for the next outage. In some plants, these post-outage spent fuel assembly moves are performed soon after the outage; then the Region I racks will typically accumulate less dose than the Region II racks. This is the case in Figure 4-3, which illustrates various scenarios for accumulating panel gamma dose over time. If instead the transfer from Region I to Region II happens later, it is possible for panels in Region I to accumulate more gamma dose than panels in Region II.

Figure 4-3 shows the gamma dose exposure history of several Boraflex panels, all located in the same pool. The racks were installed at time 0. Various spent fuel assemblies were placed around these panels over the years, typically during or just after refueling outages starting at around 270 days and then every 2 years (730) days after that: at 1,000 days, 1,730 days, and so on, based on the two-year refueling cycle. These exposure histories were selected to illustrate various dose accumulation scenarios and phenomena, as discussed in the following paragraphs.

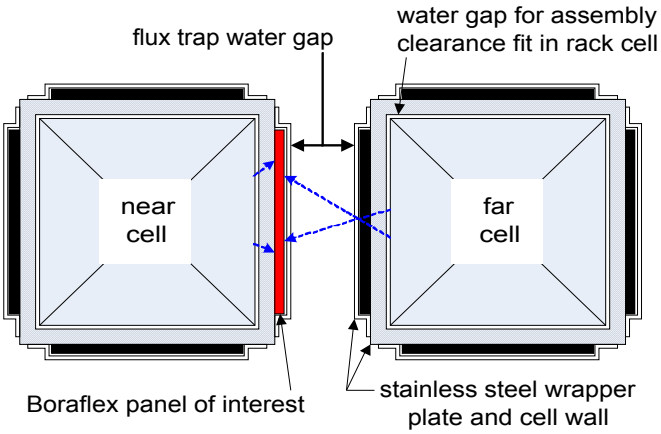

Figure 4–2. Flux-trap rack showing "near" and "far" rack cells.

Looking at the Region I plot in Figure 4-3, the dark green lines toward the bottom represent three panels arranged around a common cell. Although the three panels share a common fuel assembly discharged into the cell they surround, their doses are slightly

different because of differences in the assemblies on the opposite side of each panel. Newly discharged fuel assemblies were placed around these panels 270 days after the racks were installed. Some 90 days later, at 360 days since installation, the fuel assemblies were removed and transferred into Region II racks. (As discussed below, these transfers led to panels in Region II accumulating dose—as represented by some of the green curves in the accompanying Region II figure.) In the 90 days, these Region I panels were surrounded by freshly discharged spent fuel assemblies, each panel accumulated an average panel dose of around $2 \cdot 10^9$ rads, which is above the $5 \cdot 10^8$ rads at which the panel dissolution rate becomes significant, as noted in Table 3-1.

The light green and yellow stair-step curves in the Region I plot represent panels which had spent fuel discharged around them nearly every outage, for about 90 days each. The accumulated dose to the panel represented by the yellow curve initially lags because at the first outage (at 270 days) only one assembly was placed next to that panel (the opposite rack cell was left empty). However, at the second outage (at 1,000 days), no fuel is placed around the light-green panel, and so the yellow panel curve overtakes it. These can be contrasted with the orange and red curve panels in the Region I plot, which had spent fuel assemblies around them for an entire cycle (until just before the second refueling outage, at about 1,000 days). At that time, one of the assemblies next to the orange panel was removed while the red panel still had the same two assemblies around it.

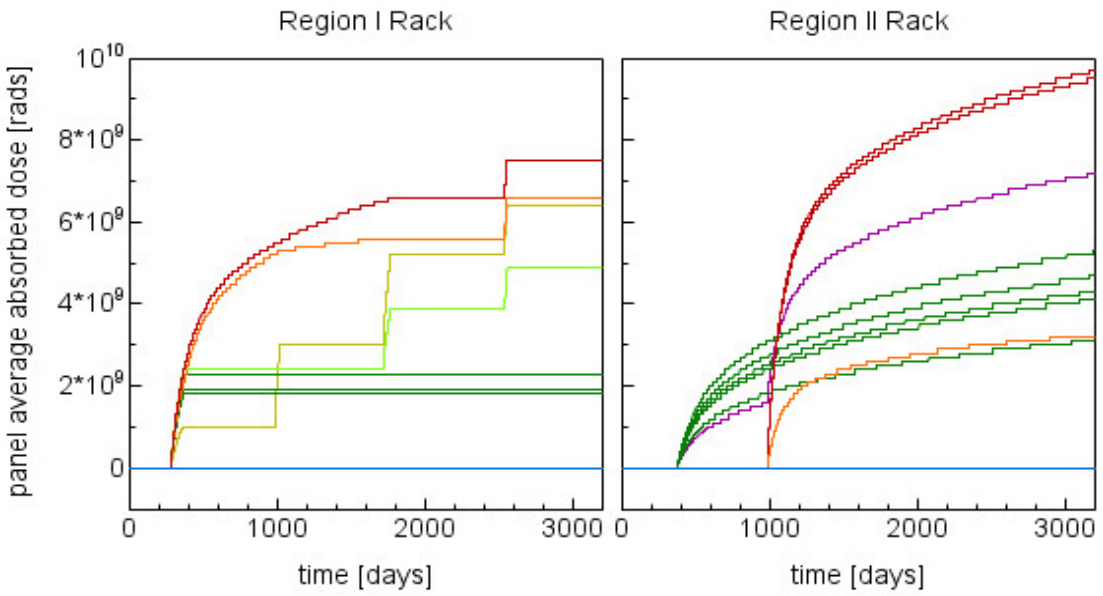

Figure 4–3. Example gamma dose exposure histories for select individual panel in PWR Region I and Region II racks.

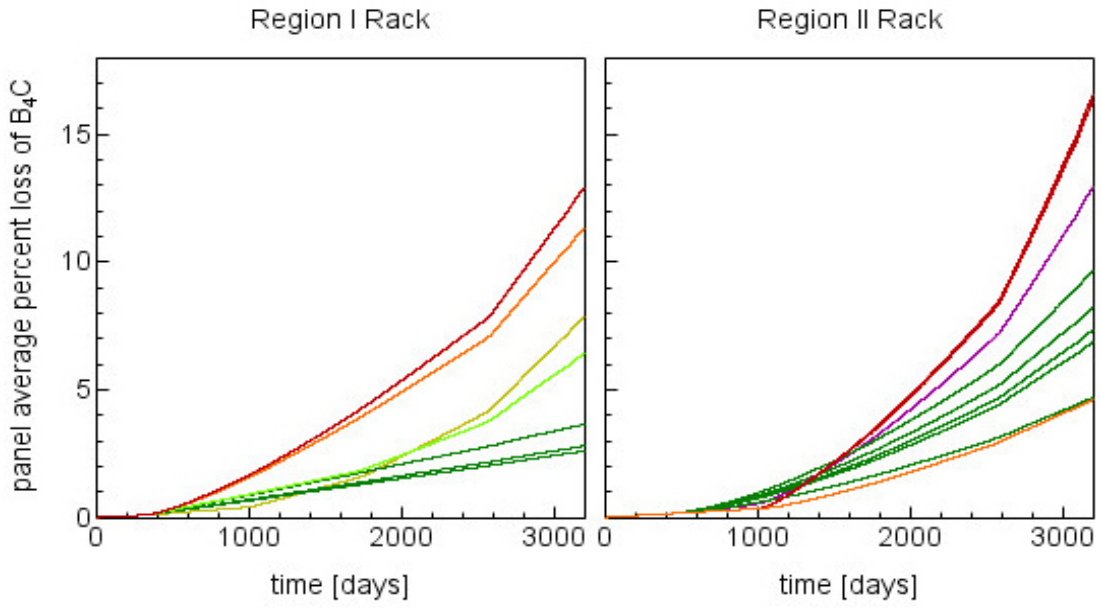

Figure 4–4. Example B_4C loss histories for the same PWR Region I and Region II panels as in figure 4.3.

Then, just before the third outage (at about 1,730 days), all of the neighboring assemblies are removed; newly discharged assemblies are placed during the fourth outage, and soon after removed. After four cycles, the highest dose panel has accumulated an average dose of almost $8 \cdot 10^9$ rads. Because of the long residency of freshly discharged assemblies around the red panel, it is possible that the red panel could have accumulated dose comparable to or greater than that accumulated by Region II panels (illustrated in the accompanying figure). However, in this simulation, the size of the flux-trap water gap was such that the geometric and material attenuation effects dominated and the accumulated dose was lower.

As noted above, the green curves in the Figure 4-3 Region II plot represent panels with assemblies on both sides that were:

- originally discharged to Region I during the first outage (at 270 days),
- remained in Region I for about 90 days,
- and were then transferred to Region II to cause the accumulation of dose shown by the green curves in the Region II figure.

Thus, per Figure 4-1 (approximately, ignoring the significant Region I versus Region II dose rate differences noted above), while it took about 90 days for a panel in Region I to accumulate $\sim 2 \cdot 10^9$ rads, it took nearly a year for the Region II rack to accumulate the same dose because of the cooling time of the assemblies; and it will take about 17 years more to double that dose. The red curves represent panels adjacent to assemblies discharged during the second refueling outage (at 1000 days) that were moved directly to the Region II racks. Because these assemblies did not cool for 90 days as the assemblies adjacent to the green panels had, the rate of dose accumulation in the red panels is much higher. Because of the much smaller geometric and material attenuation of the Region II racks, these red panels accumulate dose faster than the red Region I panel did.

The purple curve in the Region II plot shows a panel which had two assemblies moved next to it after the first outage (at 1,000 days), indicated by the corresponding increase in dose. After the second outage, those two assemblies were replaced with two freshly discharged assemblies. Again, the curve shows an increase in dose. The orange curve represents a panel that had one low power-sharing spent fuel assembly moved next to it soon after shutdown for the second outage (at 1,000 days). Observe that the orange curve soon closely follows one of the green curves—a panel surrounded by two average power sharing assemblies that had cooled for 90 days before being moved in. This shows that a simple inspection of a cumulative dose curve may not be enough to characterize the history of the assemblies around a panel. A key feature of RACKLIFE is its ability to track thousands of assembly moves over time to compute the cumulative gamma dose exposure history of each of the thousands of Boraflex panels in the spent fuel pool.

Figure 4-4 (shown previously, below Figure 4-3) illustrates the corresponding boron carbide loss from these same panels as predicted by a simulation of RACKLIFE developed for this report (discussed subsequently in Section 4.4.2). For the examples shown here, RACKLIFE would predict that panels with the most gamma exposure will show the most boron carbide loss. While

this is most often true in practice, there are numerous panels in spent fuel pools where this is not the case. To illustrate this, note the differences between the orange and yellow panel curves in the Region I plots. At 3,000 days, the orange and yellow panels have accumulated very similar gamma doses, about $6.5 \cdot 10^9$ rads. However, because the orange panel received most of its dose earlier in time, it is predicted to have an 11 percent loss of boron carbide, while the yellow panel is predicted to have only an 8 percent loss. After the next outage, the yellow panel could overtake the orange panel in accumulated dose but still lag in boron carbide loss. This is why the exposure history, and not just the current accumulated dose, is important to predicting boron carbide loss from a Boraflex panel. The theory behind the calculation of these boron carbide losses is presented next.

4.2.2 Silica Kinetics

Section 4.2.1 above has outlined how RACKLIFE calculates the gamma irradiation dose to a neutron absorber panel. Section 3.0 detailed the degradation and dissolution mechanisms that result from the irradiation of Boraflex in water. This section summarizes how these are combined to predict the magnitude and rate of Boraflex degradation. The basis of the RACKLIFE boron carbide loss prediction is the numerical solution of a system of silica mass balance equations. The mass balance incorporates a multistate and multicompartment model of silica concentration. That is, silica is present in both the reactive (monomeric) silica state and the polymerized silica state. The silica mass is distributed among many water-filled compartment locations: the N panel cavities surrounding the Boraflex panels, and the remaining water volume described as the *bulk pool*. RACKLIFE tracks the loss of silica from its measurement in the bulk pool, through its states and compartments, back to its dissolution from a panel of Boraflex. From silica loss, the associated proportional loss of neutron-absorbing boron carbide from the panel (as discussed in Sections 3.1 and 3.7) can be computed. The details of the silica mass balance model are illustrated in Figure 4-5.

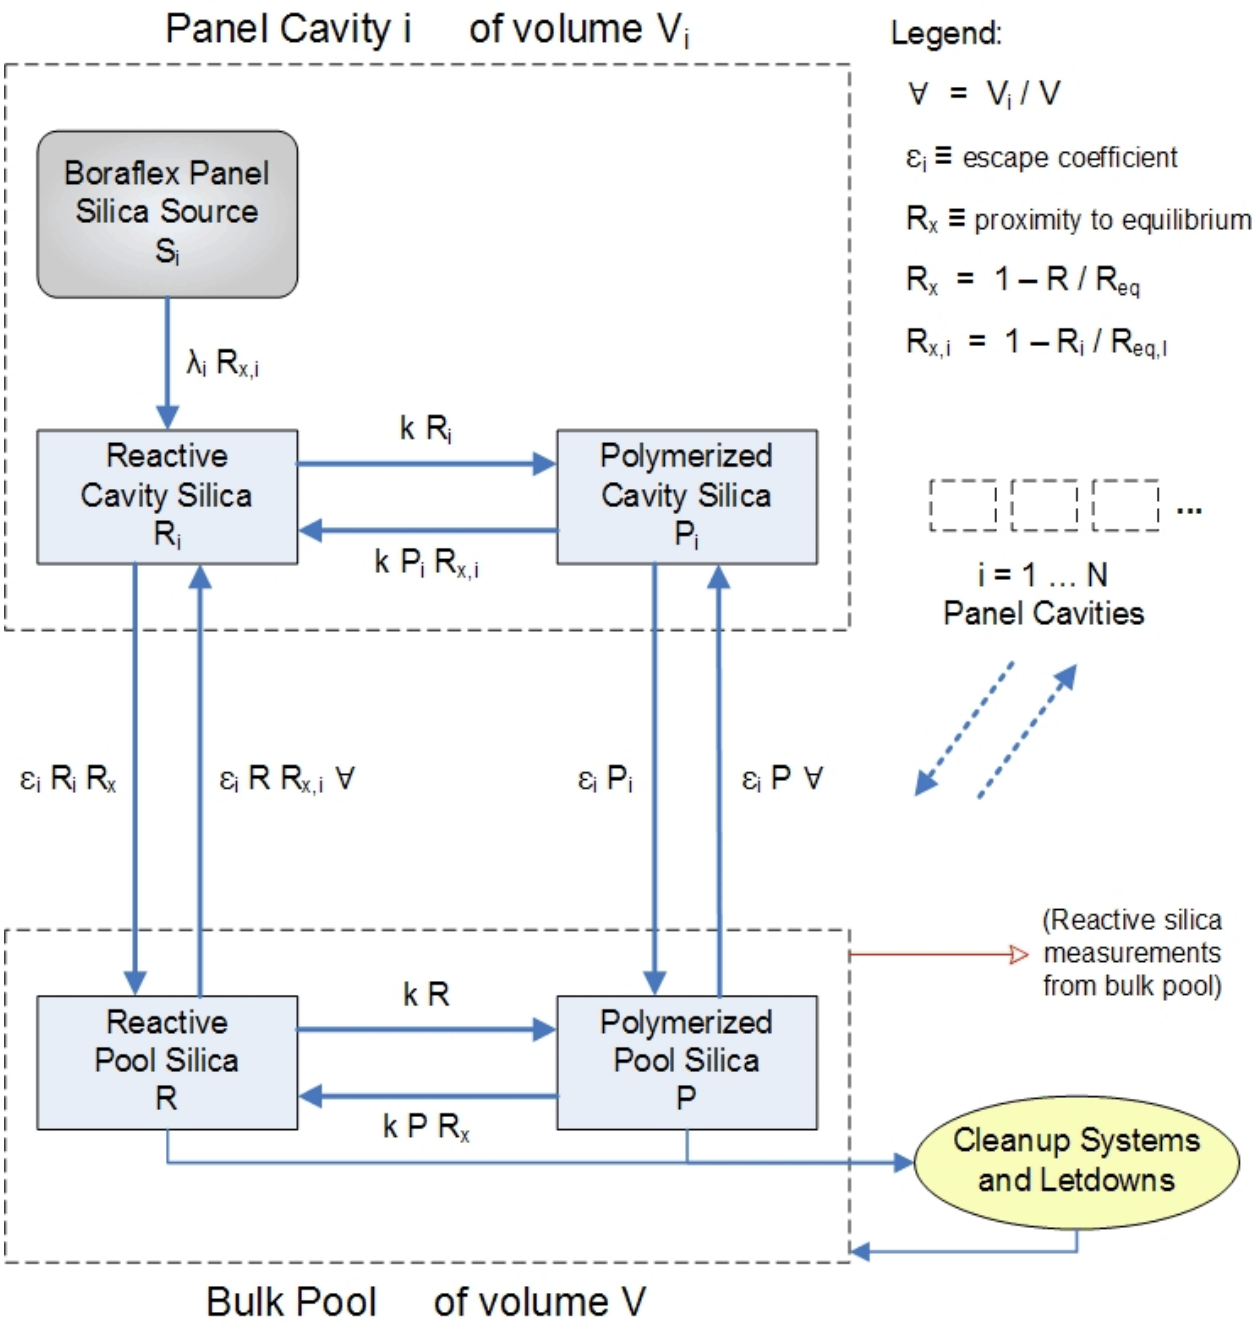

Figure 4–5. Silica mass balance kinetics model for a spent fuel pool

The model begins with the dissolution of irradiated Boraflex in the aqueous environment of its panel cavity. The rate of dissolution of silica was determined from laboratory dissolution test data involving various samples of Boraflex at various absorbed doses and temperatures. (These laboratory experiments are detailed in [EPRI 1993b] and [EPRI 1997a, Appendix A].) The silica dissolution rate per unit surface area of Boraflex is represented by the function λ'' [mg/cm^2·day]. Non-linear regression analyses of the data combined with phenomenological Arrhenius fits of the temperature dependence resulted in the following functional forms for the Boraflex silica source rate, λ_i, from the ith panel of Boraflex, in [mg / day].

$$\lambda = \begin{cases} \textit{fit below using } I_i = 5\cdot 10^8 & I_i < 5\cdot 10^8 \textit{ rads} \\ \lambda''(T_i, I_i)\, A_i & 5\cdot 10^8 < I_i < 1\cdot 10^{10} \textit{ rads} \\ \textit{extrapolation of above fit} & I_i > 1\cdot 10^{10} \textit{ rads} \end{cases} \tag{4-1}$$

where

λ'' is the laboratory measured release rate per unit surface area of Boraflex [mg/cm^2·day] for a given temperature T_i and absorbed dose I_i, as shown in Figure 4-6

A_i is the surface area of Boraflex available for dissolution

$\log(\lambda''(T_i, I_i)) = a(T_i) + b(T_i)[\log(I_i)] + c(T_i)[\log(I_i)]^2$

$a(T_i)$, $b(T_i)$, $c(T_i) = ke^{-q/T}$ (with different k and q for a, b, and c)

T_i is the temperature of the water in the panel cavity.

In the original version of RACKLIFE the panel cavities were all assumed to have the same temperature as that measured for the bulk pool. The most recent version of RACKLIFE uses a more sophisticated temperature model (see [EPRI 2009]).

Figure 4-6 shows that some level of silica dissolution occurs even from

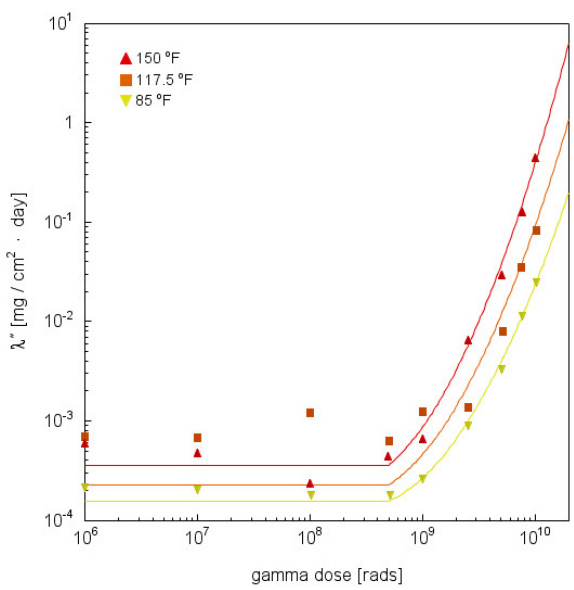

Figure 4–6. Silica release rate per unit area of Boraflex. Adapted from EPRI Report TR-107333 (1997), with permission from Electric Power Research Institute.

panels with low accumulated doses. As discussed in Section 3.2, the PDMS in Boraflex with its full complement of methyl groups is hydrophobic and so should repel hydroxide ion attack. However, along a cut edge of the Boraflex, especially where contact angles are acute (see Figure 3-15), hydroxide ion attack may be viable. It was decided that the temperature dependent regression fit dissolution rate at $5 \cdot 10^8$ rads would be used as the constant dissolution rate for all panels with lower doses because the measurement of pool silica was most accurate above that dose threshold. Because the magnitude of the low-dose dissolution rates are very small compared to Boraflex at higher doses, the uncertainty introduced into predictions of boron carbide loss and panel cavity and bulk pool silica concentrations are not significant.

The time required to steadily irradiate Boraflex in water precluded experiments beyond 10^{10} rads. While some material properties do appear to saturate around this dose (see Table 3-1), as confirmed by experiments on utility coupons irradiated to higher doses, porosity and scission degradation continue to increase without saturation. Thus, it was decided to extrapolate the dissolution rate regression curves to higher doses.

To complete the source term, a factor describing the proximity to equilibrium (introduced in Section 3.5), shown as $R_{x,i}$ in Figure 4-5, is multiplied to give

$$S_i = \lambda_i R_{x,i} = \lambda_i (1 - R_i / R_{eq,i}) \tag{4-2}$$

If there is no reactive silica in a panel cavity, then the silica source rate is the unadjusted Boraflex dissolution rate ($R_i = 0 \Rightarrow R_{x,i} = 1 \Rightarrow S_i = \lambda_i$). As the panel cavity silica concentration increases, the proximity to equilibrium factor, $R_{x,i}$, decreases from one toward zero, thereby reducing the silica source term. As discussed in Section 3.5, if the reactive silica concentration in the panel cavity water were able to reach the silica solubility limit, the source term would go to zero ($R_i \rightarrow R_{eq,i} \Rightarrow S_i \rightarrow 0$). However, as shown in Figure 4-5, there are two mechanisms to preclude this:

- some of the reactive silica will polymerize at a rate kR_i, (a linear polymerization rate constant, k, times the amount of available reactive silica in the panel cavity, R_i), as described in Section 3.5; and

- some of the reactive silica will "escape" the panel cavity and enter the bulk pool at a rate $\varepsilon_i R_i R_x$ (a linear escape coefficient, ε_i, times the amount of available reactive silica in the panel cavity, R_i, times the proximity to equilibrium factor for the bulk pool that the silica is escaping into, R_x).

Silica from a degraded Boraflex panel may participate in several other processes, including:

- depolymerization of silica within the panel cavity;

- return of reactive and polymerized silica from the bulk pool into the panel cavity (the opposite of escape);
- polymerization and depolymerization within the bulk pool; and
- escape of polymer from the cavity into the bulk pool.

This results in the following system of equations that model the silica kinetics, where each process is characterized by its own rate constant as illustrated in Figure 4-5. The system of equations represents the interaction of a single panel cavity with its Boraflex panel and the bulk pool; thousands of other panel cavities also interact with the pool.

For the ith panel cavity,

$$\frac{dR_i}{dt} = \lambda_i R_i - k\left(R_i - P_i R_{x,i}\right) - \varepsilon_i \left(R_i R_x - R R_{x,i} \forall_i\right) \tag{4-3}$$

and

$$\frac{dP_i}{dt} = k\left(R_i - P_i R_{x,i}\right) - \varepsilon_i \left(P_i - P \forall_i\right) \tag{4-4}$$

while for the bulk pool,

$$\begin{aligned}\frac{dR}{dt} = & - k\left(R - P R_x\right) \\ & + \varepsilon_i \left(R_i R_x \right. \\ & \left. - R R_{x,i} \forall_i\right) - cleanup\end{aligned} \tag{4-5}$$

and

$$\begin{aligned}\frac{dP}{dt} = & k\left(R - P R_x\right) + \varepsilon_i \left(P_i - P \forall_i\right) \\ & - cleanup\end{aligned} \tag{4-6}$$

RACKLIFE solves this system of ordinary differential equations numerically with a time step dependent on the escape coefficient and a maximum step size of 1 hour. This means that, to simulate a spent fuel pool with 5,000 Boraflex panels for 20 years will require at least 175,000 time steps, over 10 billion state point variables, and well over 100 billion computations.

The ultimate result of these calculations is the mass of silica lost, and thus the mass of boron carbide lost, from each panel of Boraflex, computed by integrating the source over the time interval of interest:

$$boron\ carbide\ lost = \frac{mass\ of\ silica\ lost \cdot \dfrac{w/o\ B_4C\ in\ Boraflex}{w/o\ silica\ in\ Boraflex}}{initial\ mass\ of\ B_4C} \tag{4-7}$$

where

$$mass\ of\ silica\ lost\ from\ panel\ i\ in\ time\ [t_0, t_1] = \int_{t_0}^{t_1} S_i\, dt = \int_{t_0}^{t_1} \lambda_i R_{x,i}\, dt \tag{4-8}$$

The results of these calculations are preserved for user output in two ways:

- state point files, and
- pool state files.

When a RACKLIFE model is created, the user creates an *initial conditions state point* that characterizes the state of the racks and the pool when the racks were first installed—unirradiated and with zero loss. At each time step, RACKLIFE updates a list of variables (e.g., reactive silica, polymerized silica, temperature, pH, water volume, gamma dose, B_4C loss, and more) for every panel and associated panel cavity in the pool, as well as for the bulk pool. (The bulk pool calculations exclude panel-specific variables such as gamma dose and B_4C loss.) As noted above, this could result in tens of billions of values for the user to store and wade through—which is simply unmanageable.

