WESTINGHOUSE POSITION ON IN-VESSEL RETENTION OF MOLTEN CORE DEBRIS FOR AP600

January 1996

WESTINGHOUSE ELECTRIC CORPORATION
Advanced Technology Business Area
P.O. Box 355
Pittxburgh, Pennsylvania 15230-355

Westinghouse Position on In-Vessel Retention of Molten Core Debris

The purpose of this paper is to recommend an approach for the closure of issues related to preventing the release of core debris from the reactor vessel during a severe accident by flooding the cavity and cooling the external surface of the vessel wall. This severe accident management strategy is called in-vessel retention of molten core debris (IVR), and increases the safety of the AP600 by eliminating the threat of ex-vessel severe accident phenomena associated with the relocation of core debris to the containment. These low frequency, high consequence phenomena, specifically ex-vessel steam explosion and molten core-concrete interaction, generally have large uncertainties associated with them and may be postulated to threaten containment integrity and produce large offsite doses. Prevention of the ex-vessel phenomena through IVR contributes to the low large release frequency of the AP600. The Westinghouse position on IVR proposes a change to the regulatory treatment of severe accident phenomenological issues.

Severe Accidents and Large Release Frequency in the AP600

The NRC safety goal (reference 1) for the large release frequency and conditional containment failure probability for advanced light water reactors are that the containment failure frequency is:

- less than 1.0×10^{-6} per reactor-year, and
- less than 10% of the core damage frequency.

The AP600 Probabilistic Risk Assessment (PRA) concludes that the core damage frequency for AP600 is 2.4×10^{-7} events per reactor-year which, by itself, is less than the large release frequency goal. The AP600 large release frequency is 1.0×10^{-8} events per reactor-year. The conditional containment failure probability is 4.1% of the core damage frequency. The AP600 therefore meets the NRC safety goals. Based on these results, AP600 provides significant protection against large release by preventing core damage accidents from occurring. The containment provides substantial defense-in-depth against the release of radiation to the environment in the event that a severe accident does occur.

Most of the containment failure probability is due to initially failed containment sequences (steam generator tube rupture which bypasses the containment and containment isolation failure). These are reliability-type failures in which the containment is failed as a result of the accident initiation. Release from an initially failed containment is approximately 4% of the core damage frequency. The mechanistic type of containment failure induced by severe accident phenomena (i.e. ex-vessel steam explosion or hydrogen combustion) represents only a small fraction (approximately 0.1%) of the core damage frequency.

The severe accident management strategy to flood the reactor cavity with IRWST water and submerge the reactor vessel is credited with preventing vessel failure in the AP600 PRA. The water cools the external surface of the vessel and prevents molten debris inside the vessel from failing the vessel wall and relocating into the containment. Retaining the debris in the reactor vessel eliminates ex-vessel severe accident phenomena, such as ex-vessel steam explosion and core-concrete interaction, from threatening the containment integrity.

The AP600 is uniquely suited to in-vessel retention of core debris because of several design considerations:

- the lower power-density fuel produces a lower heat flux through the vessel wall
- the reliable RCS depressurization produces lower stresses on vessel wall
- the "clean" vessel lower head has no vessel penetrations to provide a failure mode for the vessel other than creep failure of the wall itself
- the floodable reactor cavity can submerge the vessel to the coolant loop elevation with water intentionally drained from the in-containment refueling water storage tank,
- "IVR-friendly" reactor vessel insulation design concept to prevent the insulation from interfering with the water cooling of the vessel and the steam venting from the reactor cavity.

Basis for Low IVR Failure Probability in the PRA

The PRA estimates the upper bound failure probability for vessel integrity in the flooded cavity configuration to be 0.01. The best-estimate failure probability is 0.0001, which means that failure is virtually impossible. The ARSAP report on in-vessel retention of molten core debris (reference 2, IVR report) supports the failure probability as it is used in the AP600 PRA. The IVR report uses the Risk Oriented Accident Analysis Methodology (ROAAM), a probabilistic framework used to analyze the important "sub-phenomena" which make up IVR. Based on the testing and understanding of the sub-phenomena and their uncertainties, the report concludes that vessel failure into a fully-flooded cavity is "physically unreasonable", meaning that within the laws of physics, vessel failure is virtually impossible under the conditions and assumptions presented in the IVR report.

The IVR report has been peer-reviewed by an international panel of thermal-hydraulic, chemical and structural experts. All of the peer review comments have been included and responses provided in the IVR report. With the amount of testing and peer review, the report provides a higher quality assessment for the IVR phenomena than required for closure of severe accident issues in past PRAs (reference 3).

Contribution of IVR to Simplified Treatment of Ex-Vessel Phenomena in the PRA

At the conditional probability level for vessel failure used in the AP600 Level 2 PRA, it is possible to make simplifying, conservative assumptions in the quantification of the ex-vessel steam explosion and debris coolability since they can contribute no more than 1% of the core damage frequency to the conditional containment failure probability. This treatment is illustrated in the "mini-containment event tree" in Figure 1. The top events of the tree are defined as:

BP - containment bypass or isolation failure
IR - cavity flooding by IRWST water
VF - vessel failure
ESX - containment failure by ex-vessel steam explosion
DQ - debris coolability

The end-states on the event tree are:

IC - intact containment
CCI - long-term containment failure from core-concrete interaction

CF - early containment failure
BP - containment bypass or isolation failure

Failure probabilities for BP and IR are assigned to the tree based on approximations from the AP600 PRA quantification. As stated above, the probability of an initially failed containment (SGTR or containment isolation failure) at top event BP is 0.04. System, hardware or operator action failure probability of cavity flooding is approximately 0.02. Vessel failure in the flooded configuration is assigned an upper bounding failure probability from the PRA of 0.01. Top event ESX with the cavity flooded (path 4) is treated conservatively by assuming that the containment fails as a result of the ex-vessel steam explosion with a probability of 1 if debris is relocated to the cavity. The debris coolability issue is not significant in the flooded cavity paths since either debris is retained in the vessel or the containment is assumed to be failed by the ex-vessel steam explosion.

If the cavity is not flooded because of human error or system failure, vessel failure is guaranteed since the external surface of the vessel is not cooled. Based on the PRA treatment, there is very little containment challenge from ex-vessel steam explosion in this unflooded configuration. Also from the PRA, the probability of failure to cool the ex-vessel core debris is approximately 0.1 which is dominated by the formation of an uncoolable debris bed in the cavity. Failure to have sufficient water inventory to cool the debris may also produce core-concrete interaction but requires multiple water source failures and has a negligible failure probability in the AP600 PRA. Therefore, water inventory failure of debris cooling is not considered on this simplified tree.

Quantification of the event tree in Figure 1 results in a conditional containment failure probability of 0.051. This is higher than the PRA conditional containment failure probability (0.041) because of the scoping nature of this analysis and not crediting the vessel failure mode for reducing the failure probability of ESX on path 4. The event tree quantification demonstrates how crediting IVR at a failure probability of 0.01 or less allows conservative treatment of the ex-vessel steam explosion and debris coolability issues without exceeding the NRC conditional containment failure probability goal of 0.10.

Reduced Credit for IVR in the PRA

The NRC has not yet issued an official position on IVR or on how much credit they will accept in the PRA. They are in the process of reviewing the IVR report and cannot draw conclusions prior to completing the review. In early IVR meetings, the NRC "indicated that they saw IVR as a very positive aspect of the AP600 design that would obviate more detailed consideration of ex-vessel phenomena if the applicability of the favorable IVR experiments could be established." (reference 4).

It is possible that the NRC may not accept a failure probability less than 0.1 for vessel failure with the cavity flooded, regardless of the review. The NRC minutes from the August 17, 1995 IVR meeting (reference 5) state, "Because of the uncertainty involved with IVR and the novelty of the external reactor vessel cooling approach, the staff stated that some amount of ex-vessel severe accident work (e.g. fuel-coolant and core concrete interaction) would be required."

If a 0.1 failure probability is used for top event VF in Figure 1, and the tree is requantified with the conservative assumption for ex-vessel phenomena, the result is a conditional containment failure probability of 0.14 (Figure 2) which does not meet the 0.1 conditional containment failure probability goal. To meet the conditional containment failure probability goal, additional work must be done to

analyze ex-vessel steam explosion, debris spreading and coolability with non-realistic assumptions to provide additional defense-in-depth for the containment integrity in a severe accident configuration.

Current State of Regulation of Severe Accident Issues

The "probabilities" assigned to the failure of the IVR phenomena to prevent vessel failure are, in a sense, arbitrarily picked to either support the position that the ex-vessel phenomena are not important to the conditional containment failure probability (failure probability less than or equal to 0.01), or to support the position that defense-in-depth is required to protect against the ex-vessel phenomena in case the vessel fails (failure probability greater than or equal to 0.1). This argument cannot be won on this probabilistic battleground since there is no way to assign and defend an actual split fraction probability to such a complex phenomena.

Westinghouse believes that our estimate of the failure probability in the PRA is technically sound since it provides an upper bound of the position supported by the IVR report that vessel failure is not physically reasonable, and ex-vessel debris relocation sequences are not important contributors to the large release frequency. But because the failure probability is not readily quantifiable and defendable as a probability split fraction, we end up having to quantify and defend conditions for other equally complex ex-vessel phenomena. Based on the above, it makes sense to abandon trying to defend the 0.01 failure probability split fraction for IVR. A new strategy is needed for such low frequency, high consequence severe accident issues.

Purely deterministic analyses are not the answer, especially in severe accident space, since arbitrary bounding values for uncertain parameters typically used in such analyses can usually be chosen to produce failure. Also, bounding values are often not consistent among applications. So-called best-estimate values do not produce better results since they often do not satisfy the regulatory position that is driving the analysis and, given the uncertainty in the severe accident, the best estimate may not be readily apparent.

For example, in the issues of the ex-vessel steam explosion and debris spreadability, the key uncertainty controlling the phenomena is the vessel failure mode since it will determine the rate, quantity and composition of the debris relocation. The bounding condition for debris spreading is a slow pour which would occur from a vessel failure high in the debris pool. However, the bounding condition for a steam explosion is a failure at the bottom of the vessel that produces a rapid relocation of core debris. These bounding conditions are not consistent. The best estimate is that there is no vessel failure, but this position may not be satisfactory from a regulatory viewpoint. The "best estimate vessel failure", based on the location with the least margin-to-failure in the IVR analysis, is at a location in the metal pool at the top of the stratified debris bed in the lower head. The debris relocation resulting from this vessel failure mode is a slow pour of molten metal. The relocated mass is limited to the debris at the top of the pool. This type of relocation does not promote debris spreading, but since there is little decay heat in the metal pool, there is little threat of core-concrete interaction. The steam explosion that could be produced from such a relocation will be of limited magnitude. This scenario is not particularly challenging or bounding, so in this case, from a regulatory stand-point, it does not resolve the issues.

Proposed Change to the Regulatory Philosophy

The issue with the current treatment is that the sequences and the boundary conditions must be arbitrarily chosen. Similarly, the methodology used to perform the assessment, and the methodology to review the assessment are arbitrary, as well. The overall philosophy of this process requires revision. An approach that provides a frame that is clear, consistent and complete is needed to resolve low probability, high consequence issues and avoid the arbitrary treatment in severe accident space that has come to be called defense-in-depth.

