ORIGINAL

UNITED STATES OF AMERICA

NUCLEAR REGULATORY COMMISSION

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

SUBCOMMITTEE ON

RELIABILITY AND PROBABILISTIC ASSESSMENT

Airport Park Hotel
Jockey Club
600 Avenue of Champions
Inglewood, California

Tuesday, July 1, 1980

The meeting of the Advisory Committee's Subcommittee convened, pursuant to recess, at 8:35 a.m.

Present:

Mr. D. Okrent, Subcommittee Chairman
Mr. Gary Quittschreiber, designation federal employee
Mr. W. Kerr, ACRS
Mr. P. Shewmon, ACRS
Mr. C.P. Siess, ACRS
Mr. Carson Mark, ACRS

Mr. Lester Lave, ACRS Consultant
Mr. Paul Slovic, ACRS Consultant
Mr. Richard Wilson, ACRS Consultant
Mr. William Lowrance, ACRS Consultant
Mr. Samuel Saunders, ACRS Consultant
Mr. Ivan Catton, ACRS Consultant

Mr. Kastenberg, ACRS Fellow
Mr. Johnson, ACRS Fellow
Mr. Bessette, ACRS Fellow
Mr. Griesmeyer, ACRS Fellow

P R O C E E D I N G S

MR. OKRENT: This is the meeting of the Advisory Committee on Reactor Safeguards Subcommittee on Reliability and Probabilistic Assessment. I am D. Okrent the Subcommittee Chairman. The other ACRS members present today are Mr. W. Kerr, Mr. P. Shewmon, Mr. C.P. Siess, Mr. Carson Mark. Also in attendance are ACRS consultants Lester Lave, Paul Slovic, Richard Wilson, William Lowrance, Samuel Saunders and Ivan Catton. And we have three ACRS Fellows here, Messrs. Kastenberg, Johnson and Bessette, and we'll have one more, Mr. Griesmeyer, in a little while.

The purpose of this meeting is to continue its evaluation; that is, the ACRS evaluation, of the development of quantitative safety goals for nuclear power reactors in consideration of the actual form these goals may take and what they should accomplish. The Subcommittee will also discuss some items related to the NRC published announcement of staff's ongoing program and budget, and suggestions from the industry on specific proposals for safety goals.

This meeting is being conducted in accordance with the provisions of the Federal Advisory Committee Act, and the government and the Sunshine Act. Gary Quittschreiber is the designated federal employee for the meeting, the gentleman on my left. The rules for participation in today's meeting have been announced as part of the notice of this meeting previously published in the Federal Register on Monday, June 16, 1980. A transcript of the meeting is being kept and will be made available as stated in the Federal Register notice. It is requested that each speaker first identify himself or herself and speak with sufficient clarity and volume so that he or she can be readily heard.

We have received no written comments or requests for time to make oral statements from members of the public. We will proceed with the meeting. There are agendas available. I think we'll go right into the agenda unless the Committee members wish to have preliminary discussion.

All right. The first speaker is Mr. Zebroski of NSAC, and according to the agenda he's to disucss general attributes of criteria and a particular formulation of criteria.

MR. ZEBROSKI: Thank you, Dave, for the invitation and the opportunity and also for the fact that this is the subject now of increased interest and intensive effort on the part of the ACRS and the subcommittees. As you know, there is some view that a safety goal can be a very beneficial tool for regulation, and there's also some view that it may be a very dangerous tool if poorly formulated and poorly applied.

MR. OKRENT: Excuse me, can he be heard in the back?

(Response of no.)

MR. OKRENT: Can you turn up the volume?

MR. ZEBROSKI: I think the two things that I'm supposed to touch on -- one is the attributes of a safety goal. Obviously, one of the most important attributes is that it be workable; particularly, that the regulator and the designer or operator can both use some kind of a discipline and come to a similar conclusion.

One of the present problems in the absence of a clear safety goal is the striving for near zero risk where near zero is undefined or is a number which no one understands, like 10^{-6}. So I think I'll structure this a little bit with some viewgraphs.

So the feeling I think that everyone shares is that WASH-1400 technique is very useful, very important in setting a risk evenlope but it hasn't been a very practical tool for either design or regulation. It has, however, been very useful in a Basian decision sense in helping you to decide what is more and what is less important, what you put more effort on and what you use less effort on. In a general sort of way, both the industry and the NRC have used this in a loose way to determine priority both of regulation and of remedial actions.

However, we do see the effort in regulatory agencies generally that the measure of the diligent regulator is to drive his assigned risk to as near zero as he knows how, and the Delaney Amendment is the extreme indication of that where virtually everything in this room, given a big enough dose, would be carcinogenic and would be banned. And obviously, you can't do that, so there is some legislation proposed currently by

Congressman Ritter which basically supports one of the theses of my presentation here in that any kind of regulation must inevitably involve a relative judgment of risk; relative compared to what, is the risk to be assessed. And he proposes this as a guide for regulatory agencies generally. It's not formulated as primarily a nuclear methodology.

I think the bioethics community certainly has made the point that you cannot drive some risk to zero, particularly if that means deprivation of that particular benefit, without having some moral as well as economic concern for the alternate risks of loss or deprivation or excessive cost and social chaos, or even more in some cases.

The present existing legislation is deficient in this respect. I think there was a small step in a good direction, a letter by three of the Commissioners, which recommended that the consequences of legislation, of regulation, be taken into account. I guess the thing it implied to me was the consequences of deficiency of energy was the implication I got from that letter, and they recommended a legislative remedy.

I think one of the points that is the troublesome one is to look at directing indirect costs of the consequences, and I offer at least one perhaps extreme but nevertheless, not unrealistic view. Some of you know when you did an assessment or review of the NUREG-0660 Action Plan, the estimate developed by the industrial group was that the cost of delay associated with

Three Mile Island up to that point was about $30 billion; about 130,000 megawatt years in delay of domestic energy production, which replaced by a mixture of coal and oil, spaced energy, had a cumulative cost impact of $30 billion dollars.

Now, in one sense that is the tip of the iceberg. I think more broadly, society will not have, cannot have -- it's paying most of the social costs of nuclear energy and the regulatory process and legislative process that go with it, but it's enjoying only a fraction of the benefit, and the fraction is specifically that we are not going to build at least 60 and possibly as much as 100 gigawatts of capacity that, in principle, could have been built and were at one time in the planning process for the 1980's. And that amount of capacity, if you convert that to replacement power costs or some equivalent measure of the cost of deprivation, by my arithmetic comes out to about five or six hundred billion dollars in uninflated dollars, and in current dollars would be well over a trillion dollars. So there is a penalty to every man, woman and child in this country over his lifetime on the order of five to ten thousand dollars decrease in whatever other options that kind of resource would give you in your lifestyle, which we have paid, in a sense, for the absence of a publicly agreed upon and manageable safety guard. Again, overdrawing the picture, but at least that's one bound to the size of the issue.

So I think in that sense, if a good safety goal can be constructed, the incentive to have it and use it in some rational way is a very substantial one.

Now, what are some of the desirable attributes of a safety goal? Let me paraphrase this. It requires a definition of practical methods for design and operating decisions. That is, if the guide is vague so you have a great deal of controversy, even today when you have a bulletin order that seems rather specific, and even when we have the prescriptive ones which the industry doesn't like, even then you end up with a good deal of debate of whether you did or didn't meet that particular objective. So the clarity is important, I guess, is another way of putting that attribute.