Instead, in performing a RACKLIFE simulation, the user selects a point in time at which the "state of the system"—a *state point*—is desired, and computes from a given state point to the selected final point. For example, a user could run RACKLIFE to create a state point when a BADGER campaign occurs, so that RACKLIFE's predictions can be compared with BADGER's measurements. The given state point from which calculations are started is any state point file that predates the selected final point; the initial conditions state point predates all possible state points. Any subsequent state point can be used as the initial conditions for a simulation; there is no need to always start the calculations from the Boraflex installation date.

As a RACKLIFE simulation progresses from one state point to another, it preserves a running summary of the state of the bulk pool—a *pool state* file. This stores the RACKLIFE predicted bulk pool reactive and polymerized silica, any user input silica measurements for comparison, and average and maximum gamma doses, B_4C losses, and rates of dissolution to summarize the state of all of the panels.

The next two sections discuss the primary confounders to what might otherwise appear to be a straightforward simulation model with some coefficients based on laboratory data.

4.2.3 The Escape Coefficient

As discussed in Section 3.8, it is impossible to know the rate, ε_i, at which a particular panel cavity exchanges silica-laden water with the bulk pool. Therefore, the escape coefficient is treated as a tunable parameter. The pool state data file contains both daily RACKLIFE bulk pool silica concentration predictions and measured pool silica concentrations as input by the user (from pool chemistry silica measurement data over time). Data from the pool state file is plotted as RACKLIFE computations proceed so that the user can visually check that the predicted pool silica curve coincides with the measured data as much as possible. If the user observes that they do not coincide, the escape coefficient is adjusted and the state point calculation is repeated from the beginning. The user can also analyze the differences between the predicted and measured silica data in the pool state file outside of RACKLIFE to quantitatively optimize the escape coefficient selection (e.g. using a correlation coefficient or chi-square statistic to compare the measured and predicted data). In practice, the escape coefficient is typically adjusted to two significant figures.

When developing a RACKLIFE model, a constant average escape coefficient is typically assumed to apply to every panel cavity throughout the spent fuel pool. The value depends primarily on the rack design, and secondarily on the spent fuel pool temperature and the state of the Boraflex panels (as discussed subsequently in Section 4.4.1). As shown in Figure 2-3, rack designs vary considerably, and so do escape coefficients. Plant to plant, the constant average escape coefficient that best models the panel cavities has been observed to vary over three orders of magnitude, from below 0.1 to over 100 [day^{-1}].

If a good fit of the measured pool data can not be obtained using a constant escape coefficient for every panel cavity, the user can:

- set different escape coefficients for different rack panel cavity designs in the same pool; and
- change the escape coefficient over time.

For example, PWR Region I and Region II panel cavities may have significantly different escape coefficients based on their design. Feedback from a BADGER measurement campaign can be used to balance predicted versus measured Boraflex losses between the two regions by adjusting the escape coefficients separately. As discussed in Section 3.8, while there are degradation mechanisms that will tend to increase the exchange rate of panel cavity water over time, RACKLIFE can not predict this. It is thus left to the user to best match the available pool silica data, possibly by increasing the escape coefficient in a sequence of state point calculations. Within a state point calculation, RACKLIFE assumes the escape coefficient is constant, so the

sequence of state point calculations will have escape coefficients that resemble a staircase rather than a continuous increasing curve.

Once a best-estimate RACKLIFE model has been developed, the user can then adjust individual escape coefficients, either within RACKLIFE or by externally manipulating a text file that stores escape coefficient data, and rerun the state point calculation (or sequence of calculations). This can be done after a BADGER campaign to best match individual panel predicted losses to BADGER data of actual losses. However, this is limited to the typically dozens of panels measured in a BADGER campaign, compared to the thousands of panels in the spent fuel pool. Sensitivity studies can also be performed by stochastically varying the escape coefficients according to the distribution of individual panel cavity escape coefficients used to best match RACKLIFE predictions to BADGER measurements. The results of such Monte Carlo analyses can be input to a criticality safety analysis.

4.2.4 Cleanup Systems

As discussed in Section 3.6, spent fuel pool cleanup systems respond differently to different forms of silica in solution. RACKLIFE will recommend best-estimate cleanup efficiencies for various cleanup systems, such as filters, deep-bed resin demineralizers, precoat filter/demineralizers, and reverse osmosis systems. Based on discussions with cleanup system vendors, the actual efficiencies vary, and so are also considered tunable parameters. The flow rates through the systems are also tunable because some systems are operated continuously while others are intermittent. In summary, RACKLIFE allows the user to specify the efficiency, flow rate, and start and end times for operation of the cleanup systems. An end time far in the future is used to indicate a continuously operating system.

During refueling outages, reactor coolant is exchanged with spent fuel pool water via the fuel transfer canal. While not intended as a "cleanup," this often results in a considerable dilution of the spent fuel pool silica concentration (with the transferred silica ultimately removed by the reactor water cleanup system). This rapid drop in pool silica concentration is enhanced by the typical practice of operating both the spent fuel pool and reactor coolant system cleanup systems at full capacity during the refueling outage. When a RACKLIFE user observes a sudden drop or rapid reduction over time of the measured spent fuel pool silica concentration, the change is modeled as a *dilution*. In RACKLIFE this is also referred to as a *letdown*. The use of letdowns, their magnitude, and their duration are also tunable parameters at the discretion of the user.

4.3 Using RACKLIFE

As discussed in [EPRI 2002a], RACKLIFE has five principal uses:

- to estimate average B_4C losses and pool silica concentration;
- to manage the movement of spent fuel in the pool;
- to plan a BADGER measurement campaign;
- to corroborate BADGER measurements (e.g., during a campaign); and
- to predict future Boraflex losses – rack average losses, and specific panel losses if BADGER loss data is available for that panel.

Each of these is described in turn below.

4.3.1 Estimating Average B_4C Losses and Pool Silica Concentration

It is again emphasized here that RACKLIFE is designed to simulate an average boron carbide loss from a given Boraflex panel. As discussed in Section 3.8, every panel cavity has its own pattern of fluid "escape" (inflow and outflow) that can change over time as the Boraflex panel degrades in its own unique way. Thus, the actual distribution of loss from a panel is unpredictable. RACKLIFE predicts a panel loss that is independent of the distribution of that loss, sometimes referred to as a *panel average loss*. With respect to RACKLIFE results, this is often shortened to *panel loss*, or simply *loss*, where the average over the panel must be inferred.

The escape coefficient used when developing a RACKLIFE model is typically an average value applied to all of the Boraflex panel cavities throughout the spent fuel pool, as discussed in Section 4.2.3. This is a different kind of averaging process, whereby the maps and files of data typically produced by RACKLIFE using average escape coefficients illustrate the state of the panels as if the flow characteristics through all of the panel cavities are identical. In this case, the panel average loss is computed using an average escape coefficient. As discussed in Section 4.2.3, escape coefficients for each panel cavity can be specified individually based on additional knowledge, such as loss data from BADGER measurements, for more robust analyses. However, many analyses make use of only average escape coefficients.

Panel average losses using an average escape coefficient are nonetheless useful. Panels can be identified that merit near-term action, such as implementing loss mitigation strategies as discussed in Section 2.3 (e.g., checkerboarding), or scheduling BADGER testing, depending on the severity of what is observed. Pool maps of panel average losses are useful for developing boundaries for dividing the racks into *zones*—groups of contiguous rack cells and their associated Boraflex panels – based on common amounts of loss in each zone. Each zone can then utilize mitigation strategies specifically tailored to some bounding estimate of loss to all of the panels in the zone: for example, burnup credit such that only spent fuel with sufficient burnup may be placed in a particular zone, and/or different degrees of checkerboarding in each zone, such as 1-of-2, 1-of-4, etc. Maps of panel average losses using average escape coefficients can also help identify zones that can be depended on in the future to be relatively intact. While RACKLIFE cannot predict local dissolution effects, such as scallops, gap widening, etc., such effects have typically not been

observed in low dose panels predicted to have low loss. A dose below the $5 \cdot 10^8$ rad threshold in Table 3-1 is always considered low dose and will be predicted to have a low loss.

As discussed in Section 4.2, RACKLIFE will also predict the spent fuel pool silica concentration. Plant personnel can use these predictions to estimate the rate at which filter media and demineralizer resins will become fouled with silica so that maintenance schedules and budgets can be adjusted accordingly. Predictions of pool silica levels are also useful for planning cleanup operations prior to a refueling outage, when the spent fuel pool will be connected to the reactor coolant system and silica may both diffuse and be transported into the reactor coolant system, as discussed in Section 2.3. The quantity of crystalline silica released from Boraflex can also be calculated from predicted boron carbide losses by plant personnel to estimate when pool clarity problems might begin to arise, as discussed in Section 3.7.

4.3.2 Managing the Movement of Spent Fuel in the Pool

Section 2.3 discussed strategies that plants have followed to preserve zones and to sacrifice zones in the spent fuel pool. Figure 4-6 shows that panel cumulative gamma dose is the strongest factor in determining the silica release rate. Thus, controlling the dose from the spent fuel assemblies to the Boraflex panels can have a significant effect on panel loss. The maps of panel loss noted in Section 4.3.1 can be used to define boundaries for these preservation and sacrificial zones. Different spent fuel assembly discharge and in-rack movement strategies can also be tested.

4.3.3 Planning a BADGER Measurement Campaign

RACKLIFE can generate graphical maps and text files listing the computed absorbed dose and B_4C loss for every Boraflex panel in the spent fuel pool. RACKLIFE can list these in order of loss, dose, or panel ID. An example of a loss map is illustrated in Figure 4-7 showing the distribution of panel loss in a simulated spent fuel pool. RACKLIFE can also zoom in on particular rack modules to graphically and quantitatively identify specific panel doses and losses. In this example, the plant has typically discharged fuel from the fuel transfer canal to the closest rack modules, resulting in clearly identifiable regions of high loss (red), moderate loss (yellow), low loss (green), and near-zero loss (blue). The bounds on these color regions are selected by the RACKLIFE user each time a map is generated.

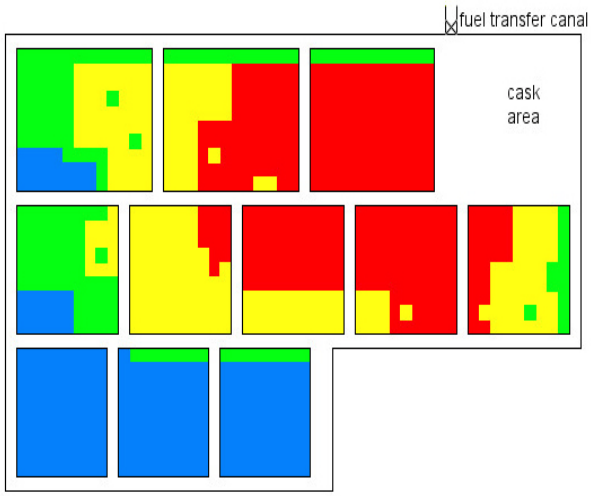

Figure 4–7. Example of a RACKLIFE simulated spent fuel pool showing regions of high (red), medium (yellow), low (green), and near-zero (blue) loss.

RACKLIFE results are also used to produce panel population dose and loss curves, which visually display the spectrum of predicted panel doses and losses for a spent fuel pool. These can be used to select panels to be measured by BADGER. Figure 4-8 illustrates an example of a population spectrum as cumulative dose and percent loss curves for a RACKLIFE simulation of a few thousand Boraflex panels in a spent fuel pool. In this example, the maximum dose panel is $1 \cdot 10^{10}$ rads, and the maximum loss panel has sustained 22 percent loss, as shown the upper right corner of each graph. Looking at the left graph (a), the red curve corresponds with the red axis of cumulative dose. The red curve is a sequence of points of each panel simulated producing a trend curve showing the fraction of panels which received a certain amount of dose. In this graph, about 18 percent of the panels in the simulated pool have received close to zero dose, and just over 50 percent of the panels have sustained $2 \cdot 10^{9}$ rads or less of dose; 100 percent of the panels have sustained 10^{10} rads or less.

Looking again at graph (a), the blue points represent the same panels as the points that make up the red line. Each blue point is the percent loss, as measured along the blue top x axis, for a panel. The degree of scatter in the blue points indicates the distribution of losses associated with a particular dose. For example, panels with doses of around $6 \cdot 10^{9}$ rads (bottom x axis to red curve) range between 5 percent and 14 percent loss (top x axis to blue points). This again illustrates a point made in regard to Figures 4-3 and 4-4 that dose is a necessary but not a sufficient predictor of loss.

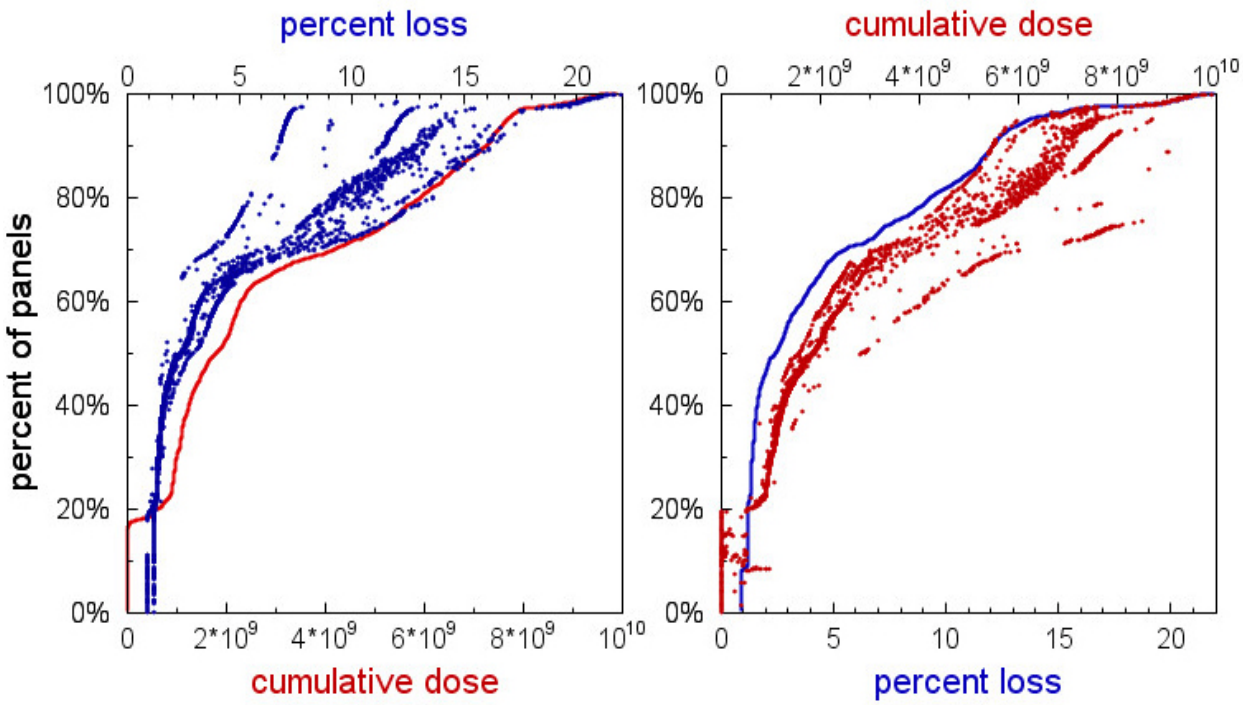

Figure 4–8. Sample cumulative dose (left) and percent loss (right) distribution curves.

Figure 4-8b shows the same information, except that the focus is reversed. For example, the blue curve shows that about 60 percent of the panels (y axis) have 4 percent loss or less (bottom x axis), and the panels with around 4 percent loss have received between about $2 \cdot 10^9$ and $4 \cdot 10^9$ rads. The range in dose comes about because the integration of dose over time is what drives panel degradation and leads to panel loss. Recall that the degradation and dissolution of Boraflex proceeds in what is effectively a two-step process—gamma radiation and free oxygen breakdown of the PDMS into silica, and dissolution of the resulting silica, with corresponding B_4C loss, into the pool water. Although a panel may have received a large cumulative dose, if that dose was very recent, then the degraded PDMS (i.e., now soluble silica) in the panel has not been exposed to the pool environment long enough to have significantly dissolved. In contrast, a panel receiving only a moderate dose dozens of years ago may have undergone substantially more silica and B_4C loss over time.

Note that, despite 18 percent of the panels having close to zero dose, the smallest loss is around 1 percent, not 0 percent. This is a consequence of uncertainty in the low-dose silica dissolution rate. Looking back at Figure 4-6, the release rate data at low doses shows considerable scatter due to difficulties in measuring the corresponding low aqueous silica concentrations (recall Section

3.6). It was decided to use the best fit release data at $5 \cdot 10^8$ rads as representative of the release rate for all doses less than that.

The actual panels measured in a BADGER campaign are a subset of these dose and loss spectra based on how many panels can be scanned in the time allotted to the measurement campaign, and how easily the spent fuel rack cells around the panels can be emptied of spent fuel. In making the limited selection from cells that can be made available, panels are selected with a wide distribution of predicted dose and loss, but with an emphasis on the higher predicted dose and loss panels. Dose is used as a selection variable because, as discussed in Sections 4.2.3 and 4.2.4, the calculation of loss is dependent on properly tuning the escape coefficient and cleanup system data, which is subject to greater uncertainty than the dose calculation, especially if no prior BADGER data has been obtained to refine the RACKLIFE model.

A mixture of panel orientations in the rack modules may also be sampled. For example, rack designs that use the "L-wrapper" cover plates illustrated in Figure 2-3 may exhibit different loss characteristics depending on the panel orientation due to fabrication asymmetries with respect to the cover plate bends and welds. A number of panels all from the same rack module may also be selected to observe the similarities and variations in the distributions of gaps and local dissolution effects between panels that presumably experienced very similar manufacturing processes. Likewise, panels from different rack modules may be selected to observe the same similarities and variations. In practice, for a given rack design, the degree of similarity and variation has typically been uniform across the racks in a given pool. For example, all of the rack modules which use L-wrappers will tend to show similar characteristics of gap and dissolution effects. Significant differences between individual rack modules of similar design have often been attributed to external events such as dropped or stuck assemblies, as noted subsequently in Section 4.3.4. In selecting panels, it is particularly important to identify some zero dose panels, or at least very low dose panels, below the $5 \cdot 10^8$ rad threshold for accelerating loss rate, to establish an intercept for the measured loss versus BADGER detector signal calibration curve. While not necessarily a truly zero loss panel, the dissolution mechanism will only attack the cut edges of the zero dose Boraflex panel as noted in Sections 3.4 and 3.8, leaving the panel substantially intact. This is discussed further in Section 5.3.

The silica kinetics and panel loss tracking in RACKLIFE is specific to the dissolution of Boraflex. However, as noted in Section 4.1, RACKLIFE could be used to plan a BADGER campaign of non-Boraflex panels using only the predicted panel absorbed dose, provided that the degradation mechanism of that panel material can be correlated to gamma dose.

4.3.4 Real-Time Corroboration of BADGER Measurements

During a BADGER measurement campaign, RACKLIFE predictions give the BADGER operator and data analyst a broad sense of what is expected and can indicate that there is a malfunction in the BADGER instrument. If RACKLIFE predicted small losses but large losses are observed

during the subsequent BADGER scan, such as gross panel loss along an edge, the first step is to examine the BADGER system to rule out experimental error. For example, disturbing the power or signal cables during a scan can cause noise spikes in the output data. Occasionally, a detector will suffer water ingress, causing spurious signals that can look like large areas of local dissolution, and this can be checked for if the results contradict expectations. A rescan can also be performed to observe if the anomaly is transient.

The next step is to question the service history of the panel (e.g., histories involving the motion of dropped or stuck assemblies that could have bent a panel cover plate such that the escape coefficient became substantially larger). Such events have correlated with BADGER observing higher losses than RACKLIFE predicts. When the information is available, anomalous service histories can be an indicator to select specific panels for BADGER testing.

4.3.5 Predicting Future Boraflex Losses

As emphasized in Section 4.3.1, RACKLIFE by itself can only predict panel average losses. Thus, future predictive calculations of loss for any but the lowest dose and lowest loss panels (as noted in Sections 4.3.3 and 4.3.4) may have too much uncertainty to be useful when used alone. Only by iteratively correcting RACKLIFE predictions against measured losses (e.g., by using in situ BADGER measurements of actual loss) can the results become useful for quantitatively understanding the current state of the panels and for predicting their future loss of boron carbide.

The evolution of a RACKLIFE prediction model is illustrated in Figure 4-9. The development of the RACKLIFE model begins with the user input of an extensive amount of data, as indicated by the gray data parallelograms across the top and as described previously in this chapter. Each data parallelogram corresponds to a separate text data file used by the RACKLIFE code. Once this data is entered, an initial conditions state point file can be created as described at the end of Section 4.2.2. This results in an initial state point corresponding to when the Boraflex racks were installed in the spent fuel pool. This state point file is indicated at the top of the model development timeline shown in Figure 4-10.

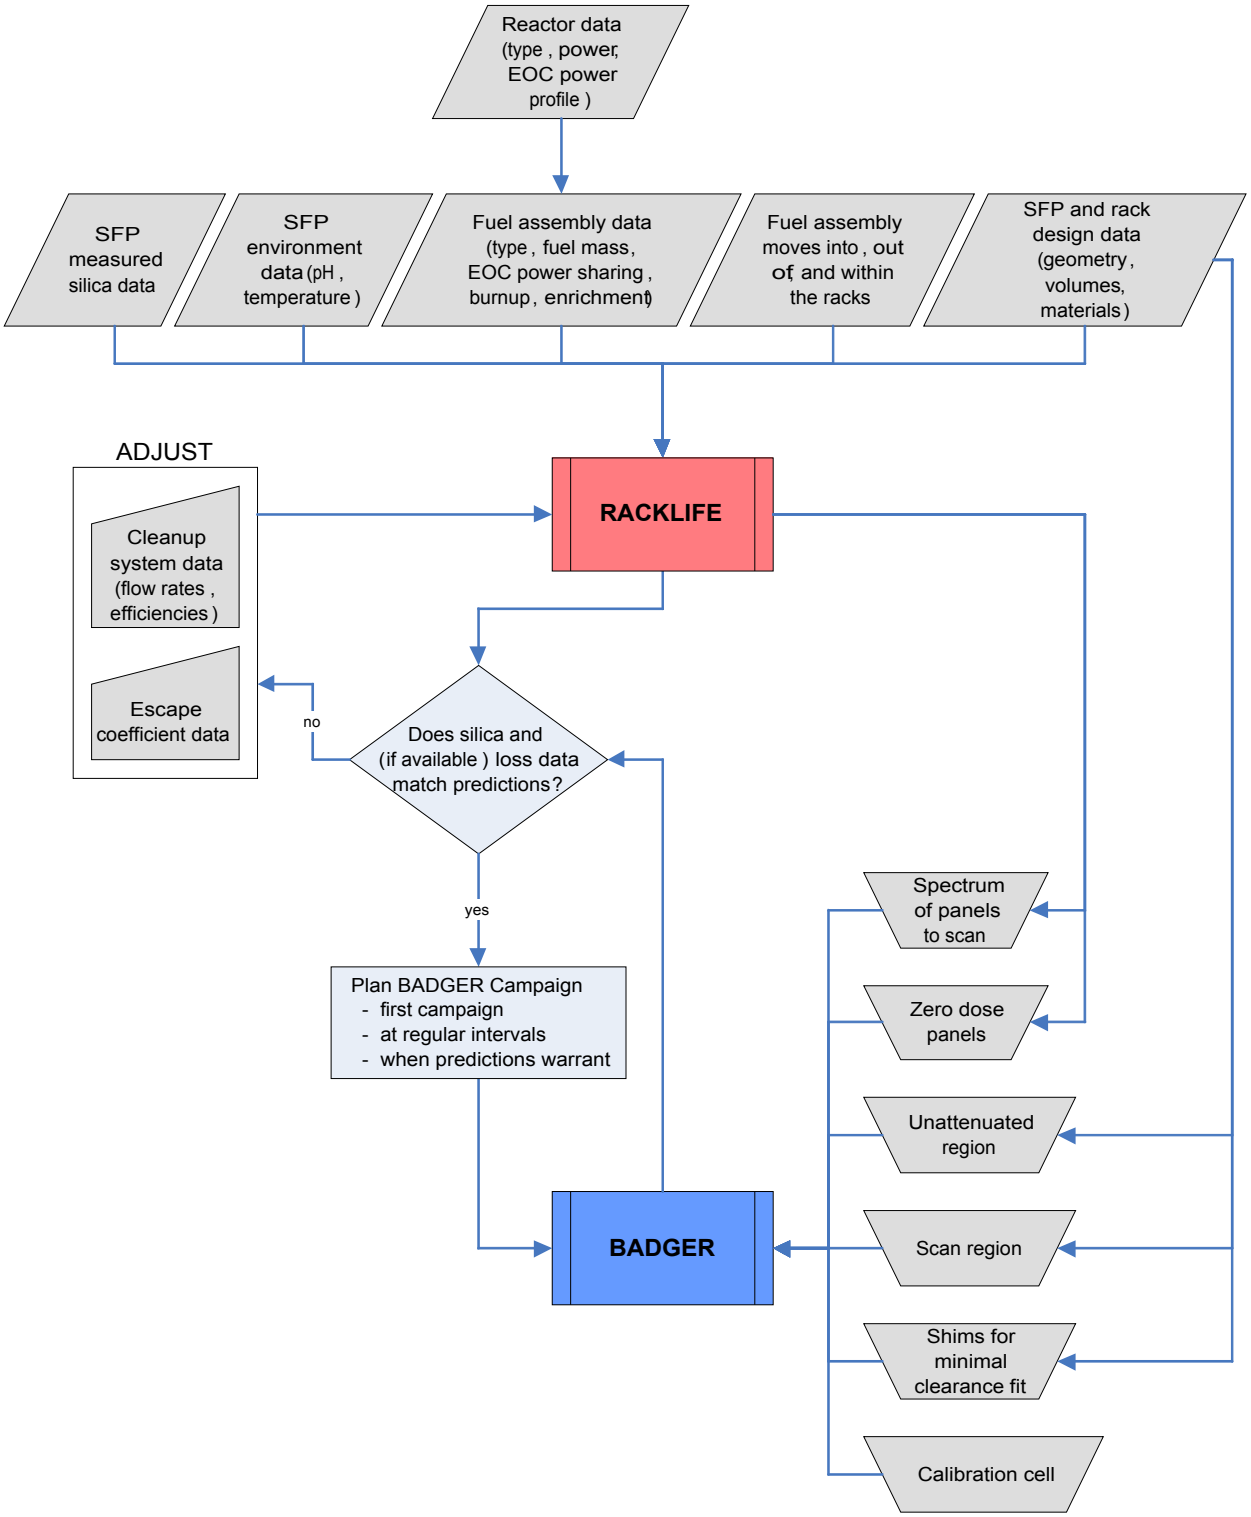

Figure 4–9. Interaction between RACKLIFE and BADGER.

