

DRAFT

**REVIEW AND CRITIQUE OF
PREVIOUS PROBABILISTIC ACCIDENT ASSESSMENTS FOR
THE SHOREHAM NUCLEAR POWER STATION**

**Report on Task 1 of the Project
"Consequence Assessment for Suffolk County Radiological Emergency Response Plan"**

VOLUME I: MAIN REPORT

**Prepared for
Radiological Emergency Response Plan
Steering Committee
Frank Jones, Chairman
Office of the County Executive
County of Suffolk
H. Lee Dennison Building
Hauppauge, New York 11788**

**Robert J. Budnitz
Principal Author**

**Peter R. Davis
Stan Fabic
Howard E. Lambert
Contributing Authors**

September 17, 1982

Future Resources Associates, Inc.

2000 Center Street Suite 418 Berkeley, CA 94704 415-526-5111

TABLE OF CONTENTS

- 1.0 Introduction
 - 1.1 Overall Conclusion
 - 1.2 Limitations
 - 1.3 Objective of the Report
 - 1.4 Organization of the Report
 - 1.5 Summary of SAI's Results

- 2.0 Are All Important Accident Sequences Considered?

- 3.0 Are the SAI Calculations of the Accident Probabilities Correct?
 - 3.1 Methodology Considerations
 - 3.2 Data Base Considerations
 - 3.3 Specific Accident Sequences -- An Eclectic Review
 - 3.4 Implications of the Review of Specific Accident Sequences

- 4.0 Are the Accident Phenomena Within the Plant Treated Correctly?
 - 4.1 Grouping of Accidents Into Classes
 - 4.2 Review of Phenomena After Meltdown
 - 4.3 Some Specific Issues Involving Meltdown and Related Phenomena
 - 4.4 Core Vulnerability vs. Core Melt: Containment Event Trees
 - 4.5 Implications of the Review of In-Plant Phenomena

TABLE OF CONTENTS

(continued)

5.0 Summary and Conclusions: Are the Probabilities and Magnitudes of the Calculated Potential Radiation Releases Correct?

6.0 Acknowledgments

7.0 References

Appendices

- A. "Scope of Work," excerpts from the Contract between Suffolk County and Future Resources Associates, Inc.
- B. "Review of Within-Containment Phenomena and Containment Challenges," report by Dynatrek, Inc. under subcontract to this project
- C. "A Selected Review of the PRA Prepared for the Shoreham Nuclear Power Station," report by P. R. Davis prepared as part of this project
- D. "Internal Flooding Analysis," prepared by H. E. Lambert as part of this project

1.0 Introduction

This report summarizes the work carried out by Future Resources Associates, Inc. (FRA) under contract to Suffolk County, New York on the project entitled "Consequence Assessment for Suffolk County Radiological Emergency Response Plan." The overall goal of the project is to provide Suffolk County with technical support in its development of an emergency response plan for the Shoreham Nuclear Power Station, in particular by providing technical input as to the probabilities, severities, and radiological dispersion characteristics of potential large accidents at Shoreham. Shoreham is a boiling water reactor (BWR 4-Mark II) in final stages of construction on the north shore of Long Island, New York, with an electrical gross power rating of 846 megawatts. The reactor was manufactured by the General Electric Company and the architect-engineering work has been done by Stone & Webster. The location of the reactor on Long Island and the site plan for the Shoreham facility itself are shown on the next two pages: these are reproduced directly from Figures 2-1 and 2-3 of the NRC's "Safety Evaluation Report" for Shoreham (Ref. 6).

This project is a joint one involving a single unified scope of work under two contracts, one with FRA and the other one with Finlayson & Associates of Cerritos, California; under the arrangement with Suffolk County, Fred C. Finlayson of Finlayson & Associates has assumed overall responsibility to

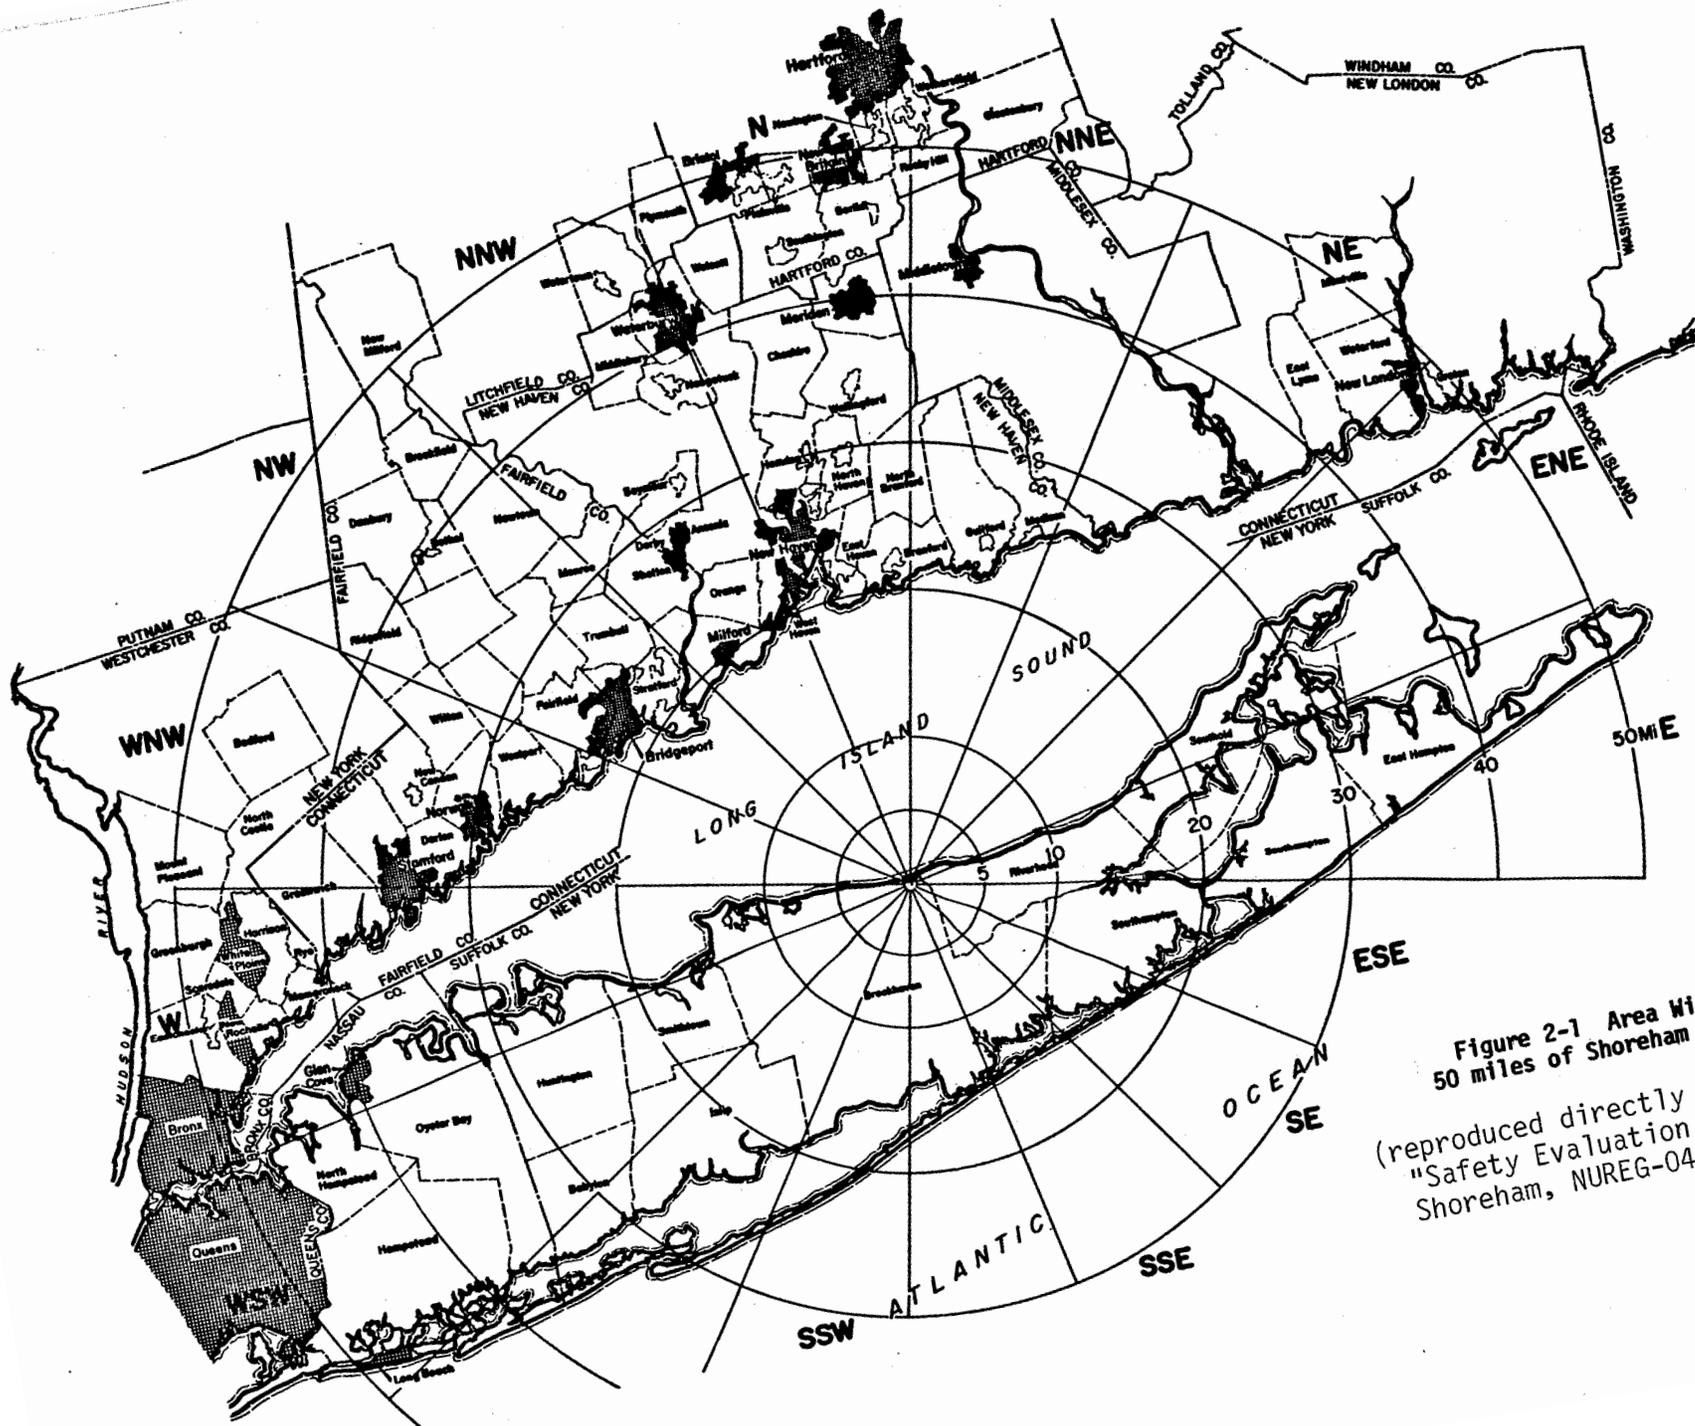


Figure 2-1 Area Within
50 miles of Shoreham Site
(reproduced directly from the NRC's
"Safety Evaluation Report" on
Shoreham, NUREG-0420)

(reproduced directly from NUREG-0420, the NRC's "Safety Evaluation Report" for the Shoreham plant)