Secondly, it must be an objective basis for agreement on how safe is safe enough. An important official of the NRC staff in the technical society meeting recently said you know, we don't regulate by rational judgments; we regulate by what we think public perception asks of us. That's why Paul Slovik is here, I suspect.

The issue of regulating by public perception in a sense says you regulate by a kind of mob rule, by a kind of what the media or what the TV tube thinks is important. That's good politics but it's maybe lousy for safety. So I think the objective basis on "safe enough" at least in some part of the intellectual community, should be something that can be somewhat objectively agreed upon.

The next one is the most controversial one in a way; there must be a non-zero goal. It's very difficult for a politician to admit that he's asking you to accept some risks, and yet, practical life experience says that all of us took some risks in getting here this morning and getting out of bed and all these things that we know about. So the non-zero goal.

The next one is perhaps even harder. That is, describable in terms of understandable and acceptable by reasonably informed and "emotionally stable" laymen. Now, I throw in that "emotionally stable" -- that, incidentally, aroused a very uncomfortable set of comments from one of the Commissioners. He thought I was pointing at him, I guess. The point on "emotionally stable" is that my friends at NIMH tell me that almost any population that you examine in any depth will run between 15% and 25% of people who have some easily measurable objectively observable neuroses, in the sense that some functional aspect of their life has been, at least for a time, crippled. They're not parenting or husbanding or teaching or working at the level that their life would expect. So there is always a constituency of neuroses available, and that's at least part of what sells newspapers, I guess. So the idea that you can satisfy an unstable situation is basically not attainable, so we must have to at least draw that distinction that we keep this part of the discussion on a rational basis with people who are willing to understand the terms of discussion.

Now, I have an observation on the importance of public perception as far as this process, but that's a much more philosophical one and if I have time I'll mention it at the end.

Ideally, -- I think this is just comment sense -- it should make use of the best available data and the best available decision process. And being prejudiced in this direction, I think the Basian decision process has not been demonstrated to be less than optimum, or we don't know of a more optimum type of process than that, so basically you're talking about a Basian decision process using the available, observables and theories.

And finally, this is the toughest one and I think this Subcommittee and ACRS generally and some of the legislature have been moving in this direction, is to take account of the alternate risks both of deprivation and alternate sources. And the risks which are most difficult to take account of and which may be dominant are the risks of risk associated with deprivation or with high costs. It is very likely and it is certainly -- let's put it this way. There are very few people I know who would consider the likelihood of highly stable supply of oil and this politically stable situation in the Mid East in the next 10 years is a very likely scenario. We're facing a situation where we can sort of see the disaster approaching and no one can do anything about it. And yet, that risk, that very explicity risk, one which may kill some of us or some of our children, is not taken into account in this process in the sense of it has some impact on the importance of domestically-based energy and resources.

Here's one fearless possible formulation of safety goal that is not endorsed by any either industrial or non-industrial group; it simply is intended to suggest the format of the issues to be covered, and you can plug in your own numbers and your own format.

The basic difference here versus the WASH-1400 -- it basically uses the same discipline as WASH-1400, but it takes recognition of the fact [illegible] the criticism of how well do you know when something is 10^{-}[illegible]? When you use that kind of a number as a criterion, you have an uncertainty band which is of the order of 1 to 1½ orders of magnitude, do you do your uncertainty analysis? So that's a fundamental and perhaps fatal defect which many people have pointed out at a session that involved Harold Lewis yesterday, and he denies -- he said that the uncertainty is two orders of magnitude, but he did say that it's certainly greater than implied in WASH-1400. So the uncertainty of an overall number of that kind is one of the basic defects in the very first attribute; that is, the ability of the regulator and the licensee to agree on the magnitude of a particular hazard.

So, by defining the issue into a frequentist and a probabilist, a frequentist and a consequence issue, this is at least one possible way of handling this conceptual difficulty, a very small number and large uncertainty.

So we have to design an operation to ensure that expected time to core damaging accidents is not less than 30 years, and a more complete formulation of it gives confidence levels on that number and I'll show you some of the mathematical formulations. That turns out to be an extremely sticky -- sounds like a simple statement but very sticky one to treat rigorously mathematically.

But just in round numbers, I would say that just as you could, in one oversimplified viewpoint, say that on an WASH-1400 basis Three Mile Island occurred not at an expected mean time but it occurred when its expectation was at about a 20% confidence level. In that same sense, this kind of a number is at about a 20% confidence level.

Next, that alone, however, certainly would not be acceptable if the implication is that when this happens you're going to kill 20,000 people. So the other necessary attribute is that the consequences be acceptable and manageable. The suggestions here are -- that the reactor containment system and design maintain assurance of not less than 99.9% probability of termination of the accident without radiation release leading to a total dose of one rem to the public. Now, I think the foreign guideline, the one that I'm familiar with is the French, is very similar to this except it takes 10 rem as plant boundary as the threshold.

This is really an observation -- use relative risk assessment methods equivalent to the conventional engineering tradeoff studies which have established the need for adequacy of design or operating improvements which establish that criteria 1 and 2 above are met. Specifically, what this means, if you take the pre-TMI2 accident statistics in a gross sense, which you shouldn't do, but just for simplification I'll take them in a gross sense, then the expected mean time to another core damaging accident if nothing else changed except the increased population at reactors would be about six and a half years; again, at this 20% confidence level. That clearly would not be a tolerable situation, and you say, how much do I need to improve that, and if this principle was used you would say clearly, you must do things which make it reasonably objectively supportable that you would improve that probability by a factor of about 5. So, going for the population of ultimately 150 reactors, the per reactor risk or hazard for probabilities on the fault trees, event trees would have to be reduced.

Then, the statistically rigorous formulation is necessary; otherwise, you don't meet the attribute of objective agreeability between regulator and licensee.

Now you say how do you pick that; why do you pick an arbitrary -- why isn't 30 or why isn't it 50 or 100 years or 500 years? Why shouldn't the probability of containment integrity be 99.999? And so on. You have to have some cut off. The suggestion is that this kind of a formulation then be periodically re-evaluated against the risk of alternates. And the difficult part of that is includin; not just the risk of coals, oil, gas what have you, but the risk of deprivation risks, the moral risk, the social chaos risk of insufficient energy or extremely high cost energy.

But at least the objective part of that is the direct health risk of the alternate sources. And I'm curiously suggesting that something below a third of the cumulative risks of the alternate sources be the at least near-term target, and that a cutoff on the other side be that if you're proposing an improvement; let's say, somebody says if I put a second and third containment shell around the building I can calculate another factor of 10 improvement and now I'm down to 1% of the risks of alternative energy sources, should you do that? And I'm suggesting that that's too much, and that the cutoff be that improvement goes to, aims at, making the risks of the total nuclear population less than one-tenth of the total risk of the alternative sources; and that such improvements be measured against cost effectiveness in a fairly rigorous sense, and that they be implemented only if they have no measurable effect on cost or availability of electricity. In other words, it doesn't preclude the idea that somebody may come up with some very clever improvements which could markedly clip off a number of unpleasant event trees, and if they're cheap obviously you should implement them, but if they involve major upheaval and design, loss of many plants, delay of many plants, for a benefit which is now out in very small statistical probability land, that they would not be implemented.

A further element of this which now goes a little bit to the issue of public perception. I think public perception is, in fact, an undeniable and necessary important force in the regulatory process and in risk-benefit process. So, what are the things that -- if we get away from the issue of manipulation of public perception, which we have a lot of, what can we do about it objectively?