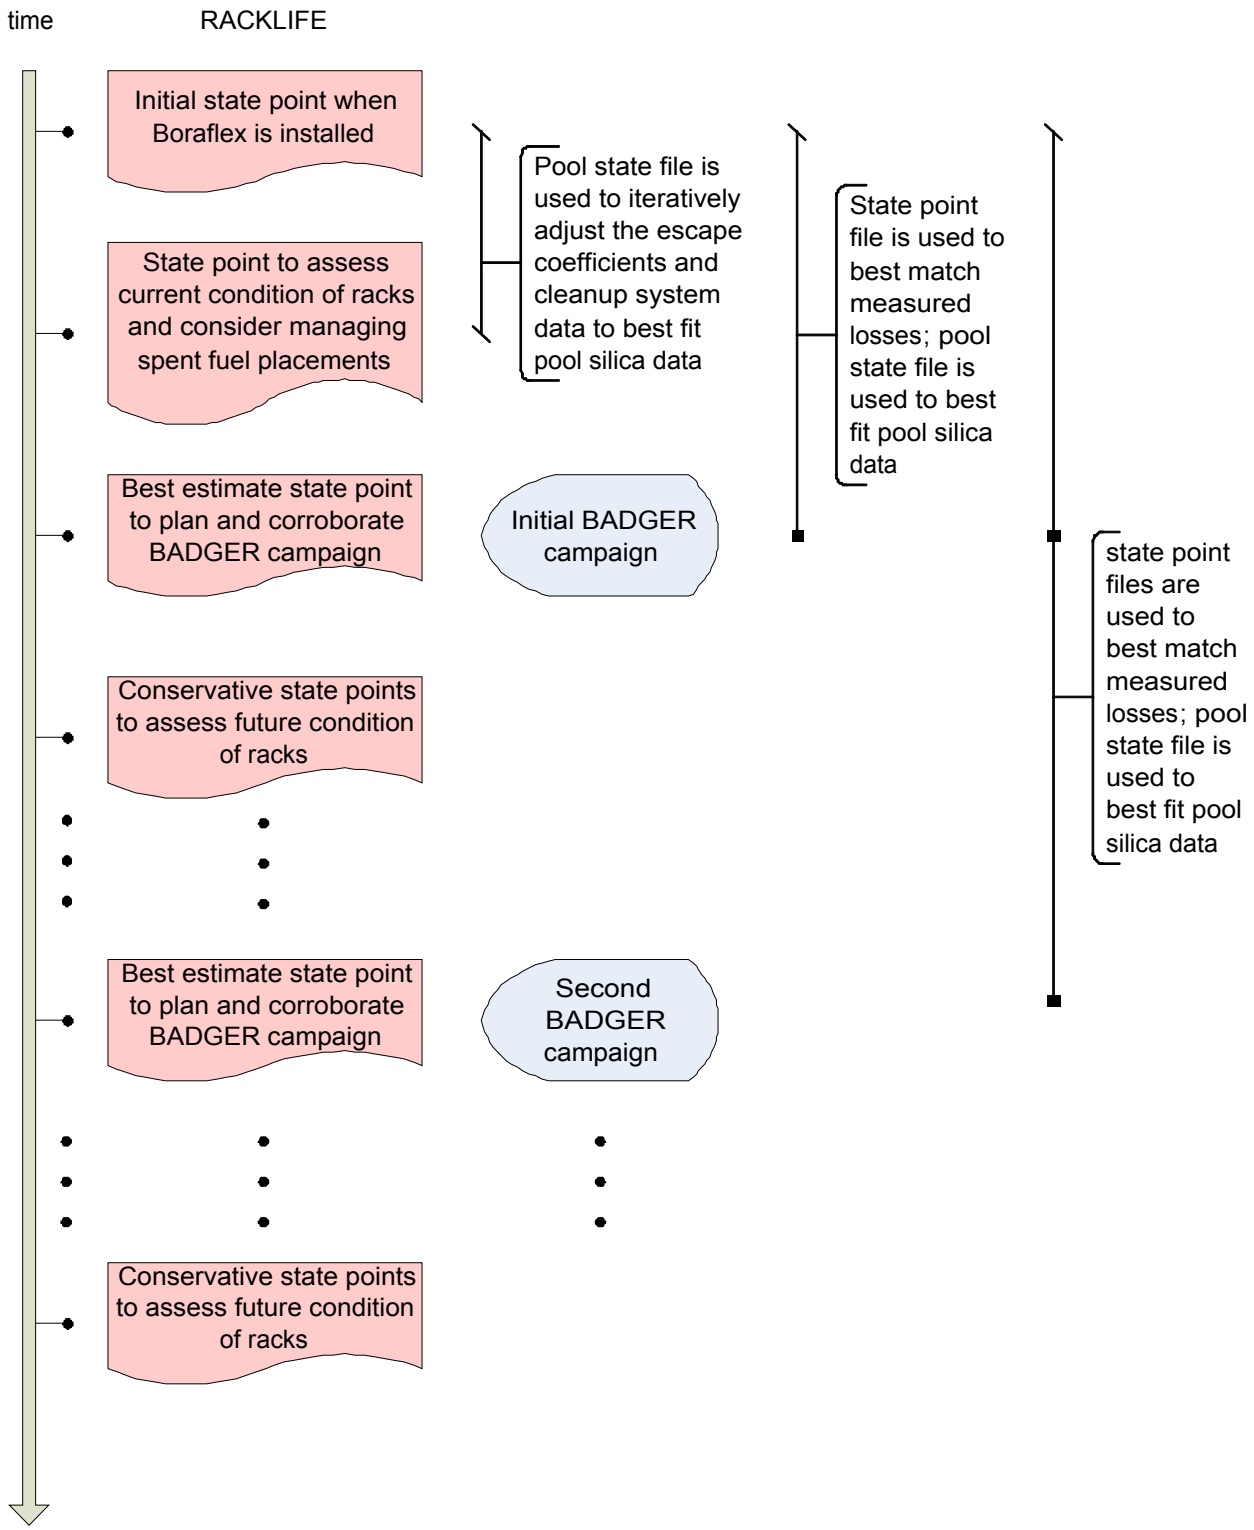

Figure 4–10. Use of RACKLIFE and BADGER over time.

RACKLIFE is then executed using an initial guess for the escape coefficient. The initial guess is typically based on the experiences of other RACKLIFE users, if any, and especially those modeling spent fuel racks of similar design. Dilution "cleanup system" events, as described in Section 4.2.4, are also included to model rapid drops observed in the silica data; these drops typically occur during refueling outages when the spent fuel pool water mixes with the reactor coolant system water. The results are plotted so that the user can compare the predicted bulk pool silica concentration over time with measured concentrations as previously input to the code. As indicated by the "no" loop through the decision diamond in Figure 4-9, the escape coefficient and cleanup system data (such as dilutions, system efficiencies, and system flow rates, as discussed in Section 4.2.4) are adjusted such that RACKLIFE predictions best match the measured silica data (as discussed in Section 4.2.3). An ideal fit is shown in Figure 4-11, which plots an example showing the measured and predicted bulk pool silica concentration as a function of time. This figure will be further discussed below.

The user can now execute a RACKLIFE state point calculation from the initial conditions state point to the current date. This will give a RACKLIFE best estimate of the current condition of the Boraflex panels. Based on these results, the user may conclude that a BADGER test campaign is warranted. The user then executes a best estimate state point calculation from the current date to the expected start date of the BADGER measurement campaign. This is the third state point (in red on the timeline) shown in Figure 4-10. From this new state point, the dose and loss spectrum of the panels can be studied to select a subset for testing as indicated in Figure 4-9 and discussed in Section 4.3.3. Figure 4-9 also shows that rack design data acquired for input to RACKLIFE will also be used to plan the BADGER scans and determine whether shims will be needed; this will be discussed further in Chapter 5.0. During the BADGER campaign, the RACKLIFE results can also be used to corroborate the BADGER data as discussed in Section 4.3.4.

Figure 4-11 shows a hypothetical idealized match between measured and predicted spent fuel pool silica. The silica data points show very little scatter in this figure; reference [EPRI 2000a] shows numerous plots of reactive silica data measured in the spent fuel pool, reactor coolant system, and reactor water storage tank versus time, which exhibit varying degrees of scatter. The scatter is typically more pronounced at low levels, particularly below 1 ppm, which is most common for BWRs, as noted in Section 2.3. At very low levels other sources of silica, such as some paints, epoxies, and resins, can also become significant. Nonetheless, most plots in the reference show very clear trends in spent fuel pool silica which can be fit visually by the user based on plots generated by RACKLIFE.

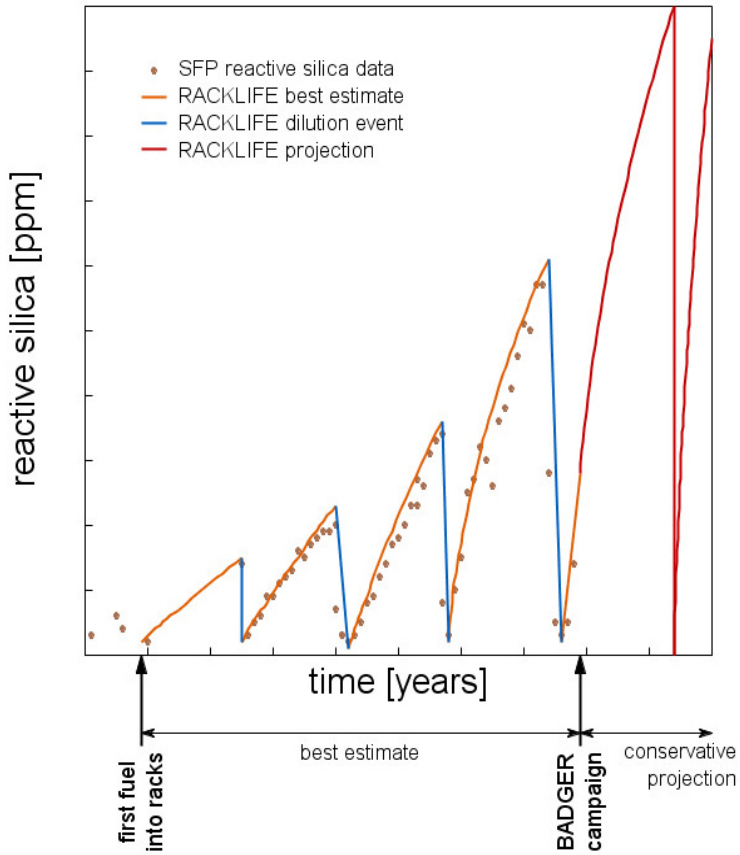

Figure 4–11. Matching silica data in RACKLIFE and projecting the future.

Figure 4-11 also shows that the trends are piecewise continuous. At each outage, the spent fuel pool silica concentration typically drops over a matter of days or weeks. This is due to some combination of:

- dilution with the low silica water from the reactor coolant system (< 1 ppm to satisfy fuel warranties), which is coupled to the spent fuel pool during fuel transfers;
- spent fuel pool cleanup systems operating before refueling to reduce the amount of silica that will be transferred to the reactor coolant system and/or to increase visibility in the pool by reducing the suspended crystalline silica; and
- dilutions from spent fuel pool makeup water storage tanks.

It is not important how much each of these factors contributes to the overall dilution; only the total magnitude of the dilution is important.

As discussed in Section 4.2.4, the dilution event is modeled in RACKLIFE by the input of a "letdown" in the cleanup system data. This will cause a ramp down in the pool silica from the current concentration to a specified concentration over a specified amount of time. The specified end concentration and duration of the drop are chosen by the user to best fit the bulk pool silica data. A letdown only affects the bulk pool silica concentration; the cavity concentrations then decrease over time at a rate governed by the escape coefficient.

Difficulty in matching the calculated silica data to measured silica data most often correlates with uncertainty about cleanup system utilization and efficiency. Those pools that do not have clearly observable trends in silica concentration are often pools with higher flow rate and/or higher efficiency cleanup systems continuously removing the silica from the pool. BWRs in particular, because of the damage silica can do to the turbines, keep their water systems cleaner of dissolved silica. As noted in Section 4.2.4, cleanup systems are more efficient in BWRs compared to PWRs, given the absence of dissolved boron in the spent fuel pool. In BWR pools typically both the escape coefficient and a combination of the cleanup system flow rate and efficiency are tunable parameters. In general, the escape coefficient is adjusted to predict the measured magnitude of pool reactive silica at the time of the current state point, while the cleanup system is adjusted to match the curvature of the silica data from the chosen initial state point to the current state point.

Figure 4-11 indicates a hypothetical initial BADGER campaign after the fourth letdown event. As discussed above, a best estimate RACKLIFE model is used to plan and corroborate the BADGER campaign. To predict future panel boron carbide loss, however, a conservative approach can be taken. For example, the temperature can be set to its highest anticipated value—based on previous measured temperatures plus a bit more to account for potential spent fuel gamma heating in the panel cavities—for maximum silica release. To maximize anticipated gamma exposure, refueling outages can be assumed to occur at the earliest anticipated date, with the instantaneous discharge of identical highest gamma source term bundles to every cell in the pool. Letdowns can be modeled as instantaneous to drop the silica concentration effectively to zero for the duration of the outage, again maximizing the release of silica from the Boraflex panels. In this manner, a conservative bounding state of the future condition of the panels can be predicted.

This conservative prediction can be used to decide when it is appropriate to schedule a followup BADGER campaign to assess the actual condition of a selection of panels as shown in Figure 4-10. For example, a current criticality analysis may assume that panels have a maximum of 40 percent loss, while BADGER measurements show a maximum of 10 percent loss so far. A subsequent BADGER campaign could be scheduled for a date before conservative projections indicate that the worst panel may be near 40 percent loss, although the best estimate maximum loss (without maximized conservatism) may be less than 20 percent loss.

In practice, RACKLIFE state point calculations are typically only performed following a refueling outage—every 18 to 24 months. However, a plant may run the RACKLIFE program more often if previous RACKLIFE calculations suggest that the average panel losses of some panels approach a limit where the panel is no longer viable for neutron absorption. Fuel moves from the recent outage are input to the code and, as shown in Figure 4-10, either a best estimate state point calculation to the next BADGER campaign or a conservative state point calculation to the next refueling outage or RACKLIFE run is executed. Because of this infrequent use, users make significant use of the user's manual and tutorials, EPRI reports that offer guidance, RACKLIFE workshops, and fellow user support to refresh their skills.

4.4 Analysis of Uncertainties in RACKLIFE

This section details an analysis of the uncertainties inherent in a RACKLIFE model. First, the fundamental silica kinetics model is examined. Then, using a program written specifically for this analysis, variables within RACKLIFE are examined for their effect on the uncertainty of the panel loss predictions. Finally, a summary table is presented.

As discussed in Section 2.3, spent fuel pool rack designs (including as-built variations such as defects) and service histories vary considerably from plant to plant. RACKLIFE uncertainties are strong functions of the rack design and service history, and so the uncertainties derived for one spent fuel pool are likely inappropriate for another. Further, global bounds on RACKLIFE uncertainties would typically be so large as to be useless. An objective of this section is to present techniques for analyzing RACKLIFE uncertainties, as well as quantitative results from applying these techniques to a select set of design and service history conditions. These uncertainties are only for illustrative purposes and to guide analysis efforts to target the largest uncertainties. The uncertainties presented here should not be considered representative of any particular spent fuel pool.

4.4.1 Silica Kinetics Model Uncertainties

The RACKLIFE silica model described in Section 4.2 is a multistate, multicompartment model that is nonlinear by its dependence on the equilibrium reactive silica, as expressed by the "proximity to equilibrium" factor, $R_x = 1 - R/R_{eq}$. Other than this factor, linear constant coefficient exchange rate kinetics between states and compartments is assumed. It is possible that, under certain conditions, the linear kinetics model is not ideal. For example, very low escape coefficient panel cavities might be better modeled by including a diffusion process. However, very low escape coefficients result in low silica loss rates, and thus low overall Boraflex panel loss. Therefore, increasing the model precision by including diffusion processes would have had an insignificant impact on predicted panel loss.

For panels that have experienced levels of degradation greater than about 40 percent, the linear exchange rate kinetics assumption may not be appropriate. As described in Section 3.8, a high-level of loss of Boraflex material will generally manifest as numerous large-scale "local" dissolution anomalies (e.g., large scallops, large gap widening, long dissolution flow paths) in addition to uniform dissolution. These anomalies may grow large enough or combine to create significant natural convection flow paths.

For example, Section 2.3 noted that the panel cavity formed by the rack wall and the wrapper plate provides a clearance fit for the contained Boraflex panel. The clearance perpendicular to the Boraflex panel face for intact Boraflex can be about 0.01 inch; as the Boraflex panel thins, this clearance will increase, to a "clearance" of about 0.135 inches if the panel is missing. Panel cavity clearances in this range are sufficiently small that the cavity flow will be a very slow fully developed laminar flow with a relatively small rate of convective heat transfer to the cavity water from the rack structure being heated by the spent fuel gamma radiation. In areas where the Boraflex dissolves from its as-built thickness to zero, the Grashoff number (characterizing the natural convection flow) and the Rayleigh number (characterizing the natural convection heat transfer) can increase from $Gr \sim 1 \cdot 10^{-3}$ to $Gr \sim 1$, and from $Ra \sim 5 \cdot 10^{-3}$ to $Ra \sim 10$, thereby increasing the amount of flow and the heat transfer to the flow by 3 to 4 orders of magnitude.

The character of the natural convection flow will strongly depend on the orifices by which the flow enters and leaves the cavity (e.g., the resistance to flow through small orifices will reduce the natural convection potential, while larger orifices, such as an aperture due to a buckled cover plate, or the natural openings in some cover plate designs, will not impede the flow). Thus, there are cases where natural convection could strongly couple the local dissolution rate to radiation induced heating from neighboring spent fuel assemblies, making the exchange rate kinetics superlinear—that is, faster than the linear assumption used in modeling the kinetics rate parameters. Such effects may have already been observed in practice for moderate (20 percent) levels of loss.

In RACKLIFE, these effects are addressed by adjusting the escape coefficient to where the calculated concentration of reactive silica matches the measured bulk pool concentration of reactive silica. Typically the escape coefficient is implemented as a constant throughout the time of the RACKLIFE run, even if the initial and final state points are years apart (i.e., the constants may be different for each panel but they do not change with time). For some pools, however, the only way to closely match the measured silica data has been to increase the escape coefficient over time. This is done in RACKLIFE by computing silica concentration stepwise over a sequence of state points, instead of one initial and one final state point. The stepwise state points are usually positioned at an outage, and the escape coefficient is increased at each state point to match the measured silica trend. For moderate levels of loss such as 20 percent, this approach has given reasonable results when compared to BADGER data. For high levels of degradation, such as above 40 percent loss, there is insufficient data available at this time to assess the adequacy of the kinetics model used in RACKLIFE.

4.4.2 Methodology to Assess Loss Variation due to Variation in Kinetics Model Variables

To quantitatively assess the uncertainties in the RACKLIFE model for low to moderate loss panels, a Fortran program was written to apply the RACKLIFE methodology as detailed in [EPRI 1997a] and summarized above in Section 4.2. The program simulates only one typical Boraflex panel in a typical spent fuel pool. The basis for the model is a PWR egg-crate rack, but because a range of dose exposure histories are examined, the results are broadly illustrative of what might be observed in BWR egg-crate and PWR flux-trap rack designs. However, as noted at the beginning of Section 4.4, the quantitative results are specific to this model and should not be considered directly applicable to any actual rack. Because only one panel is simulated, this analysis does not capture the stochastic distribution of uncertainty (e.g., in the escape coefficient) over a sample of panel cavities. Nonetheless, it provides quantitative guidance on how well model variables must be understood, and under what conditions, to have confidence in the results from RACKLIFE.

The program was instrumented to study the effects of variations in key model variables on the predicted Boraflex loss for the following *condition parameters*:

- multiple projection times
  (t = 1 to 10 years, in 1 year increments);
- low, medium, high, and very high panel absorbed doses
  ($d = 1\cdot10^9$, $6\cdot10^9$, $1\cdot10^{10}$, and $1.5\cdot10^{10}$ rads);
- very low, medium low, medium high, and very high escape coefficients
  ($e = 0.01$, 0.1, 1.0, and 10.0 day^{-1}).

Each run of the program will use a particular *triplet* of condition parameters, (t, d, e). Each triplet will stress the variable being tested in a unique manner. A total of $10 \cdot 4 \cdot 4 = 160$ (t, d, e) triplets will be used in 160 unique program runs so that a spectrum of panel and panel cavity conditions are examined for each variable tested.

Given the logarithmic nature of these condition parameters, their base-10 log transformations are used in regression analyses, such that they actually vary over the intervals $t \in [0, 1]$, $d \in [9, 10.2]$, and $e \in [-2, 1]$. Taking log transformations of these independent variables is essentially the same as plotting on graph paper with a logarithmic x axis; order of magnitude differences will appear as linear. Regression calculations are performed using the transformed variables.

For the instrumented RACKLIFE analysis performed in this report, each uncertainty variable was given both a *static* treatment and a *dynamic* treatment. In a static analysis, all other variables were held static while a deviation was applied to the uncertain variable. However, in actual practice RACKLIFE is used to obtain predictions based on a match between the calculated and measured silica data. Once all of the data is input, and RACKLIFE generates a calculated silica

concentration, the user dynamically tunes the escape coefficient (and possibly cleanup system data, though this tuning variable was not exercised in this study) until the spent fuel pool silica predictions from RACKLIFE match the measured silica data. For example, if the RACKLIFE user were to change the value of a variable prior to a state point calculation, and the change caused the calculated silica concentration to mismatch the measured silica concentration, then the user would iteratively adjust the escape coefficient (and/or cleanup system data) and rerun the state point calculation until the predicted and measured silica concentrations again matched.

In short, a static analysis changes only the variable of interest without regard to matching the predicted and measured silica concentrations. A dynamic analysis changes the variable of interest, but then iteratively tunes the escape coefficient condition parameter to compensate for any mismatch between the predicted and measured silica concentrations. Dynamic interaction can dramatically mitigate the effects of uncertainty in a variable. However, the uncertainty from a variable does not disappear in a dynamic treatment. The uncertainty is simply transferred to the mismatch between the predicted and measured silica concentrations, and then some portion of the uncertainty is "tuned away" by varying the escape coefficient.

The process of estimating the Boraflex panel average loss variation due to variation in a model variable is detailed below for two illustrative model variables – the polymerization rate and the bulk pool temperature – and then summarized for all of the variables examined.

4.4.3 Polymerization Rate Uncertainty – Static Analysis

The reactive silica concentration in the panel cavity is limited by the reactive silica equilibrium concentration. As silica release from irradiated Boraflex increases the panel cavity concentration toward the equilibrium concentration, the proximity to equilibrium factor, R_x, approaches zero. At the equilibrium concentration, the release rate then also goes to zero, as does the overall panel loss rate. Fluid exchange between the panel cavity and the bulk pool, quantified by the escape coefficient variable, typically lowers the silica concentration and serves to increase R_x, thereby allowing continued release of silica and loss of B_4C. Another variable controlling reactive silica concentration is a silica kinetics model coefficient: the polymerization rate, k (as shown in Figure 4-5 and Equations 4-3 through 4-6). This coefficient measures the rate at which reactive silica will polymerize, thus removing reactive silica from the panel cavity, increasing R_x, and again allowing continued release and loss.

As detailed in [EPRI 1997a], the polymerization rate is derived from experimental data which has evident trends, but significant scatter. Thus, there is some inherent uncertainty in this value. The effects of polymerization rate uncertainty on loss calculations can not be tested by the user of the commercial RACKLIFE software because the coefficient is hard-coded into the silica model. The instrumented RACKLIFE model applied here, however, can perform this testing.

The first phase of the testing measures the *static variation* in loss due to variation only in the polymerization rate, k. For all 160 triplets of condition parameters (time t, dose d, and escape

coefficient e), the polymerization rate k is set to 50 percent and 150 percent of its nominal value, and the panel losses computed for these perturbed conditions are compared to the nominal loss for the hard-coded value of k. This ±50 percent variation is very likely much larger than the actual combined experimental and regression error, and so can be considered bounding.

The results of this test are shown in Figure 4-12. The x-axis represents the series of combinations of condition parameter triplets. On the y-axis, the 0 line on the graph represents panel loss calculated from the nominal hard-coded value of k. The red points show the relative change in loss when the polymerization rate is increased to 150 percent of nominal, while the blue points represent the relative change in loss when the polymerization rate is decreased to 50 percent of nominal. As can been seen on the graph, for some combinations of condition parameters, varying k does not significantly change the panel loss.

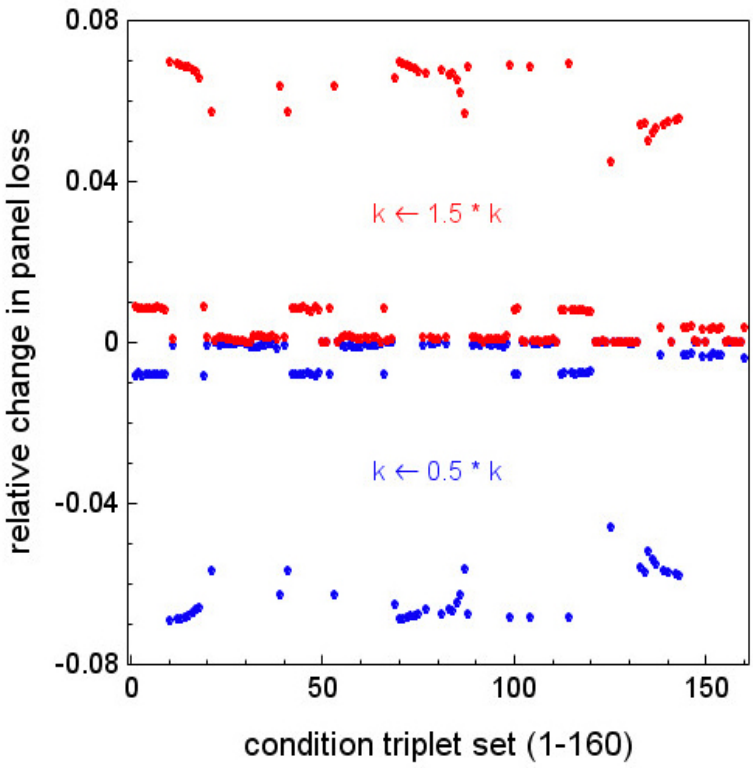

Figure 4–12. Uncertainty in panel loss due to ±50% variation in polymerization rate - static test.