The ROAAM methodology meets these requirements (reference 6) and has already been used to resolve the BWR liner melt-through and direct containment heating severe accident issues (references 7, 8, 9, and 10). The ROAAM methodology provides a framework to resolve phenomenological issues and to develop insights to prevention, mitigation and safety goals for severe accident issues. This methodology is being proposed to the NRC as a better alternative for the assessment of safety margins for severe accidents than the current safety goals.

The IVR ROAAM analysis demonstrates that building the reactor vessel and cavity to prevent debris relocation to the containment provides adequate defense-in-depth against severe accident sequences. If vessel failure can be adequately shown to violate the laws of physics if the cavity is flooded to cool the vessel, then the likelihood of containment failure from ex-vessel relocation phenomena is clearly a moot point. ROAAM systematically evaluates the uncertainties associated with the IVR phenomena and assures a consistent treatment in the analyses. The completeness is provided by a quality system-level PRA which is used to assure that all relevant sequences are addressed. Demonstrating that phenomenological failure is physically unreasonable becomes the safety goal instead of trying to assign and defend a conditional probability for complex phenomena.

Westinghouse Position

The ROAAM of the IVR analysis, which includes the evaluation of the thermal-hydraulic uncertainties of the in-vessel circulating debris pool and boiling crisis on the exterior vessel surface, has been performed and is being defended through the DOE ARSAP program by Professor Theo G. Theofanous, the developer of ROAAM. Given the strong technical basis and defendability of the IVR ROAAM assessment, Westinghouse supports the DOE ARSAP effort to change the regulatory environment and demonstrate that AP600 complies to the conditions and assumptions under which the IVR ROAAM analysis was performed. This effort entails:

- verification that the RCS depressurization system is reliable
- verification that the cavity flooding system is reliable
- demonstration that the reactor vessel reflective insulation and containment water recirculation flow paths are "IVR-friendly" and allow sufficient ingress of water and the venting of steam from the cavity
- assurance that the treatment of the lower head outside surface (painting, coatings, etc.) will not interfere with the cooling of the vessel by the water

To show that AP600 meets these conditions, Westinghouse will quantify failure probabilities for reactor and containment systems using defendable system-level PRA methods and to demonstrate the feasibility of the vessel insulation design. These are things that are done in the PRA or are needed to support IVR in any case. Once these conditions are established for AP600 and the IVR thermal-

hydraulic issues are resolved, then vessel failure into a flooded cavity is considered to be "physically unreasonable" for AP600. It is Westinghouse's position that no ex-vessel phenomena can occur, and therefore, no assessment of their impact on the containment integrity in a flooded configuration needs to be included on the AP600 docket. Any other treatment is seen as arbitrary based on speculative scenarios.

Summary

- The current severe accident issue regulatory environment results in inconsistent treatment of severe accident phenomena in an attempt to provide defense-in-depth to low-probability, high consequence accident sequences. This treatment results in having to apply probabilities and to perform analysis for phenomena that are difficult to resolve and defend.

- The ROAAM provides a frame for examining uncertainties, establishing safety goals and resolving severe accident issues by demonstrating whether or not the threat to the containment is "physically unreasonable".

- The ROAAM analysis of the IVR thermal-hydraulic issues concludes that vessel failure into a fully flooded cavity is physically unreasonable.

- It is the Westinghouse position that the IVR report provides a defendable technical basis for demonstrating vessel integrity and eliminating the consideration of ex-vessel phenomena related to debris relocation in the flooded configuration. Westinghouse will demonstrate that the conditions and assumptions in the IVR analysis are met in the AP600 design. This will be accomplished through defendable methods. This approach entails:

  - verification that the RCS depressurization system is reliable
  - verification that the cavity flooding system is reliable
  - demonstration that the reactor vessel reflective insulation is "IVR-friendly", and allows sufficient ingress of water and the venting of steam from the cavity
  - assurance that the treatment of the lower head outside surface (painting, coatings, etc.) will not interfere with the cooling of the vessel by the water.

References

1. SECY-93-087, "Policy, Technical, and Licensing Issues Pertaining to Evolutionary and Advanced Light-Water Reactor (ALWR) Designs," July 21, 1993.

2. Theofanous, T.G., C. Liu, S. Addition, S. Angelini, O. Kymalainen, T. Salmassi, In-Vessel Coolability and Retention of a Core Melt, DOE/ID-10460, July 1995.

3. Theofanous, T.G., J. H. Scobel, S. W. Sorrell, W. F. Pasedag, "Experience with Risk Analysis Methods in the Design and Certification of Advanced Passive Plants," American Nuclear Society PSA 95, Seoul, Korea, November 1995.

4. Additon, S. (Tenera), Informal Memorandum to Steven Sorrell (DOE ARSAP Program Director), April 28, 1995.

5. NRC Letter, Summary of Meeting to Discuss External Reactor Vessel Cooling for the Westinghouse AP600 Design, August 30, 1995.

6. Theofanous, T.G., "On the Proper Formulation of Safety Goals and Assessment of Safety Margins for Rare and High-Consequence Hazards," to be published in Reliability Engineering and Systems Safety.

7. Theofanous, T.G., W. H. Amarasooriya, H. Yan, U. Ratnam, "The Probability of Liner Failure in a Mark-I Containment," NUREG/CR-5423, August 1991.

8. Theofanous, T. G., H. Yan, M. Z. Podowski, C. S. Cho, D. A. Powers, T. J. Heames, J. J. Sienicki, C. C. Chu, B. W. Spenser, J. C. Castro, Y. R. Rashad, R. A. Dameron, J. S. Maxwell, D. A. Powers, "The Probability of Mark-I Containment Failure by Melt-Attack of the Liner," NUREG/CR-6025, November 1993.

9. Pilch, M.M, H. Yan, T.G. Theofanous, "The Probability of Containment Failure by Direct Containment Heating in Zion," NUREG/CR-6075, SAND93-1535, December 1994.

10. Pilch, M. M., M.D. Allen, D.L. Knudson, D.W. Stamps, E.L. Tadios, "The Probability of Containment Failure by Direct Containment Heating in Zion," NUREG/CR-6075, SAND93-1535, Supp. 1, December 1994.