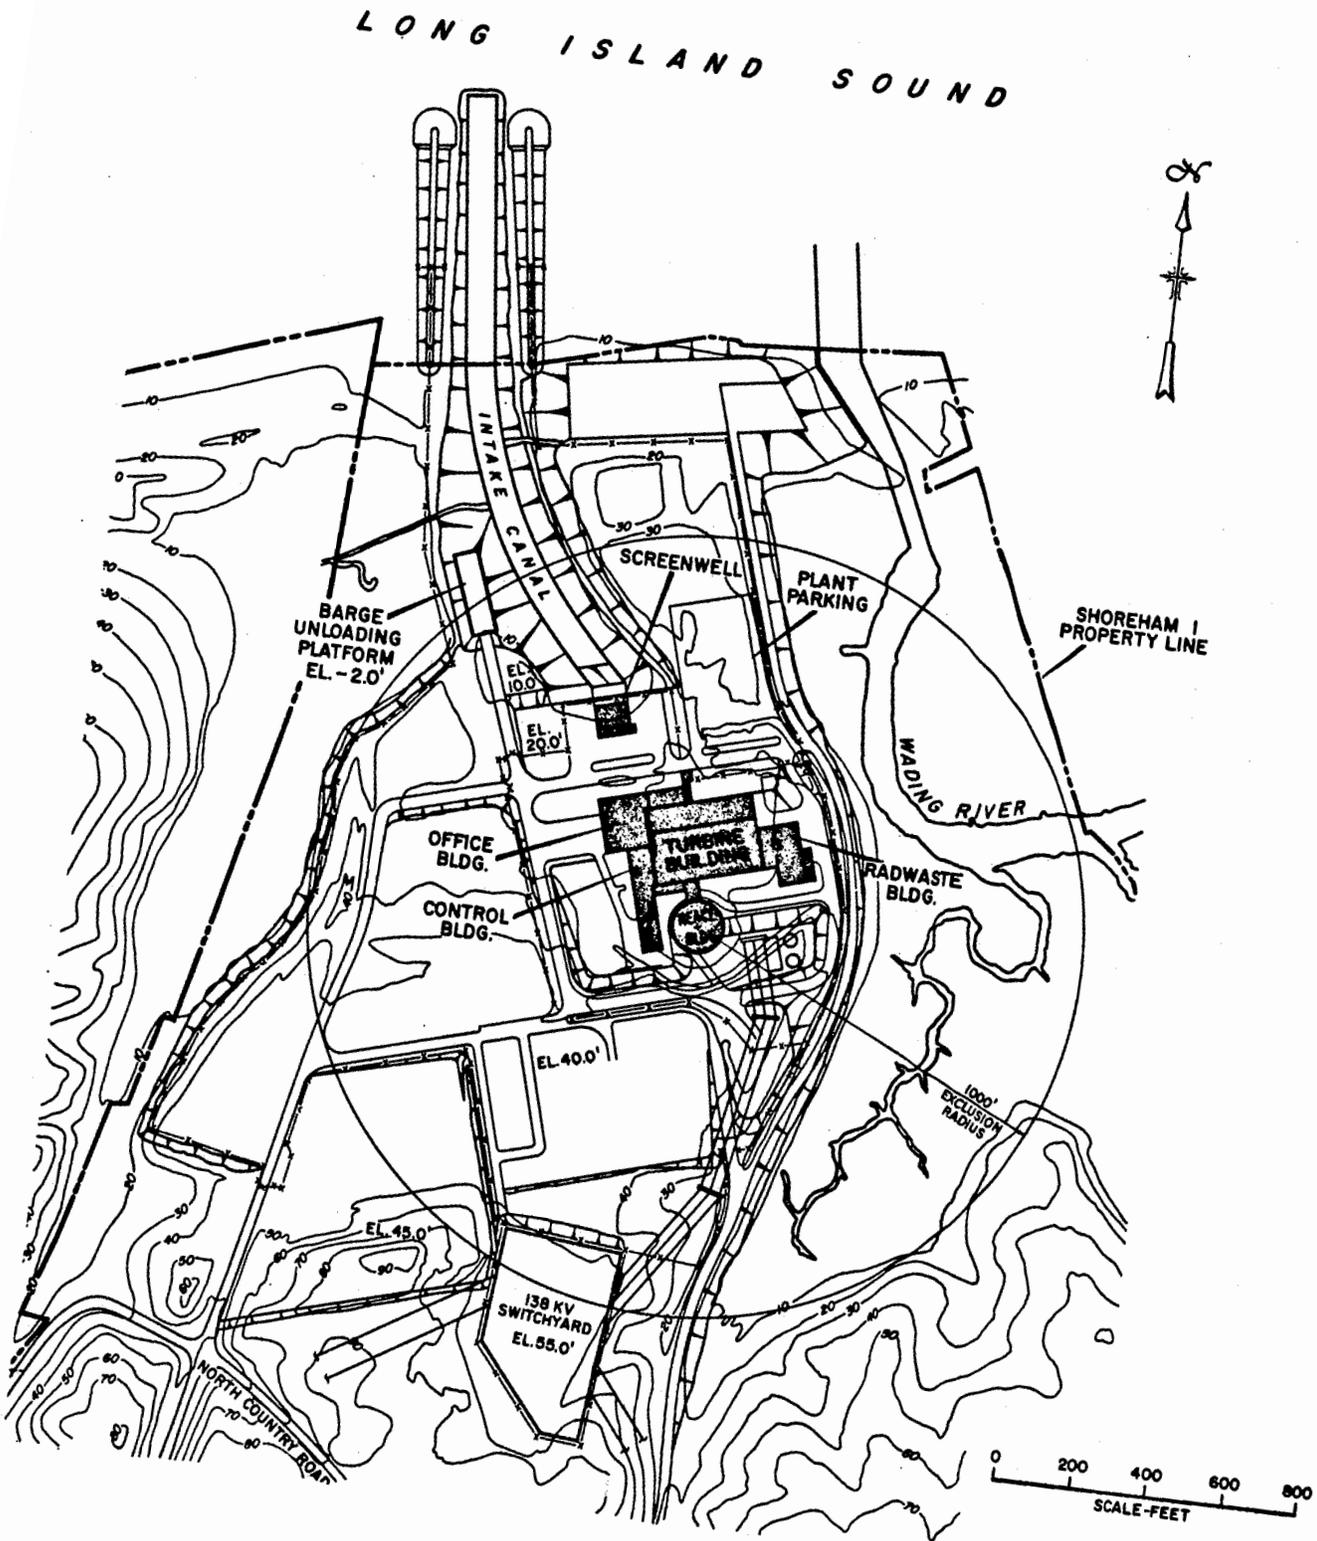


Figure 2-3 Plant Layout

coordinate the two contract efforts. The unified scope of work of the two contracts is reproduced in Appendix A of this report, and covers four tasks. FRA, under the technical direction of Robert J. Budnitz, has been principally responsible for the work under Task 1, while Finlayson & Associates is principally responsible for Tasks 2 and 3; in each task the other party has played a supporting role. Task 4, involving project management, integration, documentation, and technology transfer, has been a shared responsibility.

The contract was signed in late June 1982, and work under it has taken place predominantly in the months of June, July, and August, 1982 with this report due in draft form on September 15.

The review team made a site visit to the Shoreham plant in late June. Preliminary discussions of findings have taken place on an almost continuous basis, informally and verbally, between Dr. Budnitz and Dr. Finlayson in order to assure that both parts of the overall project are integrated effectively. Both parties agree that the integration has been successfully accomplished.

The FRA responsibility within the study has been predominantly to carry out Task 1, "Review and Critique of Previous Probabilistic Risk Analysis," and this report covers the work that FRA has accomplished in carrying out Task 1. Task 1 has consisted largely of a review of the preliminary draft report coordinate the two contract efforts. The unified scope of work of the two

entitled "Probabilistic Risk Assessment, Shoreham Nuclear Power Station," that was carried out by the San Jose, California office of Science Applications, Inc. (SAI) for Long Island Lighting Company (LILCO), the owner of the Shoreham facility. To carry out this review task, FRA has utilized a team of four individuals who possess expertise in various aspects of water reactor safety and in particular of probabilistic risk assessment (PRA). The FRA work has been led by Robert J. Budnitz, President of FRA, and has included Howard E. Lambert and Peter R. Davis, FRA consultants, and Stan Fabric of Dynatrek, Inc. (Rockville, MD), an FRA subcontractor.

The scope of the FRA work in Task 1 has been that of an independent review, which must be understood as quite different from an independent analysis of the potential accidents at the Shoreham plant. The purpose of the review has been to ascertain whether the PRA results obtained by SAI in their large and voluminous study are sufficiently reliable to form an acceptable basis for the County's emergency planning work. Because the scope of work has not included significant independent analysis, it is important to realize that its conclusion cannot be considered as a "stand-alone" conclusion. That is, it depends heavily upon the quality of the detailed work by SAI. This point was made in the original proposal to the County, where it was pointed out that "if major flaws are indentified in the earlier studies, . . . it may be necessary to devote a higher level of effort to this project." Thus the basis for the project has been an assumption that the SAI work under review is a credible effort, up to the standards of the state-of-the-art, and requiring no major upgrading. The

FRA team has attempted to challenge this assumption by carrying out a critical review with the intent to uncover facets of the SAI analysis that might contain inadequate methodology, inappropriate assumptions, errors of omission/ commission, or biases.

1.1 Overall Conclusion

The assumption that the SAI work is basically sound has turned out to be correct, in the opinion of Future Resources Associates, and we believe that the following conclusion is essentially the most important overall summary of the results of our study: FRA concludes that the overall results of the Shoreham PRA, contained in the preliminary draft report by Science Applications, represent reasonable conclusions as to the likelihood and magnitude of releases from large accidents at the Shoreham reactor.

As discussed in the text below, we believe that the likelihood of core melt accidents is somewhat higher, and the magnitude of radiation releases somewhat smaller than found by SAI. However, we believe that the differences should not cause Suffolk County to modify their emergency preparedness activities significantly compared to what they would do by using the SAI results directly as published. Subsidiary to this important conclusion are our conclusions that the methodology used in the SAI study is at the level of the state-of-the-art of reactor risk assessment at the present time, and that SAI's application of this methodology to Shoreham has generally been a competent one.

1.2 Limitations

There are limitations to our acceptance of the SAI study, and these are based on limitations within the study itself. The most important of these study limitations, from the point of view of its applicability to the County's emergency planning needs, is the absence of any analysis of accidents arising from the set of "external initiating events" that could initiate accidents from outside the plant. Most important of these external initiators is earthquakes, with high winds (hurricanes, tornadoes) and floods also of possible concern. In addition, internally-initiated fires are not treated. Because these types of accidents could be important risk contributors, their omission means that neither LILCO nor the County has a fully satisfactory set of important accidents to use as a basis for emergency planning.

There are other limitations to the SAI study, including the approximations, conservative in nature, as to the fission product source term; the way the very numerous accidents were grouped into classes for ease of treatment in the within-plant and containment phenomenological analysis; the absence of a plant-specific failure data base (since Shoreham has not yet operated); and other issues involving methodological approximations made by the study team. Fortunately, it is our view that none of these limitations, with the exception of the conservative approximations as to the fission product source term, is sufficiently important that it might significantly alter the substance of the County's emergency planning effort. In any event,

while we found a few places where the SAI report could be improved, we also believe that many of the report's limitations are not SAI's limitations per se but rather are limitations of the present state-of-the-art of PRA generally.

Our own study within Future Resources Associates, unfortunately, has had limitations of its own. The most important of these has been the relatively short time period (about 2-1/2 months) and relatively small level of effort that has been devoted to the project. Although the contract has been sizable from Suffolk County's perspective, the total level of professional effort devoted to reviewing SAI's draft PRA report has been only a few percent of the effort that SAI spent on carrying out the PRA itself and, of course, such a few percent effort cannot reasonably be expected to study every facet of the problem. This limitation is, however, balanced in part by the fact that the review team has had considerable experience in the PRA field, which has enabled the effort to be focused on what are thought to be the main issues.

Another limitation that hampered our group's work during the first month, but was cleared up in mid-July, was a restriction on us that effectively barred direct contact with the SAI PRA study team. After it was lifted, our interaction with the SAI team was a full and open technical exchange of questions and issues, at the level of professional mutual respect that we found refreshingly matter-of-fact. We wish to acknowledge SAI's full cooperation, for which we are grateful.

There are substantial uncertainties associated with the numerical conclusions that SAI quotes in its report. The uncertainties arise from several sources, including the validity of the data base, approximations in the accident sequence fault-tree/event-tree modeling, uncertainties and gaps in our understanding of within-plant accident phenomena, and incomplete understanding of the role of human error and human ingenuity in reactor accidents. Our general view is that the treatment of uncertainty within SAI's study is a reasonable one, including as it does advanced methods for estimating contributions from various sources. We have concentrated on those uncertainties that could particularly affect what Suffolk County might do or decide in the context of its emergency response plan. The discussion in the main body of this report will be in that context.

1.3 Objective of the Report

It is important to state clearly the objective of FRA's review work, which has been to provide Suffolk County with an independent technical opinion as to the probabilities and magnitudes of large potential accidents at Shoreham. FRA's work under Task 1 of its contract has concentrated on ascertaining, through independent technical review, whether the preliminary draft version of the LILCO- supported probabilistic risk assessment carried out by SAI provides a technically sound basis for emergency response planning.

1.4 Organization of the Report

The body of the report to follow will be organized generally along the lines of the scope of work (see Appendix A) for Task 1 of the overall project.

The technical issues that we have covered within the SAI report can be conveniently separated into the following:

- Are all important potential accident sequences considered?
- Are the calculations of the accident probabilities correct?
- Are the accident phenomena within the plant treated correctly?
- Are the magnitudes of the calculated potential radiation releases correct?

The treatment of environmental transport of radioactive materials after release, and their impact on populations, is the subject of Tasks 2 and 3 of this study and is not covered in this report.

1.5 Summary of SAI's Results

Before continuing with the results of FRA's review, it is useful to reproduce here the main results of the SAI study, in tabular and figure form. These, on the following pages, are reproduced directly from the SAI

draft report. The first page reproduces SAI's Table 4.1 and Figure 4.1, in which the frequencies of core vulnerable conditions are shown. Of particular interest are the uncertainty ranges of the SAI results. The second page reproduces SAI's Table 4.2, which contains the detailed numerical results for SAI's five accident classes. Core vulnerable frequencies are shown along with contingent probability of core melt, and the characteristics of the releases are also shown. The details of Table 4.2 will be discussed below in the body of our review.