The relevant value in the analysis is the maximum relative change in panel loss, which gives the upper and lower bounds to the uncertainty. In this case, varying the polymerization rate by +50

percent will increase the cumulative panel loss by a maximum of 0.07, as a fractional relative change (i.e. 7 percent). Thus, if the polymerization rate were really 50 percent higher than what is hard-coded into RACKLIFE, then a panel predicted to have 20 percent loss would actually have 20 percent · (1 + 0.07) = 21.4 percent loss. Similarly, the blue points show that decreasing the polymerization rate will decrease the cumulative panel loss by what is approximately a symmetric amount; the same panel would actually have 20 percent · (1 − 0.07) = 18.6 percent loss.

4.4.4 Identifying Condition Parameters Which Most Influence Uncertainty

Figure 4-12 shows that, while the maximum loss uncertainty for the ±50 percent variation is ±0.07, most of the uncertainties are closer to ±0.01. Thus, only certain condition triplets produce an uncertainty nearer to the 0.07 maximum value. To identify these governing conditions which maximize uncertainty and are statistically significant, a multiple regression with subset selection was performed. The process begins by postulating a fit of the form

$$\mathbf{\Delta_{loss} = b_0 + b_1\,[t] + b_2\,[d] + b_3\,[e] + b_4\,[t{\cdot}d] + b_5\,[t{\cdot}e] + b_6\,[d{\cdot}e] + b_7\,[t{\cdot}d{\cdot}e]} \qquad (4\text{-}9)$$

where

- Δ_{loss} is the relative uncertainty in the panel loss;
- t, d, and e are the base-10 logs of their respective condition parameters (time, dose, and escape coefficient); and
- $b_0 ... b_7$ are coefficients to be solved for by multiple linear least squares regression.

The above fit measures *univariate* dependencies on t, d, and e, the second order *interaction* dependencies [t·d], [t·e], and [d·e], and the third order total interaction dependence [t·d·e]. Other higher order interaction dependencies, such as $[t{\cdot}d^2]$, are not useful for this analysis; further, limited testing showed that these terms were generally not statistically significant.

There are several other possible subset models, which base the regression fit on a subset of the seven b_n coefficients, by effectively setting any unused $b_1 ... b_7$ coefficients to zero. For example, there is 1 subset model using all 7 coefficients, 7 using only 6 coefficients, 21 using only 5 coefficients, and so on, totaling 127 subset models. In the regression analysis, all 127 models were tested and ranked and further tested to check that small changes in the data set do not strongly perturb the selection by using a statistical method called a "bootstrap." This is an important check because when the same data is used both to search for the best subset model and to estimate regression coefficients for that model, the coefficients will be biased. The bias will often be negligible, but not necessarily; the bootstrap method is highly effective at detecting if this bias is significant, indicating that in this analysis the model should be rejected. After comparing all of the models, the most parsimonious model (using the fewest terms) that still fits the data well, as measured by the correlation coefficient, R^2, Mallow's C_p statistic, or other methods, and that has

negligible bias, is then selected. The governing conditions are then identified as having non-zero coefficients

For the analysis of static variation in polymerization rate, the +50 percent polymerization rate is most conservative in that it gives true losses higher than predicted. At the +50 percent rate, the selected model is

$$\mathbf{\Delta_{loss} = 0.0080 - 0.0020\ [d{\cdot}e]} \tag{4-10}$$

Among the 127 possible regression models that fit the data, the model in Equation 4-10 has the best combination of low number of independent variables, relatively high correlation coefficient, acceptable C_p, and negligible bias in the regression coefficients. Subsequent discussions of the factor [d·e] will refer to it as [dose × escape coefficient] for clarity. However, note that Equation 4-10 is derived in terms of the log-transformed condition parameters, not the actual dose or escape coefficient values.

Equation 4-10 indicates that the dose and escape coefficient combined govern the relative uncertainty in the panel loss for a +50 percent uncertainty in the polymerization rate. The negative sign means that it is the low [dose × escape coefficient] conditions that result in the high uncertainties. This is an important finding, because low [dose × escape coefficient] conditions generally result in low overall Boraflex losses. For example, although low [dose × escape coefficient] conditions will generate a larger relative uncertainty of 0.07 if the polymerization rate is 50% higher than nominal, the overall Boraflex panel will have low losses due to the same low [dose × escape coefficient] conditions. Therefore, there is little impact on the validity of RACKLIFE's projections. For higher [dose × escape coefficient] conditions, the absolute loss will be higher, but the relative uncertainty will be only 0.01 for a ±50 percent uncertainty in the polymerization rate.

The selection of negative dependence on [dose × escape coefficient] as the governing combination of condition parameters is best illustrated by the conditional effects plots in Figure 4-13. The x axes measure univariate and interaction of log-transformed condition parameters. The y axes measure the relative uncertainty in the Boraflex panel loss between the perturbed and the nominal polymerization rate. Except for the plots for [e] and [d·e], all of the plots show wide data scatter as each individual term, indicated in red in each plot, is varied along the x axis. In contrast, the lower right [e] and middle right [d·e] plots show strong trends with much less scatter. The subset selection analysis indicated that the second order [d·e] interaction term was a more robust measure of the data, and so was selected over the univariate [e] term.

Note, however, that the dependence of the uncertainty on [d·e] in Figure 4-13 is not actually linear. The [d·e] plot shows that refining the regression with an exponential fit will model the uncertainty very well. Such a fit represents unnecessary refinement to this analysis, however. It suffices to conclude that the hard-coded polymerization rate is not a major source of uncertainty in a

RACKLIFE prediction, unless the dose and escape coefficient are relatively low, in which case the uncertainty is not of concern because the panel losses are low.

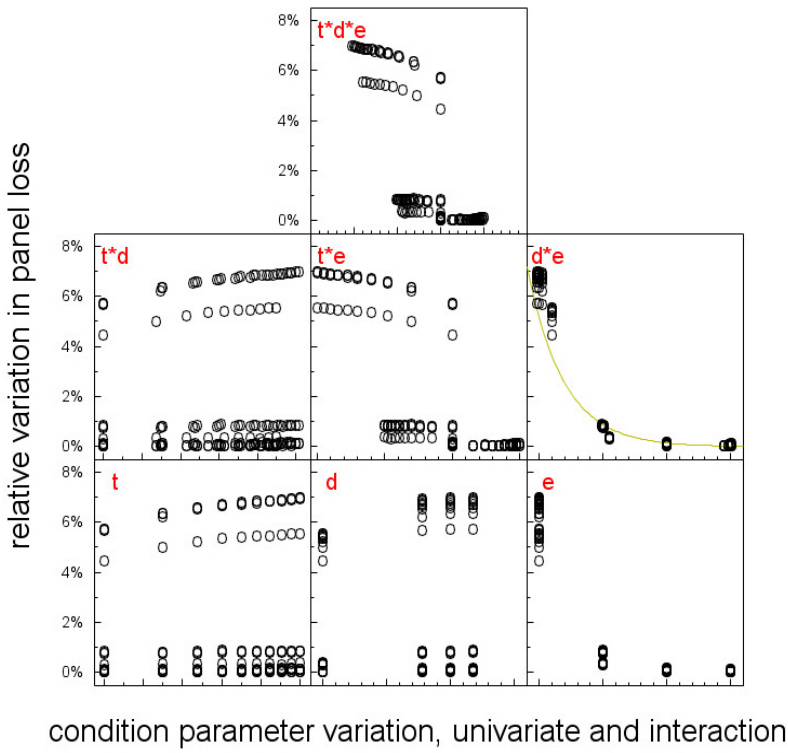

Figure 4–13. Conditional effect plot for the polymerization rate, k.

4.4.5 Polymerization Rate Uncertainty – Dynamic Analysis

Dynamic analysis of the uncertainty due to variations in the polymerization rate gives results that are not significantly different from the static results. As discussed in Section 4.4.2, in a dynamic analysis the escape coefficient is assumed to be adjusted to offset any changes in calculated pool silica caused by a change in the target variable. For changes in the polymerization rate, the highest static uncertainty occurs for low escape coefficient cavities; therefore, the reactive silica in the cavities is very close to the equilibrium concentration. The amount of polymer silica in the cavities will vary as the polymerization rate is varied, and so various quantities of polymer may

escape the panel cavity into the pool. However, this has a small effect on the pool reactive silica concentration, because, while some of the escaped polymer may depolymerize to reactive silica, the effect is relatively small compared to the effectiveness of the cleanup systems effective at removing polymer silica. Thus, regardless of the uncertainty in the polymerization rate, in a dynamic analysis the user will tune the escape coefficient to approximately the same measured reactive silica data. Therefore, the 0.07 fractional bounding relative uncertainty holds for low [dose × escape coefficient] systems, and 0.01 fractional relative uncertainty for most other conditions, under the conservative assumption that the polymerization rate could be uncertain by up to ±50 percent.

4.4.6 Temperature Uncertainty – Static Analysis

As discussed in Section 4.2.2, the silica release rate per unit area, λ, is a strong positive function of temperature, as is the equilibrium reactive silica concentration. Thus, a higher temperature will both significantly increase the rate at which silica is released, and increase the amount of silica that the cavity water can hold. Increasing the equilibrium concentration in the ith panel cavity, $R_{eq,i}$, relative to the given amount of reactive silica in the panel cavity, R_i, will move the panel cavity further from equilibrium, thereby increasing the proximity to equilibrium factor, $R_{x,i} = 1 - R_i / R_{eq,i}$ away from zero. Equation 4-2 shows that this will even further increase the silica source term. In the commercial RACKLIFE code, the temperature is input by the user as point-in-time data—a value and the date and time that the value was measured—based on measurements made in the spent fuel pool. RACKLIFE then simply linearly interpolates these values to get temperatures at all computation time-steps in the numerical integration. When computing projections into the future, the user may enter best estimate or conservative values; the last value entered is extrapolated as a constant for state point calculations with dates later than the last point-in-time data entered. Some plants record the bulk pool temperature daily, others only monthly or less. Because of the large sensitivity of the panel average loss, uncertainty in temperature could lead to significant uncertainty in loss.

In the present model, static testing again begins by examining the relative variation in loss for every possible triplet of conditions (time t, dose d, and escape coefficient e) while the escape coefficient is not adjusted to compensate. The temperature is varied by ±20 °F (±11 °C) from a nominal 96 °F (36°C) down to 76 °F (24°C) or up to 116 °F (47 °C). This range covers the majority of operating temperatures of the spent fuel pools, and the nominal temperature is higher than what most pools operate at during non-outage periods. Note, however, that the measurements are of the bulk pool temperature. The water in the panel cavities will be somewhat hotter than the bulk pool due to spent fuel radiation heating of the water and the surrounding racks. Prior to Version 2.1.1 in 2009, RACKLIFE did not include a heat transfer model for the panel cavities, so, as noted in Section 4.3.5, for conservative projections the bulk pool temperature would have been adjusted upwards.

The variation in loss for a ±20 °F variation in temperature from the 96°F nominal is shown in Figure 4-14. It shows that there is a maximum increase in relative loss of almost 80 percent for the worst case temperature increase and an over 40 percent maximum decrease in relative loss for the worst case temperature decrease. The multiple regression conditional effects plot in Figure 4-15 shows that the variation appears to depend on all of the condition parameters. The multiple regression subset selection bears this out, with all of the terms being statistically significant and five of the seven coefficients being statistically significantly different from zero. This can also be seen from the yellow lines in the conditional effects plots, which are univariate linear regressions with respect to the indicated condition parameter alone. All of these regression lines have shallow slopes and large data scatter around them. Thus, for variations in bulk pool temperature, one cannot easily localize the "worst cases" of loss uncertainty to any particular conditions of time, dose, or escape coefficient.

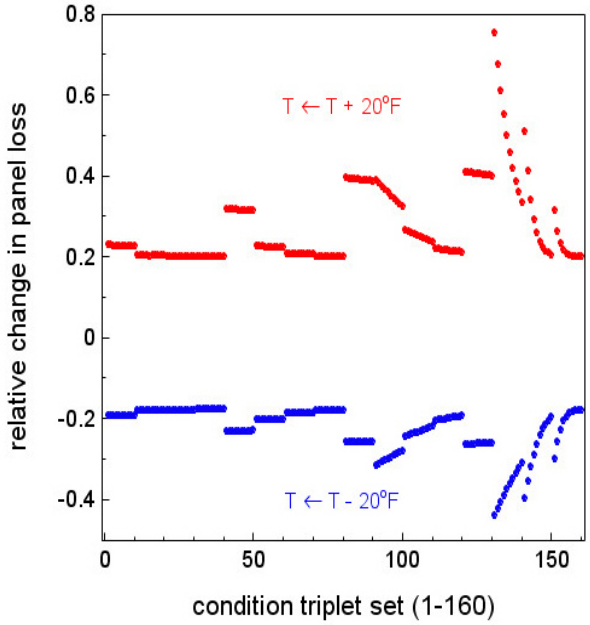

Figure 4–14. Uncertainty in panel loss due to ±20°F (11°C) variation in temperature - static test.

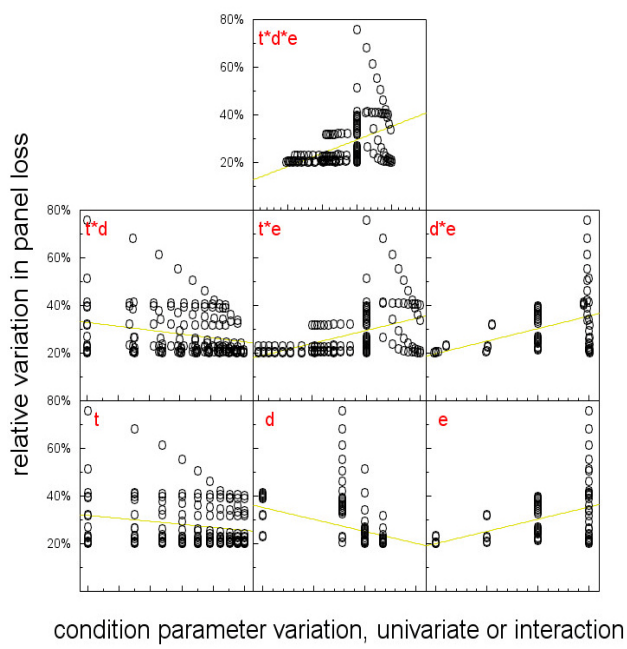

Figure 4–15. Conditional effects plot for bulk pool temperature, T.

4.4.7 Temperature Uncertainty – Dynamic Analysis

Despite the dire predictions of large uncertainty in the static analysis, the dynamic analysis clearly shows the compensating mechanism of the escape coefficient. Figure 4-16 shows the calculated pool reactive silica (bottom) and panel loss (top) over time for the condition triplet which gave the most uncertainty as temperature is varied. For this condition triplet, the panel dose $d = 1.5 \cdot 10^{10}$ rads, and the escape coefficient $e = 10.0\ d^{-1}$. As shown in the top plot, the nominal temperature of 96 °F (purple) predicted a 5.9 percent loss after 1 year and a 29.2 percent loss after 10 years. If the temperature in the panel cavity were actually +20 °F higher, the static analysis (red) predicts a 10.4 percent loss after 1 year (for a 0.75 fractional relative error) and a 39.0 percent loss after 10 years (0.33 relative error).

However, the nominal temperature case assumes that the escape coefficient had been tuned to match the calculated pool silica with the measured pool silica data, close to the purple curve in the bottom plot. When the temperature is varied +20 °F higher, the perturbed temperature case both predicts a higher percent loss from the panel (top plot red curve), and it significantly changes the calculated pool silica (bottom plot red curve). During the dynamic analysis, the escape coefficient is adjusted from $e = 10.0\ d^{-1}$ to $e = 1.8\ d^{-1}$ to better fit the measured pool silica data. The dynamic calculation with the adjusted escape coefficient results in the green silica curve (bottom plot), which is clearly closer to the nominal purple curve. The green curve in the upper plot, calculated by dynamic analysis, predicts a 5.3 percent loss after 1 year (−0.11 relative error) and a 30.8 percent loss after 10 years (0.05 relative error). Dynamic analysis in escape coefficient has thus mitigated the uncertainty from temperature variation.

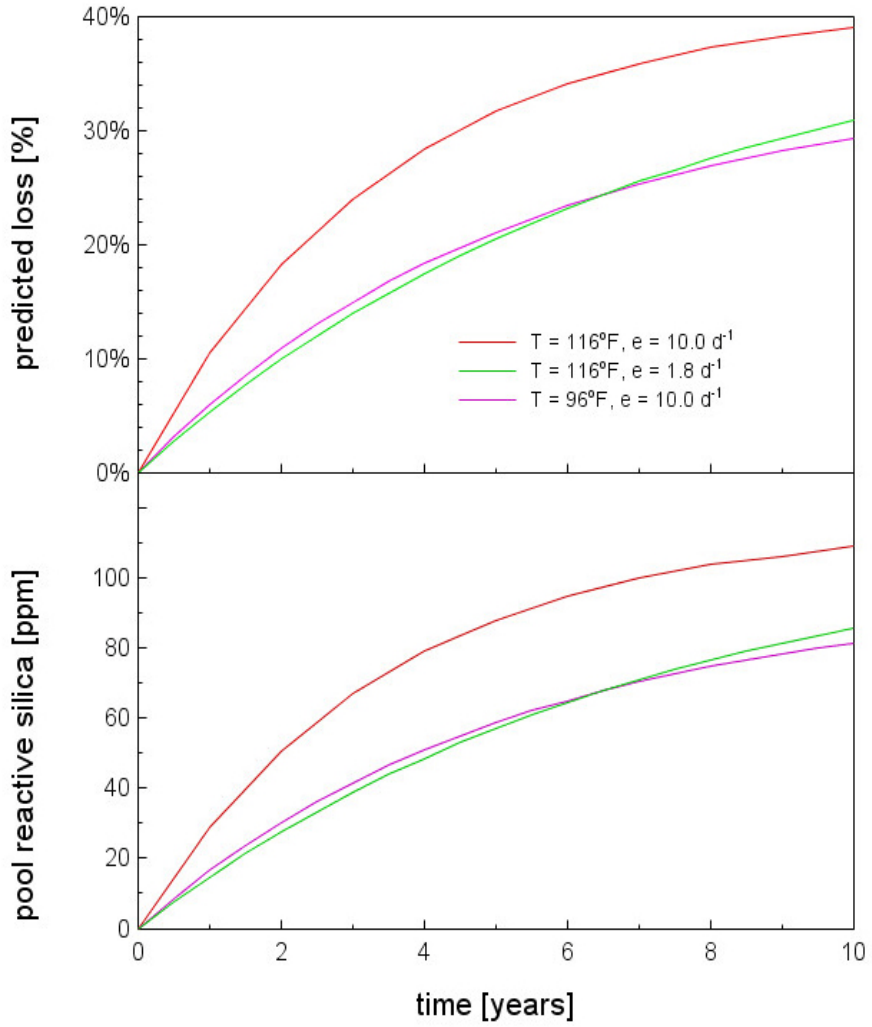

Figure 4–16. RACKLIFE predicted pool silica (lower plot) and panel loss (upper plot) for static and dynamic variation of temperature and escape coefficient.

While the dynamic analysis in this model more closely resembles the algorithm in commercial RACKLIFE and produces a more realistic measure of RACKLIFE's uncertainties in panel loss, note that the uncertainties are generally only applicable to the panel loss averaged over a zone or pool. In most RACKLIFE applications, one escape coefficient is assumed to apply to every panel cavity in the pool, or two for some PWRs with significantly different Region I and Region II rack designs. To match the silica data, this escape coefficient is tuned the same for all panel cavities. However, as discussed previously, every panel and panel cavity is unique. Imagine a pool where

most panels have long cooling time spent fuel adjacent to them, but some have freshly discharged fuel. The panels near the freshly discharged fuel may have significantly more radiation heating, leading to a very localized uncertainty due to uncertainties in the actual temperature. And yet the bulk silica data, driven by the other panels that have been dissolving for many years, may not change significantly, and so will not cause a mitigating decrease in the global escape coefficient. Thus, for identifying the most conservative uncertainty for a particular panel (as opposed to the uncertainty of a more typical panel whose percent loss is closer to the average loss of all the panels in the pool), the dynamic uncertainties are not applicable. This caveat can be extended to any source of *local* uncertainty. Because tuning the escape coefficient (dynamic analysis) is a global action, dynamic analysis may not sufficiently mitigate the uncertainty of a local effect. In such cases, the static uncertainties may be more appropriate.

Similar static and dynamic analyses were conducted for seven additional variables, as discussed below. The calculated fractional relative uncertainties for all nine variables are then summarized in Table 4-1.

4.4.8 pH Uncertainty

The pH measured in most spent fuel pools is between pH 4 and pH 8. As shown in Table 3-2, the equilibrium silica concentration, R_{eq}, decreases from pH 3 to pH 7, is constant from pH 7 to pH 8, and then strongly increases from pH 8 on up. In this analysis the nominal condition was pH 5, with variations down to pH 4 and up to pH 8, which should bound most variations. The static analysis gave a maximum relative fractional uncertainty in the percent panel loss of +0.11 for pH 4 and −0.19 for pH 8. The dynamic analysis reduced these to +0.05 for pH 4 and −0.06 for pH 8 for the worst conditions of high dose and low escape coefficient.

4.4.9 Panel Cavity Volume Uncertainty

The panel cavity volume was estimated to be "about 2/3rds of a cup of water" (0.150 liters) in Section 2.3 with the Boraflex intact, and about six cups (1.5 liters) if the Boraflex were completely dissolved. RACKLIFE calculates the increased volume as the Boraflex degrades. However, while degraded Boraflex will lose its amorphous silica to the bulk pool, some amount of boron carbide and crystalline silica may remain trapped in the cavity. Further, manufacturing tolerances and mechanical deformation during fabrication and use can change the cavity volume from its nominal value. Applying a ±20 percent variation in the volume cavity resulted in a static analysis ±0.20 fractional variation in the panel loss. The dynamic analysis reduces this to ±0.04.

4.4.10 Panel Surface Area Uncertainty

As Boraflex degrades, its porosity increases, which increases its surface area available for dissolution. In Reference [EPRI 1988; p. 5-15] the open porosity is estimated to be between 3 percent and 10 percent at 10^{10} rads. In the commercial RACKLIFE code, the effect of porosity is accounted for in the experimental data used to calculate the dissolution rate. However, the experimental data used to develop the source term only tracked one year of Boraflex degradation, resulting in only a 2 percent loss of boron carbide (see [EPRI 1997a; p. A-9]). It was simply not practical to conduct an experiment for a decade, and no accelerated testing results were applied to developing the silica source term. Further, not all Boraflex surfaces are necessarily available for dissolution, since the use of adhesives, dimensional uncertainties in the cavity clearance fit, pinching by deformed cover plates, or other anomalies, may cover significant areas of the Boraflex to greatly reduce or eliminate local water access. As discussed in Section 3.4, the dissolution mechanism is fundamentally a surface effect, and so uncertainty in the surface area will lead to uncertainty in the predicted loss. Applying a ±30 percent variation in the surface area resulted in a static analysis ±0.30 variation in the panel loss. However, the dynamic analysis dramatically reduces this to only ±0.02.

4.4.11 Panel Mass Uncertainty

As discussed in Section 3.1, sufficient information is provided in Boraflex material certification sheets to compute for each panel the initial mass of B_4C and the mass ratio of B_4C to amorphous silica from degraded PDMS, exclusive of the crystalline silica filler. As shown in Equation 4-7 in Section 4.2.2, these factors are used in RACKLIFE to calculate the mass fraction of each Boraflex panel that has been lost—and thus the amount of neutron absorbing boron carbide lost—based on the loss of silica. RACKLIFE calculates the initial mass of B_4C from user input Boraflex density, panel dimensions, and mass fraction of B_4C in Boraflex. As described in Section 3.1, the panel composition is not consistent between batches, nor even among all the panels in a particular batch, especially near the beginning and end of a manufacturing run. Users often input nominal density and dimension data, though material certification sheets can be used to enter batch-specific densities. A spreadsheet that comes with RACKLIFE can be used to calculate the mass fraction of B_4C based on data in the material certification sheets and independent chemical analysis of Boraflex (see, for example, [EPRI 1993b, Section 2]). This same spreadsheet also calculates the mass ratio of B_4C to amorphous silica from degraded PDMS, exclusive of the crystalline silica filler.

Because the B_4C mass and B_4C-to-amorphous-silica mass ratio are not used in the silica kinetics equations which contain the terms for escape coefficient, the static and dynamic uncertainties are identical. That is, changes in the mass or mass ratio have no effect on the predicted pool silica concentration, and so the escape coefficient is not adjusted to compensate for anything. The mass uncertainty thus directly translates into a loss uncertainty. For example, if RACKLIFE predicts a panel loss of 40.0 percent from a typical 2.700 kg panel, but in fact the panel was only 2.647 kg, a

difference of ~2 percent, then the actual loss is 40.8 percent, also a 2 percent difference. Similarly, a 10 percent variation in mass would result in a 10 percent variation in predicted loss. Chemical analysis (see [EPRI 1993b, p. 2-8]) suggests that the mass fractions are accurate to approximately 1.5 percent. However, the total mass has a larger uncertainty because point-to-point the panel thickness can vary from its nominal value by up to 10 percent (see [EPRI 1993b, p. 2-3]). A 5 percent variation in total panel mass is a reasonable approximation for a given panel, although higher departures from nominal have been reported. Across the spent fuel pool, the average variation should be much less.