The things we can do about it objectively, first of all, if we can say there is now high confidence that you won't see another accident that frightens people, regardless of whether it hurts them, if it doesn't even frighten them for at least 30 years or well into the 21st Century, I think that would be comforting to at least a good part of the community. If you then say even if that accident happens, it has less than one change in a thousand of hurting anybody at the site boundary, I think that's comforting. But even given that, there still won't be the issue, have you taken Class 9 accidents into account in your environmental impact statement. And perhaps one way of managing that is to set a criterion for emergency planning that you have a high assurance that you can limit the total population man rem to something less than catastrophic by your emergency and evacuation procedures. And the only thing I would make a plea for there is that that be done in an objective way. It is not now being done in an objective way; that is, the short warning times and the guarantee of constant false alarms which we have. I almost get a false alarm every day in the newspaper. You know, the alarms went off at Three Mile Island because of a faulty detector. But there is a manipulation of public fear which unwittingly, the regulatory process now contributes to daily, and I think that really requires that we show that the risk management -- okay, so if this highly improbable thing happens, how do I still limit the damage to less than -- so it isn't killing a large number of people.

We have some work in progress on trying to make some of the statistics of this thing hang together. It turns out that in the published work so far, many people -- I think we now have eight different safety goals that we're comparing and contrasting and trying to understand and measure against attributes. And I might say that I don't particularly have -- at this time I cannot make even a biased recommendation in favor of a particular formulation that I showed. I think there are elements of good ideas in almost any of them, but none of them that we've looked at, including the one I presented, meets all the attributes very adequately. So there's still some work to be done. Particularly when you try to get to the practical application level.

So, the derivation of the sensitivity of a given event tree, to the number of linked events and the time variation that the statistical distribution entails; the basic fallacy of the fault tree/event tree as a predictive tool is that if you look at one and two (?) of the tails and the distribution of the component failure probabilities, or even the failure modes over time, they cannot be regarded to be homogenous. They will change; you will have a different value, you'll have a different environment, you'll have a different operator, you'll have a different maintenance procedure, and basically the tail of a distribution can never be regarded -- the individual tails -- can never be regarded as stable in time. That's just an unpleasant fact of life and anybody who's worked with statistics of manufactured goods or operations knows. And you have to say how do I take care of that, and we have at least a therem which we hope to published for peer group review, which basically just says that if you have a fault tree probability, which is the function of many component failure probabilities -- this was an error I found, unfortunately, just a few minutes ago.

The dependence of the total probability on the stability of the individual coefficients. You're interested in a couple of questions. What is the dependence on what the value of "n" is. Intuitively you say, as "n" gets larger, the fact that one or another of these components may be a square wave instead of a bell has less and less effect. And intuitively and experientially, that's true. It's very hard to prove it mathematically.

If the theory is that "n" goes to infinity, then it fits nicely, but when "n" is typically in the range of 3 to 5, as some of the important fault trees that we know about have, then you really have to do a specific analysis of it. This is u. derway.

I guess I should leave off with one final thought on the effect of experience in increasing expoaure versus the effect of experience in narrowing the statistical uncertain event. Here again, I think it's both a theoretical and practical question that can be addressed.

The interesting thing is that we are all aware of the effect of increasing time and exposure in increasing risk probability. If I have 5 reactors instead of 1, I'm inclined to believe I have more risk.

On the other hand, if you have a well-structured situation and people are striving mightily to do that, there are four elements striving to decrease the entrophy of the situation, and these elements are, first of all, more intensive use of the re-review of the original design assumptions as a periodic exercise. This is both of interest to the utilities for operating purposes, for the manufacturer to protect his good name and finances, and obviously for the regulatory process.

Another element that I think is in the right direction is the increasingly intense review of operating procedures for both expected and rare transients, and the testing of these procedures on simulators and by logic testing. The third element is the increasingly intense review, analysis and feedback on the operating experience process itself. You've heard from Carl Michelson on his plans and work in this area. We'll be happy at some point to review what NSAC and NPO are doing in the operating experience area, but I can tell you that it is a very large effort and I think a very satisfying one. And just by the volume we have now an underground jungle telegraph between all the utilities where they can exchange information on this. The volume of information flow on those channels has doubled every month since last December. So there's a degree of instant awareness of both the event and some of the analysis and some of the remedies, which is a basic and new element in learning from the experience process. I think that the chances we have of getting the industry back on something like the Duane reliability curve where the time between failures and major overhauls increases monotonically as the log function of a cumulative operating experience. We clearly were not on that kind of a curve in the recent past. I think we are getting on that kind of a curve.

So the probabilities of many of the event chains will improve with time, even though the exposure is increasing. And finally, the use of the probabilities assessment process as a rigorous tool for bookkeeping of this logic on the plant-specific basis. As you know, we're involved in several such projects now. One that was announced just last month is the PRA with O'Honey(?), the major effort on the order of 15 man years, and similar exercises are going on with respect to Sequoyah, Yankee, Rowe, Big Rock Point and one or more of the (?) plants and perhaps several others that I don't know about. So I think the availability of the plant-specific PRA as an operating tool, not as a risk-proof tool or as a regulatory retrofit ratchet tool, but something that the operating people use day by day to make sure that the type of experience that happened in another plant in another country five years ago gets entered into their thinking and their procedures and is unforgettable and is the corporate conscience and mentality on a national and perhaps international basis. That, to me, is perhaps the most hopeful thing; that the balance between the exposure, increased exposure with time versus improved learning to reduce the hazard will be moving in a favorable direction. That's all I have.

MR. OKRENT: We have time for discussion or questions. Does anybody want to start the ball rolling?

MR. SHEWMON: With regard to this -- any member of the public who gets one rem -- fencepost is supposed to be tethered there. What's the mobility of this individual, or does that come into the calculation?

MR. ZEBROSKI: I think that has to be a fencepost column. Otherwise, it becomes too uncertain. What it really says is you've got a factor of safety there in reality that no one will actually get that one rem if you have a well-structured system.

MR. OKRENT: Could I understand a little bit more about that point? Does that Item 2 in your safety goals suggest, then, that if you look at the spectrum of core damaging accidents and I haven't a clear definition from you of what you would consider to be a core-damaging accident, but I --

MR. ZEBROSKI: I believe we have a presentation on this scheduled tomorrow.

MR. OKRENT: All right, but I assume you're talking about only accidents which involve relatively severe damage. You're not talking about a few fuel pins being --

MR. ZEBROSKI: Negative.

MR. OKRENT: And that over this spectrum you expect to have less than .1% probability of exceeding one rem at the site boundary.

MR. ZEBROSKI: Yes.

MR. OKRENT: Have you looked at this goal in terms of existing containments to see whether you think it's in the ballpark?

MR. ZEBROSKI: Yes. Fairly exhaustively for TMI-type containments, and less exhaustively -- as you know, studies are going on at the Zion and Indian Point, Sequoyah on just this issue. But the basic point, and I think Gary Collins will cover this tomorrow, is that once you grant that there are many observables for a hypothetical progression of severe damage, and once you grant that you have many ways of adding water and heat sink, once you work out those probabilities, you get very satisfying answers. And then if you back it up to a fire engine for a total blackout, total loss of water, total loss of power case, it can be very satisfying that the 99.9 is a reasonably attainable goal. And I had to restrain myself from putting another 9 on it.

MR. OKRENT: But this conditional probability begins with a core damaging accident as its premise. We're not going to terminate it before it's damaging the core. Is that right?

MR. ZEBROSKI: Yes. Say you have a big break LOCA and failure of reflux.