(reproduced directly from SAI's report on Shoreham)

DRAFT - PRELIMINARY

Table 4.1

**SUMMARY OF THE FREQUENCIES OF CORE VULNERABLE
CONDITIONS BY ACCIDENT CLASS**

GENERALIZED CLASS CLASS	CLASS	FREQUENCY OF CORE VULNERABLE (PER REACTOR YEAR)
Loss of Coolant Makeup	I	2.7E-5
Loss of Containment Heat Removal	II	1.1E-5
LOCA	III	3.6E-7
ATWS W/O Poison Injection	IV	6.1E-6
LOCA Outside Containment	V	2.0E-8
Total Core Vulnerable Frequency (Per Rx Yr)		4.4E-5

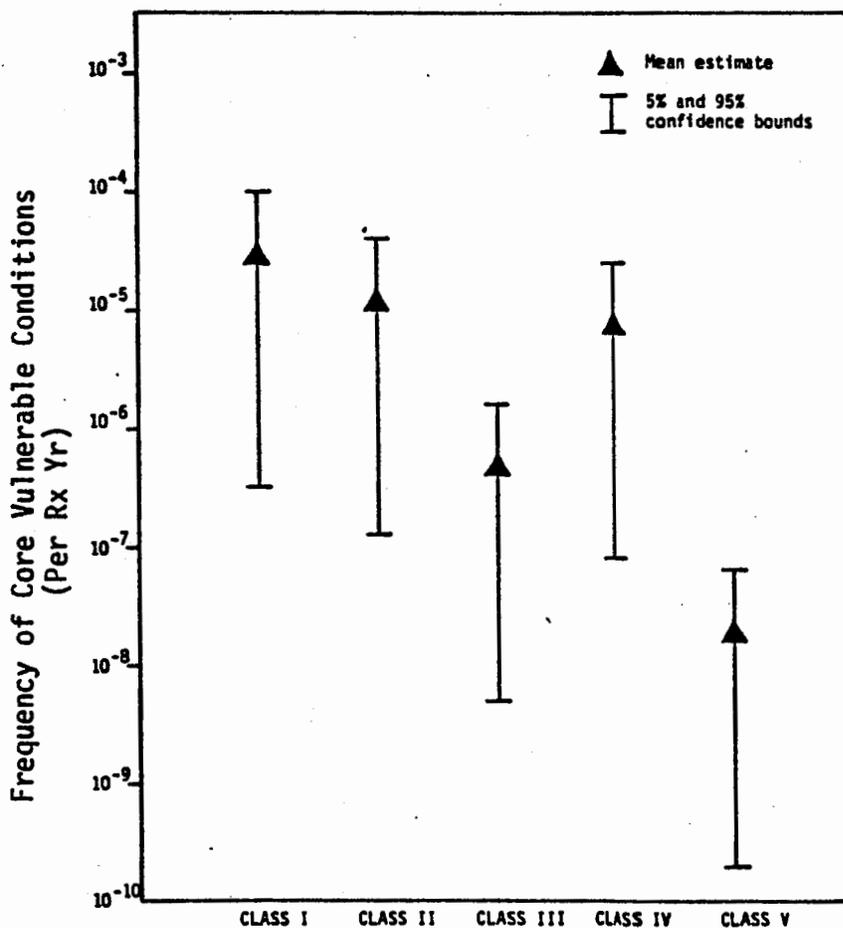


Figure 4.1 Summary of Core Vulnerable Frequencies, Including the Uncertainty Characterization

(reproduced directly from the SAI report)

DRAFT - PRELIMINARY

Table 4.2
SUMMARY OF SNPS PRA RESULTS, OVERPRESSURE EVENTS

ACCIDENT CLASS	(a) CORE VULNERABLE FREQUENCY	(b) CONDITIONAL FREQUENCY OF RELEASE	(c) RELEASE BOUNDS	(d) ACCIDENT RELEASE DESIGNATOR	(e) FRACTION	(f) TIME OF RELEASE	(g) DURATION OF RELEASE	(h) WARNING TIME FOR EVACUATION	(i) ELEV OF RELEASE	(j) AVG ENERGY OF RELEASE	RADIONUCLIDE RELEASE FRACTIONS								
											Kr-86	01 ⁻	I ⁽¹⁾	I ⁻	Cs ^(m)	Te ⁽ⁿ⁾	Sr ^(o)	Ru ^(p)	La ^(q)
1	2.7x10 ⁻⁵	8.0x10 ⁻²	UPPER	C1R,T _{1Y} ⁻	.33	2	6 ^(v)	1.5	60	10	.99	3.0x10 ⁻⁴	2.4x10 ⁻²	2.1x10 ⁻⁴	3.9x10 ⁻²	1.7x10 ⁻¹	3.9x10 ⁻³	2.2x10 ⁻²	3.6x10 ⁻³
			BASE CASE	C1R,T _{1Y} ⁻	.34	8	1	7	60	6	.98	2.97x10 ⁻³	1.8x10 ⁻²	1.5x10 ⁻⁵	5.1x10 ⁻²	2.1x10 ⁻¹	5.1x10 ⁻³	2.5x10 ⁻²	5.1x10 ⁻³
			LOWER	C1R,T _{1C} ⁻	.33	8	1	7	75	.3 ^(v)	.15	9.6x10 ⁻⁶	2.0x10 ⁻⁶	1.4x10 ⁻⁷	5.4x10 ⁻⁶	2.4x10 ⁻⁵	5.1x10 ⁻⁷	3.0x10 ⁻⁶	5.1x10 ⁻⁷
2	1.1x10 ⁻⁵	7.0x10 ⁻²	UPPER	C2R,VSFT _{1Y} ⁻	.1	25	20	2	60	60	.98	6.4x10 ⁻⁴	1.15x10 ⁻²	3.1x10 ⁻²	8.2x10 ⁻²	1.5x10 ⁻¹	3.0x10 ⁻³	1.5x10 ⁻²	2.3x10 ⁻³
			BASE CASE	C2R,T _{1Y} ⁻	.8	30	20	2	60	50	.98	6.4x10 ⁻⁴	1.14x10 ⁻²	3.1x10 ⁻²	8.1x10 ⁻²	1.5x10 ⁻¹	2.9x10 ⁻³	1.5x10 ⁻²	2.2x10 ⁻³
			LOWER	C2R,T _{1Y} ⁻	.1	30	30	2	60	.3 ^(v)	.96	6.4x10 ⁻⁴	3.1x10 ⁻³	1.4x10 ⁻²	3.1x10 ⁻²	3.5x10 ⁻²	6.4x10 ⁻⁴	2.5x10 ⁻³	3.2x10 ⁻⁴
3	3.6x10 ⁻⁷	8.4x10 ⁻¹	UPPER	C3R,T _{1Y} ⁻	.33	2	6 ^(v)	1.5	60	10	.99	6.3x10 ⁻³	8.0x10 ⁻³	4.7x10 ⁻³	4.4x10 ⁻²	1.7x10 ⁻¹	5.0x10 ⁻³	2.0x10 ⁻²	3.9x10 ⁻³
			BASE CASE	C3R,T _{1Y} ⁻	.34	8	1	7	60	6	.99	6.3x10 ⁻³	1.1x10 ⁻²	3.5x10 ⁻⁶	4.6x10 ⁻²	1.7x10 ⁻¹	4.2x10 ⁻³	2.1x10 ⁻²	6.1x10 ⁻³
			LOWER	C3R,T _{1C} ⁻	.33	8	1	7	75	.3 ^(v)	.49	3.1x10 ⁻³	1.5x10 ⁻³	2.9x10 ⁻⁵	1.0x10 ⁻²	4.3x10 ⁻²	1.1x10 ⁻³	5.2x10 ⁻³	1.0x10 ⁻³
4	6.1x10 ⁻⁶	4.3x10 ⁻¹	UPPER	C4R,VSFT _{1Y} ⁻	.1	1.5	10 ^(u)	1	60	60	.99	6.4x10 ⁻⁴	2.3x10 ⁻²	3.5x10 ⁻²	1.0x10 ⁻¹	2.1x10 ⁻¹	4.1x10 ⁻³	2.2x10 ⁻²	3.4x10 ⁻³
			BASE CASE	C4R,T _{1Y} ⁻	.8	1.5	10 ^(u)	1	60	60	.99	6.4x10 ⁻⁴	2.3x10 ⁻²	3.5x10 ⁻²	1.0x10 ⁻¹	2.1x10 ⁻¹	4.1x10 ⁻³	2.2x10 ⁻²	3.4x10 ⁻³
			LOWER	C4R,T _{1C} ⁻	.1	1.5	10 ^(u)	1	60	.3 ^(v)	.98	6.4x10 ⁻⁴	7.4x10 ⁻³	1.3x10 ⁻²	2.9x10 ⁻²	4.2x10 ⁻²	6.9x10 ⁻⁴	4.0x10 ⁻³	4.1x10 ⁻⁴
5	2.0x10 ⁻⁸	x1.0	UPPER	CSR _{1C} ^(r)	.1	1	1	0	60	6	.99	6.3x10 ⁻³	6.6x10 ⁻²	6.5x10 ⁻¹	6.9x10 ⁻¹	5.9x10 ⁻¹	1.8x10 ⁻¹	3.9x10 ⁻²	6.1x10 ⁻³
			BASE CASE	CSR _{1C} ^(s)	.45	1	1	0	50	6	.99	5.1x10 ⁻³	4.6x10 ⁻¹	6.6x10 ⁻²	6.9x10 ⁻¹	5.9x10 ⁻¹	1.8x10 ⁻¹	3.9x10 ⁻²	6.1x10 ⁻³
			LOWER	CSR _{1C} ^(t)	.45	1	2.5	0	75	.3 ^(v)	.99	2.5x10 ⁻³	3.5x10 ⁻¹	4.0x10 ⁻²	4.3x10 ⁻¹	3.0x10 ⁻¹	1.0x10 ⁻¹	1.2x10 ⁻²	2.0x10 ⁻³

FOOTNOTES

- (a) Obtained from functional event tree quantification.
- (b) Given that the core has become vulnerable, this is the conditional probability that a substantial release occurs.
- (c) Releases are treated as distributed with simple discrete model (see uncertainties). Values are shown for the best estimate and two bounding cases.
- (d) See Appendix M for discussion of nomenclature.
- (e) Treated as a discrete distribution with three values this is the conditional probability of either the point estimate or the bounds.
- (f) in hours
- (g) in hours
- (h) in hours
- (i) in meters
- (j) in 10⁶ BTU/hr
- (k) Background on the isotope groups and release mechanisms is presented in WASH-1400, Appendix VII.
- (l) Includes I and Br.
- (m) Includes Cs and Rb
- (n) Includes Te and Sb
- (o) Includes Ba and Sr
- (p) Includes Ru, Rh, Co, Mo, Tc
- (q) Includes Y, La, Zr, Nb, Ce, Pr, Nd, Hf, Pb, Am, Cm
- (r) Case where significant deflagration in secondary containment causing a reducing atmosphere to exist; difference is in chemical form of iodine.
- (s) Case where no significant deflagration in secondary containment has occurred.
- (t) Extrapolated from data for release from containment terminated after .3 to .6 hr.
- (u) Release characteristics are such that an initial puff occurs during the first hour followed by another puff towards the end of the release duration.
- (v) Low energy represents slower blowdown after a smaller containment breach.

2.0 Are All Important Accident Sequences Considered?

The answer to this question, as mentioned above in the introduction to this report, is negative. In particular, there has been no treatment of internally-initiated fires, nor of any externally-initiated events, the most important of which are earthquakes, floods, and high winds (hurricanes, tornadoes). These omissions are discussed in the SAI report (P. 1-13), and were beyond the scope of that study. The implications of these omissions will be discussed below, but first we will consider whether any important internally-initiated accidents (besides fires) seem to have been omitted.

On this latter point, we have concluded that the SAI draft report has apparently considered all of what the reactor safety community considers the important internally-initiated accidents. Specifically, we have not found any internally-initiated sequences likely to contribute significantly to the overall risks that have been omitted.

Of course, the unconscious omission of important sequences known to others in the reactor safety community would be quite unlikely in a study of this kind . . . the safety community maintains close enough and open enough communications that any new or unusual accident sequences would almost surely have come to the attention of the study team or its outside consultants. So our conclusion comes as no surprise.

This observation of the apparent completeness within SAI's draft report does not mean that there are no important sequences omitted. It only

means that we are unaware of any, nor do we believe that anyone else in the reactor safety community has thought of any. The strength of our convictions about completeness is based on the general observations that, over nearly a decade of time since the initial probabilistic accident delineation work of WASH-1400 (Ref. 1), there has been hardly any addition to the list of important internally-initiated accidents that WASH-1400 considered. But, on the other hand, who knows whether or when a new sequence will arise, either from operating experience or from analysis?

The grouping into sequences that the SAI team used is discussed in detail in Section 4.1 later in this report. The five classes are described in SAI's Table 3.3.1, which is reproduced on the next page.

It is important to recognize in this connection, however, that the approximately 1600 reactor years of commercial nuclear power operating experience worldwide without a serious accident leading to off-site consequences is a statistically significant data base providing evidence that the result for Shoreham is unlikely to underestimate the probability of serious accidents by a large factor, unless Shoreham is somehow very untypical in its risk profile of the entire group of reactors; of course, it is just this question that is being addressed by the Shoreham-specific risk assessment.