4.4.12 Panel Dose Uncertainty

As described in Section 4.2.1, the dose to a Boraflex panel is a function of a large number of variables, including reactor data, assembly data, and spent fuel rack geometry. The dose model itself is also subject to a large uncertainty. The original dose model, developed in 1991 and described in [EPRI 1993a, Section 5 and Appendix A], is referred to as the PWRDose or BWRDose model, depending on the assembly type, or generically as the LWRDose model. Subsequent experimental testing of Boraflex coupons from plant spent fuel rack surveillance programs indicated that the dose was likely being overestimated. The experiments measured the dynamic shear modulus ratio of the coupon, which had been rigorously correlated with the known dose to other experimental Boraflex samples, as described in [Lindquist et al. 1994]. It was hypothesized that the overestimate was due to changes in the spent fuel gamma spectrum as the assembly cooled. Thus, based on the experimental data, a correction was developed, called the "cooling time corrected" (or CTC) dose model. In RACKLIFE, the user selects which model to use: LWRDose or CTC; the user is advised to try both to determine which best allows them to match their pool silica data.

Subsequent verification and validation modeling, as detailed in [EPRI 1999a, pp. 9-5 through 9-14], showed that the actual dose was somewhere between the LWRDose and the CTC models. Data was extracted from the [EPRI 1999a] reference and plotted in Figures 4-17 and 4-18 to illustrate the magnitude of the variation, which can reach as high as 52 percent. It is not that any individual dose calculation varies by 52 percent—the reactor, the assemblies, and the racks are roughly the same throughout the pool—so what is called an uncertainty is more properly a bias. For example, in Section 4.4.6, on the dynamic analysis of temperature uncertainty, it was noted that individual panel cavities could be at various temperatures higher than the bulk pool temperature. Thus, there would be a distribution of temperature uncertainties across all of the panel cavities; two otherwise identical panel cavities could have different temperatures due to the random flow paths around the partially dissolved Boraflex contained in each. In contrast, two panels that have absorbed the same amount of gamma dose from identical assemblies would be in error by exactly the same amount; there would be no random distribution of doses, only a constant dose bias

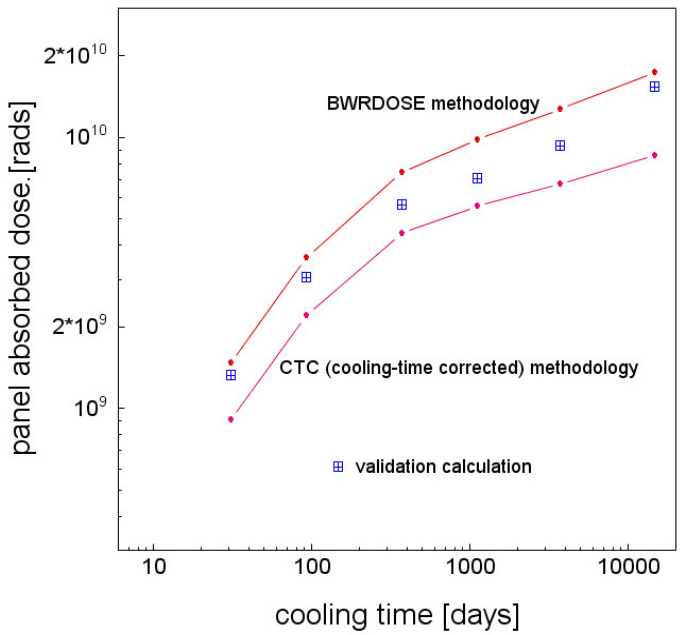

Figure 4–17. Comparison of BWR dose calculation methodologies to verification and validation data.

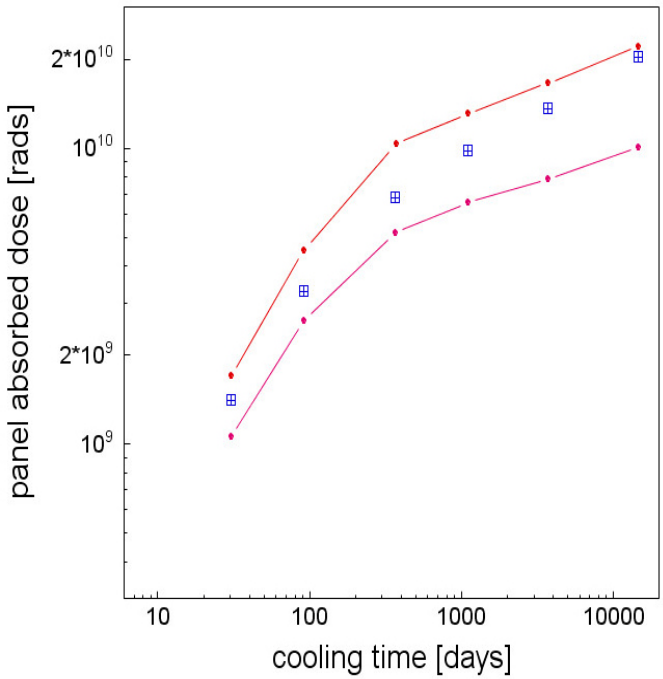

Figure 4–18. Comparison of PWR dose calculation methodologies to verification and validation data.

At face value, up to 52 percent variation in the dose may seem extreme for a degradation model which relies heavily on dose. Static testing bears this out, showing that a 52 percent variation in panel dose results in a maximum 107 percent uncertainty in the panel loss, as shown in Figure 4-19. The regression analysis in Figure 4-20 shows that uncertainty is strongly correlated with dose and escape coefficient, with moderate dose and high escape coefficient conditions showing the greatest uncertainty in loss. Despite this very large static uncertainty, however, dynamic testing shows a less than 1 percent uncertainty in loss for the worst condition triplets. For a 52 percent increase in dose, the escape coefficient is simply adjusted from 10 d^{-1} to 2.2 d^{-1} to match the silica data and the bias is corrected. This seemingly remarkable reduction in uncertainty is based on the assumption that the bias in the dose calculation methodologies is consistent across the assemblies in the pool. For local variations in dose, such as variations in an assembly's end-of-cycle power sharing, static analysis may be more appropriate than dynamic analysis. Panel-specific variations in dose are discussed in Section 4.4.14.

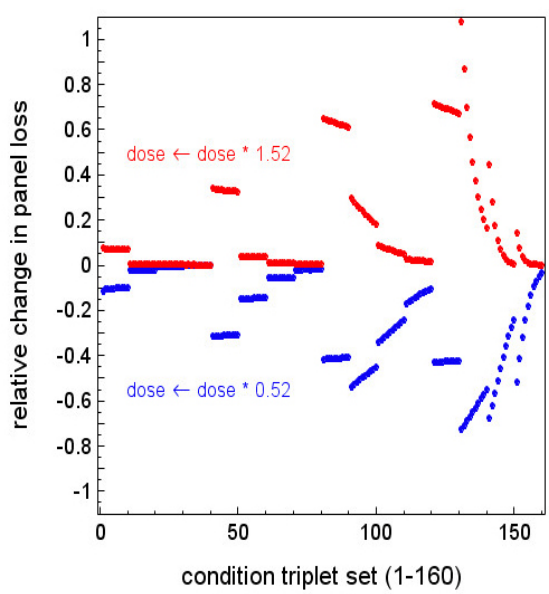

Figure 4–19. Uncertainty in panel loss due to ±52% variation in dose - static test.

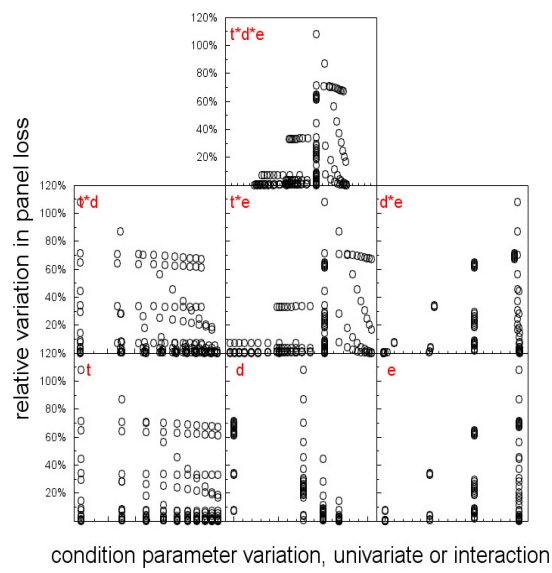

Figure 4–20. Conditional effect of panel dose.

4.4.13 Average Escape Coefficient Uncertainty

There is no dynamic uncertainty for the escape coefficient—by definition, it is the tuning of escape coefficient that makes the uncertainty analysis dynamic. Adjusting the escape coefficient is most effective at mitigating variation in variables which are closely related to the escape coefficient (e.g., for panel cavity variations in temperature or volume). Taken to its logical conclusion, a dynamic analysis on the escape coefficient, by definition, mitigates the uncertainty entirely. However, there is static uncertainty with escape coefficient variation. Most commercial users tune the escape coefficient visually based on plots of measured and predicted silica produced by RACKLIFE, and there is variation in visual matching. The difference between calculated and measured silica data can also be numerically matched (e.g., by minimizing root-mean-square variations between the two). However, the use of letdowns to simulate outages would require multiple analyses of this type, which may not reduce the variation any more effectively than visual matching. Regardless of matching methodology, there will be some variation about the optimal setting of the escape coefficient. A variation of 10 percent in the escape coefficient results in a maximum variation in the panel loss of 10 percent. The loss variation correlates positively with dose and negatively with nominal escape coefficient, so that high dose and low escape coefficient represent the worst case conditions. However, the dependence is not very strong. Figure 4-21 shows the wide distribution of variation for the condition triplets.

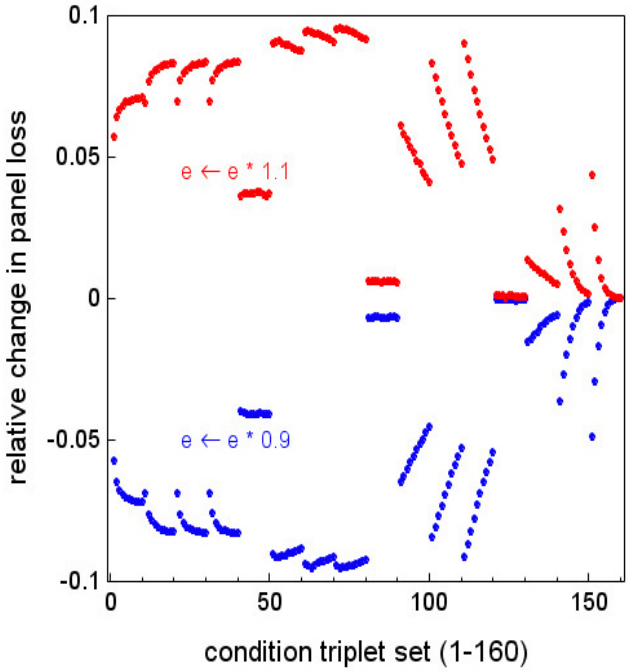

Figure 4–21. Uncertainty in panel loss due to 10% variation in escape coefficient.

4.4.14 Panel Specific Uncertainty

The above uncertainty analyses, static and dynamic, were generally developed for panel conditions averaged across the spent fuel pool. To illustrate the limitation of panel averages, suppose that a user enters the wrong assembly end-of-cycle (EOC) power sharing for a single assembly. The dose rate from this assembly is a strong function of its EOC power sharing and so the dose received by any neighboring panels will be similarly incorrect. However, because only one assembly is in error, which affects only a small number of panels in the pool, the effect on the calculated pool silica concentration will be negligible. Thus, the user will not adjust the escape coefficient to compensate. For this case, the dynamic uncertainties are not appropriate. Based on the static uncertainty analysis, a typographic error of 1.0 instead of 1.5 for the EOC power sharing could result in a 100 percent error in predicted loss—the panel loss would actually be twice what was predicted.

RACKLIFE can generate plots and tables to characterize the data input by a user. For example, for assembly EOC power sharing data, the user can view a plot of the input values by assembly index number, a histogram plot showing the distribution of the values, and a rate plot showing the difference from one value to the next, which can be helpful for detecting unusually large jumps that may indicate a typographic error. The user can also view and save a text file listing all of the data, as well as summary statistics such as maximum, minimum, average, and standard deviation. At the start of a state point calculation, prior to any results being generated, RACKLIFE will internally check all of the input data. RACKLIFE will output warnings for data that it considers unusual (e.g., unusually low or high pool temperatures), and output errors that terminate the state point calculation for obvious mistakes (e.g., temperatures below freezing). Regardless, the risk of a typographic error remains, and the static uncertainty analysis results would apply for that specific panel.

In addition, there is also natural uncertainty in the input data. For example, the EOC power sharing values likely came from a core calculation code, with its own inherent uncertainties. However, while a typographic error may affect only a single panel, uncertainties in the methods used to develop the input data will be common to all panels. If the EOC power sharing output by a core calculation code is biased, the effect is similar to the dose bias discussed in Section 4.4.12, and the dynamic analysis will largely mitigate the uncertainty. However, if the core calculation code EOC power sharing output contains a random error, then the static uncertainty analysis results would apply, scaled to the magnitude of the random error.

When BADGER data for a particular panel is available, then that individual panel cavity escape coefficient can be tuned so that the RACKLIFE predicted loss matches the BADGER measured loss for that panel. However, the remaining panels would still be controlled by a global escape coefficient tuned to the bulk pool silica data. Consistent discrepancies between the global escape coefficient and the BADGER-effected escape coefficients can indicate that the spent fuel pool cleanup systems may be more or less effective than was originally assumed. The BADGER measurements allow the escape coefficient to be tuned to better match the observed panel losses,

and the cleanup system to be tuned to better match the bulk pool silica measurements. This double-tuning process is described and applied in [EPRI 1997b, Section 4.3].

An alternative to keeping the panels untested by BADGER at some fixed global escape coefficient would be to perform a constrained Monte Carlo simulation. The RACKLIFE simulations would be constrained in that the RACKLIFE predicted pool reactive silica should still match the measured silica data. The RACKLIFE panel loss results would be stochastic in that each panel would be assigned an escape coefficient randomly selected from a distribution of escape coefficients. Because the distribution is modeled using escape coefficients which are tuned to match RACKLIFE panel loss data to BADGER measured loss, non-random and non-independent effects are accounted for. In this manner the panels could be properly simulated and variations in the escape coefficient rigorously accounted for.

4.4.15 Summary of RACKLIFE Uncertainties

Table 4-1 summarizes the above information. As discussed at the beginning of Section 4.4, the uncertainties presented here are only for illustrative purposes and to guide error analysis efforts to target the largest uncertainties. These results should not be considered representative of any particular spent fuel pool.

Table 4-1. Summary of RACKLIFE uncertainties.

variable	nominal value	variation	Static fractional relative uncertainty	Dynamic fractional relative uncertainty
polymerization rate	$1.644 \cdot 10^{-3}$ d^{-1}	±50%	0.01 (0.07 for low d×e)	0.01 (0.07 for low d×e)
all water temperature	96°F	±20°F	+0.80 / −0.40	0.06 after 1 year, 0.31 after 10 years
all water pH	pH 5	pH 4 / pH 8	+0.11 / −0.19	+0.05 / −0.06
cavity volume	0.707 liters	±20%	±0.20	±0.04
panel surface area	$1.47 \cdot 10^{4}$ cm^{2}	±30%	±0.30	±0.02
panel mass	2.68 kg	±5%	±0.05	±0.05
panel dose	various	±52%	+1.07 / −0.73	±0.01
cavity escape coefficient	various	±10%	±0.10	(not applicable)

To illustrate an example application of the tabulated uncertainties, consider a rack design where the volume of a panel cavity is not clear. For example, the L-wrapper design in Figure 2-3 may have significant variations in panel cavity volume due to variations in how the wrapper plate bends and how welds are made. Suppose, from rack fabrication drawings and engineering judgment, the user can deduce that the value input to RACKLIFE for all panel cavity volumes may be low by up to 10 percent. The table states that, for a 20 percent variation in volume, a maximum ±0.04 fractional relative uncertainty may result. Assuming that the uncertainties scale linearly, the user can estimate that a +10 percent variation in volume will result in a maximum +0.02 fractional relative uncertainty. If in some zone of the spent fuel pool RACKLIFE predicts 15 percent panel losses, then the actual loss will likely be no higher than 15 percent · (1+0.02) = 15.3 percent. That is, if the user were to input the 10 percent higher panel cavity volume into RACKLIFE, tune the escape coefficient, and rerun the model, the predicted panel losses would increase from 15 percent to 15.3 percent. A number of uncertain assumptions were made in this example analysis (e.g., that the L-wrapper design and the nominal cavity volume are the same in the example as in the analysis that produced the table and that the tabulated uncertainties scale linearly, among others). However, the user has a better understanding of what to expect and where to focus error analysis resources. Because a 0.3 percent increase in panel loss is likely easily mitigated in a criticality safety analysis, the user might simply apply the higher panel cavity value to bound the condition and move on.

Suppose instead the example user is concerned about a particular panel that is again predicted to have 15 percent loss, but was recently discovered to have sustained an impact and cover plate deformation, likely during fuel movement a few years ago. It would be inappropriate to use the dynamic analysis uncertainties because the consequences to the bulk pool silica concentration from one panel cavity will be negligible, so no compensating dynamic tuning would have been performed. In this case, the static analysis suggests a ±0.20 fractional relative change in loss for a 20 percent change in panel cavity volume. If engineering judgment determines that the deformation resulted in a 20 percent decrease in panel cavity volume, then an estimate of the actual panel loss would be 15 percent · (1 − 0.2) = 12 percent loss. Such a conclusion is may actually increase uncertainty, however. First, the decrease in panel cavity volume is a bounding estimate; the actual decrease could be much less. Second, while the change in panel volume may be observable, the change in the escape coefficient for that panel cavity is likely neither observable nor predictable, and the uncertainty due to changes in the escape coefficient will likely be larger than the uncertainty due to changes in panel cavity volume.

The range of variation in each variable in Table 4-1 is intended to be generally bounding. However, as discussed in previous sections, typographic errors, other user input errors, and variations of a specific panel or panel cavity could exceed the indicated range. When the error is local and involves a relatively small number of assemblies, panels, or panel cavities, the effect on the silica release rate may be too small to significantly change bulk pool silica concentration. In such cases, it is not appropriate to rely upon the dynamic analysis in this model because the dynamic variation of the escape coefficient would not have been utilized in practice. Proper

verification of the input data is vital when attempting to use RACKLIFE to predict individual panels.

Because RACKLIFE stores its input data in text files, they are easily accessible for the user to manipulate using an external task-specific program. For example, programs can be written to randomly assign escape coefficients to panels, perhaps based on BADGER data as discussed in Section 4.4.14, to assess the uncertainty due to the stochastic nature of this variable. Similar variations in everything from panel dimensions to temperatures to assembly powers can be performed to quantify uncertainties specific to a particular pool and its specific conditions. The uncertainties obtained will generally be much less than the static uncertainties reported in Table 4-1. The one variable not accessible to the commercial RACKLIFE user is the polymerization rate; however, the effect of this variable on panel loss was seen to be quite small.

5 BADGER

BADGER is the Boron Areal Density Gauge for Evaluating Racks: an in situ nondestructive diagnostic tool for measuring the areal density of ^{10}B in a neutron absorber as installed in fuel racks. BADGER is a multicomponent instrument comprising a source probe and detector probe, each suspended from its own pole, a computer-controlled step-drive motor system, and algorithm software for data capture. The detector probe head holds four BF_3 gas neutron detectors along a line parallel to the plane between the probes. During a BADGER scan of a neutron absorber panel, the probes are lowered into adjacent cells so that the single panel in question is located between the source head and detector head. As neutrons stream from the source to the detectors, the neutron transmission through the panel is counted. After a predetermined length of time, counting is stopped and both probes are then hoisted a short distance. The process is repeated upward for the length of the panel. Thus, in a stepwise manner, the entire panel is scanned for neutron absorption. The loss of ^{10}B is then calculated from a calibration curve for each detector, resulting in a two-dimensional representation of ^{10}B areal density. It should be noted that unlike RACKLIFE, which only calculates a panel average loss, BADGER provides measured spatial data. A typical BADGER campaign at a spent fuel pool will scan 25-60 panels over the course of about ten days.

Figures 5-1, 5-2, and 5-3 illustrate the BADGER system deployed in spent fuel racks containing Boraflex as the neutron absorber. There are two BADGER systems: one designed exclusively for PWR racks, and one designed exclusively for BWR racks. A ^{252}Cf source produces a fission spectrum of fast neutrons, which then thermalize in the source head. In the BWR design, the aluminum source head housing fills with moderating pool water when submerged. In the PWR design, the aluminum source head housing contains moderating high density polyethylene in order to displace the neutron-absorbing borated pool water. The slowed neutrons with trajectories toward the neighboring cell are then absorbed by the intervening Boraflex panel in proportion to the areal density of ^{10}B present, or may be detected by the BF_3 neutron detectors in the detector head. After a horizontal section of a panel is analyzed, both probes are typically drawn upwards for two inches, and the process is repeated. In this way the ^{10}B areal density of an entire panel can be characterized at the resolution of the 2 inch high and 5/8 inch diameter detectors. For a typical 144 inch Boraflex panel, this results in 288 data points per panel, arranged in a grid 72 cells high and 4 cells wide. (This is discussed and illustrated subsequently in Figure 5-6.)

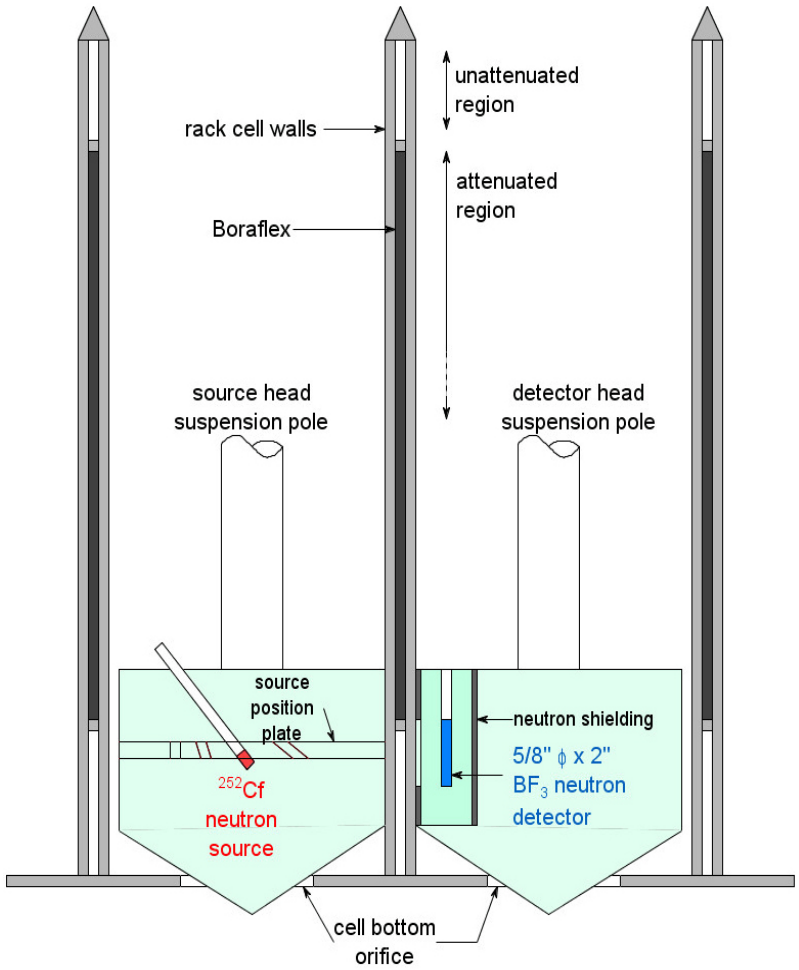

Figure 5–1. Axial view of BADGER in egg-crate racks.

Figure 5-1 shows a side view of the two BADGER probes positioned at the bottom of two adjacent rack cells, just before it begins a scan of the Boraflex panel inside the rack wall between the two cells. The source head is illustrated with a "source position plate," which was specific to the original BWR BADGER. This allowed different source positions to be tested during development to optimize the location of the source experimentally for maximum system sensitivity. Subsequent versions of BADGER use a fixed position for the source to be inserted into. The black bars within the rack walls represent Boraflex, where neutrons are absorbed and attenuated. The upper section within the rack wall, empty of neutron absorber, is labeled the unattenuated region, where neutrons may pass through, unaltered except for the intervening water and stainless steel.

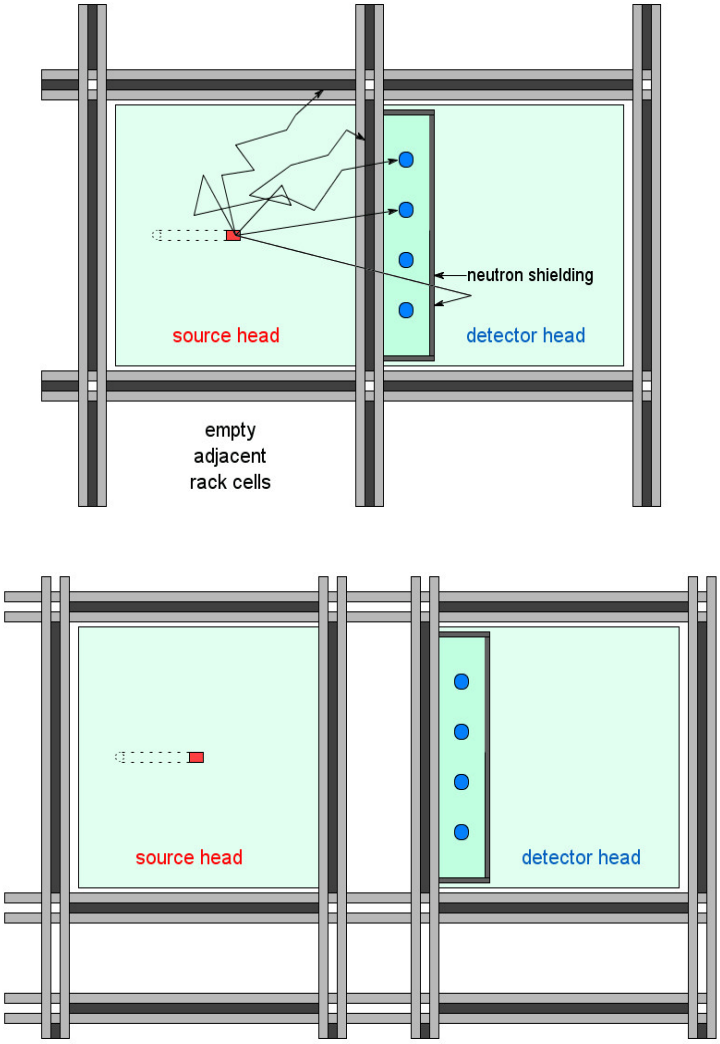

Figure 5–2. Lateral view of BADGER in egg-crate (top) and flux-trap (bottom) rack cells.