Figure 1 - Simplified CET with PRA Credit for IVR

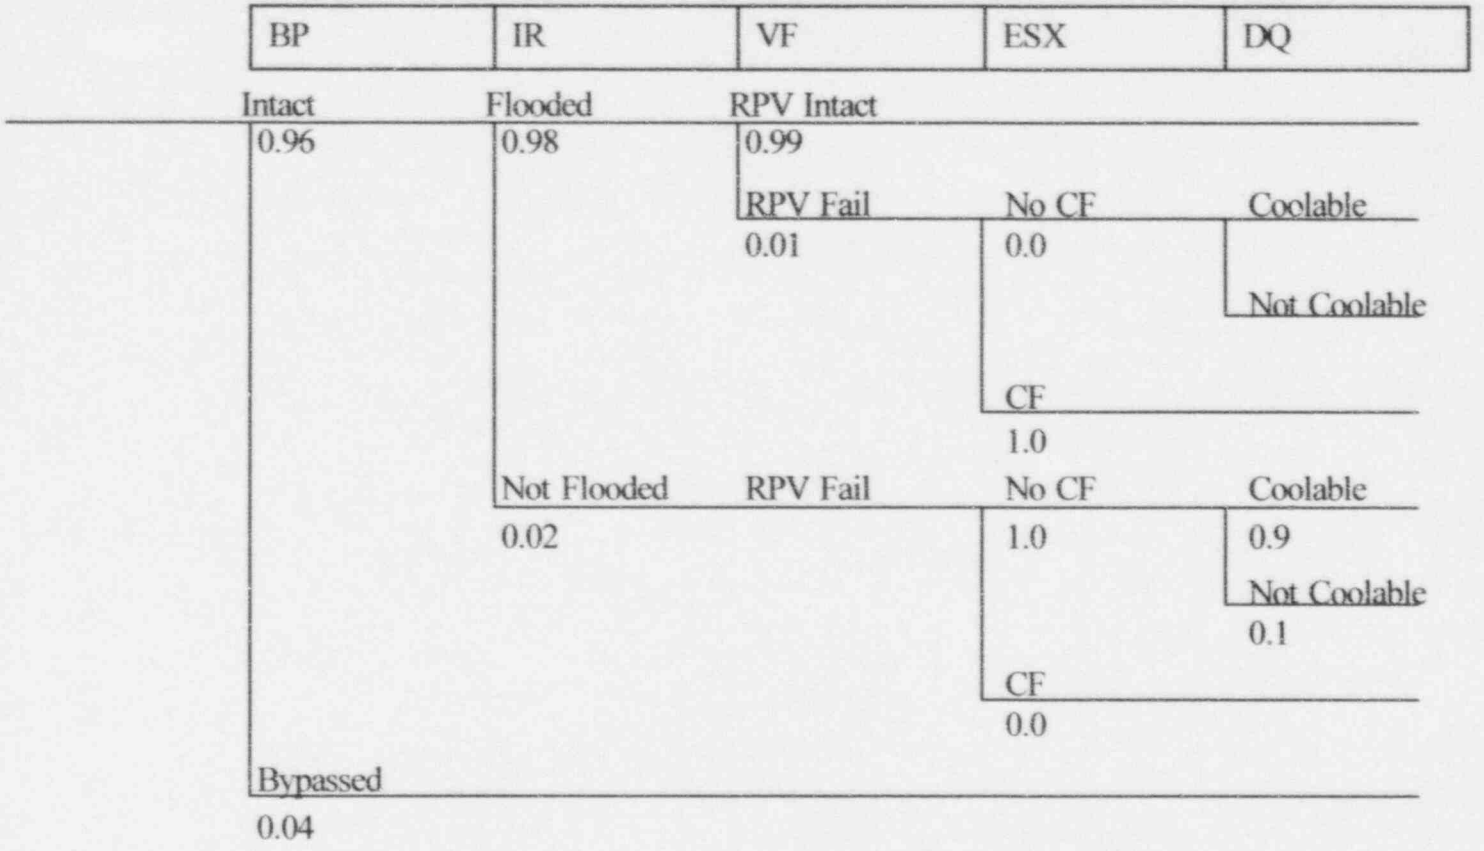

1	IC	0.9314
2	IC	0.0
3	CCI	0.0
4	CF	9.41E-3
5	IC	1.73E-2
6	CCI	1.93E-3
7	CF	0.0
8	BP	0.04

End State	Conditional Probability
IC	9.49E-1
BP	4.00E-2
CF	9.41E-3
CCI	1.93E-3

CCFP = 5.1E-2

Figure 2 - Simplified CET with Low Credit for IVR

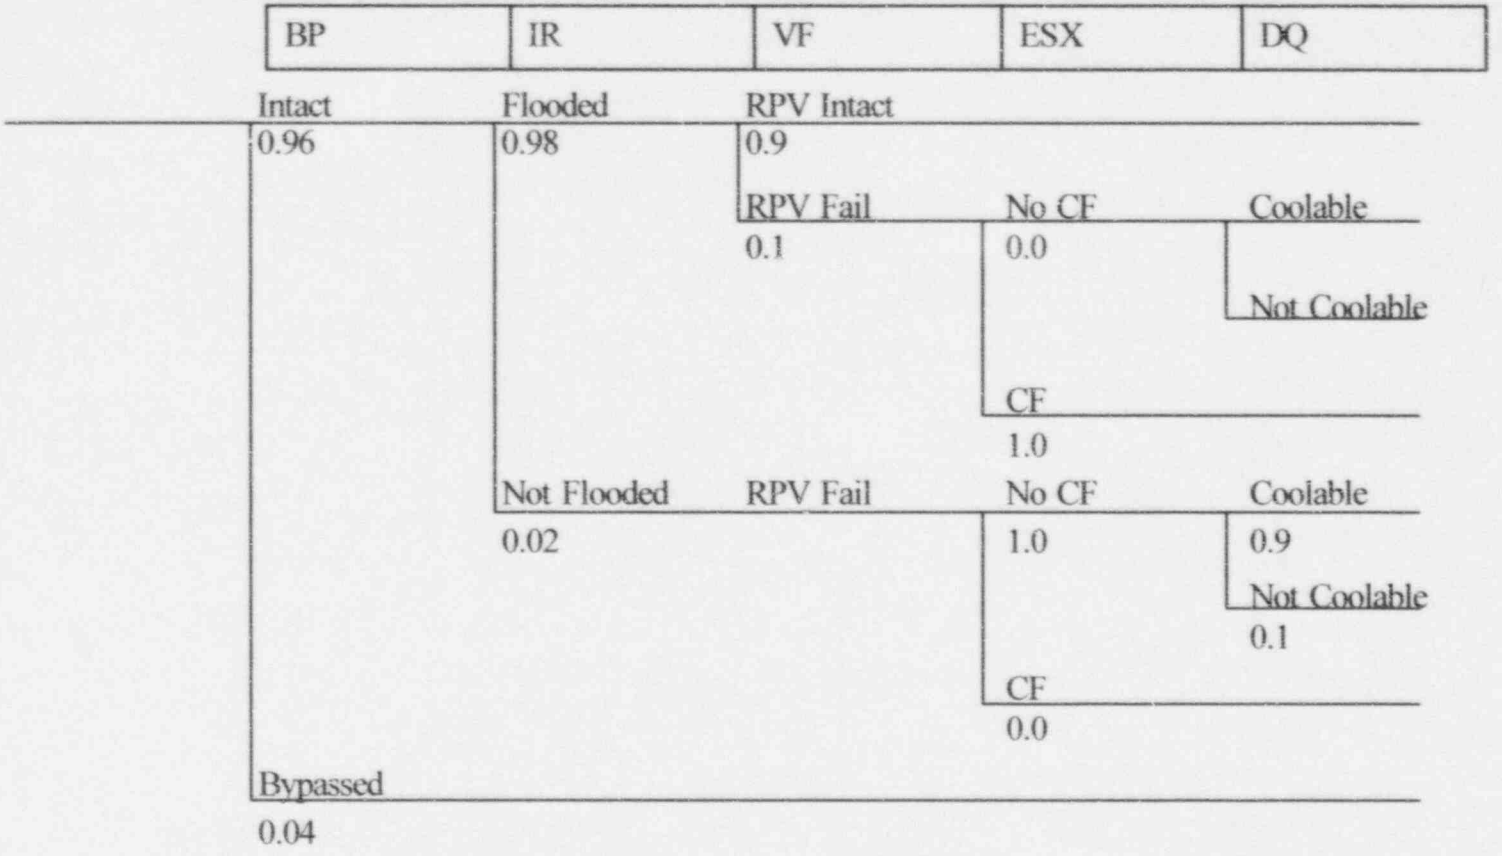

1	IC	0.847
2	IC	0.0
3	CCI	0.0
4	CF	9.41E-2
5	IC	1.73E-2
6	CCI	1.93E-3
7	CF	0.0
8	BP	0.04

End State	Conditional Probability
IC	8.64E-1
BP	4.00E-2
CF	9.41E-2
CCI	1.93E-3

CCFP = 1.4E-1

Enclosure 2 to Westinghouse
Letter NTD-NRC-96-4628

January 23, 1996

DRAFT

Invited for a special issue on "Aleatory & Systemic Uncertainty in Performance Assessment of Complex Systems" in **Reliability Engineering & Systems Safety**

ON THE PROPER FORMULATION OF SAFETY GOALS AND ASSESSMENT OF SAFETY MARGINS FOR RARE AND HIGH-CONSEQUENCE HAZARDS

by

T.G. Theofanous

ABSTRACT

The issue of "uncertainty" is addressed in the special context of assessing and managing risks from rare, high-consequence hazards. It is suggested that rather than the usual "formal treatments" on how to combine expert opinions that diverge widely, such "uncertainty" must be approached in each case as a research question that encompasses frame of assessment, approach methodology, risk management, and safety goals, with the aim of obtaining resolution in a clear, consistent, and complete manner. This, together with some basic considerations on "defense-in-depth," and certain practical aspects of communications and synergism needed for resolution (of such uncertainties), leads us to the Risk Oriented Accident Analysis Methodology (ROAAM). The purpose of this paper is to explain these views, to follow them through to the definition of the methodology and its implementation, and to indicate some of the insights gained through the several practical applications available so far.

1. INTRODUCTION

This year is the 20th anniversary of the publication of the Reactor Safety Study (RSS), the first fully integrated assessment of risk for a complex technological system. The study had a profound impact on nuclear technology but also well beyond, in every facet of human activity involving potential hazards, as "risk assessment" became a field of inquiry in its own right. While still young as a field, risk assessment is already becoming a dominant force in technology and environment and through them in economics, politics, and everyday life—indeed, as some issues are of global significance, we can visualize a potential impact on our civilization as a whole.

As the public is becoming increasingly aware of the concept of risk and, roughly, of its measure (chance, or probability for a given hazard), and as politicians and regulators are increasingly willing to take it up as a regulatory measure (as in "Risk-Based Regulation"), experts within the field are still debating the meaning of probability, while experts on the outside doubt its very veracity. Actually, the methodology itself (and underlying philosophy) evolved only slightly, as applications (actual

studies) mushroomed both in numbers (e.g., for nuclear reactors) and in topical areas (e.g., chemical hazards in industrial and environmental settings), not always with high quality results. Central to this question of quality is the proper recognition and treatment of **uncertainty**.

The problem has been recognized from the very beginning, as a criticism of the RSS (Lewis et al., 1978). It has been pursued methodologically (Kaplan and Garrick, 1981) and in the major subsequent applications (Oyster Creek, 1982; Zion, 1981; and Indian point, 1982), but it was most effectively dramatized in a revisit to the RSS about a decade later (NUREG-1150, 1990). A good reflection of the current status and the many differing points of view on it can be found in a recent workshop held especially for this purpose (NUREG/CP-0138, 1994).

At this point, we find ourselves in a rather paradoxical situation that, while risk is widely recognized as **the** preeminent, and perfectly rational, concept in addressing hazards at all scales (from the personal to the global), its utility is found to be severely hampered, especially in those situations that matter the most—those that involve rare, high-consequence (at the society level) hazards. I believe that the difficulty is, to a large extent, self-inflicted; the purpose of this paper is to explain this view and suggest an approach towards overcoming it.

My central theme is that while established risk assessment methodology has been very effective (at least in nuclear power reactor applications) in revealing and prioritizing uncertainties, it has paid insufficient attention to resolving them; indeed, it may not even be ideally equipped for this purpose. As long as they exist such major uncertainties are a real impediment to the clear interpretation of results, and hence to their utility. Worse, they may even (as they have) cast doubt on the whole enterprise; that is, the methodology itself on the one end, and the technology in question on the other. My suggestion towards a more effective approach to resolution involves two intertwining parts: safety goals and methodology of evaluation. I argue that the two must be considered **together** and be specific to the particular class of applications. The goals must be clear and communicable to the public, **expressed in a manner consistent with the nature of the uncertainties involved**, and they must entail a clear and robust sense of closure—clear to recognize when closure has been achieved, and robust enough to allow the unequivocal acceptance of a final result that meets the goal. The methodology is relatively recent (Theofanous, 1994), and it has been employed, so far, in the assessment and management of severe accidents in nuclear power reactors. There is an effort made here to generalize; however, another main message is that every resolution effort requires specificity and focus, and these should take precedence, especially in any new class of applications.

The presentation begins with a discussion of the nature of the "uncertainty" problem (Section 2). We quickly bring the focus to our main present interest, which is uncertainty in predicting sequences

of complex phenomena in the context of assessing rare and high-consequence hazards. Some ideas on the deterministic and stochastic natures of such complex physical processes are presented, and in this light some questions on current representation of epistemic uncertainty (knowledge, phenomenological, or modelling, etc.) are raised. For rare events the nature and magnitude of uncertainty must be synthesized, usually from incomplete evidence. In the PSA frame it is obtained through an expert elicitation process, and it is known as **epistemic uncertainty**. We argue that the presence of significant epistemic uncertainty requires very special consideration and care, so as not to mask potential impacts. We also argue that if it is significant ("issues") special efforts (outside the normal PSA frame) are required for resolution.

This brings us to the goals (Section 3), and the idea that when epistemic uncertainty is significant, safety goals can be only qualitatively defined. Together with the risk aversion that characterizes high consequence hazards, this implies the need for a goal that effectively communicates the idea that in a sense the perceived hazard "can't happen." This may appear, at first, contrary to the underlying principles of risk analysis; however, it is amenable to a rational deduction that marries the probabilistic (risk) and deterministic approaches, through the concepts of "residual risk," "screening frequency," and "defence in depth." For both the goal and methodology we require, as basic principles, **clarity**, **consistency**, and **completeness**. In layman's, or primitive terms, this goal is expressed as: the perceived hazard is "**physically unreasonable**" under "**any circumstances**" leading up to it in a "**physically meaningful**" context.

The methodology (Section 4) provides a working interpretation of these terms, and an approach for implementation. The key aspects of it are a physically-based decomposition that allows transparency in separating out the essential portions of epistemic uncertainty, and the use of multiple scenarios in combination with the conservative estimates of epistemic uncertainty, so as to obtain convincingly conservative results. The purpose, of course, is to show compliance with the goal, and to thus achieve closure; or, in case the goal is not met, to determine risk management actions (procedures, hardware changes) necessary to achieve this purpose.

As a matter of high level perspective, let us say at this point, that as much as PSA emphasizes comprehensive, system-level treatment, the present methodology emphasizes focus on controlling physics that underlie specific issues; this is done at a much greater level of detail than can be afforded within the normal scope of a PSA. Put another way, while PSA systematically evaluates all potential outcomes from given sets of initiators, the present methodology begins from the failure mode (or consequence) that is to be prevented and systematically seeks to eliminate all physically meaningful causal paths that could lead to it—thus we have found it to be a very convenient and effective tool for accident management purposes. Finally, the two are really complementary, for once an issue is resolved it can be reflected back to a PSA, while, as we will see further below, the

present methodology requires consideration and use of system status information (as determined in a Level 1 PSA, for example) from accident initiator and equipment availability standpoints.