Returning to the known "external" omissions (fires, earthquakes, high winds, floods), the SAI team acknowledged in their report (P. 1-13) that

(reproduced directly from the Shoreham PRA by SAI)

Table 3.3.1
GENERIC ACCIDENT SEQUENCE CLASSES

GENERIC ACCIDENT SEQUENCE DESIGNATOR	PHYSICAL BASIS FOR CLASSIFICATION	SYSTEM LEVEL CONTRIBUTING EVENT SEQUENCE	REPRESENTATIVE SEQUENCE FOR CLASS
Class I (C1)	Relatively fast core melt; containment intact at core melt and at low pressure	Transients involving loss of inventory makeup; small-small LOCA events involving loss of inventory makeup; transients involving loss of scram function and inability to provide sufficient coolant makeup	Transient with loss of high and low pressure coolant makeup
Class II (C2)	Relatively slow core melt due to lower decay heat power; containment failed prior to core melt	Transients or LOCAs involving loss of containment heat removal; inadvertent SRV opening accidents with inadequate heat removal capability	Transient with loss of residual heat removal
Class III (C3)	Relatively fast core melt; containment intact at core melt, but at high internal pressure	Large LOCAs with insufficient coolant makeup; transients with loss of heat removal and long-term loss of inventory makeup; RPV failures with insufficient coolant makeup	Large LOCA with loss of low pressure ECCS
Class IV (C4)	Relatively fast core melt; containment fails prior to core melt due to overpressure	Transients involving loss of scram function and loss of containment heat removal or all reactivity control; transients with loss of scram function followed by actuated depressurization	Transient with failure of RPS and failure of SLCS
Class V (C5)	Relatively fast core melt; containment failed from initiation of accident due to equipment failures	LOCAs outside containment with insufficient coolant makeup to core; RPV failures which result in immediate containment failure	LOCA in main steam lines with failure of MSIV closure and loss of ECCS

these were consciously and specifically excluded. The reasons for the omission are probably a combination of two things: first is the fact that treating external events and fires in a PRA is considerably different than treating internal initiators, requiring a different methodology, large additional manpower resources different in character from the rest of the study, and yielding results of even greater uncertainty than the uncertainty in the rest of the PRA analysis. These large uncertainties are because the methodologies are immature and the data base weak. Second, the SAI study team believed when the project began in 1980 that these external events and fires were not as important contributors to risk as internally-initiated accidents, nor as amenable to cost-beneficial risk reduction (Ref. 2).

In the intervening two years, our ability to analyze the risk contribution from fires and externally-initiated events has advanced considerably. Benchmarks of this advance include the recent publication by Pickard, Lowe, & Garrick of seismic and fire PRAs at both the Zion and Indian Point reactor stations (each a two-unit station with Westinghouse PWR's); the completion of an important NRC-sponsored seismic methodology development effort, the Seismic Safety Margin Research Program at Lawrence Livermore National Laboratory; and the inclusion by consensus within the draft PRA Procedures Guide (Ref. 3) of an "acceptable" methodology for PRA analysis of earthquake- and fire-initiated accidents. With these methodological advances have come the first probability-based insights into the quantitative contribution to risk from these sources, albeit with very

large uncertainties and important conservatisms in the analyses. To the surprise of some in the reactor community, the Zion and Indian Point studies have told us that neither earthquakes nor fires can be neglected as contributors at those power stations to overall public risk. Their apparent preeminence at Zion and Indian Point is partially because internally-initiated accidents were found to be much less important than was found in WASH-1400.

Fortunately for the purposes of this report, one member of the PRA review team (Budnitz) has recently been reviewing the Indian Point PRA specifically from the perspective of the risk posed by earthquakes and high winds (but not fires). His basic conclusion vis-a-vis earthquake-initiated and wind-initiated accidents at Indian Point is that the methodology is clearly adequate to tell us what types of accident sequences are likely to be of most concern, and whether these initiators pose important safety problems (at Indian Point, they do). But the methodologies do not seem to be mature enough to provide reliable quantitative calculations of the probability of core melt . . . thus any numerical comparison of internally-initiated core melt frequencies or public risks with those from earthquakes and winds is of little value.

An important observation from the Indian Point study is that accidents initiated by earthquakes and winds seem to involve phenomena quite similar to those involved in the ensemble of internally-initiated accidents: that is, the accident phenomena themselves do not seem to form a different set

of phenomena that must be considered separately in the sense of comprising different types of releases. If this observation is generally valid, then the main impact on overall PRA results would be to increase the probability of releases already treated in the analysis, with less impact on the spectrum of releases and consequences.

What insights about Shoreham can be obtained by transfer to Shoreham of our recent increased understanding of externally-initiated accidents at other plants? Unfortunately, not much. This is in part because there is not yet available any external-events PRA analysis for any BWR reactor (the PWR's analyzed to date are in detail not remotely similar to the Shoreham design), and in part because the accident-initiating events found to be important at Zion and Indian Point are quite site- and design-specific in detail, involving features that are unlikely to be reflected at another plant.

Much thought has been given by the PRA review team as to why useful insights applicable to Shoreham cannot be reliably gained from studying external-events PRA analyses at other reactors. Our negative conclusion arises basically from our belief that if earthquakes, high winds, or fires give rise to important accident sequences at Shoreham, the sequences themselves are likely to be idiosyncratic to Shoreham, or in some cases possibly generic to BWR's (or BWR Mark II reactors) as a class. Absent any specific analysis, we conclude that the contribution of these sources to residual public risk from Shoreham is simply not quantitatively known, in terms of either probability or character of consequences.

The insight (at the PWR's studied elsewhere) that the accident phenomena are not qualitatively different in kind from those arising from internal sources is, of course, a reasonable one consistent with the intuition of most students of the problem. If this insight were to hold at Shoreham, then emergency response plans based on accident scenarios from internal initiators would likely afford reasonable protection from these other types of accidents as well, provided that the special external circumstances surrounding a large earthquake or hurricane are adequately included in any response plans, but the "if" in this proposition could be a weak reed.

3.0 Are the SAI Calculations of the Accident Probabilities Correct?

To answer this question, the FRA team reviewed the methodology used by the SAI analysts, considered the validity of the numerical data base, and repeated selected calculations to ascertain how sensitive some of the results were. Because the numerical quantification of literally hundreds of sequences, within many different event trees, was beyond our capability, we are not able to affirm the specific validity of each SAI accident sequence. However, we believe that such a detailed review has not been necessary to satisfy Suffolk County's objective.

3.1 Methodology Considerations

The methodology of the SAI analysis includes numerous advances over the WASH-1400 analysis of Peach Bottom in 1973-1974. We concur in the judgments of the SAI team that the use of these advances improves the analysis. An example of the different approach taken at Shoreham is the incorporation of certain important support systems (such as instrument air, AC power, DC battery power) consistently within the fault trees rather than having some of these in the event trees, as in WASH- 1400. This approach allows for easier analysis, but carries with it an increased risk that common dependencies might be missed through oversight if the analyst gets sloppy or forgetful. We believe that SAI's approach is a valid one, which can produce valid results if executed with care.

Another advance is the differentiation in the Shoreham study between accidents leading to a "core-vulnerable" condition and those that proceed further to "core melting." In the WASH-1400 approach, any accident proceeding as far as what SAI here calls a "core-vulnerable" condition was assumed to proceed all the way to "core melt"; that is, there was no differentiation in WASH-1400 between the two conditions and the more conservative approach was taken. The reason for this was that in 1973-74 there was no existing methodology for consistently calculating the differentiation between these two conditions, and the WASH-1400 study team was not able to develop one within the resources and technical knowledge of the time. Now, eight years later, this differentiation is feasible, and the SAI analysts have developed it further and applied it in the Shoreham analysis.

In our opinion, this differentiation is an advance: it was understood in the WASH-1400 period that the conservative core melt assumption was not correct, and the Three Mile Island accident told us this as well. (The TMI accident, if treated with the WASH-1400 methodology, is called a full "core melt," although it was certainly not.) For purposes of comparison, the probability of "core-vulnerable" conditions at Shoreham is what should be compared to "core melt" in WASH-1400. The numerical results of the Shoreham analysis show that the fraction of "core-vulnerable" sequences that proceed to "core melt" ranges from several percent (for what are called Class I and II events) to 100% (for the Class V category). In Section 4.4 below, we will discuss in detail the question of whether SAI's quantitative conclusions in this area are valid.

Still another methodological advance is in the way uncertainties are estimated. Since the WASH-1400 team pioneered with their analysis, a variety of different methods have arisen for getting a numerical handle on these uncertainties. The approach taken by SAI in the Shoreham study is to use Monte Carlo methods to model how much difference in the final results would arise from various changes at the front end, or in the data base. This seems to be a reasonable approach, and its conclusions are also quite reasonable (see Table 3.8.1 on Page 3-172 of the SAI main report), but we did not review it in detail because of our judgment that the conclusions were reasonable.

Of course, there are still methodological problems that temper the confidence we might have as to the validity of PRA results for frequency of, say, core-vulnerable or core melt states. Among these are a variety of issues in the arena of human factors, and these methodological issues stand apart from issues of human error quantification, about which we will say more below as part of our discussion of the error data base.

A key example of methodological inadequacies is the failure of PRA methods generally to consider adequately the ability of reactor operators to cope, through improvisation and ingenuity outside of standard procedures, with accident sequences as they develop in real time. The coping must surely allow operators to terminate some sequences successfully that otherwise would develop further into high-risk accidents. Yet we do not know how much difference this inadequacy makes to the final results. On the other

side of the same coin is the possibility that even well-trained operators, being fallible humans like the rest of us, might significantly aggravate a sequence that in the standard PRA analysis is terminated successfully without damage.

Another methodological limitation in a related arena is the present inability to model control system failures well enough. The event-tree/fault-tree approach intrinsically views accidents through the concept of system functional failures caused by underlying component and support system failures or unavailabilities. In its present state of maturity, the methodology can only incorporate control system failures through ad hoc analysis of when, and how, multiple component system failures might arise from control system failures. While this ad hoc analysis is probably adequate for most sequences, we do not know whether it does acceptably well when "time is of the essence," that is when rapidly developing events, especially in the early stages of some accident classes, could severely tax the operators' capabilities or the resilience of some components.

Time sequencing issues also underlie uncertainties in event tree delineation (and, to a lesser degree, fault-tree delineation). The event trees are written down in a time-ordered fashion, but issues can arise of "which failure occurs first" and of whether recovery might occur later in a sequence. Again, whether this set of issues makes a significant difference to the final results is not known, although our experience leads us to expect that differences are unlikely to be large.

As described in the NRC's PRA Procedures Guide, quality assurance is very important in the generation and analysis of fault trees. For example, if the same event appearing in two or more places in a fault tree were mis-typed when entered into a computer, one can make an error as large as a factor of 100 (usually non-conservatively) in computing accident probabilities. (Such numerical errors would probably show up, of course, in the final analysis.) It is important to note that we did not check SAI's fault trees, nor did we run an independent computer analysis of their fault trees.

None of these methodological limitations is special to the Shoreham PRA analysis by SAI: all are generally shared by other PRA studies on other reactors. Indeed, progress is gradually improving our confidence that these limitations are not important enough to invalidate our confidence in the conclusions of PRA . . . but they do temper our confidence, in an unquantifiable way.

The tenor of the above discussion reveals our overall conclusion about the methodology used by SAI in the Shoreham analysis. That conclusion is that the SAI methodology, though it suffers from some generic limitations common to all reactor safety PRA's to date, nevertheless represents the present state-of-the-art; it includes important methodological advances in some areas. We believe that it is an acceptable basis for estimating the probabilities of the important internally-initiated accident sequences at Shoreham.

3.2 Data Base Considerations

The validity of the failure data base is vital to the validity of any PRA analysis. A couple of issues in our present case represent limitations that might be important. First and foremost, the Shoreham reactor is still in final stages of construction, so we have no data on failures at Shoreham: the failure data must generally come from industry-wide sources. But operating experience has told us that some failure and unavailability rates can vary widely from plant to plant, and we don't know whether Shoreham will be above or below average.

The other side of this coin is that significant advances have been made recently in identifying below-average design, maintenance, and test practices, by study of LERs and plant-specific attention to determining root causes of system and component failures. To the extent that these activities represent improvements, a new plant such as Shoreham can take advantage of this experience to improve its performance over the average performance of plants already running.

Human error quantification is the other arena where data base issues seem to represent a limitation. The SAI study team has used the accepted industry-wide reference for most of its human factors failure data (Ref. 4), but these are widely understood to contain large uncertainties. In particular, the failure data are lumped into broad categories whose applicability to specific sequence situations at Shoreham must surely be only approximate. Also, how Shoreham's operators will behave, compared to industry averages, is of course not known.

Despite the limitations just discussed, it is our conclusion that the Shoreham PRA analysis under review has used state-of-the art data bases generally. We have found some specific cases (see below) where we do not agree with certain specific numerical values, but we conclude that the data used are a generally acceptable basis for estimating accident probabilities at Shoreham.