Figure 5-2 shows an overhead view of the two BADGER probes inserted into two adjacent rack cells. The top of Figure 5-2 shows BADGER in an egg-crate rack along with illustrative neutron trajectories. Neutrons thermalized in the source head will encounter the surrounding Boraflex panels and either be absorbed by ^{10}B or will pass through the panel, in proportion to the amount of ^{10}B present. Neutrons in the source head that pass through cell walls other than the one with the Boraflex panel being measured (i.e., neutrons that migrate toward rack cells other than the one containing the detector head) are very unlikely to reach the detectors because of the amount of interposing water and the presence of the other panels of Boraflex. Fast neutrons will easily pass through the Boraflex panel being measured but are unlikely to interact with the detectors before passing into the water in the back of the detector head and thermalizing there. To reduce the

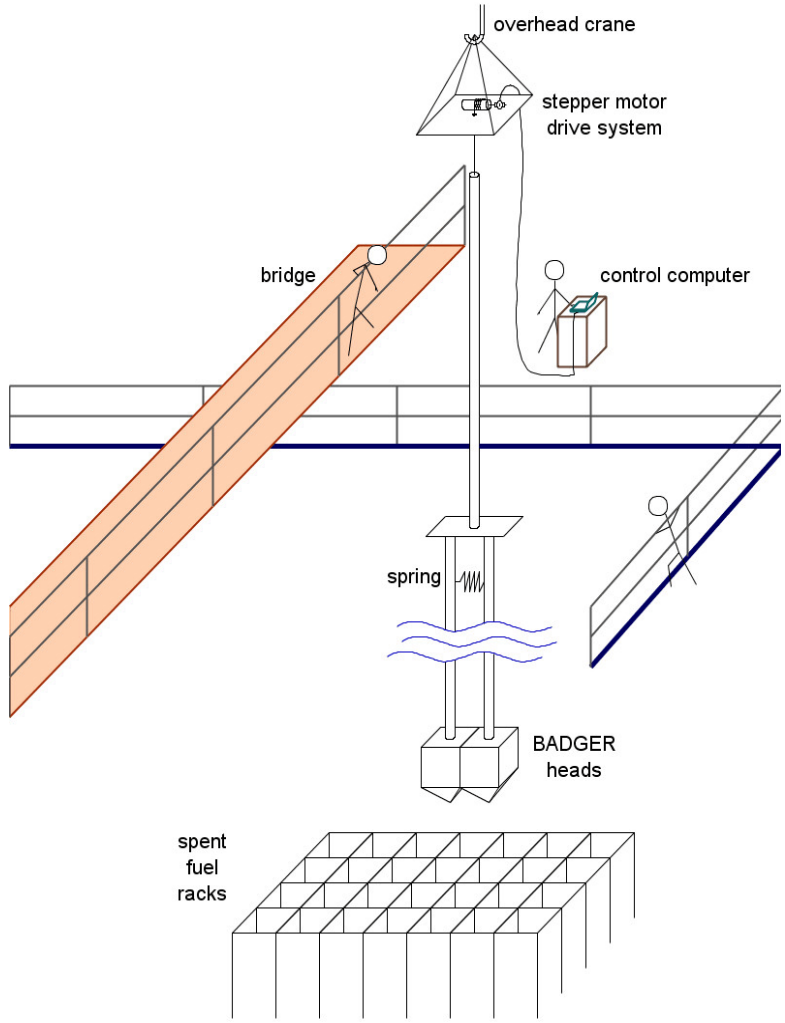

Figure 5–3. BADGER in operation.

detection of neutrons backscattered from the detector head water, thermal neutron shielding surrounds the detector block except for a 2-inch-high full width window facing the panel being scanned, where neutrons enter the detector volume. These features greatly increase the proportion of neutrons detected which passed through the panel as a consequence of loss of ^{10}B.

The bottom of Figure 5-2 shows BADGER in a flux-trap rack. Experimental and numerical testing (see [EPRI 1998a, Section 5.2] for details on the numerical testing of a flux-trap rack) has shown that the system is almost exclusively sensitive to the absorber panel nearest to the detectors. Thermal neutrons that emerge from the face of the first panel will be attenuated in the flux trap water gap. Fast neutrons that pass through the first panel encounter the pool water in the flux trap water gap, which efficiently thermalizes the neutrons prior to encountering the thermal neutron absorbing ^{10}B in the second panel. The overall neutron signal to the detectors is greatly reduced,

however, especially due to the dissolved boron in the PWR spent fuel pool that will attenuate the neutrons thermalizing in the flux trap water gap. The net effect is a nearly order of magnitude reduction in the neutron count rate for flux-trap racks. Thus, while BADGER properly measures only a single panel of Boraflex, those measurements are much less efficient than for egg-crate racks, thereby increasing the necessary count time, as discussed subsequently.

Figures 5-1 and 5-2 show that the BADGER probe heads have a clearance fit with respect to the rack walls. On the sides the clearance is nominally around 3/8ths inch; shims can be used to reduce the clearance. Differences in spent fuel rack designs make this clearance vary in size from plant to plant, and even region to region in multi-region pools with different rack designs for each region. Some clearance is necessary to allow the movement of the probe heads upward during a BADGER scan of a panel, but it also allows the heads to become misaligned by lateral offset, twist, and tilt, and to be separated if not held flush to the cell walls. These misalignments will be discussed in detail in Section 5.4.

Further description of the BADGER components can be found in [EPRI 1997b] and [EPRI 1998a]. The following sections describe BADGER operation in the spent fuel pool (Section 5.1), the theory behind the use of BADGER to measure Boraflex panel degradation and dissolution (Section 5.2), the calibration of the BADGER system to measure ^{10}B areal density (Section 5.3), and the uncertainties inherent in its use (Section 5.4).

5.1 Using BADGER

Figure 5-3 illustrates the BADGER system in operation. A BADGER operator controls the system from a computer interfaced to a programmable digital controller. From here the BADGER operator controls the demand elevation of the probe heads, the axial length of a panel scan, and the counting time for each step during the scan. The BADGER operator can also view an axial trace from each detector in real time. A sample trace is illustrated in Figure 5-4, based on a trace shown in [EPRI 1997b, p. 4-15]. Figure 5-4 shows, from bottom to top, all four detectors starting below the bottom of the Boraflex panel. As the source and detector heads move into the *attenuated region*,

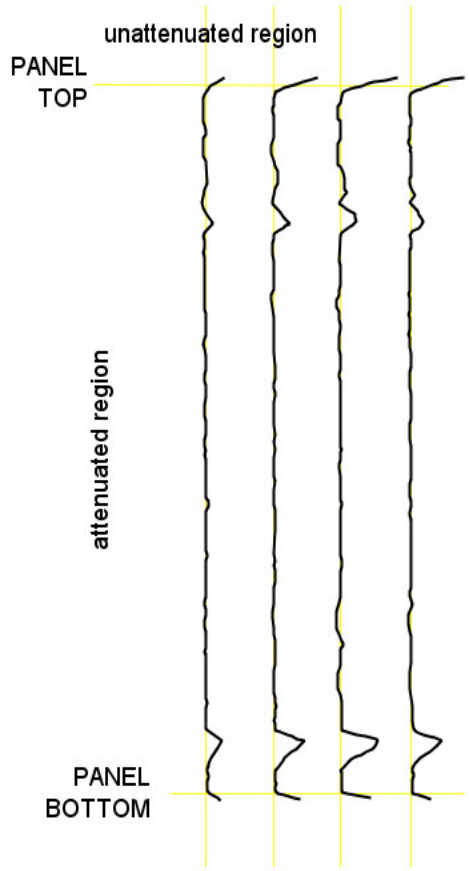

Figure 5–4. A BADGER scan trace. Adapted from EPRI Report TR-107335 (1997), with permission from Electric Power Research Institute.

concurrent with the Boraflex panel as shown in Figure 5-1, neutron flux drops as the Boraflex panel absorbs neutrons. Moving upward, the trace shows small peaks as BADGER detects Boraflex gaps and regions of dissolution, which will be discussed in detail in Section 5.2. Moving out of the attenuated region and into the *unattenuated region*, above the Boraflex panel, neutron flux again increases.

Either as an extended part of the attenuated region scan, or as a separate scan, the unattenuated region above each Boraflex panel is scanned as a quality assurance measure for the BADGER instrument. The rack geometry is often the same in the attenuated and unattenuated regions, except that the Boraflex is replaced by pool water. In some wrapper plate designs, the very thin wrapper will also not be present, a minor difference. However, the length of the probe heads ensures that the lower part of the heads will still contact the top of the wrapper plate so that the heads remain separated the same distance in the lower part of the unattenuated region as in the attenuated region.

In theory, the results from the unattenuated region scans should be the same for every identically manufactured cell. However, rack wall and cover plate tolerances and drift in the electronics will cause unavoidable differences. Misalignment of the heads within the clearance space will also cause differences, though these are to some extent correctable through data analysis. Considerable statistical analysis can be brought to bear on the scan data to quantify electronic drift, cell wall to cell wall differences, detector to detector differences, and to assess the degree of head misalignment, as discussed in Section 5.4.

At the beginning of a BADGER campaign, the instrument is calibrated. The calibration process is described in detail in Section 5.3. The Boraflex panel being scanned is between the source and detector heads. As shown in Figure 5-2, the heads are held lightly against the interposing rack cell walls by a spring between the suspension poles above the top of the racks, as seen in Figure 5-3. BADGER is designed to achieve a clearance fit in the vast majority of BWR and PWR spent fuel rack cell designs. When the lateral dimensions of the rack cell are substantially larger than the BADGER heads, stainless steel shims can be used to reduce the clearance to reduce head misalignment.

Lateral motions of BADGER to the desired rack cells are effected by the refueling bridge auxiliary crane to the precision of the crane operator, as aided by viewers stationed along the sides of the pool or on the bridge. In Figure 5-3, the BADGER operator is shown facing the BADGER suspension poles from the side of the spent fuel pool. During lateral motion, the heads at the ends of the suspension poles are subject to several types of misalignment—tilt, separation, offset, and twist—which are described in detail in Section 5.4.1 and illustrated in Figure 5-9. The BADGER operator and other viewers can detect tilt misalignment by viewing the orientation of the suspension poles with respect to any vertical sight line in the fuel handling room. Ideally, sight lines concurrent with the suspension poles and perpendicular to the plane of the suspension poles would be used to detect tilt misalignment visually. Separation problems—where the heads are not flush against the cell walls, which would change the character of the system further from the

geometry of the calibration cell by allowing additional water between the heads—may be observed from the perpendicular side or by viewers on the bridge. Offset misalignment, where the head centerlines do not coincide, can be observed from the concurrent side by noting if the near pole eclipses the far pole. Twist misalignment, where the heads are rotated with respect to the cell walls, is most often caused by tension from the power and signal cables; the BADGER operator will adjust the cables to minimize the tension.

The axial drive system platform is suspended from the overhead crane and holds the stepper motor, rotary encoder, cable winch winding a 0.25-inch vinyl-clad steel cable, and cable load sensor. Excess load on the cables indicates that the system is wedged in the cell for any reason (e.g., tilt misalignment, or deformation of the panel covers due to Boraflex off-gassing or mechanical impact during fuel moves). The axial drive raises and lowers BADGER to precise axial elevations based on demand signals from the BADGER software and as measured by the shaft encoder. While the digital rotary encoder is very precise (600 steps per revolution), there is an axial elevation bias due to cable strain. The bias changes with the length of cable deployed; the maximum bias typically ranges from 0.5 to 2.5 inches.

The elevation bias is measured for each panel by identifying the motion limits of BADGER in the rack cell being scanned. Prior to a scan, the BADGER probes are lowered together by their common connection plate to the bottom of the rack cell. As shown in Figure 5-1, the lead-in tapers of the heads will protrude through the cell bottom orifices. Geometric calculations at this position can determine the vertical distance between the BADGER detectors and the bottom of the Boraflex panel. As long as the detectors can be lowered below the bottom of the Boraflex panel, these calculations can be verified by examining the traces for unirradiated (i.e., undegraded) panels to mark the initial drop in neutron flux, from which the elevation of the panel bottom can be calculated based on calibration scan data for transitions from unattenuated to attenuated regions. At the end of a scan, the heads are visually positioned to be flush with the top of the racks. This point can also be compared with geometric calculations to establish an elevation at the top, again as verified by traces for unirradiated panels. These known bottom and top points determine a linear correlation between the BADGER software demand elevations and true elevations, including the linear effects of cable strain, throughout the range of axial motion in the rack cell.

Figure 5-1 shows the source and detector heads moved to the bottom of the cell such that the tapered lead-ins penetrate the cell bottom orifice. Typically this positions the detectors below the bottom of the Boraflex panel. However, in certain rack designs it is not possible to begin the scan below the panel bottom, so that scan verification of the geometric calculations is not possible. On very rare occasions, there have been cells that BADGER could not descend to the bottom of because of rack deformation; the scan then proceeded from the lowest elevation BADGER could achieve, with the recognition that differences between the demand and true elevations in that panel scan could not be fully compensated for; instead, comparisons to elevation data from other panels were used.

From the bottom of the cell, the BADGER system moves up in two inch steps—equal to the height of the detectors. The original BWR BADGER used 1-inch high detectors, but this was later modified to use 2 inch detectors. At each step elevation, a count is taken for a specified amount of time; the same amount of time is used for each count in a given scan. Typically, this same count time is used throughout a campaign for a given rack design, though count times will differ between different rack designs in the same pool (e.g., when scanning a Region I versus a Region II rack). A typical 144-inch-high panel, including the unattenuated region, can be completed within a 156-inch scan height comprising one scan at the bottom plus one scan every two inches of elevation resulting in 79 total elevations at which a count is performed. In typical egg-crate racks, with only one interposing panel, each count may last on the order of 10 seconds to obtain attenuated counts in the thousands. This results in count rates on the order of a hundred counts per second. The counting scan will have taken 13 minutes (79 elevations at 10 seconds each). However, combined with lateral positioning to the desired rack cell, insertion, elevation bias measurements, and scan, a typical measurement of an average egg-crate rack panel will take almost 1 half hour. Batches of panels with higher nominal ^{10}B areal densities require more time as the panel will absorb more neutrons and, therefore, a longer count time is necessary to achieve sufficient statistical resolution.

In a flux-trap rack, with two interposing panels, dissolved boron in the flux-trap water gap, and a larger distance between source and detector, a typical count is on the order of 1 minute, a typical count rate is on the order of two dozen counts per second, and one panel measurement can last over an hour and a half. This low count rate indicates the low absolute efficiency of the BADGER system. The optical efficiency of the detectors—the solid angle from the source that they subtend—is on the order of 1/1000. For a typical undergraded panel of Boraflex, about 2 million source neutrons are necessary to average a single count in a detector.

At each elevation, the demand elevation and the total neutron count observed by each detector during the count time is recorded to a text file. This file can then be sent to an analyst for analysis. This analysis is discussed in the next section.

5.2 Analysis of BADGER Data

Figure 5-4 illustrated a BADGER trace for a panel that had been exposed to a dose of about $1 \cdot 10^{10}$ rads accumulated over at least a decade. The faint superimposed horizontal yellow lines indicated the actual top and bottom of the as-built panel, as correlated with the change in neutron counts. Figure 5-5 shows the same trace, this time annotated to highlight the panel features observed. These annotations are hypothetical and for illustration only, based on observing the sketched trace made from the original figure in [EPRI 1997b, p. 4-15]; they are not meant to be representative of the actual panel scanned. The vertical yellow lines show the expected flat trace of the panel if uniform dissolution were the only loss mechanism. These lines are based on the average trace for each detector, exclusive of statistically significant deviations. Significant deviations from the vertical yellow lines indicate areas to scrutinize, whether large deviations such as the easily

observed peaks, or small deviations that may not be noticeable to the eye but are nonetheless statistically significant. Only the application of proprietary statistical techniques to correlate observations between detectors and with calibration scans can make strong conclusions about what a deviation represents, as discussed in the remainder of this section. The calibration process will be discussed in Section 5.3.

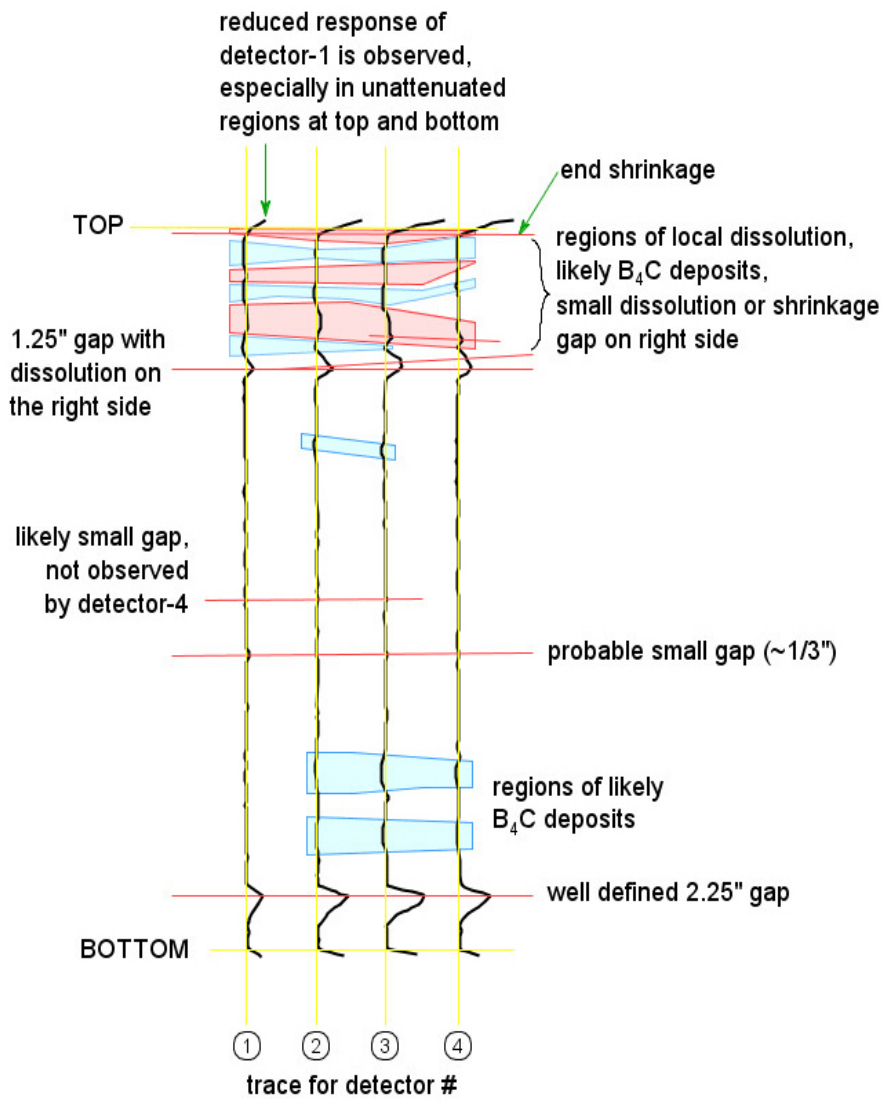

Figure 5–5. A hypothetical annotated BADGER scan trace. Adapted from EPRI Report TR-107335 (1997), with permission from Electric Power Research Institute.

The red horizontal lines in Figure 5-5 indicate missing Boraflex in the form of a gap. Two of the gaps, 1.25 inches in size near the top and 2.25 inches in size near the bottom, are well defined. Since the detectors move in 2-inch steps, the asymmetry of the peaks defining the gaps is to be expected: it is not likely that the detector center will coincide with the gap center to present a symmetric peak. Thus, to characterize a gap requires an analysis of both the height and the shape of the peak. The 1.25-inch gap shows a double peak in Detectors 3 and 4, as indicated by the bifurcated red line. This may be due to some combination of local dissolution that has widened the gap, the breaking of highly thinned sections of Boraflex that settled into the gap, and the settling of residual boron carbide particles into the gap area.

About halfway up the height of the panel are two probable gaps. The lower of the two is indicated as "~1/3" inch. Although the peaks would not individually be considered statistically significant, the fact that all four detectors show the same small peak at the same elevation indicates that the increase in counts can be identified as a gap. The actual height of the peaks is smaller than what a 1/3-inch gap 1 inch off-center from the detectors would show, but a properly conservative analysis would round the peak size up to 1/3 inch, the lower limit in reliably identifying gaps. Just above this "probable small gap" is another anomaly, but here only three of four detectors are statistically significantly correlated—Detector 4 does not observe it. There are several potential hypotheses to account for this data. This could be another gap smaller than 1/3 inch. Detector 4 may not observe it because of normal statistical variations in its observed count rate. Or the gap may be slightly tilted from the horizontal, so that one end of the gap was not near enough to the Detector 4 window to detect it. In general, BADGER can reliably detect gaps which are 1/3 inch in size or larger. Smaller gaps may or may not be detected.

Except for the double-peaked dissolution edge of the 1.25 inch gap, all of the gaps in this trace appear to be horizontal, perpendicular to the panel height. The horizontal nature of the gaps is due to tensile stresses resulting from the shrinkage of Boraflex near points of attachment (as discussed in Section 3.3). Peaks indicating gaps are also expected to occur with about 80 percent or more thinning loss, though the shape of the peak will typically be broader since nearby areas are also likely to exhibit some degree of local dissolution leading up to the $\gtrsim$80 percent area. Panels experience dose gradients based on the dose profile from adjacent spent fuel assemblies. Panels also experience the aqueous environment of the pool differently in different locations (as discussed in Section 3.8). These gradients will lead to local shear stresses, which can manifest themselves as cracks or gaps at angles. While horizontal gaps are by far the most commonly observed, some gaps have been observed at angles. This is often the case for cracks that do not penetrate across the entire panel. Angled gaps will more readily support a natural convection flow, and so are generally more prone to local dissolution effects.

Just below the top of the panel, a horizontal red line indicates some combination of end shrinkage and dissolution across the top. A zone of local dissolution is shown by the very top red area, as compared to the horizontal yellow line which indicates the as-built top. This top dissolution zone is larger in the middle, which is expected—this rack design has a small off-gas hole through the rack wall at the top center. This hole allows Boraflex panel off-gasses to easily vent from the

panel cavity but also greatly increases the exchange of panel cavity water there, which allows for accelerated local dissolution of silica. End shrinkage is common in unrestrained panels but, because of the hole, the amount of end shrinkage versus local dissolution for this panel is uncertain. Two other larger zones of local dissolution are also shown. In the middle of these three zones, the trace shows irregular broad plateaus of higher than normal count rates but without a defined peak. The lowest of the three red zones shows the same irregular broad plateaus, but Detectors 3 and 4 also show a small peak, indicated by the red line. This is likely a small area of $\gtrsim$80 percent dissolution, such as a scallop as illustrated in Figure 3-17, where the Boraflex has completely or almost completely dissolved. The shape of the significant peaks observed by Detectors 3 and 4, combined with the absence of peaks for Detectors 1 and 2, indicated this to be a dissolution effect rather than a shrinkage-induced gap.

Below each of the zones of local dissolution are zones, shaded blue, where the count rate is lower than nominal in broad correlated areas. Without correlation across neighboring scan elevations and detectors, many of these counts would be considered within the statistical uncertainty of the counting process. These are areas where the panel may have been protected from the uniform dissolution experienced by the rest of the panel. These may be areas where the Boraflex had swelled or the cover plate was deformed inward, so that the panel clearance fit in the panel cavity was eliminated, leaving an area of panel-wall contact that excludes the dissolving water. It is also possible that these are areas where insoluble boron carbide particles, released from the PDMS matrix in the dissolution areas above (as discussed in Section 3.7), settled by gravity and became trapped in areas where, again, the Boraflex had swelled or the cover plate was deformed inward. This effect has been observed in coupon and panel inspections.

It is important to recognize that the features discussed above—end shrinkage, large gaps, small gaps, zones of various degrees of local dissolution, gaps widened by dissolution, scallops, and B_4C deposits—are often difficult to characterize in practice. It is important to have well-characterized calibration data to compare against that simulate as many of these features as is practical. The statistical methods that the analyst can apply in the analysis, either heuristically or quantitatively, are also important. Most important of all is the experience of the analyst.

Figure 5-6 shows how Figure 5-5 might be interpreted as a grid. For an example of a BADGER grid, see, e.g., [IP2 2001; PDF p. 56 of 131]. Note that Figure 5.6 is again a hypothetical interpretation for illustration only; although the grid was adapted from data in Figure 5-5, no correspondence to an actual panel is suggested. The left column represents the panel elevation in 2-inch increments, corresponding to the step size of the scan. The elevation is at the axial midplane of the detectors; in this idealized scan, each 2-inch step moves the detector midplane from elevation 1 inch to 3 inch, and so on up the panel. Each column represents $1/4^{th}$ of the width of the panel, each corresponding to one of the four detectors.

The gaps are shown in dark gray, with numbers indicating the number of thirds-of-an-inch of gap in the grid cell, per Reference [IP2 2001]. Thus, at 117-inch elevation is a 4/3-inch gap, denoted by the row of 4's; this was rounded up from the 1.25-inch gap actually observed. Because the 2.25-inch gap at the 7-inch elevation has a larger size than the scan step size of 2 inches, the remaining 0.25 inch is applied to the neighboring cell. The algorithm for conservatively corresponding trace peaks to grid cells is discussed in part in [IP2 2001, Section 4.2], but further details are proprietary. The 3/3 inch = 1 inch of end shrinkage and loss is shown at the top, and the possible 1/3-inch gaps are also included.