MR. OKRENT: Okay. And again, just to understand your items, in Item 6 you say emergency plans provide 99% assurance of total population dose less than 5000 man rem, even if containment failure were to occur after a core-damaging accident which is not safely terminated. Now, in this case --

MR. ZEBROSKI: It's again a conditional probability that first of all, you have the core damage, and then all the termination measures, which we'll say is at the 99.9 level, -- that's roughly equivalent to this 10^{-6} worst event in the WASH-1400.

MR. OKRENT: But again, is this over some spectrum of containment failure mode, then, or are you assuming that it's all of these failure modes involve a release to the environment directly, not through the ground?

MR. ZEBROSKI: This is taking the worst case of a partly-ruptured containment releasing essentially all of the gases and a physically realistic fraction of iodine and non-volatiles.

MR. OKRENT: Okay, but you haven't washed those all out with the core spray before the rupture or something like that?

MR. ZEBROSKI: No. I think what we are trying to take explicit account of is the fact that the iodine transport is less than historically assumed, and that the non-volatile transport will probably be far less than WASH-740. The iodine transport will be less than WASH-1400 by quite a large factor.

MR. OKRENT: Are you excluding from consideration accidents which might lead to larger amounts of iodine and cesium, or are you going to try to calculate the probability of different kinds of releases and get an expected value?

MR. ZEBROSKI: Yes, I think we're just saying that the physically realistic probable case is a great deal less aggressive or catastrophic than -- far so than the WASH-740, even considerably so than WASH-1400. Now, if you say given that I have together a sufficient number of hypothetical assumptions which are contrary to physical reality, the battleship can fly, but I think that the -- say at least out to one sigma or two sigma, you'll find much lower releases than the historical assumptions when you look at the physically realistic thing. People are applying this now explicity for iodine, and I'm very hopeful this will also be true for the non-volatiles. There are some codes now which enable you to calculate this.

They are not yet, to my knowledge, being explicitly taken into account in either the siting or the Class 9 rulemaking, that is degraded core rulemaking, or the emergency planning process. So I think this is something hopefully that can be taken into account.

MR. OKRENT: I sort of feel like I'm being taken back to 1963, 1964, when people were proposing the LOFT experiment in its original form. The P&L people didn't seem to feel that one could make a good case for plate out and so forth, or at least some of the accidents involved --

MR. ZEBROSKI: I think paper proof of these things, which are, shall we say, techno-emotional issues now -- paper proof is certainly not manageable; they have to be experimental.

MR. OKRENT: That's a low number, 5000 man rems. If you get any substantial amount of cesium out, I assume you can't meet anything like it, unless --

MR. ZEBROSKI: Again, you have to assume -- you have to say how many assumptions do I allow to operate. Do I allow the assumptions to operate that nobody does anything for X days? At some point, you can stop one of these accidents as far as major release is concerned with a garden hose. If you want to insist that there are no garden hoses, then you can -- you know, the battleship flies.

MR. OKRENT: But I'm just trying to understand what you have in mind by this one. Again, 5000 man rem is a small number compared to numbers which have been calculated like a factor of 1000 --

MR. ZEBROSKI: Let me say why I think that's an attainable number. It comes from the perception that the time to get severe releases, even on the Jerry Rosen probability basis, are rather long -- tens to hundreds of hours for the scenarios that we've looked at. And given that, evacuation may be relatively effective.

MR. OKRENT: But my recollection of the calculations, and I may be wrong because I don't do these calculations myself, is that a large part of it comes from the regions far enough away that, infact, you don't even take restrictive measures. People live there and they get some fraction of an r in a year, or whatever it is, but there are a lot of people. Now, if you're including those doses, then I think you end up saying I can't afford to have anything like cesium or this sort get out; anything with a long half-life or a reasonably long half-life.

So that's why I'm getting back -- are you assuming that there's, ineffect, a 99% or higher probability that effectively no cesium will get out, even in a core melt?

MR. ZEBROSKI: No, I think it's saying that given the perception that the times are rather long, given the perception that ingenuity you do actions to minimize the exposure. I think given that if you're living in an environment where you're getting one r per year from the ground, I think you could find ways to reduce that. So the do-nothing assumption is the one that is perhaps to be challenged here.

MR. SHEWMON: Is that do nothing assumption with regard to cesium which did get out and

MR. ZEBROSKI: Yes, I think at the low levels. At the levels that Dave is talking about I think -- you know, washing, and people have talked about even silly things like you get a lot less dose from radon if there's snow on the ground.

Given that there's incentive, that there's cleverness, that there's modern capability and the times are relatively long -- this I think is the main perception that is comforting to me. There are relatively long times involved to manage these things, so even if you don't have a formal plan to take care of them in advance, people will improvise. That's where I get hopeful.

MR. MARK: I was a little concerned also with your possible item number 6, not intending to point out that something was wrong but just slightly nervous because the 5000 is a surprisingly small number. It's of the same order as the number estimated for TMI.

MR. ZEBROSKI: Maybe I should comment on that. Our analysis of TMI is bringing that number lower and lower. I think the four agency report now -- not unreasonable to believe it's more than a factor of 10 higher than the .

MR. MARK: Well, I was just quoting the first --

MR. ZEBROSKI: In fact, Livermore Lab is one of the strong supporters of that number.

MR. MARK: The Japanese have concluded it was 10 times higher than --

MR. ZEBROSKI: But the Hatashi report is terribly flawed in data.

MR. MARK: It seems to me it's lower perhaps than it need be, and that a lot of care would -- while a criterion of this sort has almost certainly got to be part of the picture, and you are not saying it must be this one, it's one that I guess I would only say deserves probably at least as much thought as the other. You can no doubt more easily come on the fencepost cow number a lot easier than this one, and things like deciding in what way you're going to consider this integral; whether you limited it to 10^{-3} roentgens or something of that sort -- you'll surely get a lot of discussion on it.

MR. ZEBROSKI: It may fail. I was really expecting to throw out this one from the public perception standpoint, where you'd expect to get a good deal of argument that the 5000 is way too high. If you go at it from the standpoint of the, say, just to take another aribtrary number, a fifth of the cumulative risk of ernate sources, that's probably clearly too low.

MR. MARK: I didn't suppose that you meant when you were talking about comparing it with alternate sources that you were going to compare them on a rem basis.

MR. ZEBROSKI: No, it has to be societal impact; that's where the public perception problem is severe. That equivalence is by no means transparent even to the technical community.

MR. MARK: I didn't mean to say I didn't rather like your pattern, but I was hoping and thinking, of course, that an enormous amount of discussion has to go into those numbers.

MR. WILSON: I would like to ask for your comments on it. My impression of this is I have the horrible fear that lacking now something like safety goal 3, that those words at some time will appear in the Federal Register as the NRC safety goals, without any of the conversations that you put forward and the clarifications you just made. And I think that would be another disaster.

Because immediately, every one of these can be interpreted by someone in a completely different way. So I get problems with writing them down in this way without discussion such as -- Paul Slovic insisted at one time at an early meeting we start the decision theory at the beginning, what do we decide to decide. And then go on and then Griesmeyer and Okrent had a procedure for discussing some risk levels. And this, then, is quite farther down the chain of what we're trying to do; the specific methods of addressing some of these levels. So without any of the calculation procedures which you're discussing, I think this could be -- it's almost meaningless, and if you leave it open people are going to consign almost any meaning to some of it. Such as, it has to be qualified with what realism one puts in the calculation procedure.

So I wondered if -- what I want to ask then is that in almost all of these, were you intending this as a summary of the procedure, or just a summary for discussion here? How would you imagine one would, in fact, write it down in the Federal Register?

MR. KERR: Let me ask, are you discussing Roman Numeral III?