Implementation (Section 5) carries several special provisions which are an integral, and crucial, part of the methodology. These provisions include a comprehensive review effort that involves essentially all experts in the field, through an iterative and fully documented process towards resolution. Other key provisions include guarding against premature closure, and some guidelines in communicating an overall measure on the inherent quality of evaluation. But the real nature of implementation, its quality, and indeed a complete understanding of the methodology itself, can only be appreciated through the complete documentation available from each actual application. An overview of these, with appropriate citations, is given in the next and last paragraph of this introduction. Some global perspectives on the overall process in each case, the timing, and the number of experts involved are provided in Section 6. As an illustration the probabilistic framework from one of these applications is provided in the appendix.

The methodology is known as the Risk Oriented Accident Analysis Methodology or ROAAM (Theofanous, 1994), and it has been successfully employed in the resolution of the major severe accident issues relating to early containment failure; namely, Mark-I Liner Attack (Theofanous et al., 1991, 1993) and Direct Containment Heating (Pilch et al., 1994a, 1994b). In its earliest version ROAAM was also applied and made decisive steps towards the resolution of the α-mode (steam explosion induced) containment failure (Theofanous et al., 1987; Theofanous and Yuen, 1994; Turland et al., 1994). Currently, another application to In-Vessel Retention as a severe accident management scheme in Advanced Light Water Reactors (the AP600 in particular), is near completion (Theofanous et al., 1995b), and a complementary effort addressing lower head integrity under steam explosion loads is in the offing (Theofanous et al., 1995a). Finally, ROAAM played a key role in design, assessment, and implementation of severe accident management in the Loviisa NPP in Finland (Tuomisto and Theofanous, 1994, 1995). The underlying philosophical bases of the methodology evolved significantly through these experiences, and particularly so in the cases of Loviisa (Tuomisto and Theofanous, 1995) and AP600 (Theofanous et al., 1995c).

2. THE NATURE OF THE "UNCERTAINTY" PROBLEM

The need for risk analysis begins with the perception of a potential hazard, and the goal of those that practice it is to provide a quantitative measure of likelihood and the magnitude of corresponding impacts. These results can be used to make decisions such as: requiring enhanced safety measures and/or emergency planning, accepting or rejecting a certain technological endeavor, or selecting among various alternatives. There is a lot of judgment involved in making such decisions, but those that make them (i.e., the public and politicians) are least qualified to exercise such judgment—

they are inherently limited from appreciating the key technical nuances imbedded in the results expressing these likelihoods. For hazards that actually materialize (observable frequencies) the situation is self-correcting, attaining a long-term, mutually acceptable equilibrium. This is not the case for rare, high magnitude hazards, and this is the crux of the matter conveniently expressed by the collective term "uncertainty." Clearly, there is not much opportunity for an approach to equilibrium here, empirical evidence is inherently lacking, and even the concept of "likelihood" may not be open to straightforward interpretation. Communication gaps now develop even among experts, and the usual gaps with the public now become real chasms. **In this very special but critical context we wish to address "uncertainty," here.**

The level of representation of physics (degree of decomposition in the model) and level of understanding (and verification) that goes with it, dictates the nature and magnitude of uncertainty in prediction. For rare events, absence of direct evidence requires that the nature and magnitude of uncertainty must be synthesized on the basis of partial, and relevant to varying degrees, information. In PSA frame this is accomplished through an expert elicitation process, and it is known as epistemic uncertainty (for those that render it are επιστημῶνες, Greek for scientist or expert). Beyond epistemic uncertainty there are inherent limits to prediction when physical instabilities are involved, e.g., bifurcations, chaos, threshold behavior such as failures. Such behavior can be specified only on the basis of data (statistical inference), and it is known as aleatory or stochastic uncertainty. Further, this quantification is subject to epistemic uncertainty. Although efforts are made to make epistemic uncertainty coherent with available information, in the final analysis it is subjective. An important contribution of NUREG-1150 is that it demonstrated that subjectivity can lead to major incoherencies among experts, even though the knowledge base is the same. It led to a plethora of formal methods to combine and reflect such widely varying expert inputs, but perhaps not unexpectedly the "issues" persisted. Several important questions emerge from this experience.

(a) Since not every PSA can take on the scope of NUREG-1150, how do you identify such areas where evidence is insufficient to allow a reliable and coherent estimation?

(b) Since lack of knowledge can be expressed in a way to maximize uncertainty in the input (this is the preferred PSA way—i.e., maximum entropy principle), and not only as optimistic or pessimistic renderings (this happens quite often, too—see 1150), how do you prevent masking out important areas of uncertainty?

(c) And last but not least, how do you resolve "issues" sufficiently well as to allow the intended utility of results?

It is our view that addressing these questions requires very special and intense efforts. The most difficult is question (a). In a well-explored field, as is that of nuclear risks, we can rely on

the ground-breaking studies notes in the introduction and efforts such as the Containment Loads Working Group (NUREG/CR-xxxx). In other fields, depending on the magnitude of the potential hazards, one may require studies of a similarly major scope. Still, there are concerns of applicability, and the potential for surprise in any new application. Ultimately, we have to rely on real expertise at the regulatory level, while pitfalls due to widespread availability and use of PSA tools (systems level codes, PSA software) entails major pitfalls for the originators (of the PSA) and the reviewers (the regulatory) alike. Question (b) can be addressed by appropriate sensitivity studies within the PSA frame, but care and deep experience are crucial, too. Finally, for question (c), we propose special efforts, as described in the following sections.

3. ON SAFETY GOALS

We begin with the philosophical position that **when epistemic uncertainty is significant** safety goals can be only qualitatively defined. Really, this is a consequence of the subjective nature of the evaluation (to be performed against the goal), where the primitive variable is "meaning" (or idea) rather than "numbers." Then, for a high-consequence hazard, this inherent "softness" requires that we take measures that in effect ensure that it "can't happen." For example, in the case of severe nuclear accidents, this "can't happen" refers to containment failure.

Parenthetically, compare this to current NRC goals, specifying a conditional containment failure probability of 10^{-1}. It is rather easy to see that in the presence of significant epistemic uncertainty, such a goal is not very meaningful, nor does it provide an adequate defense in depth. Smaller numerical goals could imply better defense in depth, but they would imply increasingly more dubious (in both scrutability and communicability) evaluations.

Thus, we envision a marriage of probabilistic and deterministic approaches, which is possible through the concepts of "residual risk," "screening frequency," and "defense in depth." Further, the goals must be **clear**, imply **completeness**, and be subject to **consistency**. In layman's, or primitive, terms this goal can be expressed as: the perceived hazard is "**physically unreasonable**" under "**any circumstances**" leading up to it in a "**physically meaningful context**."

The methodology described in the next two sections provides a working interpretation of these terms, and an approach to implementation. As noted above, goals and methodology must be very closely intertwined.

4. THE RISK ORIENTED ACCIDENT ANALYSIS METHODOOGY (ROAAM)

To begin with, we must emphasize again that for present purposes we restrict attention to high consequence hazardous situations, whose likelihood has been made extremely low by preventive

measures. Besides the basic physics, which normally can be considered known (i.e., as belonging within the design basis domain), the success to prevent depends to the greatest extent on equipment and human reliability, including all sorts of systems' interactions. Failures outside this prevention envelope have the potential to create the perceived hazard(s), they are expected to occur at a certain finite (non-zero) rate, and it is the hazard that we wish to mitigate in a defense-in-depth approach. As stated above, we wish to show that this mitigation is effectively accomplished, or if not, to determine the appropriate factors available at our disposal so that this be the case. The central thrust of our approach is that at this stage we wish to place ultimate reliance on the basic laws of physics in a simple, clear, and demonstrable manner.

Let us characterize each major class of such rare realizations of "failure to prevent" by the corresponding set of equipment/human failures; that is, the system damage states symbolically expressed as D_i, and the corresponding predicted rates, or frequencies, as $f_i(D_i)$ (i takes on integral values in a finite, normally small interval). By definition, the f_is are smaller than some target value, f_t, achieved as the prevention goal, and we are interested in the subset bounded at the other end by some screening frequency, f_s, below which the failure likelihood may be considered negligible ("residual risk"). That is, we are interested in each major damage state D_i, for which $f_t > f_i(D_i) > f_s$. The screening frequency is an essential concept in establishing **consistency** between the physics-based approach adopted here and the probabilistic treatment of the domain to which it is to be applied. The consistency derives from the idea that normal design procedures ensure that events with a frequency below f_s are in a sense physically unreasonable. In other words, f_s is an inherent cornerstone in the design, and the consideration of failure states as described above is to ensure **a similar level of protection**. For example, in a nuclear setting f_s would describe the estimated rupture frequency of the reactor pressure vessel. Note that as $f_t \rightarrow f_s$ the $\{D_i\}$ approaches a null set, and in effect we obtain a deterministic design. In this case defense-in-depth is replaced by well-established design margins. Again in the nuclear setting, this does not seem to be the case even for the advanced passive plants, **hence the need for defense-in-depth for severe accidents at the containment level**. Moreover, considering uncertainties in estimating f_t for a complex system it is really doubtful whether this can ever be reasonably expected. To the extent that, as noted above, such damage states are dominated by equipment/human reliability, the uncertainties are basically stochastic—that is, they are subject to statistical inference. A quality treatment requires that they be supported by appropriate/applicable data bases, and that estimation of system response includes thorough consideration of common cause failures. These plant damage states describe the "**any circumstances**" mentioned above, or the general setting for consideration of the efficacy of the mitigation measures.

Each such circumstance gives rise to a physical evolution, or a sequence of events that in general is impossible to predict in its entirety and all detail. **Completeness** requires that all such sequences that cannot be excluded on physical grounds should be considered. Thus, the "**physically meaningful context**" mentioned above refers to the exclusion of scenarios found to be otherwise, in a process that while allowing significant reduction in the space of evaluation, it maintains a clearly understood enveloping character. This methodology of reduction differs crucially from assigning relative likelihoods, or weights, to parameters that affect scenarios through a system model, or to the scenarios themselves. It constitutes the crux of our approach in addressing uncertainty in this physically-oriented aspect of the evaluation, and it arises naturally, as described in the following.

The starting point is our interest in knowing the likelihood of realizing a certain hazard, H_k, given a set of system damage states $\{D_i\}$. Formally, this likelihood can be written as a function of the system response, and the various parameters that characterize it, as

$$L_i(H_k) = \mathcal{G}_i(p_1, p_2, \ldots p_\ell) \quad \text{given} \quad D_i \tag{1}$$

The usual approach is to approximate this function by a mathematical model, the so-called system or integral model, and uncertainty is to reflect how good this approximation is. Much of the effort actually goes to determine parameter uncertainty. Experience shows that there are grave difficulties with this approach, not the least of which is assessing model uncertainty. In fact, often enough the key physics are relegated to input parameters, to be determined by one or more individuals, depending on the resources available in any given application. In the presence of significant epistemic uncertainty, the end results lack transparency and are inherently non-convincing, as revealed by actual experience so far.