The losses due to local dissolution are shown in the light gray cells as a relative percent loss, rounded up to the nearest 5 percent, per Reference [IP2 2001], as compared to the uniform dissolution state of the panel. In a purely mathematical representation, the possible B_4C deposits would be shown as negative numbers; however, conservatively no credit is taken for possible deposits. As discussed in Section 3.8, independent of the local degradation and dissolution effects, the panel surface will have some degree of uniform dissolution. This uniform loss state of the panel can be quantified by averaging the BADGER results exclusive of these local effects, and comparing the results to a reference zero-dose panel while normalizing for differences between the nominal as-built ^{10}B areal density of a zero-dose panel and the nominal as-built ^{10}B areal density of the test panel. The local losses are then computed by comparing the measured ^{10}B areal density of a particular grid cell to the average ^{10}B areal density of the panel exclusive of the end shrinkage, gap, and local loss cells. Quantifying the measured ^{10}B areal density is the subject of the next section.

Elevation [inches]	Detector 1	Detector 2	Detector 3	Detector 4
143	3	3	3	3
141		30%	30%	
139				
137				
135				
133	10%	10%	10%	10%
131			10%	
129				
127	15%	25%	5%	
125	30%	30%	30%	10%
123		25%	1	1
121				20%
119			1	1
117	4	4	4	4
115				
113				
111				
109				
107				
105				
103				
101				
99				
97				
95				
93				
91				
89				
87				
85				
83				
81				
79				
77				
75				
73				
71				
69				
67	1	1	1	1
65				
63				
61				
59				
57	1	1	1	1
55				
53				
51				
49				
47				
45				
43				
41				
39				
37				
35				
33				
31				
29				
27				
25				
23				
21				
19				
17				
15				
13				
11				
9				
7	6	6	6	6
5	1	1	1	1
3				
1				

Figure 5–6. Grid view of the degraded panel shown in Figure 5-5. Adapted from [IP2 2001 p56] to correspond with Figure 5-5.

5.3 BADGER Calibration

Qualitatively, the BADGER scan traces provide an indication of the condition of a panel. The number, size, and location of gaps can be correlated with exposure history and rack design, and regions of significant dissolution can be observed. However, to quantify the local and global degradation of a panel requires comparison with calibration standards. Each BADGER system includes a calibration cell assembly. The BWR calibration cell includes a representative rack cell wall for a BWR egg-crate cell. The PWR calibration cell includes representative cell walls for both egg-crate and flux-trap racks (see [EPRI 1998a, Figure 4-2]). The calibration cells are not exact representations of any particular rack geometry but are quite close to most typical designs. The rack walls within the calibration cell contain Boraflex panels of various known ^{10}B area density and well-characterized gaps. The ^{10}B areal density of calibrated Boraflex panels are known from material certification sheets and verified by chemical assay and neutron attenuation testing; the gaps are all precisely measured.

Originally, BADGER was measured against the calibration cells at the beginning and end of every shift day during a BADGER campaign, in part to measure electronics drift during the shift. When scans of the unattenuated region of a rack wall were found to be adequate for detecting and compensating for any drift, BADGER was calibrated with the calibration cell only once or twice during a campaign—at the beginning and sometimes at the end.

The calculation of areal density in a particular grid cell postulates a linear ($y = mx + b$) relationship between the actual ^{10}B areal density, ρ_A, of the panel being scanned, and the transmission, T, of neutrons through the panel:

$$\boldsymbol{\rho_{A,ij} = \rho_{A,0} + m \cdot \ell n\{T_{ij}\}} \qquad (5\text{-}1)$$

where

- i,j are panel grid cell locations by row i and column j, as illustrated in Figure 5-6, which correspond to elevation i and detector j;
- $\rho_{A,ij}$ is the ^{10}B areal density as grams of ^{10}B per cm^2 of Boraflex panel face in grid cell i,j (or, more properly, the grid cell constituent density of ^{10}B in Boraflex in g ^{10}B / cm^3 Boraflex, multiplied by the grid cell Boraflex panel thickness in cm);
- $\rho_{A,0}$ is the known ^{10}B areal density of a zero-dose (non-degraded) panel, ideally based on its material certification sheet;
- m is the local slope of the ^{10}B areal density versus ℓn(transmission) calibration curve, as determined from scans of the calibration cell; and
- T_{ij} is the neutron transmission measured through the panel (detailed below).

The transmission ratio, T_{ij}, is computed as the ratio of the measured grid cell i,j count rate, C_{ij}, to the average count rate measured in the zero-dose panel for detector j, $C_{0,j}$:

$$\mathbf{T_{ij} \quad = \quad C_{ij} \quad / \quad C_{0,j}} \tag{5-2}$$

For example, for a given detector j, if the count rate measured for the grid cell is equal to the count rate measured for the zero-dose panel, then Equation 5-2 shows that the transmission ratio is 1. Since $\ell n\{1\} = 0$, then Equation 5-1 shows that the average ^{10}B areal density in the small grid cell of the measured panel is equal to the average ^{10}B areal density of the entire zero-dose panel: $\rho_{A,ij} = \rho_{A,0}$. If instead for a given detector j the count rate in the panel grid cell is greater than the average count rate in the zero-dose panel, then Equation 5-2 shows that the transmission ratio is greater than 1 so that $\ell n\{T_{ij}\} > 0$. As will be discussed subsequently, the slope, *m*, is negative so that Equation 5-1 shows that the average ^{10}B areal density in the panel grid cell is less than the average ^{10}B areal density of the zero-dose panel: $\rho_{A,ij} < \rho_{A,0}$; that is, a higher count rate indicates that the measured panel has lost ^{10}B areal density relative to the zero dose panel.

The transmission ratio in Equation 5-2 can be multiplied by a *drift correction factor*. This is the ratio of the count rate in the unattenuated region of the panel being measured to the count rate in the unattenuated region of the zero-dose panel. This correction factor accounts for drift in the electronics between scans with respect to the time when the zero dose panel was scanned. It also accounts for rack wall, panel, and cover plate thickness variations between the cell containing the panel being measured and the cell containing the zero dose panel. In general, this factor is within a couple of percent of 1.0.

The slope, *m*, is computed from data acquired during a calibration cell scan. Because loss of ^{10}B increases the transmission of neutrons through the panel, *m* is negative. Figure 5-7 illustrates a hypothetical calibration curve. Because calibration data are not available in the public literature, this curve was developed from a simple Monte Carlo neutron transport shielding model with a geometry somewhat similar to a PWR Region 1 rack cell. While the calibration standard areal densities, shown as blue points, correspond to areal densities in Reference [EPRI 1998a; Table 1-2, p. 1-2], the geometry does not correspond to any actual BADGER calibration cell; this figure is for illustration only.

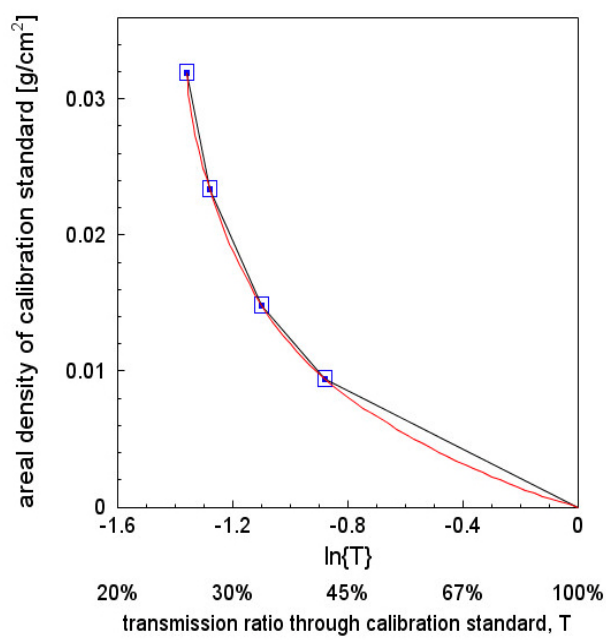

Figure 5–7. A hypothetical simulated calibration curve for a hypothetical calibration cell, adapted from [EPRI 1998a; Table 1-2].

Figure 5-7 shows an actual (hypothetical) response curve in red relating the ^{10}B areal density of a calibration standard to the transmission ratio through that calibration standard, T. Note that the transmission ratios T (in Figure 5-7) and T_{ij} (in Equation 5-2) are not the same ratios. T is the ratio of the count rate through a calibration standard to the count rate in an unattenuated area of the calibration cell. Within the calibration cell, a transmission ratio of T = 1 (100 percent) corresponds to a complete absence of ^{10}B, and therefore a complete absence of Boraflex. As Boraflex calibration standards with higher ^{10}B areal densities are scanned, the count rate decreases as does the transmission ratio as shown. The slope for 100 percent transmission is a relatively small negative value, and the slope becomes increasingly negative as the transmission ratio decreases. This negative slope is the slope *m* in Equation 5-1.

In theory, if the calibration cell exactly matched the spent fuel rack cells, then a ratio of the measured panel attenuated count rate for a grid cell to the unattenuated count rate above the panel would give a transmission ratio which, from this calibration curve, would specify the areal density. However, in practice, there is significant error due to differences between the calibration cell and the rack cell. However, whether in the spent fuel rack or in the calibration cell, for a given ^{10}B areal density of neutron absorber, the rate of change in areal density per unit change in transmission ratio is approximately the same, making the slope from the calibration cell applicable to the racks.

The local slope, *m*, is theoretically the tangent to the red curve for a given areal density. Various methods for computing *m*, from derivatives of regression fits to weighted secants to a simple linear interpolation between data points, as shown by the black lines, can be used. Because the slope is based on the areal density, calculating the ^{10}B areal density is an iterative procedure. For example, the slope may be initially calculated from the calibration curve based on the as-built ^{10}B areal density of the panel being measured. This slope is then used to calculate the areal density of the degraded panel by Equation 5-1. This new areal density can then be used to calculate a new slope, and that used to update the ^{10}B areal density.

The intercept term in Equation 5-1, $\rho_{A,0}$, is obtained by measuring a zero-dose reference panel in service in the pool. For the reference panel, the known ^{10}B areal density is the same as the nondegraded areal density, taken either from a nominal value or ideally from a material certification sheet. The relative error in $\rho_{A,0}$ produced by the differences in geometry between the calibration cell and the rack cell being measured is relatively large. Therefore, $\rho_{A,0}$ is measured on a zero dose panel of the same design as the panel being scanned. This reduces the geometry error to the small manufacturing variations between rack cell walls and panels. As noted above, even this difference can be reduced by applying the drift correction factor, which accounts for differences between the unattenuated region scans of the zero-dose panel and the panel being measured. For an even more robust assessment of measurement uncertainty, $\rho_{A,0}$ may be measured for more than one zero-dose panel and the results averaged. Section 4.3.3 discussed how RACKLIFE can be used to identify zero-dose panels.

Once the calibration curve is determined and the zero-dose panel measured, the ^{10}B areal density of any measured grid cell on a degraded panel (as illustrated in Figure 5-6) can then be computed from the detector count rates. It should be noted that the thickness of the measured panel—which can vary by up to 10 percent as noted in Section 2.3—is not necessary in that the areal density is not computed directly from the panel; it is instead computed by comparison with the calibration curve and the zero-dose panel.

Once the absolute areal density of a panel is computed, the local percent loss in a grid cell can be computed by comparing the measured ^{10}B areal density to the as-built ^{10}B areal density specified by material certification sheets. Integrating these losses across the entire panel gives a panel average loss, which can be compared to RACKLIFE panel average losses, as described in Section 4.3. In particular, as discussed in Section 4.3.4, the RACKLIFE results can be used to corroborate the BADGER results. This is best understood in the context of the uncertainties in the BADGER measurements, which is the subject of the next section.

5.4 Uncertainties in BADGER

For the purposes of this report, uncertainties in applying the BADGER system to calculate ^{10}B areal densities were categorized into five principal areas:

- scan repeatability;
- the calibration cell scan and calibration curve slope;
- the zero-dose panel scan, including uncertainty in the actual areal density of the zero-dose panel;
- the unattenuated scan of the panel being measured; and
- the attenuated scan of the panel being measured.

The first, scan repeatability, is an uncertainty common to all of the scans subsequently discussed. To quantify some of these uncertainties, the simple low-fidelity simulation model that was developed to illustrate the typical calibration curve in Section 5.3 was used. Similar to the uncertainty analysis for RACKLIFE in Section 4.4, an objective of this section is to present techniques for analyzing BADGER uncertainties as well as quantitative results from applying these techniques to a highly idealized model of a spent fuel rack containing Boraflex. These uncertainties are only for illustrative purposes and to guide analysis efforts to target the largest uncertainties. The uncertainties presented here should not be considered representative of any particular spent fuel rack design.

5.4.1 Scan Repeatability

Scan repeatability is effectively the ability to draw the same conclusions about the degradation magnitude and morphometry of a panel from two independent scans of the same panel or

calibration cell, either within a campaign or between campaigns. This is more than just the repeatability of a scan trace; because of variations in such factors as elevation, alignment, and drift between scans, the traces will be different. The issue is, would the same conclusions be made from these two different traces. There are four factors that confound repeatability:

A. variations in detector response and channel electronics;
B. detector/source head misalignments;
C. changes to racks over time (between campaigns); and
D. unavailability of cells in a subsequent campaign.

These are discussed in turn below.

A. Variations in Detector Response and Channel Electronics

Each of the four detectors and associated channel electronics respond differently. As illustrated in Figure 5-8, differences between the "inner" and "outer" detectors are to be expected—the path length from the point source is different, and neutrons further from the centerline from source to detector block are more likely to be absorbed by the detector block neutron shielding or by other panels of Boraflex around the rack cell (see Figure 5-2). Simulations show that when measuring high ^{10}B areal density panels, the inner and outer count rates will generally be within the typical counting statistics uncertainty of ~1-2 percent. This difference increases as the ^{10}B areal density of the panel decreases until a difference of ~30 percent is observed during the scan of the unattenuated region.

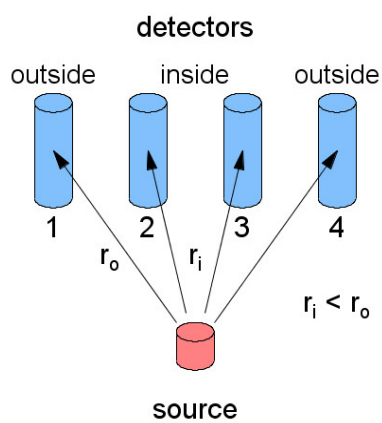

Figure 5–8. BADGER *inner* and *outer* detectors.

In early BADGER campaigns, it was not known how much electronics drift might occur from the beginning to the end of an 8-10 hour shift. As experience with BADGER was gained, it was observed that drift was generally not detectable relative to other measurement uncertainties. However, spurious signals would occasionally be noted. Spurious signals were associated either with disturbances to the electrical cables between the BADGER operator and the BADGER drive platform, or with external plant systems switching on and off on the same power circuit. The spurious signals did not affect all channels equally—sometimes only a single channel trace would show anomalous signals. Thus, it was important for the BADGER operator to watch the real time traces individually and to note any anomalous behavior to possibly correlate it with events taking place.

Occasionally a detector would fail, usually due to water ingress. If a spare detector was available, the faulty detector would be replaced with the spare, thus changing the performance characteristics of that detector location. Changing a detector always requires a new calibration scan and a new zero dose panel scan. Thus, while electronic drift is generally not a factor in repeatability, noise and detector change-out, while rare, will be a factor.

Individual detectors also vary, regardless of position or operational status. For example, Figure 5-5 notes at the top that the detector-1 response, in the uniform size unattenuated regions and the well-defined and likely uniform 2.25 inch gap, is lower than the response of what should be a symmetric detector-4. The result is that each detector must be calibrated independently. The detectors can of course be compared as they will be correlated with the presence of gaps and other cross-panel anomalies, but absolute count rate differences on the order of ~3 percent or more between the two inner or the two outer detectors (beyond counting statistics) should be expected.

B. Detector/Source Head Misalignments

Figure 5-9 illustrates three principal misalignments of BADGER in a rack cell—offset, twist, and tilt.

i. Offset occurs when the head centerlines are not kept collinear. This is a function of the amount of clearance fit available to insert BADGER into various rack cells. This is often the most difficult misalignment to correct when positioning and inserting BADGER into the cells, but it can be the easiest to detect. As noted in Section 5.1, sighting along the insertion poles can give a visual indication of an offset misalignment. A scan trace with "one side high, one side low" for the two outside detectors (as compared to a previous trace) can also be an indicator. This gradient is generally observable by the BADGER operator on the data console so that corrective action could be taken immediately. Shims on the heads to reduce the clearance can be used as long as there are no "tight" cells due to cover plate deformations (e.g., from manufacturing, off-gas buildup, or fuel assembly impacts during movement).

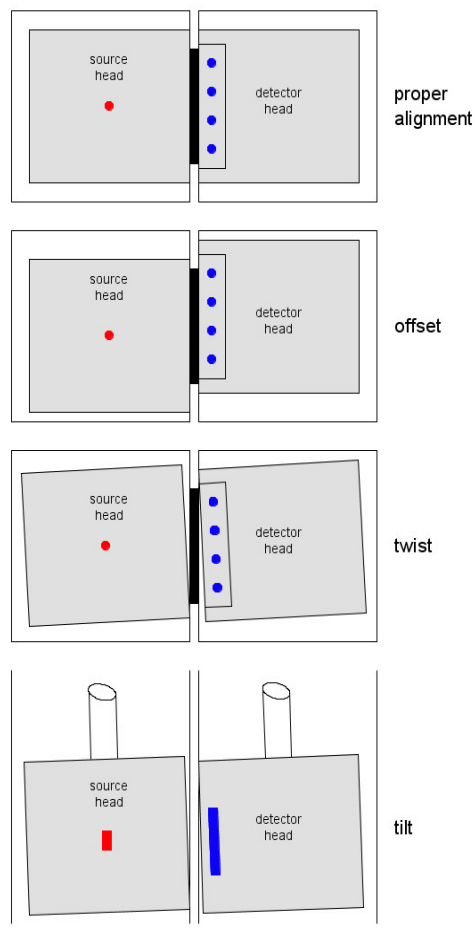

Figure 5–9. BADGER source / detector head misalignments.

ii. Twist misalignment generally occurs due to torques applied by attached electronics cables that hang above the pool from the drive system plate to the control computer, as shown in Figure 5-3. A cable swivel mount and ensuring that the system is precisely perpendicular to the panel to be scanned usually eliminates this problem.

iii. Tilt misalignment occurs when the overhead crane does not precisely position the two probe heads directly over the Boraflex panel between the two cells. The inability to insert the heads completely into the cells or visual sighting from two directions will usually detect this problem, which is easily eliminated by adjusting the position of the crane. In the earliest BADGER campaign, scans started at the top and proceeded down. However, tilt problems led to the practice of first inserting BADGER fully into the cell, aligning as necessary, and then scanning upward.

"Offset" in (i), above, refers to the lateral offset of the heads. Axial offset is prevented by the rigid suspension poles extending from the top of the heads and coupled above the top of the rack cells. As long as there is no distortion of the length of the poles, the two heads will remain axially aligned. As discussed in Section 5.1, there can be axial demand elevation "offset" (bias), but the source and detector heads will still be axially aligned, with the same degree of demand elevation bias for each head.

Separation offset is prevented by a spring at the top of these poles that draws the poles toward each other as shown in Figure 5-3. Thus, the heads are kept in light contact with the intervening cell wall and cover plate. However, surface variations, such as cover plate bumps and wrinkles due to manufacturing or off-gassing, can change the source-detector distance as the BADGER heads pass over them. In a repeat scan of the same panel, the bottom elevations of the heads (the head offset) during the second scan is almost equal to the bottom elevation of the first scan (i.e., the original bias does not change). This is because the BADGER heads are first lowered to the rigid bottom of the cell before they are lifted out in discrete 2-inch steps. However, because the bottoming process is not exact, the surface bumps may be sampled differently during a repeat scan, adding uncertainty to the repeatability.

It is nearly impossible to generally quantify the uncertainty due to offset, twist, and tilt. As noted in Section 2.3, there is a wide variety of rack designs, and the amount of clearance between the BADGER source and detector heads and the rack inside walls will vary—even cell to cell from manufacturing tolerances, in-service deformations, panel cavity swelling, or other deformations. Limited simulations of small variations in offset produced count rate variations up to ~4 percent beyond counting statistics. However, this uncertainty is sensitive to the actual clearance and how the head positions correlate within this clearance, and thus could vary widely between designs. A rigorous modeling based on *a priori* geometric considerations may indicate higher uncertainties. The best method for quantifying this uncertainty would be to perform repeatability experiments in the racks being tested to account for specific geometries and materials. Deliberate maximum misalignments could even be injected into the test to bound the uncertainties.

C. Changes to Racks Over Time (Between Campaigns)

As spent fuel is moved into and out of cells, local deformations to the thin Boraflex cover plates may occur. These may change the source-to-detector spacing between campaigns. Even within a campaign, the gentle pressure of BADGER on the cover plates to prevent separation offset can be enough to squeeze out bubbles (as discussed in Section 3.2), indicating that a slight change in the local source-to-detector distance and material composition may have occurred. A re-scan could potentially squeeze out more, further confounding repeatability within the campaign.

D. Unavailability of Cells in a Subsequent Campaign

BADGER requires that spent fuel in cells around or near (typically within 3 cells) of a panel to be scanned be moved away. This is to prevent gamma interference to the neutron detectors; however, no data are publicly available to quantify this interference. In subsequent campaigns, the cells that were once available for testing may no longer be available, so that a repeat scan is not possible.

5.4.2 The Calibration Cell Scan and Calibration Curve Slope

The importance of a thorough verification and validation of the BADGER system was understood early in its demonstration process, both in general and specific to the calibration process. As discussed in Section 2.2, RACKLIFE and BADGER were first deployed at about the same time, around 1996. For some early BADGER campaigns it was not well understood how to tune a RACKLIFE model using the escape coefficient, especially for multiregion pools. The result was that RACKLIFE could significantly underpredict the actual loss in racks, making it ineffective in its corroboration role (as discussed in Section 4.3.4). The disparity between RACKLIFE and BADGER results was noted in Reference [EPRI 1998a; pp. 1-1 and 7-1] for a PWR flux-trap rack that had shown 25-30 percent Boraflex degradation. In addition to questioning the results from RACKLIFE, it was also hypothesized that the calibration cell might not sufficiently reflect the racks being measured, leading to increased uncertainty in the results. The BADGER verification and validation process, detailed in [EPRI 1998a], addressed these uncertainties in particular.

In the verification and validation report, several data tables [EPRI 1998a; Tables 4-3, 4-4, 5-3, 5-4, 6-1, and 6-2] compare the count rate ratio for both racks and calibration cells. Counts were compiled from both experimental BADGER data and from numerical MCNP (a Monte Carlo neutron transport code) calculations. BADGER data and MCNP calculations were compared for three specific designs—a PWR Region I rack, a PWR Region II rack, and a BWR rack as well as their corresponding calibration cells.

The data from these tables are plotted in Figures 5-10, 5-11, and 5-12 to assess the uncertainties associated with the differences between the inner and outer detectors, the differences between racks and calibration cells, and the differences between BADGER experiments and MCNP calculations. The count rate ratio is the ratio of the count rate observed for a given ^{10}B areal density to the count rate observed for a reference ^{10}B areal density. The reference ^{10}B areal density in each case was the nominal as-built ^{10}B areal density of the Boraflex rack panels. These ^{10}B areal densities were known at the batch average level, but panel specific as-fabricated ^{10}B areal densities had to be inferred. Regardless, the most important result is the trend in differences with loss, as captured by the count rate ratio, rather than specific values of loss. Given the different response of the inner and outer detectors (per Section 5.4.1.A) in each figure, they are plotted separately—the inner detectors are plotted in red and the outer detectors are plotted in blue. The solid lines are linear fits between the experimental data points for a rack, and the dashed lines are linear fits between the experimental data points for the corresponding calibration cell. The letters show the numerical results for the rack, R, and calibration cell, C.

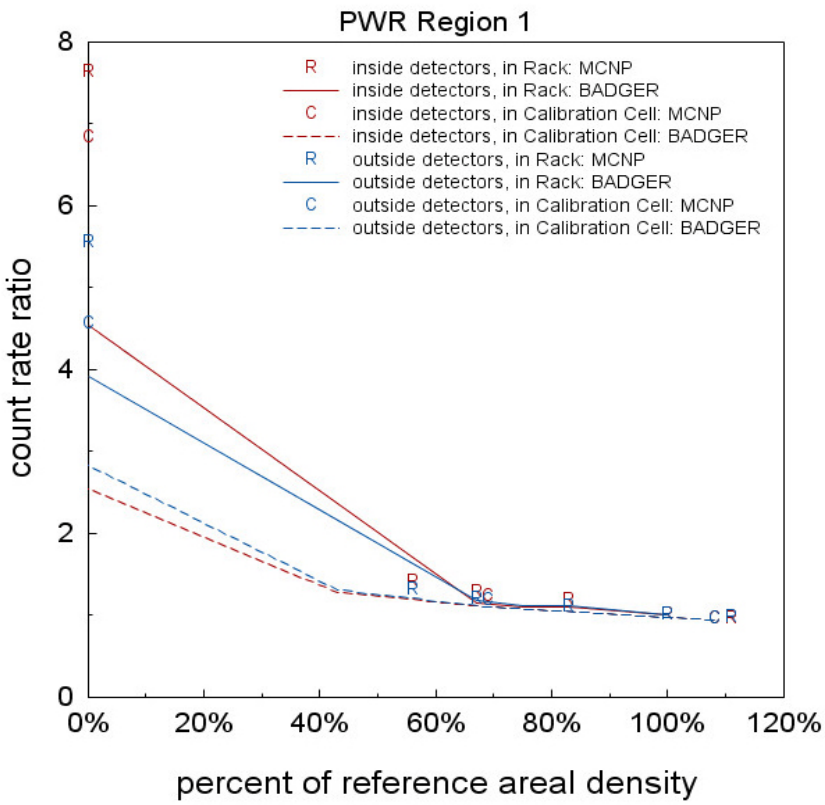

Figure 5–10. PWR Region I - rack versus calibration cell and experimental vs. numerical.