3.3 Specific Accident Sequences -- An Eclectic Critique

To review quantitatively all of the important accident sequences that SAI quantified in its Shoreham analysis was not possible (and in our opinion not necessary) within the scope of this report. We took the approach of studying what SAI wrote down in their event-tree delineation, and using our experience and judgment to ascertain whether the approach and results were roughly congruent with our expectations. This is what characterized our accident-sequence review.

This activity consumed a reasonable fraction of all the effort spent on this project. In the course of it we found that most of what we studied was unarguably consistent with our experience and understanding, while there were only a few cases to the contrary (albeit important cases). For example, we have already mentioned above that we examined the list of internal accident initiators to ascertain whether SAI had left out any from the list that we would have used, and we indicated that they had not

(except fires). We also attempted to look at some of the numerical values that SAI used in order to see whether they agreed with our collective experience. In the course of this review, we found ourselves comfortable with almost all of the numbers we examined, the exceptions being noted below.

Perhaps the most important part of the PRA analysis where we find that our review does not agree with the SAI draft report concerns their analysis of internal flooding (their section 3.4.4.1 and Appendix G). In particular, consider the situation in which portions of crucial safety systems are disassembled for routine maintenance during plant operation. If during this disassembly the valved-off component is accidentally reconnected to its pipe (such as accidentally opening a motor-operated valve that has been closed to allow the maintenance), then release of water through the opened valve will occur. If the mistake is not promptly corrected by re-closing the opened valve, an internal flood can result; such a flood at "Level 8" inside the reactor building can quickly inundate several pieces of critical safety equipment, and its analysis as a special issue was deemed so important that an entire Appendix G is devoted to it by SAI.

A detailed numerical analysis of this issue, by H. E. Lambert of our review team, is presented in Appendix D to this report. We will briefly discuss here the main issues and conclusion. We believe that SAI's analysis contains some errors that underestimate the internal flood frequency, the correction of which would raise the core-vulnerable frequencies for Classes I and II.

The events leading up to a disabling flood are as follows:

- Event A: On-line maintenance of some critical system
- Event B: During maintenance, the system is disassembled
- Event C: Inadvertent opening of isolation valve, causing the flood
- Event D: Failure to reclose valve within specified time period
- Event E: Operator erroneously isolates the power conversion system during flooding causing the accident because the heat sink is lost

We believe that the initiating event is Event C: the occurrence together of Events A and B defines a vulnerable system state that permits C to initiate the flood. Because Event C is an initiating event, we must compute its frequency, and also the failure on demand of Event D. Event C's frequency (in units of events per year, or the like) must include the pre-existing presence of the vulnerability-inducing states A and B. The units of the events, in the calculation, should be maintenance acts per year (Event A); probability of system disassembly given maintenance (Event B); conditional probability of inadvertent isolation valve opening given maintenance with disassembly (Event C); conditional probability of failure on demand to reclose the opened valve (Event D); and conditional probability of erroneous operator isolation of the power conversion system (Event E), which then would initiate the accident. We believe that SAI has incorrectly used system unavailability for Event A, and that their calculated result is about a factor of 100 lower than the value we obtain, all of which is explained in more detail in Appendix D to this report.

In addition, we believe that the SAI team has used an incorrect value for the Event D failure (probability to reclose the accidentally open valve): SAI uses 0.05 for this probability whereas their primary reference source (Ref. 4) uses 0.25, if one accepts that there will be highly stressful conditions during the period when the operator action will be required (see Appendix D for details).

An important assumption made by SAI is that flooding to the six-foot level will not result in automatic closure of the MSIV's. (SAI does assume, however, that reactor trip will occur.) It is important to verify that the assumption regarding automatic MSIV closure is true. Otherwise the power conversion system is lost and the only normally available coolant makeup system is the condensate system. In this case the accident frequency caused by flooding would increase by an additional factor of about 10, and a design change might be necessary to overcome the problem just discussed.

If our flooding analysis is correct, then the internal flooding accident frequencies are at least of the order of a few times 10^{-5} per reactor year, and as described in our Appendix D could be much higher, depending upon how human error rates are quantified. These accident sequences then become dominant for both Class I and Class II type accidents.

There is another part of the SAI analysis where our review team disagrees with the SAI work. Our difference of opinion enters critically into a key class of accident sequences, the so-called ATWS group (Anticipated

Transients Without SCRAM), which comprise Class IV of the five classes into which SAI grouped the Shoreham accidents. These sequences arise when one of several anticipated transients cannot be properly controlled because the "SCRAM" system (which inserts the control rods) fails to shut down the chain reaction. If this event were to occur, back-up engineered functions are brought into action to bring the reactor to safe shutdown: these are discussed on pages 3-94 ff. of the SAI report and include the alternate rod insertion system; the standby liquid control system to inject boron poison into the core; trip of the recirculation pumps; and operator procedures. Of course, the critical failure that drives this sequence is the failure to SCRAM on demand, for which there are essentially no empirical data upon which to base an analysis.

The SAI report selects as its value for failure to SCRAM on demand the likelihood 3×10^{-5} (about one failure in 33,000 demands). In its discussion (see p. B-111 ff), the report points out that "The calculation of Scram system reliability has been an issue which has taken on both technical and philosophical aspects over the last seven years." This is because within the regulatory arena the issue of whether reactors are adequately safe against ATWS events has been one of the most hotly-contested issues in recent years, whose regulatory resolution is still not complete. Unfortunately, there is also no consensus in the safety community about how to go about calculating the value for this failure nor about which value should be used in analyses such as the Shoreham case.

While admitting that the situation requires a good deal of judgment, we believe that the value chosen by SAI is too high, by a factor in the range of 3 to 10 . . . that is, we would have chosen a value of 1×10^{-5} to 3×10^{-6} instead of their 3×10^{-5} . This judgment, based upon nearly a decade of study of the issue, cannot be defended analytically; but we do believe that the 3×10^{-5} value, which SAI seems to state (p. B-115) is chosen with a conservative bias, is too high. (Of course, we are aware that certain industry calculations fall in the range of 10^{-6} , or sometimes much smaller.)

The impact of this judgment on SCRAM failure probability is practically linear in the results of core-vulnerable and core-melt frequency for Class IV. However, even though our best judgment would lower the likelihood of Class IV accidents, we believe that it is prudent for Suffolk County to base its emergency planning on the Class IV core melt values as quoted in the SAI draft report . . . this approach takes cognizance of the observation that our difference of opinion with SAI is a purely judgmental one.

There were other areas that we reviewed which are important to the SAI analysis. For example, as part of our review we studied the discussion on success criteria found on pp. 1-22 to 1-27 of the SAI draft report. It is clear that in any analysis of accident sequences, whether the outcome is "successful" (that is, leads to an adequately cooled core without vulnerability to core melt) depends crucially on whether the design

requires, say, all 4 low-pressure pumps or, say, only 1 out of 4 to cope with a large LOCA. The success criteria chosen by the SAI analysis team have been taken, according to their report, from manufacturer's (GE's) reports and analyses, and the report comments (page 1-22) that "it is believed that the success criteria so defined tend to be conservative." We were not able to assess this claim independently, although it makes some sense to us because the postulated conditions in the GE analysis are limited in a conservatively biased direction. We wish to call attention to this observation as a possible source of uncertainty in the SAI results, which if corrected would lower their calculated core melt values.

We also looked at the way human errors are incorporated into the component and system failures. Here our understanding is that the SAI team has basically relied on the standard data source, the work by Swain and Guttman (Ref.4), for the human failure data. Since this is the standard work and there is no better source at the present time, we have no quarrel with the SAI team's decision here, but we feel it important to point out that there are very large uncertainties in the values for human failures that Swain and Guttman quote. Even in the situations in which these values are applicable, there are large variations in the error rates from one human to another, but the stickier problem is that the actual environments in which the human errors will occurⁱⁿ reactor accidents are not those in which the error rates used are necessarily directly applicable. Of course, this is a problem that is generic to reactor PRAs as a class, and not special to the Shoreham effort under review here. But again we

wish to call attention to this issue as a source of possible high uncertainty in the SAI results.

Certain other assumptions, necessary in the SAI analysis in most cases for want of any better approach, merit brief discussion. On page 1-30, for example, the SAI report says that "The failure of display of information to the operator is treated as a random independent failure or set of failures and is not dependent on the accident sequence." The limitations in this assumption are clear: that is, independence is obviously not fully true for some sequences, but we understand (and concur with) the reasons why the assumption is made. Also (p. 1-28), "The maintenance contributions in the fault tree model are modeled as mutually exclusive among certain systems consistent with the Limiting Conditions of Operation. An example of this would be that HPCI is not allowed in maintenance if RCIC is unavailable." While this seems reasonable, it is clearly non-conservative to some degree, since LCOs are not always obeyed. Another guideline used in the SAI analysis (p. 1-28) is that "Plant components are presumed to meet all performance requirements consistent with licensing." This assumption implies that, although a component or system might be unavailable, if it is operating it will perform essentially at least as well as the licensing requirement. Again, this reasonable approach contains some (unquantifiable) error.

3.4 Implications of the Review of Specific Accident Sequences

Our review of specific accident sequences described in the SAI draft PRA report has resulted in several specific comments that have been discussed in sections 3.1 - 3.3 just above. Here we will consider the implications of these findings for Suffolk County's emergency planning activities.

Our most important quantitative finding is an apparent error in the SAI calculation of the probability of severe internal flooding: we have discussed this in section 3.3 and in Appendix D of this report. If our interpretation of the frequency of flooding at elevation 8 of the reactor building is correct, then the frequency of core-vulnerable conditions in accident Classes I and II is raised considerably: mostly in Class I in our opinion, with some effect on Class II as well. The SAI report gives as its best-estimate for core-vulnerable frequency in Classes I and II the following:

Class I	2.7×10^{-5} per year
Class II	1.1×10^{-5} per year

The conditional probabilities of core melt are given as 7% and 8% for these two classes, respectively.

If our analysis is correct, these core-vulnerable frequencies would be considerably higher. Because we have not been able to carry out a complete systems analysis, we believe that our choosing a number and sticking to it

as "our value" is inappropriate: we believe it is more appropriate that LILCO and SAI carry out a proper analysis. We not only believe that the core-vulnerable numbers will rise, but also that the conditional probability of core melt might increase: to be precise, if flooding dominates Classes I and II, then the flooding sequences should be used as the basis for computing the conditional probabilities, with the possibility that core-melt might be more probable than 7-8% given the amount of equipment taken out of service by a flood at elevation 8.

With these provisos, we believe that it is prudent for Suffolk County to use a number in the range of 10^{-4} per year as the likelihood of core melt from both Class I and Class II: use of this value as input to the planning basis for the County emergency planning activities will prudently allow the County to protect against accidents at Shoreham within Classes I II. Also, because no core-vulnerable/core-melt analysis has been done for these internal flooding sequences, we believe it is prudent to ignore this factor for the time being in these classes. (As it turns out, using these higher values probably makes almost no difference to the Suffolk County emergency planning activities, because Classes I and II dominate the accident planning basis one way or the other.)

For the other accident Classes, the only one where we believe the SAI results are probably not correctly representative of the real accident frequency is Class IV (the ATWS accident group). However, as discussed earlier, even though our best estimate would lower the SAI values for

core-vulnerable frequency by a factor of about 3 to 10, we recommend that Suffolk County use the SAI values as a prudent planning basis.

Our review has revealed some important qualitative conclusions that should be mentioned. Most important of these is that we have found the SAI analysis of accident sequences and their probabilities to be on the whole a fully competent, state-of-the-art analysis. Also, we believe that there are major uncertainties in the results of these analyses today, not only the analyses by SAI but analyses by all teams that carry out reactor PRAs: we have touched on some of the underlying reasons for these in earlier discussion. Nevertheless, we believe that the SAI group has taken some care to understand and quote the ranges of their uncertainties, and we find their uncertainty analysis for the accident sequences acceptable.

4.0 Are the Accident Phenomena Within the Plant Treated Correctly?

The approach taken by our review team in reviewing the accident phenomena parts of the SAI draft PRA was to ascertain whether the analytical methods employed are reasonably consistent with the present state-of-the-art methods now used for these calculations. Members of the FRA team are quite familiar with the level of understanding within the broader reactor safety community about these phenomena, including a recognition of the significant limitations to our present understanding.

The basic conceptual problem being addressed here is as follows: we postulate the unlikely situation that the Shoreham reactor has reached a physical state in which a core-degradation accident is underway. The accident sequence that has led to this physical state is assumed to be understood, in the sense that the event-tree/fault-tree analysis has revealed to us a specific sequence of equipment failures and human errors leading up to the onset of core degradation. (It is not necessary for the purposes of this part of the analysis to know the probability of the sequence.) The problem is to attain acceptable understanding of the physical sequence of events following the equipment failures modeled in the event-tree/fault-tree analysis, beginning with the onset of core degradation and continuing through until the reactor core is either securely cooled, or has melted (wholly or partially), releasing a fraction of its fission products to the reactor vessel, the containment structures, and possibly the environment.

In order to carry out quantitative calculations of the phenomena, one requires a variety of data, some conceptual understanding of the physical and chemical events, and a calculational model embodying this understanding. The calculational model should incorporate not only the physical reality of the reactor plant itself, but also the functioning (correct or degraded) of various engineered systems within the plant.