In ROAAM, the approach is quite orthogonal to that described above. First, we take the position that it is rather futile, and in fact unnecessary, to realize a **defensible** approximation to functions $\mathcal{G}_i$. Accordingly, we alter the goal to a much less demanding one—rather than seeking the likelihood L_i, we wish to establish that it is (or can be made by appropriate choices) low enough as to regard the hazard H_k as **physically unreasonable**. To the extent that conservative treatments are always possible, this may appear as a trivial undertaking. In fact, the essence of the approach is to find the middle ground of avoiding excess conservatism while still remaining convincing. Our experience (see references in the introduction) is that this is feasible and effective. Second, we wish to make a clear separation between aspects of the systems response that can be stated as well-posed physical problems (these are really the **controlling** physics), and other aspects that are subject to inherently variable behavior. The former we call "causal relations," the latter "intangibles." The structure of separation and hence the manner according to which these parts are to be combined

is called "probabilistic framework." Each framework refers to a particular "scenario," S_j, and the art in the decomposition is to envelop the behavior through the coherent use of "intangibles" and respective "scenarios." This is not a task for a computer. It has to be created by human intelligence and must be understandable (and scrutable) to human intelligence as well.

Each "causal relation" requires an in-depth and demonstrable understanding of the controlling physics; "scenarios" and "intangibles" are to fill in the gaps whenever this is not possible. Uncertainty in causal relations is well-behaved and can be made relatively small (this is a consequence of the "well-posedness"). Uncertainty in intangibles can only be qualitatively approached, but it can always be bounded. The adequacy of scenarios can be determined according to the completeness of the logical structures used in deriving them. The process of integration through the probabilistic framework is effected by introducing a scale (Table 1) for the **temporary** quantification of intangibles, and the results are rendered in qualitative terms by applying this scale in reverse. We emphasize that this scale is completely arbitrary except for the fact that it provides the definition of a physically unreasonable process as one involving the independent combination of an end-of-spectrum with one expected to be outside but cannot be positively excluded. Finally, our task can be stated as showing that

$$P_{ij}(H_k) < P_s \quad \text{given} \quad \{D_i\} \quad \text{for all} \quad \{S_{ij}\} \tag{2}$$

where S_{ij} is the jth scenario in the set belonging to the ith damage state. Because these scenarios are to independently satisfy Eq. (2), they are called "splinter scenarios." The "probabilities" are computed from the probabilistic framework, representing a map of the parameters in the causal relations, $\{d_i\}$, and the intangibles $\{i_i\}$; that is,

$$P_{ij}(H_k) = \mathcal{F}(d_1, d_2, \ldots i_1, i_2, \ldots) \tag{3}$$

The $\{i_i\}$ and the uncertainties in $\{d_i\}$ are, as distributions (pdf's), quantified according to the probability scale in Table 1. The P_s corresponds to the physically unreasonable level in Table 1.

It may be interesting to note that the above structure in effect separates out epistemic (from aleatory) uncertainty. This is intentional, and indeed motivated by the distinct manner necessary to judge residual risk, as explained already above; that is, screening frequency for aleatory, and the physically unreasonable concept for epistemic. Any stochastic behavior not already included in the definition of the severe accident window (the plant damage states to be considered) can be taken up in the definition of scenarios and intangibles, since they would be expected to dominate the uncertainty in any case. If necessary, however, stochastic parameters, or even processes, can appear explicitly in Equation 3. A similar separation can be effected in this case, too, by simply

Table 1. Definition of Probability Levels	
Process Likelihood	Process Characteristics
1/10	Behavior is within known trends but obtainable only at the edge-of-spectrum parameters.
1/100	Behavior cannot be positively excluded, but it is outside the spectrum of reason.
1/1000	Behavior is **physically unreasonable** and violates well-known reality. Its occurrence can be argued against positively.

finding the total probability in each frequency range, and applying the same criteria for judging the results—but now these frequencies should be combined with the respective plant damage state frequencies.

5. ROAAM IMPLEMENTATION

From the description above it should be clear already that to a large extent the essence of ROAAM lies in implementation. The basic premise is that once the **whole** community of experts in a given problem area is convinced (that the demonstration has been effected), the problem may be considered solved, in a robust and final way. A key aspect of the implementation, therefore, is to involve, to the extent possible, all the experts. Procedurally, this is effected by making available resources for review and interaction with those who lead the effort (technically). There is a **process** involved here that takes time, usually one or more years. This process must remain traceable and scrutable, and open. In particular, the reviews are eponymous (signed by individual experts), and they are included together with the responses in the publication. This individual-based approach also protects to a large extent against organization interests, which can be detracting, even if not intentionally, which is not always the case. Still, constant care in this area is necessary. Eventually the complete reaching of all experts is effected by publication in the technical literature, with additional iterations thereof if necessary.

For the whole process to succeed the originating document must be basically sound. This is because the creative part of the work involves a close interaction between analysis and synthesis,

which is amenable to a close-knit group, but not to a panel or committee setting. Even when this is the case, questions abound. However, the decomposition afforded by the probabilistic framework allows transparency, manageability, and focus. It also allows finding and involving a much larger set of experts, especially in the examination of causal relations, from respective fields. In this manner the effort acquires a strong fundamental thrust. This provides for enrichment, refinement, synergism, and eventually for convergence in quantification. For resolution, we must have agreement on the framework and convergence on each of its parts. This eliminates the highly dubious but not uncommon situation where experts agree on a conclusion by using widely varying arguments.

By quantification we mean the assignment of probability distributions to the $\{d_i\}$ and $\{i_i\}$ sets for each scenario S_{ij}. By synthesis, we mean the combination of these numbers according to the rules that define the respective probabilistic framework. The arithmetic can be done using either discrete probability distributions or Monte Carlo (the choice is a matter of convenience, and depends upon the number of variables and the structure of the probabilistic framework). For each damage state and each scenario we have one result which must be expressed in qualitative terms using the scale shown in Table 1. Each quantification is unique but amenable to adjustments as appropriate through the review process. In addition to this Base Quantification, supplementary quantifications are conducted as sensitivity and parametric calculations. They involve primarily changes in the distributions of the intangibles. These are ad hoc and aimed primarily at showing the margins to negating the goals expressed in Eq. (2). These are helpful in the originating document in providing some perspective, and as a rule further are requested by the reviewers for additional depth in this perspective. It is important, however, that the base quantification remains as the primary focus.

The quantification is to be supported by demonstrably-applicable evidence. For the causal relations this means properly scaled experiments and/or appropriate theories. Definitions of scenarios and quantifications of intangibles are strongly dependent on logical structures and technical reasoning at a basic principle level. Clearly the exercise of judgment is inevitable here, but robustness is gained by the conservative attitude that is to guide these judgments. The ultimate safeguard is in wide and coherent expert participation and synergism. Experience has shown that a great deal of the existence of an "issue" lies in confusion created by partial approaches to the problem and miscommunication. Tackling both is in the heart of the ROAAM implementation—the process has to take its course and the frame must be complete. Even before producing the final result this process is very useful in focusing the issues and maximizing effectiveness in addressing them—including appropriate research efforts. Indeed, this can be the first step in coming up with a focused and rational research plan. For communication purposes it is expedient to explicitly indicate the position of an issue on the "learning (or resolution) curve." This can be constructed, in form of Table 2, such

as to also safeguard against a premature closure. The initiating ROAAM document will normally span Phases II and III, while the ROAAM review process can take a significant portion of Phase IV. The actual completion of Phase IV would take place outside and after resolution has been achieved, while at the same time being reinforced by other related applications in Phase V. That is, Phases IV and V can be regarded as confirmatory. Finally, the need for Phase VI can only be judged on a case-by-case basis, especially with regard to "completeness." This is also confirmatory and defines a defense-in-depth in evaluation.

Table 2. Definition of "Phases of Development"

Phase I	Scoping	Simple Models, Parametric Results
Phase II	Analytical	Detailed Models, Pre-experimental
Phase III	Quantitative	Experiments, Scaling, Model Comparisons
Phase IV	Maturation	Verified Results from Multiple Origins
Phase V	Extension	Other Related Applications
Phase VI	Defense-in-Depth	Further Exploration of Intangibles

Finally, as another key aspect of properly communicating the results, we need to recognize the qualitative differences due to scenario dependence. This also provides for deeper compensation in quality of evaluation and the whole process as scenario dependence increases. Thus, we define three quality grades as summarized in Table 3. Note that Grade A corresponds to what is understood as the "deterministic approach." It is normally found to be impractical, or in converse, if it can be employed the question of "uncertainty" becomes mute. Also note that Table 3 does not include a grade based on integral system "simulations" (as in standard PSA Level II for severe nuclear reactor accidents), as this is not considered to be acceptable, **in the present context**.

6. PERSPECTIVES ON APPLICATIONS

The first application of ROAAM was on the α-mode containment failure (Theofanous et al., 1987). The work pointed to the key areas of premixing and lower head failure (as a mitigative measure), which have become major subjects of investigation, internationally, in the past few years. Further advancement towards closure were made by Theofanous and Yuen (1994), and specifically for the Sizewell plant in the UK, again within the ROAAM frame, by Turland et al. (1994). Orderly closure is within reach now, by continuing the ROAAM process to incorporate the new research results.

	Table 3. Definition of Quality Grades
Grade A	Framework characterized by a single, limiting process, evaluated on basic physical laws, with appropriate bounding inputs. No scenario dependence.
Grade B	Framework involves a single complex process evaluated at a high confidence level. There may be slight scenario dependence compensated by appropriate quantification of intangibles.
Grade C	Framework involves sequence of processes. Significant scenario dependence compensated by appropriate choice and quantification of intangibles.

The next two applications were on **the two major issues** in the NRC's Severe Accident Research plan: the "Mark-I Liner Attack" and the "Direct Containment Heating" (DCH) problems. In the Mark-I case the original document (Theofanous et al., 1991) proved to be quite controversial among the experts, in a manifestation of the preexisting polarization (as evidence on the NUREG-1150 elicitation). Follow-up work (audits) was carried out at RPI (on melt release scenarios), ANL (on spreading), SNL (on corium concrete interactions), and ANATECH (on liner structural failure, by creep), and the final document, with the complete convergence, was published a little more than a year later (Theofanous et al., 1993). There were 19 experts involved in this work. By way of illustration, the problem and the framework are summarized in the Appendix. On the DCH, the original document (Pilch et al., 1994) was again met with extensive critical comments (but many favorable as well). Further work over a period of about a year, however, produced convergence and resolution (Pilch et al., 1994). This work also proposed a basis for extrapolation of the results to all PWRs with broadly-similar Zion-like features. A subsequent study of application to Surry went through the review rather quickly with positive results. There were 17 experts involved in this work. In both, the Mark-I and DCH, splintering, in combination with conservative evaluation of the intangibles, played a key role in bringing the experts together.