Figure 5-10 illustrates the PWR Region I racks and calibration cell. It shows that for ^{10}B areal densities within a range of +10 percent and −30 percent of the nominal (100 percent) ^{10}B areal density, the count rate ratios for the racks and the calibration cell correspond very closely. Further, the experimental BADGER and numerical MCNP results are also in close agreement. The small differences indicate that, should an experimental calibration scan at the beginning of a campaign suffer from gross experimental error (e.g., due to excessive misalignment) but go undetected until after the campaign, the data collected during the campaign may still be used. (Recall that a BADGER scan of the calibration cell is typically conducted only once in a campaign.) A numerical simulation can provide the necessary calibration curve slopes to compute ^{10}B areal densities per Section 5.3.

If panel ^{10}B areal densities are less than 70 percent of nominal, however, the agreement between rack cell and calibration cell, experimental BADGER and numerical MCNP, and inner and outer detectors ends and uncertainties between the pairs increase. The actual disagreement in the 10 to 50 percent range is likely overemphasized in these plots. Because data points (the symbols R and C) are available only for panels at 67 percent and 0 percent of nominal, the linear fit over this large range changes the curvature of the graph. It is more likely that the curvature of the graph would more closely match the curvature of the simulated calibration curve in Figure 5-7. Therefore, the gap between the rack and calibration data may not manifest until ^{10}B areal densities of lower than 50 percent of nominal.

However, looking at the 0 percent (unattenuated) results on the y axis of Figure 5-10, there is clearly a large disagreement between the racks and the calibration cell, and between the experimental and numerical results. Reference [EPRI 1998a] discusses some hypotheses to account for these differences (e.g., neutron spectrum effects on detector response, per [EPRI 1998a; pp. 7-1 and 7-2]), but the source of this disagreement is unknown. Note that the figure also overemphasizes the actual nature of the disagreement. The calibration cell does not set the y-intercept of the ^{10}B areal density versus count rate calculation; the zero-dose panel does. Further, the unattenuated count rate in the transmission ratio is due to neutron absorption of the in-service rack materials and not from the calibration cell. Only the slope of the calibration curve is important, and how its slope relates to that of the racks determines the actual uncertainty in the results. For these reasons, the R and C symbols on the y axis are not valid representations of pool panel data and, therefore, should not be used to calculate the uncertainty in panel areal density.

Figure 5-11 is for the PWR Region II racks and calibration cell. Again, the disagreement in the lower areal density range is overemphasized because of so few data points for the calibration cell. But again, there is clear significant disagreement for the unattenuated results. For BWR count rates, the reference [EPRI 1998a] did not supply calibration cell data for comparison. However, the same magnitude of difference in the numerical results at 0 percent ^{10}B areal density, as shown in Figure 5-12, suggests that the disagreement will be the same.

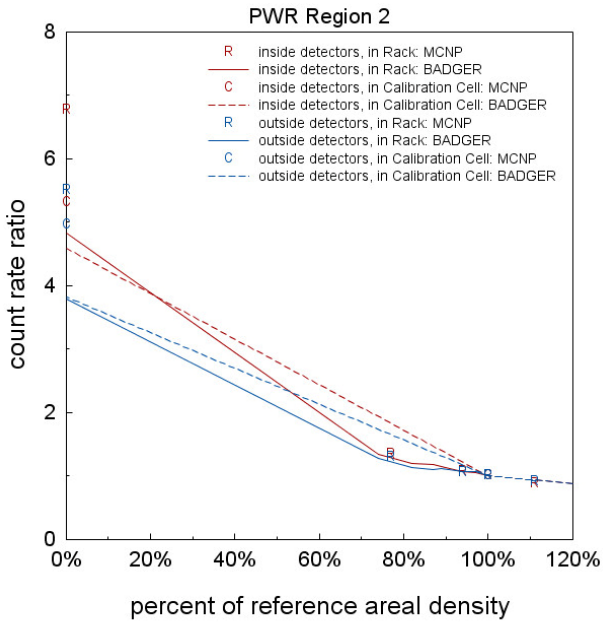

Figure 5–11. PWR Region II - rack versus calibration cell and experimental versus numerical.

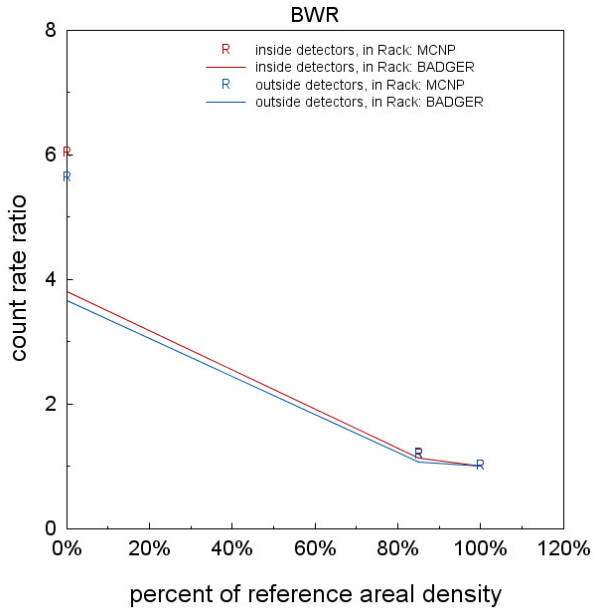

Figure 5–12. BWR - rack versus calibration cell and experimental versus numerical.

One conclusion from these figures is that Figure 5-10, a PWR Region I flux-trap configuration, appears to bound the uncertainty for all three graphs. Figure 5-10 is likely to be reasonably accurate for ^{10}B areal densities within +10 percent and −30 percent of nominal. There are insufficient data to make confident calculations of uncertainties due to the calibration cell, but the following is observed. A comparison of fits of the available PWR Region I data shows that the calibration cell slope could be greater than the rack slope by a maximum of about +10 percent in the −30 percent range. Applying this +10 percent slope to a ^{10}B areal density calculation shows that a 30 percent loss predicted by the calibration cell slope would actually be a 27.3 percent loss if the actual slope from the rack was used. Thus, a large 10 percent variation in slope translates into a less than 4 percent variation in ^{10}B areal density for panels which retain 70 percent of their ^{10}B areal density, i.e. 30 percent loss. At ^{10}B areal density losses larger than 30 percent, the gap between the rack and calibration curve increases significantly and, therefore, the uncertainty will be greater though the magnitude cannot be determined from the available data. In addition, larger variations in calibration slope will lead to greater uncertainty in *m* and therefore a greater uncertainty in B-10 areal density.

A final point is to note from Figure 5-7 that the slope of the calibration curve becomes increasingly negative with increasing ^{10}B areal density. High ^{10}B areal density panels are significantly more sensitive to uncertainty in the transmission ratio, and this is exacerbated by uncertainties in the calibration slope itself. Thus, the higher the ^{10}B areal density of the panel, the more variation will be introduced from natural variations in the count rate just due to counting statistics. This makes physical sense, as the uncertainty in detection increases when fewer neutrons are detected during the scan of a high-density panel.

5.4.3 The Zero Dose Panel Scan

As discussed in Section 5.4.1, the best method for assessing the uncertainty of BADGER is to perform repeatability experiments in the actual racks. A zero dose panel is a natural choice for one such experiment because, in addition to being representative of all panel scans, it provides the intercept to the ^{10}B areal density calculation made from the BADGER measured count rates, as discussed in Section 5.3. Any observed uncertainty can then be applied directly to the intercept.

To determine the uncertainty of the ^{10}B areal density of the zero-dose panel, one approach is to measure several zero dose panels and combine the measured ^{10}B areal densities in an uncertainty-weighted average. This is particularly useful if material certification records are unavailable or insufficient to know the ^{10}B areal density of a particular zero-dose panel. It has been the case where Boraflex manufacturing batch ^{10}B areal densities are known, but which panel in the racks came from which batch is not known. By characterizing multiple zero dose panels from multiple batches, the distribution can be understood and applied in using the intercept. Note, however, that scans of different panels, even if all zero dose, does not address the previous issue of repeatability.

Typically, material certifications are accurate to ±0.0001 g/cm^2 in ^{10}B areal density. This translates to ±1.3 percent for a low 0.008 g/cm^2 ^{10}B areal density panel and ±0.3 percent for a high 0.030 g/cm^2 ^{10}B areal density panel. If the batch that the zero dose panel came from is known, then the ^{10}B areal densities computed from BADGER data will have these corresponding uncertainties. However, a bounding example is where the specific batch is not known, and only one panel is measured. In that case, batch variation of ±0.001 for a typical nominal ^{10}B areal density of 0.02 g/cm^2 will produce an uncertainty of about ±5 percent, which may then be applied as a bias to the ^{10}B areal densities calculated. Greater batch variations will produce greater uncertainties.

A final uncertainty regarding the zero-dose panel is the question of whether a nominally zero-dose panel has actually received no gamma exposure. In Section 4.2.2, Figure 4-6 showed that panels below $5 \cdot 10^8$ rads are equivalent to zero-dose panels, regardless of actual dose. This suggests that a wide range of very low dose panels are all acceptable as negligible-loss zero-dose panels. Using RACKLIFE to identify zero-dose panels, per Section 4.3.3, will rely on RACKLIFE's dose calculation, not the silica kinetics loss calculation, to identify actual zero-dose panels. However, panel selection can be difficult if there is reluctance to move fuel to make the selected panel accessible to BADGER. Further, there have been cases (e.g., [EPRI 1997b, p. 3-27]) where fuel movements input to RACKLIFE did not necessarily account for all of the actual moves made. In the reference cited, temporary fuel moves during rerack operations may not have been input. Finally, RACKLIFE only tracks spent fuel gamma sources, not other gamma sources such as discharged reactor control elements.

5.4.4 The Unattenuated Scan

Section 5.1 described the unattenuated scan, and Section 5.3 described how it is used in computing the drift correction factor. All aspects of the uncertainty in the unattenuated region scan are the same as for the attenuated scan, described next, except the following. The slope of the calibration curve at the point of 100 percent loss (giving a transmission ratio of 100 percent, which is effectively the unattenuated region) has a relatively shallow slope. Thus, the unattenuated scan is much less sensitive to counting uncertainty and uncertainty in the calibration curve slope. This makes it well suited for detecting rack cell manufacturing differences, anomalous misalignment, and electronics drift in making comparisons between rack cells. Regarding misalignment, however, having the top of the source and detector heads at or above the top of the rack cell makes them more prone to tilt and twist misalignment. For an 8 percent misalignment variation in count rate (twice the value estimated in Section 5.4.1.B to account for being at the cell top), the simple low-fidelity simulation model suggests that the drift correction factor may vary by up to 3 percent beyond the counting statistics variation. This is comparable to the up to 2 percent variation typically attributed to drift and manufacturing variations in practice.

5.4.5 The Attenuated Scan

A 4-percent variation in count rate due to misalignment, as estimated in Section 5.4.1.B, can result in a 1.5 percent or less variation in the computed ^{10}B areal density, beyond the typical 1-2 percent counting statistics uncertainty. When measuring the calibration cell, it is typical to observe variations of 2-4 percent in repeat scans, which is 2-3 percent greater than what the counting statistics alone would predict. If this is due to varying degrees of misalignment, then this same higher variation is likely to be occurring in the racks as well. Again, the only definitive assessment of this uncertainty is to perform a repeatability experiment on a panel being analyzed.

Another aspect of repeatability is the ability to reliably detect panel gaps in repeat scans. While a repeatability analysis based on the zero dose panel has the advantage of easily factoring in its results via the intercept in the ^{10}B areal density calculation, by its nature the zero dose panel will not have any gaps. A repeatability analysis of a known degraded panel would quantify the uncertainty in characterizing the size of a gap and the reliability of observing small gaps – those close to the limit of 1/3 inch.

In analyzing the uncertainty in an attenuated scan, it must be recalled from Section 2.3 that the variation in thickness of any one Boraflex panel is typically ~5 percent and can be as large as 10 percent. To ensure that each panel contains the minimum certified areal density, each panel is manufactured such that its ^{10}B areal density is typically 15 percent higher than the minimum certified ^{10}B areal density. Thus, a 4 percent uncertainty in a particular grid cell measurement of ^{10}B areal density due to misalignment will be comparable to or smaller than the variation in areal density due to natural variations in thickness. Further, if material certification sheets are unavailable, the uncertainty will be much smaller than differences between the unknown as-built and known minimum certified conditions. Only by observing correlated responses across multiple points can the effects of misalignments be analyzed.

Dissolution effects are generally quite smooth, especially in comparison to the size of a grid cell observed by a detector. Thus, BADGER will only observe large scale changes in areal density. Applying propagation of error techniques to Equations 5-1 and 5-2 for typical count rates and calibration curve slopes obtained from the low-fidelity simulation model ^{10}B areal density, the uncertainty in ^{10}B areal density for a given grid cell without any local degradation or dissolution effects is around ±15 percent. Assuming this applies to any portion of the panel that has undergone only uniform dissolution, then averaging over these grid cells by excluding those with local effects, the average panel ^{10}B areal density uncertainty should be just over 1 percent when the increase in surface area mitigates the uncertainty in counting statistics In contrast, reference [EPRI 1997b; p. 3-8 and Fig. 3-7 on p. 3-11] identifies the overall uncertainty in average panel ^{10}B areal density as a standard deviation of 8 percent. This indicates there are mechanisms beyond counting statistics that affect the uncertainty including as-built variations in areal density, calibration slope uncertainty, zero-dose panel uncertainty, and misalignment.

5.4.6 Summary of BADGER Uncertainties

The uncertainties described in this chapter are summarized in Table 5-1. Except for uncertainty due to misalignment errors, the panel scan repeatability uncertainties are impossible to quantify without repeatability testing data; these data are not available in public documents. Other sources of error are quantified to illustrate their magnitude based in part on the low-fidelity simulation model and estimates of material certification uncertainties. Table 5-1 also shows the difference between local ^{10}B areal density uncertainty and panel average ^{10}B areal density uncertainty discussed above. Again, as discussed at the beginning of Section 5.4, the uncertainties presented here are only for illustrative purposes and to guide error analysis efforts to target the largest uncertainties. These results should not be considered representative of any particular spent fuel rack.

In principle, the above analyses apply to a BADGER scan of any neutron absorber, not just Boraflex. The uncertainties inherent in scan repeatability, including misalignment uncertainty, are features of the BADGER system and not the type of neutron absorber. These uncertainties would, within the limitations of the low-fidelity model used here, apply to other neutron absorbers such as Boral™ and Metamic™. If BADGER testing is conducted on a non-Boraflex material, then the slopes of the calibration curve will necessarily change. However, at least for the areas where ^{10}B areal densities are closer to that of an undegraded panel (i.e., losses less than 30 percent), the calibration curve could likely be simulated numerically. Four issues bearing on uncertainty were identified in [EPRI 1998a, p. 7-2] for consideration prior to using BADGER for non-Boraflex applications:

- neutron detector model fidelity in the numerical (MCNP) model;
- numerical geometric and material model fidelity;
- neutron detector characterization; and
- independent numerical model validation.

A final note of caution is warranted for non-Boraflex applications. As discussed in Section 5.2, the experience of the operator and the analyst is likely a strong determinant of repeatability. BADGER operators and analysts are attuned to look for gaps, scallops, and areas of local dissolution. For other materials, the operator and the analyst would have to attune themselves to the degradation mechanism of the particular material.

Table 5-1. Summary of BADGER uncertainties.

source of uncertainty		estimated uncertainty	notes on uncertainty
scan repeatability			can only be known via repeatability testing
	variations in detector response and electronic drift	–	largely mitigated by the drift correction factor, as described in Section 5.3; large uncertainty specific to events (e.g., induced noise)
	detector / source head misalignments	4%	from very limited simulations; may be higher with *a priori* geometric modeling, highly variable in racks in service
the calibration cell scan and calibration curve slope		4%	at 30 percent loss and a 10 percent variation in calibration slope; increases quickly at higher losses or high variation in calibration slope
zero dose panel scan		5%	bias is dependent on batch variation, e.g., ±0.001 g $^{10}B/cm^2$ uncertainty for batch average 0.02 g $^{10}B/cm^2$ areal density will produce 5 percent uncertainty in ^{10}B areal density
the unattenuated scan		–	contributes to drift correction factor, and thus only indirectly to ^{10}B areal density uncertainty
the attenuated scan			
	specific panel grid cell	15%	based on propagation of error from counting statistics
	panel average	~1%	based on averaging uniform dissolution grid cells only
	panel average – reported	8%	uncertainty reported in [EPRI 1997b]

.

6 CONCLUSIONS

When the neutron absorber material Boraflex is exposed to gamma radiation from spent nuclear fuel, it will degrade nonhomogenously in the aqueous environment of the spent fuel pool, especially under conditions of high cumulative dose and high pool temperature. The resulting loss of neutron absorber, typically over several decades, can erode the spent fuel rack's design neutron absorption margin to approach or potentially exceed the *regulatory* requirements for subcriticality. The rate and distribution of the degradation is highly dependent on the design of the high-density spent fuel racks that utilize panels of Boraflex between fuel assemblies for criticality control.

RACKLIFE is a commercial software program developed to calculate the average degradation of a Boraflex panel and to predict the future degradation of Boraflex expressed as a percent loss. The RACKLIFE code is based on the specific degradation mechanism of Boraflex and combines physical models of gamma radiation absorption and dissolution kinetics with empirically derived model rate coefficients. BADGER is a commercial in situ nondestructive diagnostic tool for measuring the areal density of ^{10}B in a neutron absorber as installed in fuel racks. Overall, RACKLIFE and BADGER approach the loss of ^{10}B areal density from Boraflex from opposite sides—RACKLIFE calculates the amount of ^{10}B that is lost while BADGER measures the amount of ^{10}B that remains. Together they complement and supplement each other. The combined RACKLIFE/BADGER method is a tool—albeit with significant uncertainties—for quantifying the loss of Boraflex in spent fuel pool racks.

For this report, a computer program to simulate RACKLIFE was written. The simulation predicts that variations from nominal in accumulated panel gamma dose and local water temperature will produce the highest uncertainties in percent degradation results calculated by RACKLIFE. It should be noted that this simulation applies only to a single Boraflex panel in one specific rack design. Therefore, the uncertainties calculated from this model are necessarily only for illustrative purposes, and serve to guide analysis efforts to identify the most significant sources of uncertainty. While the specific uncertainties calculated in this example are of limited applicability, the stochastic techniques introduced for calculating the uncertainty in RACKLIFE predictions of Boraflex loss may be implemented by commercial RACKLIFE users to rigorously quantify uncertainties in their own particular models.

Similarly, for this report, a simple low-fidelity model of a BADGER system was developed in order to simulate a BADGER scan of a panel in one specific rack design. Several BADGER uncertainties were simulated by the model, and it was determined that the greatest uncertainty arises from variations in detection and counting statistics, and, as the panel degrades further, in the calibration method.

7 REFERENCES

7.1 NRC References

The following references, in chronological order, are available from the NRC public library via their Web-based ADAMS (http://wba.nrc.gov:8080/ves/). Each reference includes its ADAMS accession number in brackets at the end of each citation.

[NRC GL 78-11] NRC Generic Letter 78-11: *Guidance on spent fuel pool modifications, entitled "Review and Acceptance of Spent Fuel Storage and Handling Application", to all Power Reactor Licensees*. United States Nuclear Regulatory Commission, Office of Nuclear Reactor Regulation: Washington D.C.; 14 April 1978. [ADAMS ML031280383]

[NRC IN 87-43] NRC Information Notice No. 87-43: *Gaps in Neutron-Absorbing Material in High-Density Spent Fuel Storage Racks*. United States Nuclear Regulatory Commission, Office of Nuclear Reactor Regulation: Washington D.C.; 8 September 1987. [ADAMS ML031130349]

[NRC IN 93-70] NRC Information Notice 93-70: *Degradation of Boraflex Neutron Absorber Coupons*. United States Nuclear Regulatory Commission, Office of Nuclear Reactor Regulation: Washington D.C.; 10 September 1993. [ADAMS ML031070107]

[NRC IN 95-38] NRC Information Notice 95-38: *Degradation of Boraflex Neutron Absorber in Spent Fuel Storage Racks*. United States Nuclear Regulatory Commission, Office of Nuclear Reactor Regulation: Washington, D.C.; 8 September 1995. [ADAMS ML031060277]

[NRC GL 96-04] NRC Generic Letter 96-04: *Boraflex Degradation in Spent Fuel Pool Storage Racks*. United States Nuclear Regulatory Commission, Office of Nuclear Reactor Regulation: Washington, DC; 26 June 1996. [ADAMS ML031110008]

[IP2 2001] *Indian Point, Unit 2 – License Amendment Request (LAR 01-010) for Spent Fuel Storage Pit Rack Criticality Analysis with Soluble Boron Credit*. 20 September 2001. [ADAMS ML012680336]

[Walker 2010] Walker, J. Samuel, and Thomas R. Wellock. *A Short History of Nuclear Regulation, 1946–2009*. NUREG/BR-0175 Rev. 2. U.S. Nuclear Regulatory Commission: Washington, D.C.; October 2010. [ADAMS ML102940169]

7.2 EPRI References

The references below are EPRI documents, in chronological order. Most are available as PDF files from the EPRI Web site (www.epri.com) by using the search box there. Search for the report number or report title to get to a download page. Those indicated with an asterisk were not available for download from the EPRI Web site when this report was being prepared.

[EPRI 1988] *An Assessment of Boraflex performance in spent nuclear fuel storage racks*. Report NP-6159. EPRI: Palo Alto, CA; December 1988.

[EPRI 1993a] *Boraflex Test Results and Evaluation*. Report TR-101986. EPRI: Palo Alto, California; February 1993.

[EPRI 1993b] *Guidelines for Boraflex Use in Spent-Fuel Storage Racks*. Report TR-103300. EPRI: Palo Alto, California; December 1993.

[EPRI 1995]* *Inspection and Testing of Boraflex from the Fort Calhoun Spent Fuel Racks*. NETCO Report NET-092-02, prepared for EPRI under RP-3907-01. Northeast Technology Corp.: Kingston, New York; March 1995.

[EPRI 1997a] *The Boraflex Rack Life Extension Computer Code – RACKLIFE: Theory and Numerics*. Report TR-107333. EPRI: Palo Alto, California; September 1997.

[EPRI 1997b] *BADGER, a Probe for Nondestructive Testing of Residual Boron-10 Absorber Density in Spent-Fuel Storage Racks: Development and Demonstration*. Report TR-107335. EPRI: Palo Alto, California; October 1997.

[EPRI 1997c] *A Synopsis of the Technology Developed to Address the Boraflex Degradation Issue*. Report TR-108761. EPRI: Palo Alto, CA; November 1997.

[EPRI 1998a] *MCNP Validation of BADGER*. Report GC-110539. EPRI: Palo Alto, California; May 1998.

[EPRI 1998b] *The Performance of Irradiated Boraflex under Seismic Conditions*. Report TR-109927. EPRI: Palo Alto, California; September 1998.

[EPRI 1999a] *The Boraflex Rack Life Extension Computer Code – RACKLIFE: Verification and Validation*. Report TR-109926. EPRI: Palo Alto, California; March 1999.

[EPRI 1999b] *The surface composition and solubility of irradiated boraflex and silica treated in metal ion solutions*. Report TE-114126. EPRI: Palo Alto, CA; November 1999.

[EPRI 2000a] *LWR Plant Water Silica Database, Revision 4*. EPRI: Palo Alto, California; February 2000.

[EPRI 2000b] *Elevated Silica Project at Palisades—Final Report*. Report 1000991. EPRI: Palo Alto, CA; November 2000.

[EPRI 2002a] *Guidance and Recommended Procedures for Maintaining and Using RACKLIFE Version 1.10.* Report TP-1003413. EPRI: Palo Alto, CA; April 2002.

[EPRI 2002b] *Sixth Inspection – EPRI Boraflex Surveillance Assembly*. Report 1003414. EPRI: Palo Alto, CA; July 2002.

[EPRI 2004] *Feasibility Assessment: Centrifugal Clarification for Control of Crystalline Silica in Spent Fuel Pools*. Report 1009697. EPRI: Palo Alto, CA; November 2004.

[EPRI 2008]* *RACKLIFE Version 2.1 Escape Coefficient Study*. Report 1016640. EPRI: Palo Alto, CA; November 2008.

[EPRI 2009]* *RACKLIFE Version 2.1.1*. Product ID 1020265. EPRI: Palo Alto, CA; September 2009.

[EPRI 2010]* *Inspection and Testing of Boraflex Surveillance Coupons Nos. 212, 221, and 224 from the Seabrook Nuclear Power Station*. Report 1021451. EPRI: Palo Alto, CA; July 2010.

7.3 Other References

[Dove et al. 2008] Dove, Patricia, Nizhou Han, Adam F. Wallace, and James J. De Yoreo. *Kinetics of amorphous silica dissolution and the paradox of the silica polymorphs*. Proceedings of the National Academy of Sciences: Volume 105, Number 29; 22 July 2008.

[Graubner et al. 2004] Graubner, Vera-Maria, Rainer Jordan, Oskar Nuyken, Bernhard Schnyder, Thomas Lippert, Rüdiger Kötz, and Alexander Wokaun. *Photochemical Modification of Cross-Linked Poly(dimethylsiloxane) by Irradiation at 172 nm*. Macromolecules: Volume 37, Number 16, pp. 5936-5943; 17 July 2004.

[Iler 1979] Iler, Ralph K. The Chemistry of Silica: Solubility, Polymerization, Colloid and Surface Properties, and Biochemistry. John Wiley & Sons: New York, NY; 1979.

[Lindquist et al. 1994] Lindquist, Ken, Donald E. Kline, and Ray Lambert. *Radiation-induced changes in the physical properties of Boraflex, a neutron absorber material for nuclear applications.* Journal of Nuclear Materials: Vol. 217, p. 223-228; 1994.

[O'Donnell 1989] O'Donnell, James H. *Radiation Chemistry of Polymers*. Chapter 1 in The Effects of Radiation on High-Technology Polymers, Reichmanis et al., editors; ACS Symposium Series. American Chemical Society: Washington, D.C.; 1989.

[Tsao et al. 2006] Tsao, Cheng-Si, Hung-Fa Shyu, Hsin-Fa Fang, Ming-Churng Hsieh and Shih-Chung Cheng. *Instrumental Design and Verification of a Nondestructive Testing with Neutron Backscattering for Boron-based Material Characteristics*. Journal of Nuclear Science and Technology: Vol. 43, No. 12, pp. 1517–1521; 2006.