MR. WILSON: Roman Numeral III, right. Because I feel that's the important issue.

MR. ZEBROSKI: Dick, I'd like to come back to my original caveat. I think we've tried to make a compare and contrast analysis of the many different formulations which we're aware of. Everyone who is here today has written one formulation or another. And I'll repeat that none of them, including the one that I'm presenting, really meet all the attribute requirements that I suggested. So that's why I started out saying the poorly-formulated one is probably more damaging than beneficial; I just agree with that. So I think the structure of a workable formulation is the principal intellectual challenge for the whole nuclear community worldwide. I don't think it's just in this room, but the whole world is struggling with this issue.

In fact, one of the thoughts we've had is that we perhaps should have a more international workshop on this subject.

I recently talked with people in Japan and Taiwan and Germany, France and Swedon on this, and they're all struggling with this question.

MR. SAUNDERS: Some of your comments in explanation preceding safety goal 3 give me even more pause. The phrase that the Committee is not to give itself over to uninformed mob opinion and to disregard the 20% of those people who are emotionally unstable -- if that appears in the Federal Register it will be as well received as the proverbial object and a fence pole. But those were even more alarming to me.

MR. ZEBROSKI: Yes, but I'm suggesting a process. I think if you tried to manage public risk perception while getting a statistical confidence level, you'll do neither. So you have to do the first. I'm suggesting the procedure that you first of all get an agreement which is at least understandable to a sympathetic community before you float it to an unsympathetic community. If you can't get the former, don't try the latter.

MR. KERR: I had the same problem because I thought you were going to tell us next that at least half of the population is below average.

(Laughter.)

MR. SIESS: I'm not sure this is a question, but on your item 6, the 5000 man rem, and public acceptance of that, as I recall from Three Mile Island in the final environmental impact statement which is a supposed to be a realistic estimate, the population dose of a design basis accident was on the order of 5000. Is that right?

MR. ZEBROSKI: As of what vintage?

MR. SIESS: The environmental impact statement for TMI2 --

MR. ZEBROSKI: That's on the order of 5 or 8 years ago. I don't think that has any of the current analyses in it. As a round number I'd agree with you, but I just don't have it.

MR. OKRENT: Thanks, Ed. I'll invite you to participate in the continuing discussion. I think we'd better go on. Chris Whipple is next. It's clear that we get very bad feedback if he turns it up. Maybe the thing to do is if you can't hear in the back wave your arm and I'll ask people to speak up.

MR. WHIPPLE: Holler if you can't hear me.

Following Dick Wilson's last point about the starting with the general formulation before you get to the specific numbers, I'll mention at the outset that I don't intend to recommend actual numbers but talk instead about the attributes of the approach.

I think starting with kind of a base point, what I see as the principal advantage of moving to the kind of quantitative criteria that the elements of all the approaches discussed have is the extent to which they separate the questions of fact from value relative to the existing procedure, and I think the strains this places on the industry and the NRC staff and on intervenors to deal with these issues on a somewhat more direct basis is a principal advantage.

I also see in the literature four types of criteria that have been mentioned and I'll mention them very briefly. The first is just an individual risk level, the fencepost number, and the occupational vel. Those I think we have procedures for dealing with and I don't think those are the really substantial issues on nuclear risk that have caused public concern. I don't think that's why were're here.

The second issue is a societal risk ceiling; some sort of farmer limit line or a modern extension of that that is very much an issue. The third point is a cost effectiveness test for safety opportunities below those existing risk ceilings. And a fourth frequently mentioned kind of criteria is a noise level lower bound, which is considered in the interest of operational efficiency of the regulatory agency rather than on any other principle of it being right to ignore low risks.

Well, the lower bounds, it's often suggested, are usually given simply as a probability, somewhere on the order of 10^{-6} or so. At least, that's the number I hear for FDA's current operations concerning which drug food additives to take action on. But I'm very hesitant to see something like that applied to nuclear because the extremely low probability accidents are precisely the ones that have caused public fears about nuclear, and the separability of consequence from probability at very low probabilities is an issue for nuclear power. That is, I don't think that the most severe nuclear accidents can be dismissed regardless of their probability. So if any sort of a lower bound is instituted, I think for the nuclear power case it should incorporate considerations of both frequency and consequence.

In so doing, I think it would be desirable if the faults of the current system in which the catastrophic accidents are treated distinct from the moderate accidents; that is, Class 9 versus design base accidents, as being separate points when, in fact, they're on the same spectrum. They were just in different parts of the spectrum.

I see some attraction for dealing with that as the continuum it is.

The big issue is, of course, the frequence and magnitude relationship for acceptable risk. Ed Zebroski went through a number of things he wanted to consider in arriving at that criteria, and notably, comparisons with alternative energy sources and risks of non-energy. And while I think that's true in a total societal sense, I'm not sure that the NRC is the place to do that. It might more readily be the role of the Congress to do that. And the Congress is on record as having somewhat endorsed nuclear power and established NRC to see that nuclear plants get built subject to numerous criteria about the public safety of those plants.

So, the levels of risk aversion that the NRC might adopt I don't think have to be strongly based on the full decision aspects of all the energy alternatives. I think implicitly the Congress has already done this in telling the NRC to go ahead and develop criteria for licensing reactors. But there are still very difficult issues involved in developing those levels.

I think the basic place to start is the performance of the existing plants, and the possibilities seen by the people developing the technology. What is achievable? A second issue is what's achievable in fairly close competitive technolgoies? For example, dams have frequency-magnitude relationships that can be studied, and that's a very worthwhile study to make in that they are truly alternatives in many cases.

The issue of risk aversion is one that I think needs to be faced explicitly, and I would recommend a modest degree of risk aversion, but that's very much my own value system and I suspect that that's a 100% value issue, and given that the NRC Commissioners have been confirmed by Congress and have some endorsement and legitimacy for choosing something of that sort, I think it's fully within their power to do so.

Now, in developing the frequency-magnitude relationship, there's really three issues. The first is -- excuse me, let me interrupt myself. We have to decide whether we're interested in a risk criteria or a plant criteria; that is, are we interested in developing criteria for probability versus curies released, or are we interested in a probability fatality, probability health effects kind of relationship. The latter is more difficult to do because it involves the three steps I started to mention. The first is the frequency of release probability; second is the exposure given the release and that involves a number of difficult factors, probabilistic factors, such as weather and transport and demography and evacuation assumptions and so forth. And finally, the also uncertain area of dose response.

I noticed Ed chose to put his formulation in a plant model in terms of specifying frequencies and for types of accidents, but then population exposure is conditional upon those accidents. It's kind of a mixed approach.

It's a difficult issue because there's kind of a mis-match between the natural units of what's desired. From a licensing viewpoint, I can see great attraction in having an old farmer curve, a frequency probability of releasing so many curies curve, as a fairly simple scientific question, fairly difficult scientific question, I suppose, but from to use those kind of criteria ignores a number of opportunities for trading off costs, for example, in siting. And it would probably result in non-uniform marginal control cost per risk by going to standard frequency curies curves.

The alternative of specifying risk levels based on a probabilistic risk level or risk curve, is that it's very difficult to do and to work backwards through all the uncertainties involved to a power plant design. And in so doing, you again bring a lot of values into the process; for example, what's your philosophy regarding errors? That is, are you going to work on most likely at every point in the chain are you going to carry uncertainty throughout the chain, and when you work back you find, as Ed pointed out, perhaps greater than one order of magnitude, perhaps two orders of magnitude, perhaps more.