In an ideal analysis, it would be desirable to calculate the sequence of phenomena for every important accident sequence that follows core degradation. Unfortunately, though every accident sequence is different from every other one, such a massive effort would be beyond the calculational ability of any team of analysts today. Also, because our understanding of these phenomena is limited it would not make much sense to model each specific sequence separately: differences among similar sequences are far less significant than uncertainties in our understanding.

For these reasons, PRA analysts generally have grouped the numerous accident sequences beyond core degradation into categories, each category being treated separately in the calculation of core degradation/melt phenomena. This grouping allows the analysts to treat a tractable set of issues, with the limitation that errors and uncertainties are introduced because of the approximate nature of the categorization process.

Each group of accidents is characterized by a set of conditions: examples include high or low pressure within the reactor vessel; containment

integrity intact or already failed before the onset of core degradation; time elapsed since SCRAM; various engineered safety systems either available or failed; and so on.

The challenge is to develop an acceptably accurate quantitative analysis of what happens to the core and its fission products, and how, and in what sequence, and under what physical-chemical conditions. The result of the analysis should be a quantification of fission product releases from the plant in terms of species, quantities, times of release, physical-chemical conditions of the release (energy of the release, physical form, etc.), and some measure of the confidence of one's conclusions.

FRA's purpose in this review has been to develop an understanding of what the SAI analysis team has done in the course of their work, so as to ascertain whether their results form an acceptable basis for the calculation of offsite consequences: if the SAI calculations do form such an acceptable basis, then Suffolk County's emergency response plan can utilize them.

The approach taken by the SAI analysis team was to utilize a well-known, widely used, and well-documented set of computer-based models known as the MARCH-CORRAL code package. These codes, originally developed under AEC/NRC sponsorship for the 1973-74 WASH-1400 analysis, have been the main method of analysis of these phenomena ever since. However, it is widely known that these codes contain important approximations, omissions, and other

limitations that make their modeling of physical reality less than precisely correct, and indeed the limitations of MARCH and CORRAL have been a continuing subject of research, analysis, and discussion in the reactor safety community for much of the last decade. In addition to the MARCH-CORRAL codes, the calculations of the accident phenomena require various types of input data, the most important of which are the fission-product release fractions or partition fractions in various stages of the postulated accidents. Examples include information about fission product releases from the fuel pins during fuel melting; fission product partitioning in a water/steam environment as a function of pressure, temperature, and other conditions; and fission product plateout, transport, and the like in the airborne state. The MARCH-CORRAL code requires data such as these essentially as parametric input to its modeled calculations.

The SAI analysis team made modifications to parts of MARCH and CORRAL in order to tailor the code to the Shoreham reactor, and to incorporate some recent insights about various phenomena. These modifications do not substantially modify the basic operating philosophy of the codes, nor their major limitations.

4.1 Grouping of Accidents into Classes

One important question that we have looked at is whether the grouping of accidents into five classes by the SAI team is reasonable. The five Classes can be briefly characterized as:

Class I	Loss of Coolant Makeup Accidents
Class II	Loss of Containment Heat Removal Accidents
Class III	Large LOCAs
Class IV	ATWS (Anticipated Transients Without Scram) without Poison Injection
Class V	LOCAs Outside Containment (including Interfacing Systems Accidents)

All accidents leading to and beyond core-vulnerable conditions are placed in one of these five Classes, and the five Classes are dealt with separately in the subsequent (MARCH-CORRAL type) analysis. We have studied the validity of this approximation, and find the grouping by SAI to be fully acceptable. We believe that some differences exist among accidents that are grouped by SAI into a common Class, an example being the variety of accidents grouped under Class II in which containment residual heat removal fails: the very long time that it takes for the Class II accidents to evolve is characteristic of the Class as a whole, but differences in detail emerge depending on which safety systems fail in which order. However, in this case (as generally), we believe that the grouping is a reasonable approach. Our reasoning is that the uncertainties in the phenomena are sufficiently great as to overwhelm any additional "information" that might have been gained by a grouping into more numerous Classes; and that the choice of specific sequences within each Class to serve as the model for each Class is also reasonable.

An important issue is the time duration between initiation of the accidents and the eventual release of fission products to the environment: this will differ from Class to Class substantially. We have studied the general

grouping from this perspective, because of a concern that perhaps there could be mis-classification from the perspective of time evolution. We tend to affirm that the time evolution of the Classes is reasonably consistent with our experience: for example, the Class II accidents evolve very slowly compared to the other Classes, and we believe that all of the accidents grouped into Class II behave in this way. This conclusion allows us to deal with all of these accidents together in discussing their emergency-planning implications in the context of Suffolk County's needs.

4.2 Review of Phenomena After Meltdown

The approach we have taken for ascertaining whether the SAI meltdown analysis of each accident Class seemed reasonable was in part the use of subjective judgment based on the experience and knowledge of the reviewers, and in part some independent calculations of a few of the phenomena involved. These independent calculations, carried out by Dr. Fabric, tend to support the validity of the SAI conclusions for many of the important post-meltdown phenomena, with some exceptions that will be discussed next.

A general elevation-view drawing of the Shoreham reactor and containment building is shown on the next page, reproduced directly from the NRC Safety Evaluation Report for Shoreham (Ref. 6). The main important features to

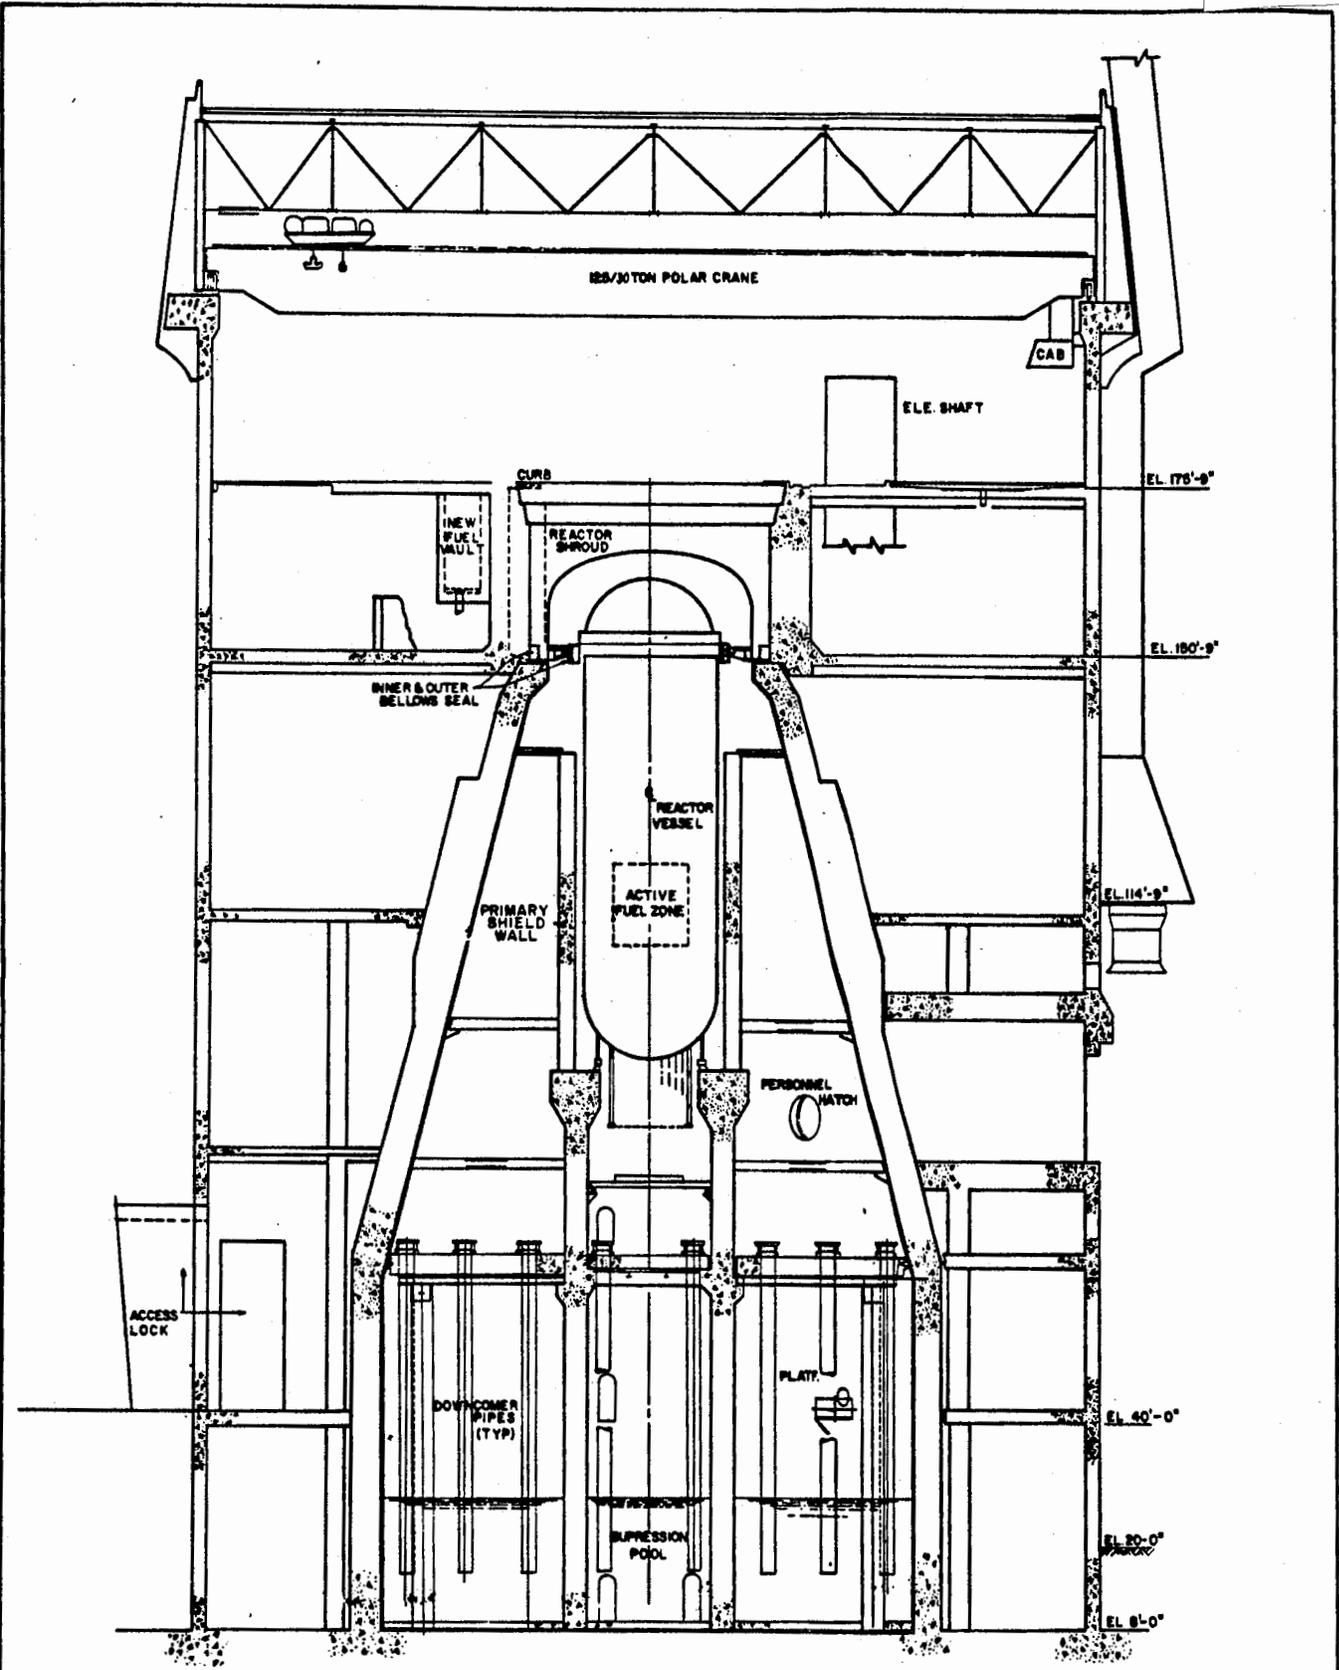


Figure 6-1

Shoreham Mark II Containment

(reproduced directly from the Shoreham "Safety Evaluation Report" by NRC, NUREG-0420)

concentrate on for our discussion below are the location of the suppression pool at the bottom of the building, and the downcomers penetrating the drywell floor below the reactor itself.

That the MARCH code (which models the highly complex processes involved in a meltdown) has deficiencies is, as mentioned above, well known. In Appendix B, some of these are discussed by Dr. Fabric, and his discussion will be excerpted here, as follows:

"In the course of the SASA (Severe Accident Sequence Analysis) calculations being performed by the Oak Ridge National Laboratory (ORNL) for the USNRC as applied to BWRs, the ORNL engineers have also come to the conclusion that MARCH has many deficiencies. Out of 21 listed inadequacies, we shall extract only a few:

- * Modeling of heat transfer to upper and lower BWR vessel internals
- * Modeling of core collapse
- * Failure of bottom head via control rod drive tube penetration not considered
- * Suppression pool and wetwell/drywell interaction
- * Rod-to-rod radiation heat transfer not included
- * Vessel water level calculation does not include variable flow areas
- * Fuel pin melt/slump/freeze phenomena are not mechanistically modeled."