Currently, ROAAM is being applied to the advanced passive plants for DOE's ARSAP program. A severe accident management scheme proposed for the AP600 (and incidentally for the Loviisa plant in Finland, too) is to externally flood the reactor vessel, and thus prevent the melt from melting through the lower head. The approach is too new to call it an issue yet, but the reviewers' comments to the initiating document (Theofanous et al., 1995b) indicate that in effect it is—there were about 300 questions from 17 reviewers, and the process so far has taken about one year. It is expected

that resolution will be obtained in the very near future. Currently under preparation also are an application to lower head integrity under steam explosion loads, and a little later an application to the structural integrity at the boundary of the lower drywell in SBWR (Theofanous et al., 1995a).

There is substantial experience with ROAAM at this point to allow us important insights on severe accident issue resolution, and methodologically perhaps even more broadly. The most important lesson is that our experience confirms that issues persist because of two methodological missing links: lack of a "basis" for fruitful and synergistic collaboration among the experts on an issue, and a lack of a "frame" to facilitate such communication. Also, we find that the naturally present somewhat competitive/adversarial atmosphere is actually helpful in the search for areas of difficulty with the proposed solution. This helps immensely on the "completeness" aspect of the work. There is some art involved, as well, in identifying the key physics and finding effective means of representing them. This requires flexibility, and certainly is contrary to generalized approaches that have as central elements large systems simulation codes. The idea is to not encumber the solution with any unnecessary "baggage," but rather begin and remain focused on the relevant physics. The decomposition affords the focused and reliable understanding of these physics by special purpose tools (analyses, experiments, or combinations), and the most effective utilization of external expertise, that usually is available in narrow aspects of the problem.

Finally, we find that the focusing and synergism obtained by each such resolution exercise are valuable more generally. Visualized as growing islands of deeper knowledge, collectively they will eventually form a firm coherent basis for severe accident management and resolution at large.

7. CONCLUDING REMARKS

The "uncertainty" addressed here is one of very special interest in the societal/regulatory context. It can cause qualitative changes in the perceptions of risk, it gives rise to major "issues," and related decisions or indecisions are likely to produce huge impacts. The critical role of such uncertainties has been recognized in the past, but this recognition and related efforts to address it have produced confusion and delays. These efforts have been principally oriented to formal definitions and supposedly quantifications aiming toward improved utility of the results. Our position is that such efforts are actually misdirected—**as long as uncertainties remain significant**, any efforts to "quantify" are bound to remain futile.

Rather, we recommend an approach that is based on defense-in-depth, and is focused at reducing the uncertainty sufficiently to recognize that such a defense-in-depth has been achieved. This solution-oriented approach has been shown to work by being both efficient and robust. Also, this approach implies a radical change in paradigm of how problems are approached and how related

work is conducted in this specialized area of risk analysis. There are also major implications in the regulatory process, including safety goals and review procedures. It may be that areas other than nuclear power can be benefitted as well from this experience.

ACKNOWLEDGEMENTS

The development of ROAAM began under NRC support, evolved under DOE, and was further benefitted by work done for IVO International of Finland. The author is grateful for this support, as well as to all his co-workers, co-authors, and participating experts in the several applications listed under references. Finally, this work would not have been possible without the interest, cooperation, and in some cases collaboration (in chronological order) of H.S. Isbin, T.P. Speis, H. Tuomisto, E. Beckjord, B.R. Sehgal, W. Pasedag, S. Additon and S. Sorrell. Special thanks also go to J. Garrick and S. Kaplan, from whom the author learned what he knows about risk analysis.

REFERENCES

1. Containment Loads Working Group Report (xxxx).
2. "Indian Point Probabilistic Safety Study," (1982) Pickard, Lowe and Garrick, Inc., Westinghouse Electric Corporation, and Fauske & Associates, Inc., prepared for Consolidated Edison Company of New york, Inc., and the New York Power Authority, March 1982.
3. Kaplan, S. and B.J. Garrick (1981) "On the Quantitative Definition of Risk," *Risk Analysis*, **1**, 11–27.
4. Lewis, H.W. et al. (1978) "Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission," NUREG/CR-0400, September 1978.
5. NUREG-1150 (1990) "Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants," U.S. Nuclear Regulatory Commission.
6. NUREG/CP-0138 (1994) Proceedings, Workshop I in Advanced Topics in Reliability and risk Analysis, Annapolis, MD, October 20-22, 1993 (October 1994).
7. "Oyster Creek Probabilistic Safety Analysis,"(1982) Pickard, Lowe and Garrick, Inc., prepared for GPU Nuclear Corporation, December 1982.
8. Pilch, M.M., H. Yan and T.G. Theofanous (1994a) "The Probability of Containment Failure by Direct Containment Heating in Zion," NUREG/CR-6075, SAND93-1535, December 1994.
9. Pilch, M.M., M.D. Allen, D.L. Knudson, D.W. Stamps, and E.L. Tadios (1994b) "The Probability of Containment Failure by Direct Containment Heating in Zion," NUREG/CR-6075, SAND93-1535, Supp. 1, December 1994.

10. Theofanous, T.G., W.H. Amarasooriya, B. Najafi, M.A. Abolfadl, G.E. Lucas and E. Rumble (1989) "An Assessment of Steam-Explosion-Induced Containment Failure," NUREG/CR-5030, February 1989. [Also Theofanous, T.G., B. Najafi and E. Rumble (1987) "An Assessment of Steam-Explosion-Induced Containment Failure. Part I: Probabilistic Aspects," *Nuclear Science and Engineering*, **97**, 259-281; Abolfadl, M.A. and T.G. Theofanous (1987) "An Assessment of Steam-Explosion-Induced Containment Failure. Part II: Premixing Limits," *Nuclear Science and Engineering*, **97**, 282-295; Amarasooriya, W. H. and T.G. Theofanous, "An Assessment of Steam-Explosion-Induced Containment Failure. Part III: Expansion and Energy Partition," *Nuclear Science and Engineering*, **97**, 296-315; Lucas, G.E., W.H. Amarasooriya and T.G. Theofanous (1987) "An Assessment of Steam-Explosion-Induced Containment Failure. Part IV: Impact Mechanics, Dissipation and Vessel Head Failure," *Nuclear Science and Engineering*, **97**, 316-326.]

110. Theofanous, T.G., W.H. Amarasooriya, H. Yan and U. Ratnam (1991) "The Probability of Liner Failure in a Mark-I Containment," NUREG/CR-5423, August 1991. [Also Theofanous, T.G., H. Yan and F. Eltawila (1993) "The Probability of Liner Failure in a Mark-I Containment, Part I: Probabilistic Framework and Results," *Nuclear Technology* **101** No. 3, 299-331.

12. Theofanous, T.G., H. Yan/UCSB; M.Z. Podowski, C.S. Cho/RPI; D.A. Powers, T.J. Heames/SNL; J.J. Sienicki, C.C. Chu, B.W. Spencer/ANL; J.C. Castro, Y.R. Rashid, R.A. Dameron, J.S. Maxwell, D.A. Powers/ANATECH (1993)"The Probability of Mark-I Containment Failure by Melt-Attack of the Liner," NUREG/CR-6025, November 1993.

13. Theofanous,T.G. (1994) "Dealing with Phenomenological Uncertainty in Risk Analysis," *Workshop I in Advanced Topics in Reliability and Risk Analysis*, Annapolis, MD, October 20-22, 1993. NUREG/CP-0138, October 1994.

14. Theofanous, T.G. and W.W. Yuen (1994) "The Probability of Alpha-Mode Containment Failure Updated," Proceedings CSNI Specialists Meeting on Fuel-Coolant Interactions, Santa Barbara, CA, January 5-8, 1993, NUREG/CP-0127, March 1994, 331-342.

15. Theofanous, T.G., W.W. Yuen, S. Angelini and X. Chen (1995a) "The Study of Steam Explosions in Nuclear Systems," DOE/ID-10489, January 1995.

16. Theofanous, T.G., C. Liu, S. Additon, S. Angelini, O. Kymäläinen and T. Salmassi (1995b) "In-Vessel Coolability and Retention of a Core Melt," DOE/ID-10460, July 1995 (peer review version).

17. Theofanous, T.G., J.H. Scobel, S.W. Sorrell and W.F. Pasedag (1995c) "Experience with Risk Analysis Methods in the Design and Certification of Advanced Passive Plants," PSA'95 –

International Conference on Probabilistic Safety Assessment Methodology and Applications, Seoul, Korea, November 26-30, 1995.

18. Tuomisto, H. and T.G. Theofanous (1994) "A Consistent Approach to Severe Accident Management," *Nuclear Engineering and Design*, 148 171-183.

19. Tuomisto, H. and T.G. Theofanous (1995) "Application of Risk-Oriented Methodology to Complete Severe Accident Assessment and Management," PSA'95 -International Conference on Probabilistic Safety Assessment Methodology and Applications, Seoul, Korea, November 26-30, 1995.

20. Turland, B.D., D.F. Fletcher, K.I. Hodges and G.J. Attwood (1994) "Quantification of the Probability of Containment Failure Caused by an In-Vessel Steam Explosion for the Sizewell B PWR," Proceedings, CSNI Sepcialists Meeting on Fuel-Coolant Interactions, Santa Barbara, CA, January 5-8, 1993; NUREG/CP-0127, March 1994.

21. "Zion Probabilistic Safety Study," (1981) Pickard, Lowe and Garrick, Inc., Westinghouse Electric Corporation, and Fauske & Associates, Inc., prepared for Commonwealth Edison Company, September 1981.

APPENDIX A

THE MARK-I LINER ATTACK PROBLEM AND THE PROBABILISTIC FRAMEWORK FROM THE ROAAM APPLICATION TO IT

1. Introduction and Background

The material in this appendix is an adaptation of the first two chapters from the document that initiated the first full demonstration of ROAAM (Theofanous et al., 1991) and is provided here for the purpose of illustration. It explains the problem and describes the framework used in addressing it. Unfortunately, there was no way to abbreviate the quantification, nor the expert involvement. These are essential to obtain the full sense of the approach, and the reader is referred to the two NUREG reports cited for them.

The problem is relevant to the management of core melt accidents in boiling water reactors (BWRs) with a Mark-I containment configuration. The issue arises because the tight containment floor geometry relative to the large corium inventory can produce direct exposure of the liner (in the vicinity of the floor) to the melt as illustrated in Figure A1. Such exposure, even in the presence of water, provides an obvious potential mechanism for liner meltthrough, which is equivalent to containment failure. More importantly, the timing is such that this failure can be classified as "early failure" with obvious implications on outside consequences.

The issue was first raised by the Containment Loads Working Group and became known as the Mark-I Liner Attack. In the late 1980s it led to a flurry of activity including a special workshop in Baltimore (1988) and a special session in the 16th Water Reactor Safety Information Meeting (October 1988). Furthermore, it has figured prominently as a risk contributor in NUREG-1150, and it is a topic of focus in the NRC's Revised Severe Accident Plan published in 1989.