I have kind of a middle of the road suggestion which might be to specify dose release limit curves that have some flexibility built into them depending upon population density. It seems quite common sense to me to have more stringent requirements for plants in more populated areas, or conversely, to provide incentives to locate plants in less populated areas. And this would be one way of doing it.

And also is has the advantage of making implicit statements of risk acceptability rather than explicit statements, and this is a difficult enough field that if the NRC is to adopt something in this area, I suspect that the operational feasibility of an implicit approach is greater.

Finally is what I consider to be the most difficult issue of all, and that is, how do you determine whether the criteria are being met? And how do you establish standards for proof. I consider this to be separable from the issue of setting the standards themselves because the degree of allowable uncertainty will clearly affect the achievable criteria.

Unfortunately, in thinking about it I keep coming down to the fact that ultimately, the probabilistic estimates are going to be made subjectively by the NRC staff; at least the ones that stick. And the difficulty there is that both the industry and the intervenors want to have a shot at reviewing and critiquing the process. Within that framework, it's very difficult to see how that can happen.

Again, I'd like to see some sort of a middle of the road approach that balances the subjective estimates of probability by the staff with the fullblown complexity of trying to conduct a WASH-1400 type study for each power plant. The kind of middle ground that might come out would be to have a kind of generic safety study for each vendor or for each major plant type with separate analyses required where there's some substantial variations from existing or proposed plants from the base plant. But that's a very difficult issue, I believe. I think that's the central issue in going to probabilistic criteria.

Within that comment, there's a need for standardized risk calculation codes as obvious, I think. And they would be of enormous benefit if there was an industry standard for evaluating the probability of a reactor having a release of certain amount of activity. I think that should be a kind of high R&D item.

Now, I think that the -- let me just conclude. I also initially was talking about the legitimacy of separating facts from values. I think that could we attack the issue of how do you measure reactor risk, and hopefully always in the absence of data, that that does do a good job of separating the public debate as well as the institutional hearings into issues of fact and issues of value much better than has been done in the past. My suspicion is that the disagreements over issues of fact will outweigh those disagreements over the issue of value. But if, in fact, that's the root of the public conflict, we're better off addressing it directly than through subterfuges. Thank you.

MR. WILSON: I have a couple of questions. You said you want to separate probability and consequence, and I would agree with that. But you gave the implication that even if the probability is very small, the consequence is to people. I would like to ask, do you really mean if the probability is calculated to be very small, or do you mean -- because I really feel that we want to separate them simply because -- that is where the belief in the actual value s calculated -- people just don't believe that that calculated consequence is that low. And for that reason, you want to separate the probability and consequence and try and get both down.

MR. WHIPPLE: That's quite right. In fact, I think the frustration of intervenors in nuclear issues on the fact that many of the most catastrophic accidents have, in many hearings, been considered off limits because they've been considered resolved a priori, is a great sense of frustration that this would resolve.

Now, it also opens the issue for debate. But you're quite right. The public issue, I believe, is over whether, in fact, the catastrophic accident probability is very small or extraordinarily small -- 10^{-4} versus 10^{-6}, perhaps.

MR. WILSON: Again, when you mention -- you don't mention the question of the procedures -- of the question of uncertainties and what you would do about -- this brings out the question of how would you bring in uncertainties into such a calculation. How do you bring in earthquakes and sabotage, which can't be calculated by the Rasmussen method?

MR. WHIPPLE: No, I didn't discuss those because I don't know.

MR. ZEBROSKI: I think that those equations exhibited at the end really address this issue very explicitly, and that is, if you use -- the virtue of the relative risk assessment is that if you take a particular line of design where you have 100 or 200 or 300 years of operating experience, that in a sense gives, at some confidence level, a limit to some of these probabilities, and then you can ratchet from this by saying, I have cut off some of the event chains which contribute to that cumulative risk, and improved them by some factor. So I think this, at least to me, is the only manageable way of addressing this issue of the unthought-of alternates of sabotage, war and so on, and seismic. Because you have to start from the experience base and then say what am I doing to improve it from there. And that's why the stability of those distributions gets to be mathematically the key question; how stable do they have to be before you can use it as a practical thing.

MR. WHIPPLE: Yes, that is an area that I have looked at to some extent. The point is, what criteria do you establish for the validity of probability estimates, in the absence of what I would call really relevant data. That is, if comparable sized power plants exist or comparable sized earthquakes, for example.

What Ed is mentioning I guess you would call extrapolation of a frequency-magnitude relationship, and I think that's a very important tool if it can be demonstrated to have validity. The conditions under which that method can be applied I think is an issue that would be well addressed, because it is, as Ed pointed out, quite valuable.

MR. KERR: Mr. Whipple, I thought I understood you to say that you did not believe that NRC should do too much comparison of risks of other sources of energy because Congress had already indicated approval of nuclear energy in setting up the energy program, or something like that.

MR.WHIPPLE: I think in talking about the legitimacy --

I don't think it's within the purview of the Nuclear Regulatory Commission to make decisions of whether the U.S. is going to go with coal power exclusively, whether to go with nuclear --

MR. KERR: No, I thought you were discussing risk comparisons in an effort to arrive at the goal.

MR. WHIPPLF: I think in a sense it's different than that even. Ed was thinking of, as I understood him --

MR. KERR: I'm not trying to get you to comment on what Ed said, but rather to comment on what you said. But I'm not sure that I know what you said. I thought you said you did not believe it was appropriate for the Nuclear Regulatory Commission to make comparisons with other sources of energy.

MR. WHIPPLE: No, I did not mean to say that. And I understand how you got that impression. I did point out that, for example, I think dams would be a very useful alternative technology to compare developing risk criteria. Let me try to clarify this. What I think is not a good idea for the NRC to try to do is to make a power plant selection decision process study; that is, they should not try to select what is the preferred alternative technology, necessarily. I'm not sure that's within their jurisdiction.

MR. SHEWMON: But to calculate the risks that do -- to establish rather firmly the risks that do come from alternate technology you feel is part of their purview, or could be.

MR. WHIPPLE Yes. As a demonstration of what technologies have been acceptable in the past or what risk levels have been acceptable in the past, that's very clearly an important issue in establishing nuclear power risk levels.

MR. KERR: I think one almost has to separate, whether it can be done or not, considerations of risk from considerations of environmental impact. But in a sense it seems to me the licensing process of the NRC, at leas in the Commission's interpretation of its responsibility, forces it to choose between alternate technologies each time it considers a new power plant license.

MR. WHIPPLE: Yes, on a siting basis, that's true as I understand it.

MR. KERR: And one of the important parameters in environmental impact is that of risk. That is not the only one. Are you saying, in a sense, that it's appropriate for the Commission to choose between technologies on an environmental basis but not appropriate for it to choose between technologies on a risk basis?

MR. WHIPPLE: No, I think what I'm saying is that it's appropriate for the Commission to choose on an individual basis but not a generic basis. I think the generic decision --

MR. KERR: But it is already choosing on an individual plant basis and environmental impact, I think.

MR. WHIPPLE Yes. But now in terms of developing criteria for reactors that apply generically to all reactors, I think the issue of whether you choose criteria against which reactors can actually be built or cannot be built is one that I think is more appropriate for the Congress to consider. That is, the general question of whether the U.S. is going to have nuclear power plants. Again, that's very much a personal thing as opposed to something directly related to the process that we're talking about.

MR. OKRENT: I see a hand in the audience.