Dr. Fabric (Appendix B) also discussed the reasons why the MARCH code contains extensive simplifications:

"The extensive simplifications in the code were introduced for two reasons: (1) to produce a tool useful for probabilistic risk assessment which requires many computer runs for exploration of consequences of various bounding assumptions. Hence, computational economy must have been one of the principal goals; (2) to produce a tool amenable to accepting some of the major uncertainties as input assumptions that could be changed from run to run. These stem from the relatively poor knowledge of various thermohydraulic processes that involve melt propagation and the attendant heat and mass transfer in complex geometries. The pertinent empirical base lags far behind the existing empirical base collected in the course of reactor safety research for situations that do not involve a degraded core. However, the word 'uncertainty' is also used here to imply simplifications one needs to make to intentionally bypass the detailed calculations that would cause the code to become long running."

Despite these limitations within MARCH, the conclusion of the FRA review team is that the code has been applied by the SAI team in a competent and conscientious manner. Our analysis of their description of their MARCH runs, the way they have thought about issues involved in MARCH, and their

conclusions has left us with both a confident feeling about what the SAI analysts have done and a heightened appreciation of the approximations inherent in the MARCH modeling.

One of the important areas of disagreement concerns the mechanism of melt-through of the molten core as it penetrates the bottom head of the reactor pressure vessel (RPV). Again we quote from Appendix B, Dr. Fabric's report (the abbreviation "CRD" means "control rod drives"):

"We do not believe that the RPV bottom head would fail in either of the two ways described in the Shoreham PRA: structural failure of the vessel wall due to elevated pressure and temperature or the wall melt-through when the RPV pressure is low.

"Instead, we agree with the scenario described by R. Henry of Fauske Associates, Inc. (see item 8 in Section 1.1) wherein the relatively thin metal that seals the CRD tube, or the CRD tube itself, is much more likely to fail first upon contact with the melt.

"The melt is very unlikely to reach the bottom head in a coherent fashion (gross slumping). It is more plausible to consider downward streaming of melt around the vessel axis. That melt will not attack the CRD tube as long as there is some water left in the lower plenum. If the amount of steam generation caused by quenching of melt is insufficient to stop further melting (a likely case),

increased amount of debris will accumulate on the bottom portion of the RPV and it would eventually remelt the fraction of debris that was frozen by the liquid--which by now has evaporated. This whole process could be delayed significantly if the CRD cooling water were continuously supplied.

"Eventually, one or more CRD tubes would fail and the debris discharge into the CRD room would commence."

Following the failure of the CRD tubes, another key question is how fast the melt will pour or stream through these CRD channels. Dr. Fabric calculates (Appendix B) that under the conditions likely to be present, only a very few seconds (perhaps less than 5 seconds) would elapse before as much as 50% of the melted metal would flow out of the lower vessel openings. In this scenario, there is no outright structural failure of the vessel wall, nor wall melt-through per se. However, as the melt penetrates the CRD tubes, significant enlargement of these penetrations will occur by ablation. If this is true, then only a small fraction of the melted core will remain in the bottom head, and that fraction will be the core material that has not melted, or has re-frozen at the bottom. Whether the flow through the CRD openings is pressure-driven or gravity-driven will make a difference, but should not make a very large difference to the phenomena in this process, which will be quite fast.

The next issue involves where the core debris streaming through the CRD openings will go. With a large amount of CRD hardware in the way (essentially hanging down in large massive metallic drive channels below the lower head), and with such a fast outflow, it seems unlikely that the CRD hardware will melt before the core material has passed through, which in turn indicates that the core material will reach the drywell floor only after being significantly broken up and "sprayed about" by the CRD hardware. However, an important insight (see Appendix B) is the conclusion that most of the core will go down into the drywell area just below the reactor pressure vessel (RPV) head, with rather little of it going into the annular outside region of the drywell floor. Once the level of core melt debris on the drywell floor exceeds the height above the floor (about 3.5 inches) of the downcomer top flanges, the molten debris will start pouring down the four downcomers that are located within the area of drywell floor just below the RPV: we find that it is unlikely that much of the melt will go down the numerous downcomer channels in the outer annular region, since the melt will go down these four inner downcomers promptly. Our analysis indicates that this will occur rather quickly, most likely within about 5 minutes after RPV breach.

Once the debris begins pouring down the downcomers into the deep wetwell, we studied whether steam produced by the molten material would produce enough countercurrent flow upward to impede further melt from penetrating downward. Our calculations (see Appendix B) seem to indicate that this phenomenon will not occur.

In the scenario described above, almost all of the melt will end up in the wetwell almost all of the time. Appendix B discusses why we are reasonably confident of this conclusion. It appears that there is only a limited range of conditions where melt will not mostly be quenched in the wetwell. The implications of this quenching are, of course, that significant removal of fission products in the wetwell water would occur, lowering by large factors the amount of fission products other than noble gases available in the gaseous phase for atmospheric release through an ultimate containment breach.

Perhaps the most important conclusions from Dr. Fabric's analysis in Appendix B are that vessel failure will probably take place through the CRD channels; that discharge from the vessel to the drywell will be very rapid (less than 1/2 minute), and discharge from the drywell to the suppression pool will also be fast (less than, say, 10 minutes); that core melt material will nearly always end up in the wetwell, where significant removal of fission products will occur; and that there is little time for debris-concrete interaction on the drywell floor, and too low a temperature for that interaction to occur on the wetwell floor.

All in all, the picture painted by our analysis is that the SAI results on core-melt behavior and fission-product release are likely to be conservative (that is, too high values for release, too much core-concrete interaction) compared to what we believe to be the real phenomena.

In particular, we expect that for two of the classes of accidents (Class I* and Class III**), there will be significant delay in the containment rupture time, and perhaps no containment rupture at all, depending on specific containment heat-leakage properties that we have not specifically addressed and that are both complicated and scenario-specific. Also, there will be quite large decontamination factors for fission products released in the wetwell pool, leaving the gaseous fission products released directly from the melted fuel pins as the principal airborne species available for release from the containment after breach.

Unfortunately, we are not able to quantify the extent of this conservatism. The phenomena that we believe will occur differ from those modeled in the SAI draft report in several areas, but almost always in the "conservative direction," meaning that if our analysis is a better representation, then the accidents will be less severe. We believe that much remains to be done before a reliable calculation of these phenomena can be carried out: some research on specific phenomena must be performed; some improved modeling of the sequence and character of the events must be accomplished; and some better data on specific coefficients must be obtained, for a variety of physical and chemical interactions.

* Class I involves accidents where failure of core cooling with the RPV at high pressure after a transient or small-break LOCA results in core degradation, with an intact containment until after core melt.

** Class III involves accidents such as a large-break LOCA inside containment where RPV depressurization occurs prior to core degradation, with an intact containment until after core melt.

4.3 Some Specific Issues Involving Meltdown and Related Phenomena

In the course of our review of the meltdown part of the draft PRA on Shoreham, a number of specific issues arose that we believe are worth discussing within this report. These are technical issues that could affect the overall "bottom line" conclusions of the SAI study. For each we will indicate what our own conclusions are as to the impact of the issue vis-a-vis Suffolk County's application of the PRA results for emergency response planning.

The first issue is the magnitude of fission product releases from the fuel. In Table D-3 (page D-9) of the SAI draft report, a list is given of release fractions from the fuel to the primary reactor system during core meltdown. The table is reproduced on the next page. Shown in the table are two values for the release fraction for each isotope: one is the value used in the 1973-74 WASH-1400 analysis (Ref. 1), and the other is the more recent value found in NRC report NUREG-0772 (Ref. 5), which contains a discussion of recent understanding of fission product behavior in core-melt accidents. The SAI report states (p. D-7) that they have used an average of the two model estimates (WASH-1400 vs. NUREG-0772), in view of uncertainties in our knowledge. While this averaging does not make very much difference to most of the release values, the value for tellurium is quite different in the two models: WASH-1400 assumed that only 15% of the tellurium escapes from the fuel, while NUREG-0772 believes the correct value is 100%. While we do

(reproduced directly from the Shoreham PRA by SAI)

Table D.3
RELEASE FRACTIONS TO REACTOR PRIMARY SYSTEM DURING CORE MELTDOWN

Fission Product Group	RSS	NUREG-0772
Xe, Kr	.9	1.0
I, Br	.9	1.0
Cs, Rb	.81	1.0
Te	.15	1.0
Sr	.10	.3
Ba	.10	.5
Ru**	.03	.02
Zr, Nb	.003	.03
La [†]	.003	.003

* Includes Sb, Se

** Includes Mo, Pd, Rh, Tc

† Includes Nd, Eu, Y, Ce, Pr, Pm, Sm, Np, Pu

NOTE: "RSS" is "Reactor Safety Study", Report WASH-1400

not know which is correct (that awaits further experiments that have not yet been done), we believe after study of the NUREG-0772 arguments that the lower figure is more likely to be correct for tellurium, which would slightly decrease the offsite releases and doses. The effect is a small one generally, but not so small as to be negligible.

SAI's draft analysis also states (page D-7) that 90% of the iodine is released from the primary system as CsI (cesium iodide), which is a major departure from the WASH-1400 assumption that iodine was released 100% in elemental form. This makes a significant difference because the CsI is soluble while the elemental iodine, in gaseous form, was a key contributor to offsite doses in WASH-1400. We tend to agree with the recent arguments that elemental iodine is not a likely form for that element in these accidents, and believe that SAI's assumption is a reasonable one. However, there are still important uncertainties in our knowledge of what precisely happens to the iodine during the accidents under consideration. The reactor safety community will require some experiments, some modeling, and considerable discussion (all of which is underway) before this important issue is resolved and a consensus reached on it.

Another issue involves whether there would be zirconium-water reactions within the reactor vessel after the fuel has melted but before the melt penetrates the bottom head: SAI assumed (page C-23) that in a melted state the Zr-water reaction would not occur, and has explained their assumption (Ref. 2) as being due to the presence of a molten pool at the bottom of the

lower vessel head, covered by a crust at the interface between molten and solid fuel. This crust would inhibit Zr-water interactions, according to the SAI opinion. While this scenario seems reasonable, the generation of zero Zr-water reaction seems too extreme, because surely some of the zirconium will be in contact with water or steam during parts of the meltdown process. This issue needs more analysis, in our view. We have no way of quantifying whether the Zr-water reaction is miniscule or only "small," but we do concur that it is unlikely to be "large" in the sense of comprising any reasonable fraction of the total zirconium in the fuel cladding. Thus we do not think this issue should have any important impact on the overall results of SAI's analysis.

Another issue that we have studied is whether there will be rapid degradation of the concrete walls of the drywell during core-melt accidents, which might lead to carbon-dioxide evolution, high pressures building up inside, and ultimate pressure failures of containment. First, we should point out that in our own analysis, almost all of the melt goes directly from the bottom vessel head to the CRD room floor to the suppression pool; almost none goes to the annular drywell floor (see Appendix B). However, for this discussion we will assume that we are considering melted core on the drywell floor. When our concern with the issue of drywell concrete degradation was raised with the SAI analysis team, their response (Ref. 2) was that only for the Class III sequences would a sufficiently high temperature occur in the drywell, and these sequences produced very short thermal transients. If this is true, the

decomposition temperature would be reached only for a very thin layer near the surface of the concrete. We have not been able ourselves to calculate any sufficiently detailed numerical values for the temperature transients, nor can we find them in the draft SAI report. However, based on our experience with this issue, and our understanding of the phenomena that will occur during melting and core transport from vessel to drywell to wetwell, we believe that this concern is not likely to be a major one: specifically, we believe that even if significant core melt were to reach the drywell annulus, it is quite unlikely that this concrete-degradation effect will produce enough non-condensable gas to challenge containment earlier than it is challenged in the scenarios set down by SAI in their report.

We have also been concerned that containment integrity might be compromised by pre-existing containment leaks, that might not have been considered properly in the SAI analysis. Also, we were concerned with failures of penetrations at high temperatures. While SAI confirmed (Ref. 2) that this latter problem was not explicitly considered in the PRA, they believe that the former is adequately analyzed within the containment event trees. We are not convinced that the analysis incorporates either of these issues quantitatively; however, we believe that containment leakage is not as important an issue for the Shoreham BWR as it appears to be for many other reactors (in particular, for many of the PWRs), because the effective fission product removal action of the suppression pool provides such large reductions for the most likely accident classes. Thus the issue probably does not introduce major additional pathways to the environment even if it has not been thoroughly considered in the SAI analysis.