Early analyses by IDCOR claimed that liner failure would be highly unlikely even in the absence of water, while a subsequent NRC task force concluded that failure would be virtually

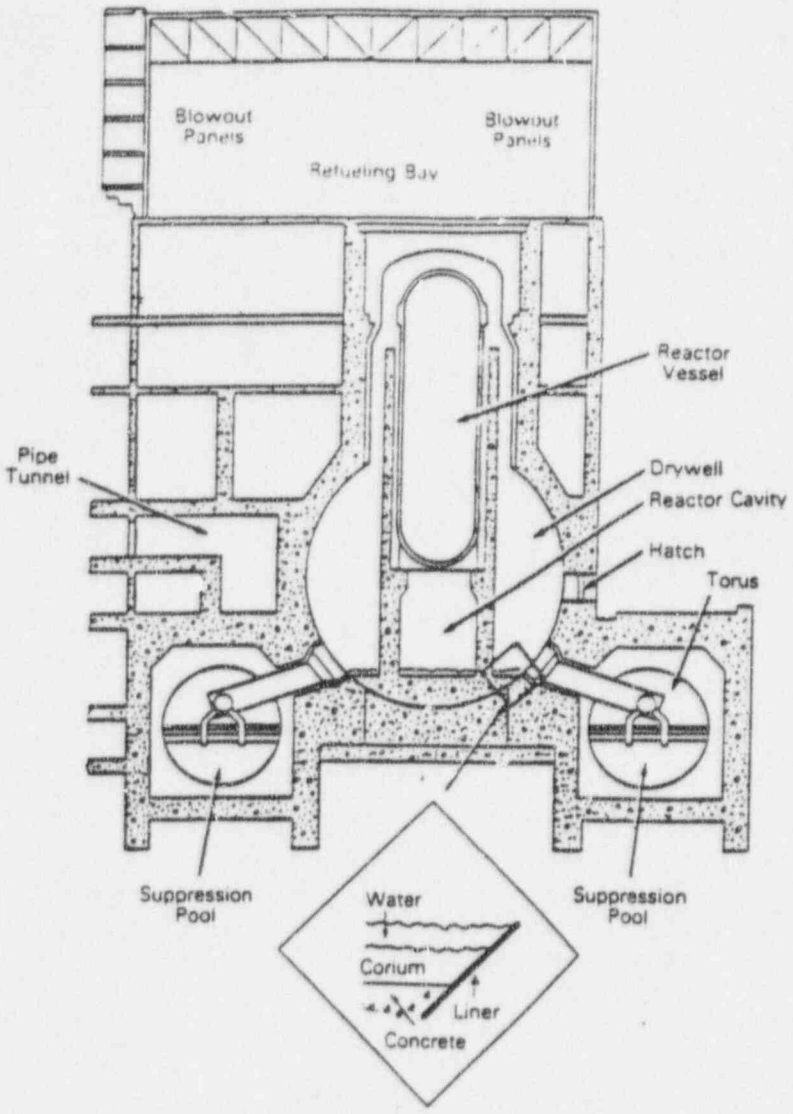

Figure A1. The liner attack configuration in a Mark-I containment.

assured even in the presence of water. It is fairly obvious that the dry drywell case would always lead to failure; however, the flooded drywell case was clearly highly contested and controversial.

We begin by recognizing that the process of liner attack is quite complex on its own; more importantly, it is preceded by a very complex sequence of events (phenomenology), each one of which can influence the outcome significantly. This phenomenology is inadequately understood at this time, and mechanistic modeling with traditional uncertainty analysis is out of the question. An important element of our approach is issue decomposition and assessment of each component separately on strictly technical grounds. A first level decomposition for the problem at hand would

involve melt release phenomena to the containment (that is, the rates, quantities, compositions and temperatures of corium exiting the reactor vessel), melt spreading phenomena on the containment floor (in particular, looking for liner submergence histories and respective melt pool temperatures and compositions), and liner attack, or thermal loading phenomena (in particular, looking for peak liner temperatures). Evidently, these three first-level components are linked as the release affects the spreading process and thus liner submergence which, in turn, and jointly with melt temperature they affect the thermal loading.

Quantitatively, this linking is effected by means of the probabilistic framework, which is presented in the next section. Each component (release, spreading, attack) was addressed separately in technical detail in Chapters 3, 4, and 6 (of NUREG/CR-5423, as is the case for all other chapters referred to in this paragraph), respectively, and a quantification appropriate to the respective probabilistic framework input was given together with the bases for it. Corium-concrete interactions underlie both spreading and liner attack phenomena; hence they have been addressed separately in Chapter 5. Previous work in each of these areas was discussed in respective chapters, and, to the extent possible, some of it was used, as appropriate, in the quantification. The results of the probabilistic synthesis of these inputs were presented in Chapter 7 together with additional calculations to test sensitivity to these inputs. The conclusion that liner failure in the presence of water is physically unreasonable was confirmed after complete resolution, as documented in NUREG/CR-6025.

2. Probabilistic Framework

Because of major uncertainties in the core melt progression, or the in-vessel portion of the accident, we chose to distinguish three separate types of melt release scenarios; namely, an immediate release of a significant portion of the melt followed by a gradual release of the remaining quantity (Scenario I), a gradual release over an extended period (Scenario II), and a large release of mainly solidified debris (Scenario III). This is a "*splinter*," and each scenario will have to be considered independently of the others.

Scenario II is intended to represent behavior dominated by early local failure(s) (instrument tubes), and release of core materials through such, as they become molten in the lower plenum. Scenario I is intended to represent the behavior where core plate failure is delayed, melt accumulates prior to its release into the lower plenum, and a local lower head failure occurs soon after this release. Finally, Scenario III accounts for the possibility that the lower head may suffer creep rupture prior to significant melting of the fuel debris on it. We believe that these scenarios envelope the range of behavior that can reasonably be expected. Following release from the reactor vessel, the melt will "flood" the in-pedestal region, spreading through the doorway toward the liner across from it, and laterally beyond, as illustrated in Figure A2. A first order decomposition, in terms of key events/parameters, is shown in Figure A3. This figure is keyed to, and can help clarify, the corresponding probabilistic framework depicted in Figure A4. As seen, the structure of this framework involves, as inputs, three intangibles (the *pdf*'s) and three causal relations (the *CR*'s). The idea is that liner failure potential is determined by three main parameters; namely, molten pool depth (H) against the liner, its initial superheat (ΔT_o),* and the time duration of the superheat transient (τ). This relationship is qualitatively illustrated as component CR3 (for Causal Relation #3) in Figures A3 and A4. That is, for each height H_i and superheat duration τ_i there is a critical value of the initial superheat, ΔT_{cr}, above which the liner will fail. On the other hand, we expect that for each height there will be a likely range of superheats and respective durations as determined by melt composition (pdf3), the initial superheat (pdf4), and the process of corium-concrete interaction (CR2). This relationship is illustrated, qualitatively, as a "combination" of CR2, with the probability distributions of the zirconium (metal) content of the melt (pdf3) and of the initial superheat (pdf4). The result can be further "combined" with the probability distribution of pool depth (pdf2) to yield pdf5 which in "combination" with CR3 can produce the total failure probability. The depth distribution itself can be obtained by recognizing that depth is primarily related to the quantity of melt released, and by "combining" this causal relation (CR1) with the probability distribution of the

* [In NUREG/CR-5423, we show that the initial melt superheat and the time duration of superheat transient (during corium-concrete interactions) are adequate to characterize the thermal behavior of the corium pool with regards to the thermal loading of the liner.]

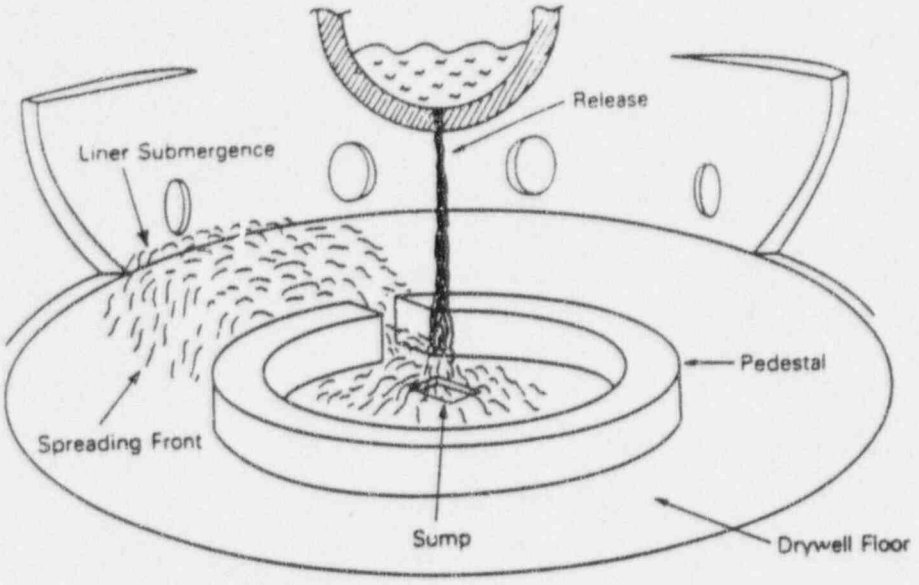

Figure A2. Illustration of melt release and spreading phenomena.

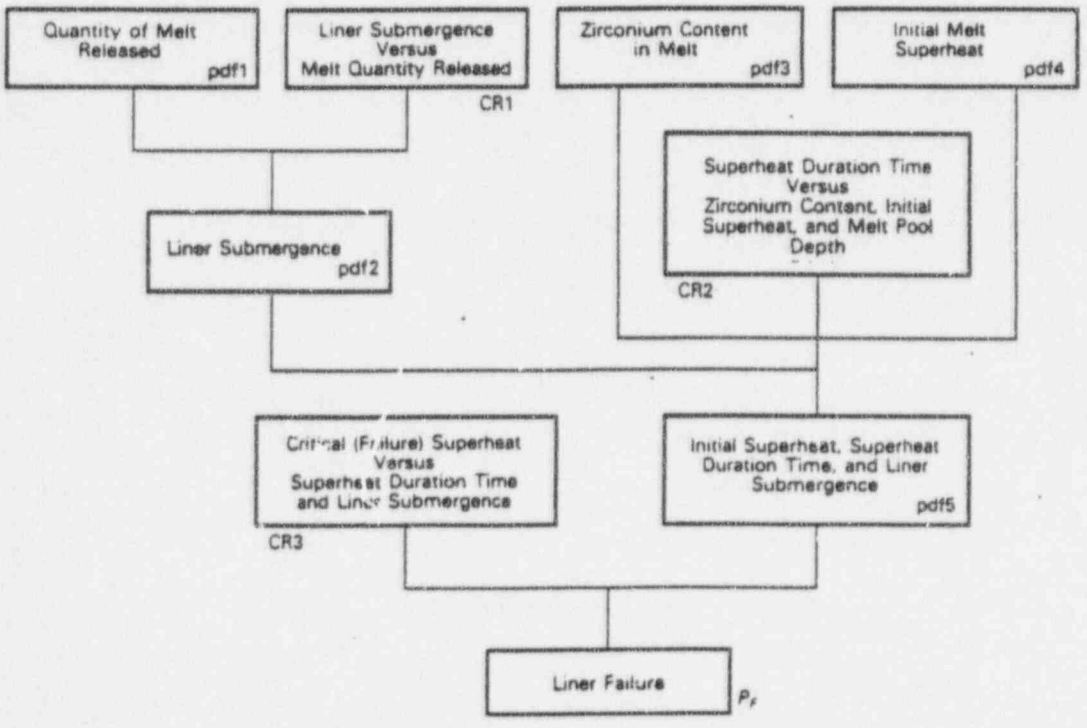

Figure A3. Schematic of the decomposition of the Mark-I liner attack issue. "Liner Submergence" and "Melt Pool Depth" are related by a factor of x1.4, and they are used synonymously.

quantity of melt released (pdf1). The details of how the arithmetic is carried out will help clarify the above and they are given below.