MR. SPANG: I'm Miller Spang with the Nuclear Regulatory Commission. We are in the throes of grappling with this very problem of what our role should be, and the special inquiry group recommended that the NRC not provide the final answer; that this should be the Congress and the President, but that we should have a leadership role in assembling and synthesizing views from a wide body of people to present to those who have the final answer. And as it's been pointed out, the NEPA required us to do, on a case by case basis, an analysis of alternative energy; we cannot escape that unless the NEPA is amended to eliminate that.

Generically, we have stayed away from making a comparative analysis saying what is an optimal mix between coal and nuclear. I think it's pretty clear that the Second National Energy Plan says we need both of these for the remainder of the Century as the mainstays of baseload electricity generation so we can quote that as a given. But then to try to find, too, what is an optimal mix of nuclear, coal and solar, geothermal or whatever, is not for us to do, as we see it. But I do think that we are grappling with this and possibly we'll be coming up with some sort of a statement indicating some acceptance of a leadership role, at least in focusing debate on this subject.

MR. OKRENT: Okay. I guess we'd better move along. We're now behind agenda, as I anticipated we would be. I hope everybody in the audience anticipates that there may be a little bit of slippage in their present agenda. This is an interesting topic and it's worthwhile having discussion.

The next speaker is Mr. O'Donnell from the Atomic Industrial Forum.

MR. O'DONNELL: Good morning, Dr. Okrent and members of the Subcommittee. My name is Ed O'Donnell, I'm a Division Vice President with Ebasco Services, and I also serve as the Chairman of the Atomic Industrial Forum Committee on Probabilistic Risk Assessment for PRA.

Our committee consists of representatives of utilities and SSS vendors, architect engineers, consultants and also EPRI and NSAC representatives.

Our current efforts have been aimed at developing proposed policy statements on the use of PRA in the regulatory process and also on quantitative safety goals. On June 2nd, we issued a statement of policy that was sent to NRC staff, and also to the ACRS, setting forth our views on these subjects. And today, I welcome the opportunity on behalf of AIF to elaborate on our views in these areas.

There are a number of our subcommittee in the audience here as well as AIF staff members that can respond to questions today.

I'd like to summarize for the ACRS Subcommittee our June 2nd statement, and also focus in particular on our views on the quantification of safety goals. And in view of the limited time available, I think I will embark on that subject first and leave the discussion of how PRA should be used for the end of the presentation if time permits.

We believe that the establishment of quantitative safety goals is a very complex undertaking, and involves essentially at least four major steps. First of all is to decide on the basic principles on which you are going to establish the safety goals. Number two is then determine exactly what it is that you want to quantify; that is, what are the parameters that you wish to attach numerical numbers to. Third is to develop those numbers; and fourth, to determine how you intend to apply them in the regulatory process.

With respect to the basic principles involved, we believe that the quantitative safety goals that are applied to nuclear plants should be generically applicable to all technologies, particularly those related to energy production. And, as a matter of fact, they should also apply to any human activity that involves risk.

Number two is that the goals should embody the principles that the acceptable societal risk; that is, the population exposure, should reflect in some way the benefits that society derives from that technology. And third and one of the primary principles, is that no individual in the public should bear an inordinate burden of that risk. And fourth, that the goals should promote an optimum allocation of societal resources in reducing societal risk.

Using these elements, after much discussion we have evolved a framework for setting quantitative safety goals that would involve four elements. And these are essentially -- provide what we feel is a simple and direct framework for setting quantitative safety goals that will apply to nuclear power plants.

The first of these is setting a limit on the individual health risk. That is, to the maximum exposed individual in the public.

The second involves a limit on population health effect, which somehow recognizes the societal benefits as derived from nuclear power. The third is a cost-benefit criterion that would apply to reductions in residual risk once you've satisfied the first two. And the fourth criterion is somewhat subordinate to the others in that there is a benefit to be derived from setting a goal on limiting the probability of events that would result in core damage such as TMI.

I think what you see here is essentially a hierarchy in terms of the individual elements in the risk criteria; that is, the individual risk criteria is probably the foremost and most absolute, and the individual must be protected at all costs from undue risk for the benefit of society.

And for that criterion, we would propose that there would be a statement of essentially a qualitative nature that would be of this nature, that the incremental risks of adverse health effects to the maximally exposed individual in the vicinity of a nuclear plant site should not result in a significant increase in annual mortality rates or in significant shortening of expected statistical lifespan.

Essentially, this criterion is aimed at protecting the individual, primarily the individual that is at the site boundary and it does not take into account the concept of balancing societal benefits or economic costs.

The numbers that I'm going to propose here are essentially our preliminary thinking on these subjects, and they should not be taken as absolute proposals by the Atomic Industrial Forum because I think there's a great deal of work that has to be done in refining these numbers. And what we're proposing this morning is essentially some numbers that we think look to be reasonable at a first cut, and they are subject to further revision or refinement as we look more deeply into this subject.

But with that said, it appears that a number such as

10^{-5} per year as an individual mortality risk; that is, to the maximum exposed individual, appears to be in the range of reasonable in value. And the basis for that is essentially that it reflects a number that is a small fraction of existing background risk. That is, if you take the average risk to the individual, which is about 10^{-2} per year, a number such as 10^{-5} would represent .1% increase in total mortality risk. And if you look at accident risk, it's about 1% or 2% of that number. And it compares favorably to other types of accidental risks that the individual is exposed to in normal everyday life. Such as the risk of motor vehicle accident, violence, fires, air travel, the risk of death from falling objects and also the risk of death from electrocution. And that last point is interesting because it is related somehow to the fact that the risk being imposed here is to essentially provide electricity to society and the individual.

The average to the individual from electrocution is on the order of 10^{-5}. Here we have six times 10^{-6}.

In addition, if we look at the expected risk to the individual from background radiation, assuming the linear dose model, the added risk to the maximum exposed individual at the site boundary of a nuclear plant would be on the order of risk that he is already exposed to from normal background sources. So we are, in effect, saying that if an individual lives near a nuclear power plant, his added risk is about equivalent to the added risk that he would be exposed to if he moved from, say, New York to

Denver, which is certainly a risk that no one at this point has gotten too excited about.

In addition, it is somewhat unfair to compare these average risk numbers with a maximum risk number for nuclear plants because if you looked at the maximum risk to an individual from any of these sources, it would most probably be significantly higher than these numbers down here.

In addition, some preliminary study indicates that the maximum risk to an individual near such installations as chemical plants and hydroelectric installations and coal plants is significantly higher than the 10^{-5} number that we're suggesting here of an individual risk for the maximum exposed individual near a nuclear plant.

MR. OKRENT: Could I ask just a question of clarification. The 10^{-5} per year individual, that would include both acute and latent effects?

O'DONNELL: That's right, that would be both acute and latent, and for the individuals nearest the plant we would expect that the early fatalities would dominate this risk.

To give you some basis for comparison of the suggested number that we're giving, we did some research into what other people have been suggesting. For instance, Mr. Vesely at the April presentation before this group, proposed as a tentative or draft number, a number of 10^{-5} which would be an unacceptable goal, and numbers of 10^{-6} and 10^{-5} as numbers in which you would be in a warning range. That is, you would conduct a case-by-case evaluation on the individual risks.

Dr. Wilson I believe has proposed as a tentative proposal that a number of 10^{-5} for the near-site occupant or individual which may be acceptable, and a number like 10^{-6} for the next township. Now, Dr. Wilson, as I understand it, makes that distinction on the basis of tax benefits that the individual near site would gain.