One of the key issues that has been an outstanding concern for many years in consideration of core-melt accidents is "debris-bed coolability," which is a shorthand phrase for the question of whether, if a bed of core rubble were to form from molten-then-solidified core material, that bed can be cooled adequately. Another question is whether steam, generated in and around the debris bed as it is cooled, can provide enough additional pressure to pose a challenge to containment. First, we must clarify that we are concerned here with cooling a debris bed on the drywell floor; we have no concern with cooling the later debris bed under many feet of suppression pool water, although ultimately after many days the pool water would boil away and the issue would arise again. The SAI response to our inquiry on this subject (see Appendix C) is reasonable, and is congruent with our own judgment on the matter. However, it must be emphasized that while our judgment and that of the SAI group agree, the entire subject of debris beds is still one where everybody's conclusions are highly speculative; in our view the safety community as a whole doesn't have enough experimental data, nor modeling talent, to put this issue to rest at this time. Fortunately, the issue is not believed by our group to pose significant additional risk potential, for the following reasons which are quoted here from Mr. Davis' discussion in Appendix C:

"Upon further consideration of this issue, it is considered not to be significant in terms of the potential for risk increase for the following reasons:

"1.If containment failure has not occurred at the time debris bed cooling occurs, most of the fission products (which are released during the initial meltdown phase) will be securely trapped in the suppression pool water. Thus, the possibility of containment failure from a steam pressure surge upon debris bed cooling would not cause a large fission product release.

"2.If containment failure occurs before debris bed cooling, the major consequences of the accident would be underway, and the added fission product release would likely not be significant.

"3. A steam pressure surge sufficient to challenge containment integrity requires a large amount of water delivered to the bed and intimate mixing. It is not likely that a high volume water source would be available since most such sources within the containment must previously have been assumed failed or degraded to cause the accident to progress through core meltdown."

4.4 Core Vulnerability vs. Core Melt: Containment Event Trees

In section 3.1 (above), we commented favorably on the methodological advance employed by SAI in the Shoreham study in which an accident that proceeds to a "core-vulnerable" state is differentiated from one that proceeds further to a "core melt" state. As discussed earlier, this differentiation was not made in the WASH-1400 analysis, and yet it is

clearly important to recognize that only a fraction of "core-vulnerable" accidents proceed to "core melt," an example of one that stopped short being the Three Mile Island accident of 1979.

The SAI report finds, for the five Classes of accidents considered, the following conditional probability that a core melt will follow:

<u>Class</u>	<u>Conditional Probability of Melt, Given Core Vulnerable State</u>
I	8%
II	7%
III	84%
IV	43%
V	100%

The way these conditional probabilities were calculated in the Shoreham PRA analysis represents an advance over earlier probabilistic reactor analyses. The effort consisted of the development of very complex containment event trees (CETs) that considered the large variety of engineered safety features and phenomena that are brought into play during the time period after a core-vulnerable condition is reached. Of course, some assumptions are necessary to simplify the problem, and the SAI analysis team made several of them. We have studied them in the SAI draft report (see the list on p. 1-19, pages H-11 to H-13, H-48 to H-56, and Table H.7), and find them generally reasonable, although the analytical basis for the specific numbers is weak.

We believe that there are major uncertainties in the CET analysis, but our study of the SAI discussion leads us to believe that the SAI analysts were aware of these and handled them acceptably. For example, a key limitation is that there are essentially no data on the relative likelihoods of containment failure modes of different sizes and types that would result

overpressure (see page H-50). Another example is that, although the analysis correctly differentiates between containment overpressure failures near the top as opposed to near the bottom of the containment (a failure near the bottom, although very unlikely, could compromise the suppression pool), their treatment is obviously an approximate one.

We believe that the point of departure of the SAI treatment is proper. They start by differentiating among three situations: those in which the core is vulnerable to melting in an intact containment; those in which containment may be vulnerable first while the core is still adequately cooled; and those involving containment bypass. These three topologically different states are then treated separately in the quantification process.

We also endorse the approach taken in the draft SAI report to consolidation of the release end-states. This consolidation has the effect of grouping numerous different accidents and treating each group singly, which inevitably implies loss of detail in the interest of calculational tractability. The SAI report acknowledges this issue (p. H-51), and states that the approximations made are conservative in nature. We affirm that this is probably correct but have not reviewed the details sufficiently to have an independent view of whether this is always true (see Table H.8, p. H-106 for details).

To mention other technical issues, we believe that the treatment of the steam explosion issue is a reasonable one. Also, the discussion on

quantification of the CETs recognizes explicitly some concerns that clearly affect the quality of the results (see Table H.6, p. H-57 to H-60). The SAI analysts have explicitly differentiated among various qualities of their supporting data, have given not only their best point estimates but 10%-90% bounds for their results, and have documented their main assumptions. Again, however, as elsewhere in this review there was no way that our review effort could examine the (literally, hundreds of) detailed numbers in the CET quantification.

Our failure to review these conditional probabilities quantitatively is not very troublesome to us, because we are of the opinion (a qualitative opinion, however) that the values quoted are reasonably within the range that we would expect. Furthermore, we believe that the inclusion or exclusion of the SAI conditional probability factors does not make any important difference to what Suffolk County will do in its emergency planning activities. Therefore, and because we are uncertain as to the quantitative validity of the results, we believe that it is prudent for the County to ignore these factors in its planning, and to take the core-vulnerable figures, as modified by our recommendations for Classes I and II, as the planning basis for emergency preparedness.

Our rationale for recommending the exclusion of these factors is as follows: for Classes I and II, we have already recommended that the County assume that the core-melt values lie in the range of about 10^{-4} /year, taking into account our improved analysis of internal flooding; for the internal-flooding sequences there does not exist a valid

core-vulnerable/core-melt analysis, and we believe that our recommended numbers are in any event only a rough estimate. For Classes III and V, the SAI factors are so close to unity (84% and 100%, respectively) that there is no difference. For Class IV, the SAI factor is 43%, within about a factor of 2 of unity, but for this class we believe that the core-vulnerable number could be too high by a factor of 3 to 10 because of SAI's use of too high a value for failure of SCRAM on demand: so for Class IV the inclusion or exclusion of a 43% factor is practically like splitting hairs.

For all these reasons, we believe that omitting the core-vulnerable/core-melt factors from Suffolk County's planning basis is the prudent choice: the factors have no reasonable basis for the internal-flooding sequences and make essentially no difference in the other Classes.

4.5 Implications of the Review of In-Plant Phenomena

Our review of in-plant accident phenomena described in the SAI draft PRA report has resulted in a collection of specific comments and remarks that are covered in the earlier sections (4.1 - 4.4) of this chapter. It is important to describe the context in which these comments are to be understood. We believe that the present state-of-the-art of probabilistic analysis of in-plant phenomena is not very well advanced. In particular, our underlying understanding is inadequate for some phenomena (for

example, core melting itself, core penetration of the vessel, debris bed formation and coolability, aerosol plateout, core-concrete interactions); for the performance of some key systems (for example, containment failure mechanisms, effectiveness of suppression pool heat removal mechanisms, efficacy of active aerosol removal systems); and for the time sequence and duration of some events during the accident (for example, duration of meltdown itself, pressure buildup).

Given the paucity of experimental information, only limited applicability of the data that do exist, and calculational intractability of models complex enough to contain detailed differential effects, it is no surprise that differences of opinion exist within the professional community. Some of these differences are reflected in our comments above. What is important to leave with the reader of this report is that we do not believe that there are important differences between what SAI has accomplished and documented and our own view about release magnitudes: differences of interpretation, differences in level of detail, differences in selecting experimental data or modeling approximations seem in every case to produce effects on the final PRA results that are within the quoted uncertainties of the SAI analysis.

If there is one overriding impression that our review team is left with in the aftermath of this in-plant phenomena review effort, it is that the magnitudes of the radioactive releases are likely, in actual accidents, to

be smaller than are calculated in the SAI draft report. This is due to the inherent introduction of numerous conservatisms in the analysis. Paramount among these are conservatisms in the description of removal mechanisms that will keep important fission products within the reactor building, for which incomplete credit has been taken in the analysis. Examples include the likelihood that bottom-head penetration may not even occur at all for some scenarios, either because melting will be incomplete, or heat transfer larger, or CRD cooling water available (see our Appendix B); the analysis of suppression pool decontamination factors (see our Appendix C); and the various assumptions that SAI has made on containment failure mechanisms themselves.

We are unable to quantify the extent of these conservatisms in the in-plant phenomena analysis; indeed, we believe that it will be several years from now before enough research has been carried out to enable a consensus to be reached on these issues. For this reason, we believe it imprudent to take account of them for the purposes of advising Suffolk County's emergency preparedness effort. However, it is important that the reader understand our view that the SAI radiation release results, taken at face value and considering their large quoted uncertainties, are more likely to be too high (that is, to represent accidents more severe than actual) than too low.

5.0 Summary and Conclusions: Are the Probabilities and Magnitudes of the Calculated Radiation Releases Correct?

In earlier sections we have already given our summary conclusions as to the correctness of the probabilities of the accident sequences (Section 3.4) and the magnitudes of the calculated releases (Section 4.5). Here we will summarize our findings, and discuss their implications when applied to Suffolk County's emergency planning activities.

Concerning the SAI calculation of accident probabilities, we have found two important differences between our analysis and the SAI analysis. They were noted in Section 3.4, as follows: First, we believe that due to an error in the SAI treatment of internal flooding, the contribution of internal flooding to Class I and II accidents has been underestimated. While we are unable to provide our own analysis of the internal flooding accidents, we recommend that Suffolk County use, as a basis for their emergency planning effort, values of about 10^{-4} per reactor year for the frequency of core melt for both Classes I and II. Second, we believe that Class IV accidents will occur less frequently than the SAI report claims, because we believe the SAI report has used too high a value for the failure on demand of the SCRAM system. However, we recommend that the County should use the SAI results for Class IV as their planning basis. We recommend that the County use all other values for core-vulnerable and core-melt frequencies directly as found in the SAI draft report.

Concerning the SAI calculation of magnitudes of radiation releases from Shoreham, our review has left us with the strong impression, which is supported by the discussion in our Chapter 4, that the releases calculated in SAI's study are more likely to be too high than too low (that is, we believe that their calculation has conservatisms within it). We have been unable to quantify the magnitude of these conservatisms, and in fact we do not believe that it is possible to do so at the present time, because of insufficient information about the phenomena that characterize accidents of the type under discussion. However, we believe that the conservatisms could amount to large factors of reduction for some accident types, particularly those important accidents (albeit, in an absolute sense quite rare accidents) where the melted core would pass into the Shoreham pressure suppression pool. In our opinion the suppression pool, and the way the Shoreham design provides for prompt passage of melted core material to it in these unlikely accidents, will be very effective in fission-product removal. The substance of our remarks here is our opinion that the removal effectiveness of the pool may be even greater than the SAI analysts have used in their calculations.

Thus we arrive at the following overview of our conclusions as to the probabilities and magnitudes of SAI's calculated radiation releases from accidents at Shoreham. We believe that the Class I and II accident groups have higher probabilities than are found in the SAI report, because of internal flooding accidents; and we believe that the magnitudes of the releases are likely to be lower than SAI has calculated.

The implications for Suffolk County's emergency planning effort will depend upon what the County planners decide to do with our findings. We believe that the implications for the County of higher core-melt probabilities in Classes I and II would be minor because these classes already dominate the composite offsite doses from Shoreham, even if SAI's lower numbers are used. If the County planners use the composite dose-distance curves as their planning basis, increasing the absolute probabilities of these curves changes nothing that the County will do or decide. Concerning the magnitudes of the releases, use of substantially smaller releases would have a major effect on the County's planning effort; however, we recommend that the County utilize the SAI draft results as their planning basis because we cannot quantify the degree of conservatism in them.

6.0 ACKNOWLEDGMENTS

The authors wish to acknowledge the advice and assistance of Fred C. Finlayson, our collaborator on this project without whose help the work could not have been accomplished. Fred has also served in a liaison capacity with Suffolk County throughout this work. We also acknowledge Christopher McMurray, counsel to Suffolk County, for much assistance. We have benefitted from important interactions with Robert Erdmann and Edward Burns of Science Applications, Inc. and Voyin Joksimovich of NUS Corporation, who have explained parts of the SAI report to us. Finally, our interactions with Suffolk County itself, through Frank Jones and Patricia Dempsey, have been smooth and successful throughout.

7.0 References

The primary reference used in this report is the probabilistic risk assessment carried out by Science Applications, Inc. on the Shoreham Plant.

The reference is:

Science Applications, Inc. (San Jose, California),
"Probabilistic Risk Assessment, Shoreham Nuclear Power Station,
Long Island Lighting Company," Preliminary Draft Report
SAI-001-83-SJ, March 1982

The following are referred to by number in the text:

1. U. S. Nuclear Regulatory Commission, "Reactor Safety Study," Report WASH-1400, October 1975
2. Letter, E. T. Burns (of Science Applications) to R. J. Budnitz (of Future Resources Associates), dated 2 August 1982, on the subject "Responses to questions generated during July 16, 1982 meeting"
3. "PRA Procedures Guide, A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants," prepared by the American Nuclear Society and the Institute of Electrical and Electronic Engineers, Report NUREG/CR-2300, Review Draft, dated April 5, 1982
4. A. D. Swain and H. E. Guttman, "Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications," Report NUREG/CR-1278, April 1980
5. U. S. Nuclear Regulatory Commission, "Technical Bases for Estimating Fission Product Behavior During LWR Accidents," Report NUREG-0772, June 1981
6. U. S. Nuclear Regulatory Commission, "Safety Evaluation Report Related to the Operation of Shoreham Nuclear Power Station, Unit No. 1," Report NUREG-0420, April 1981