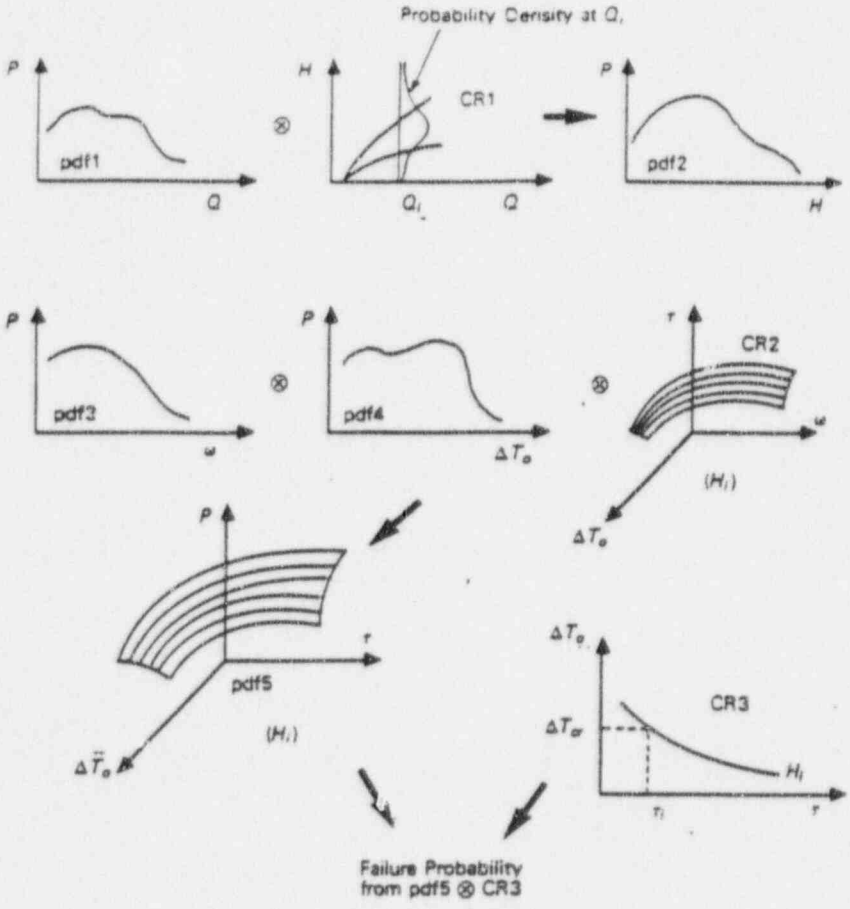

Figure A4. Probabilistic framework for Mark-I liner attack. *pdf* stands for "probability density function" and *CR*, for "Causal Relation." ⊗ denotes various operations explained in the text.

The above decomposition is not unique, nor does it consider explicitly all uncertainties. For example, in CR3 each critical superheat line could be replaced by a range corresponding to the uncertainty range of the pool-to-liner heat transfer coefficient and/or liner failure criteria. Similarly, each line in CR2 could be replaced by a range reflecting uncertainties in the corium-concrete interaction process (i.e., heat transfer coefficients to concrete and water). However, such and other effects can be easily folded in as bounding estimates during the preparation of each input (i.e., a similar process at a more detailed level) as necessary. [An illustration of an expanded probabilistic

framework including the above and its relation to the treatment adopted for quantification here may be found in Appendix I of NUREG/CR-5423.] This hierarchical approach, we believe, is desirable because it offers breadth and flexibility of application. This, in turn, is essential to the overall methodology since it allows the use of widely different phenomenological treatments. An appreciation of the basis of the particular structure proposed here, and its limitations thereof, can be obtained from the detailed discussion provided in the preparation for each input.

Lastly, we turn our attention to the arithmetic of carrying out the operations among the various distributions according to the probabilistic framework. In doing so, we will also have the opportunity to clarify the nature of the information contained in each distribution. That is, a probability density function (pdf) of a quantity, say x, is represented by a finite set of probability numbers, P_i,

$$P_i\{x_i\} \qquad x = 1, \ldots N \tag{1}$$

such that the probability that the quantity lies within $x_i - \frac{\delta}{2}$ and $x_i + \frac{\delta'}{2}$ (δ, δ' represent the intervals $x_i - x_{i-1}$ and $x_{i+1} - x_i$, respectively) is P_i and

$$\sum_{i=1}^{N} P_i = 1. \tag{2}$$

In this sense the meaning of pdf1, pdf2, pdf3, and pdf4 of Figures A3 and A4 is straightforward. The purpose of CR1 is to represent the causal relation between the quantity of melt released (Q) and liner submergence (H) obtained, which accounting for uncertainties is given as a band. With the usual notation for conditional probabilities it can be expressed as $P_i(H_i|Q_j)$; that is, for each value of Q we have a pdf for the distribution of H. The upper and lower boundaries of this band are taken to correspond to the 95 and 5% limits of a normal distribution fitted in for each value of Q as illustrated. The combination of pdf1 and CR1 may be thought of as a set of probabilities

$$P_i\{H_i\} \qquad i = 1 \ldots I, \tag{3}$$

obtained by:

$$P_i(H_i) = \sum_{j=1}^{J} P_j(Q_j) P_i(H_i|Q_j) \qquad i = 1 \ldots I. \tag{4}$$

Similarly, CR2 is the two-dimensional analogue of CR1. Here for any height H_i and any combination of initial superheat (ΔT_o) and zirconium fraction (ω), there is a range of probable values for the superheat duration time (τ) during contact with the liner. It is convenient here, and conservative, to consider the upper bound only rather than the actual range itself. The combination of pdf3, pdf4, and CR2 may be obtained as a set of probabilities, *for each height* H_i,

$$P_{k\ell}\{\Delta T_k, \tau_\ell | H_i\} \qquad k = 1 \ldots K, \qquad \ell = 1 \ldots L, \tag{5}$$

from

$$P_{k\ell}(\Delta T_k, \tau_\ell | H_i) = P_m(\Delta T_{om}) P_n(\omega_n), \tag{6}$$

under the deterministic mapping (CR2)

$$\Delta T_k = F(\Delta T_{om}, \omega_n, H_i) \qquad \text{and} \qquad \tau_\ell = G(\Delta T_{om}, \omega_n, H_i). \tag{7}$$

The meaning of discrete probability and intervals in this case is analogous to the one-dimensional case discussed above. The total failure probability is obtained by summing up the probabilities $P_{ik\ell}$ of the triplets $\{H_i, \Delta T_k, \tau_\ell\}$ that produce failure by comparison to the map represented by CR3. This $P_{ik\ell}$ probability (pdf5)is calculated by

$$P_{ik\ell}(H_i, \Delta T_k, \tau_\ell) = P_i(H_i) P_{k\ell}(\Delta T_k, \tau_\ell \mid H_i). \tag{8}$$

The comparison with the failure map (CR3) can be expressed by

$$P_{FT} = \sum_{\ell=1}^{L} \sum_{k=1}^{K} \sum_{i=1}^{I} P_{ik\ell}(H_i, \Delta T_k, \tau_\ell) H[\Delta T_k - \Delta T_{cr,i\ell}], \tag{9}$$

where H is the Heavyside function that becomes unity for positive values of the argument and zero otherwise, and P_{FT} is the total liner failure probability.

As an overview of the quantification (specification of inputs):

- pdf1 is to embody and codify the in-vessel portion of the meltdown sequence, lower head failure, and melt release,

- CR1 is to represent the spreading dynamics including the role of freezing and possibly remelting,
- CR2 is to convey the outcome of corium-concrete interactions in the neighborhood of the liner and especially the role of chemical reactions in heating the melt layer,
- pdf3 is to be based on the in-vessel behavior and lower head failure (the scenario),
- pdf4 is to take into consideration the melt release scenario and the heat transfer aspects during the spreading processes preceding liner contact, and
- CR3 is to map out the heat transfer process within the liner, as defined by the thermal loading potential (superheat), the time scale for such aggressive attack, and the thermal resistance for heat rejection to the water layer above (proportional to submergence) and to the embedded portion of the liner and the concrete below.

In the above treatment the input pdf's representing the quantity of melt released (pdf1), the zirconium content of the melt (pdf3), and the melt initial superheat (pdf4) are specified, and treated, as independent of each other. Since the question of dependence often arises it may be appropriate to explore this aspect in some more detail. We begin by recognizing that in a fully mechanistic treatment (calculation) of core heatup, meltdown, vessel-breech sequence one would end up with all three above melt characteristics quantified simultaneously, a fact that indicates co-dependence, if not interdependence. However, small changes in modelling approach can lead to bifurcations in the sequence, and these bifurcations can pile up significantly on each other in a manner that in effect destroys these internal connecting, mechanistic links. This results in an effective independence. We have seen already one major such bifurcation arising from the transportability, or not, of core materials into the lower plenum as it becomes molten. It gave rise to Scenarios I and II. Within Scenario I the quantity of melt released in related primarily to the holding capacity of the core "crusible" and the mode of its failure, circonium content is related primarily to the degree of oxidation obtained during uncovery and early stages of meltdown, and the initial melt superheat would depend not only on the dynamics of incorporating solid core debris (by melting in) into the

corium melt (held in the crucible) but also on heat losses during the massive relocation into the lower plenum (losses to water and control rod guides). We see that totally different mechanisms dominate each of these three key melt parameters, thus for purposes of quantification they can be considered independent. Another way of expressing the same thought is that the state of the art in the systemic modelling of respective phenomena allows enough latitude that independent variations in each of these melt parameters can be produced under *nominally* the same scenario. The situation for Scenario II is even more explicit in this regard. Here the quantity of melt is controlled by the debris heatup in the lower plenum (decay heat) and the detailed geometry of the failure location, the zirconium content is primarily tied into the oxidation behavior during the earlier core melt-down/drain-down phase of the sequence, and the initial melt superheat is connected primarily to the convective heat transfer effected between the point of melt formation and the vessel breech. Again, an independent quantification is quite appropriate. The flip side of the above argumentation, however, points also to the difficulties associated with such quantification. Certainly one cannot approach it from a "best estimate" standpoint nor can one hope to account for statistical variability for nominally the same scenarios or for differences in behavior among various classes of scenarios. Rather, our approach is to provide reasonably bounding distributions of these parameters generically grouped into the type I and II scenarios.