Dr. Okrent in your April article in Science I believe had suggested numbers that I've indicated here, basically distinguishing between the benefit of an activity in terms of whether it's essential, beneficial or peripheral to society, and has suggested individual risk numbers that are not too far out of line with the numbers that we're suggesting. The difference, and it may be significant, is that you've suggested that this risk be assessed at 90% confidence level. Essentially, we're suggesting that the risk should be assessed at mean value or 50% confidence. So whether that's a significant factor depends on the uncertainties involved.

Mr. Corkerton and his associates in England have suggested a number of 10^{-5} as a maximum risk to the public, and a number of 10^{-4} for the work force. If one goes into WASH 1400 and extracts the --

(Tape change.)

Studies in England have suggested a number of 10^{-5} as a maximum risk to the public, and a number of 10^{-4} for the worker. If one goes into WASH-1400 and extracts the maximum risk to the individual for site boundary you come up with a number of about 8 times 10^{-7}, and the German risk study implies a risk of about that magnitude, also.

So our value of 10^{-5} is really not totally out of line with what others have been suggesting in this area, and which indicates that if nothing else, there is strength in numbers.

With respect to the population health effects criteria, the basis for this number is that there should be some relationship between societal and benefit with respect to a technology. And to reflect that in this criterion, we've tied in the cumulative risk to the population to the capacity of generation of the plant. That is, the incremental cumulative risk of adverse health effects to the exposed population per thousand megawatts of nuclear plant capacity. Considering the probability and consequences of events integrated over the spectrum of potential accidents, it should be no more than a small fraction of the average background incidence of health effects.

And there are a number of important concepts embodied in that criterion. First of all, is that you should measure societal or population health effects against the plant capacity and that reflects the fact that a thousand megawatts of electricity is in some way a measure of societal benefit. That is, we don't build nuclear power plants solely for the purpose of building nuclear power plants. That is that they are needed in some way and will provide some benefit to society.

And as the capacity grows, so should the risk be enabled to grow.

The other important concept here is in that we are setting forth the concept that the population risk is essentially the integrated risk under the accident probability consequence curve, and we're not proposing that we should apply a penalty for high consequence, low probability events.

The third thing that this criterion would do, is it would provide some measure or some consideration of population density. That is, if you site a 1000 megawatt plant in a 1000 people per square mile population zone, the population risk will be 10 times that if you sited in a 100 person per square mile area. So this criterion does provide some incentive to site plants in lower population areas.

It, as would the individual risk criterion, considers both early and latent fatality risk, and in this case of population risk, the latent fatalities would dominate the risk. And it also facilitates comparisons with alternative energy sources and their risks.

The number that we're suggesting as a goal is on the order of .1 fatality per year per 1000 megawatt capacity. And again, the basis for this number is that it does reflect a very small fraction or increase in the existing background risk. In fact, if one looks at total mortality risk in the U.S., it's about .0 to 1% of existing risk. Since we're concerned primarily with cancer risk, it's appropriate to compare it with the background cancer risk. And there also, it's about .005% of the total existing cancer risk in background.

And again, if .e look at comparisons with other risk contributors -- I should point out that these percentage numbers that I've provided are essentially based on a total capacity of 200,000 megawatts.

If we compare these numbers to other existing individual risk contributors, again you find out that it would be a small fraction of existing risk in the public. The risk of motor vehicle accidents is one in which we're experiencing 50,000 early fatalities per year, which represents about 2½% of total fatality risk. The number for violence is about the same, and as we go down the list, it's interesting to note that the number for electrocution; that is, the risk of the in-use product for nuclear power, that is electricity, results in about 1100?fatalities per year, which represents about .05% of total risk.

Now, on this basis, it would c-- ainly assume to be reasonable that the added risk of generating or producing the power that results in this sort of end use risk, which I don't believe people are willing to forego, is a very small fraction of the end use. And in fact, if one took the total capacity in the U.S., which is about 400,000 megawatts per year, and looked at the risk of electrocution, this would work out to be about 2.5 fatalities per year per 1000 megawatts.

So certainly, the number we're proposing here of .1 is a very small increment with respect to even that very pertinent specific for electricity.

To go further, if you look at the population risk from such sources as coal and hydroelectric power, I believe you would find out that the population risk per 1000 megawatts is quite a bit higher than numbers that we're suggesting here.

To provide some basis for judging the proposal that we're putting forth here, Saul Levine, in a draft paper that he's written -- I'm not sure if he ever gave it or not -- was proposing a risk curve that essentially implied a risk of about .2 fatalities per year. If one looks at WASH-1400 and the integrated area under the risk curve, we're talking about .02 fatalities per year. The German risk study indicates that for the 25 plants they evaluated, there would be a statistical average of 10 fatalities per year, which works out to be about .4. So the number that we're suggesting is reasonable is, again, in this range of probability or within the numbers that other people have suggested or that have been implied by studies.

The third criterion essentially deals with cost-benefit.

MR. KERR: Excuse me. The set of numbers you gave on the last slide -- is that associated with a 1000 megawatt electric plant?

MR. O'DONNELL: These are normalized, I believe, to a 1000-megawatt, yes. Maybe 800.

MR. KERR: You're talking about one plant, not the total number of plants.

MR. O'DONNELL: Yes. I'm sorry. Yes. This is per 1000-megawatt.

The third criterion deals with cost benefit, and essentially this would be a criterion that would be applied after you've satisfied the individual and population risk criteria. That is, they would be applied as threshold values before which you would not consider cost-benefit criteria. Once you've satisfied the 10^{-5} individual risk criterion and the .1 fatality per year per 1000 megawatt, then in evaluating further reductions in residual risk, one would apply a cost-benefit basis. And this would be explicit in terms of dollars per man rem.

And as stated in words, this would be that the benefit in terms of population risk reduction afforded by a change in plant design or operating procedure should be comparable to that which is generally achievable through alternate investment of the cost of the change in other areas of public risk reduction. And this is aimed at the principle that we've espoused; that is, the quantitative safety goals should promote the optimum allocation of public resources in reducing public risk. That once you've protected the individual and you've satisfied yourselves that the risk to the public balanced against the benefits to the public are not outside the realm of reasonable values, one should then look at balancing cost and benefit and achieving reductions in residual risk.

We are, in terms of benefit, talking about the population risk reduction; that is, if you make a change in plant design you should evaluate both the probability and consequences effect in terms of reducing the integrated risk under the curve. That is, you will reduce the risk in terms of man rem per year.

The cost involved is the total cost of making that change to the plant.

In this respect, it's important to recognize that the quantitative goal for this value should not be "conservative" with respect to other values because if we set a level that is too high in terms of cost-benefit, thepublic will be disadvantaged in terms of misallocation of public funds in areas where they're not getting a commensurate benefit for the investment.

And also in this respect, the cost-benefit criterion also promotes or encourages small population density siting. That is, if you have two identical plants, one on a 1000 person per square mile site and one on a 100, an individual change in that plant design will bring a cost-benefit value of 10 times higher for the higher population density than for the lower. So it does promote, again, the concept of low population density siting, and would encourage utilities to site plants in lower population areas.

The number that we're proposing as a suggested goal is on the order of $100 per man rem, and the basis for that essentially is that it is equivalent, using a linear dose fatality coefficient of 1 death per 10,000 man rem, of $1 million per life saved. We're not proposing that we assign a value to human life. The important point here is that we're trying to gain the optimum allocation of public resources, and at that level we believe it does reflect a value that's somewhere in the median of what is achievable in terms of public risk reduction for the expenditure.

And I'd like to illustrate that with this slide. It tabulates the cost-benefit ratio in terms of millions of dollars per life saved for a variety of industries and regulatory approaches for a variety of public risk areas. And as you can see, in the nuclear power plant area, for instance, the rad waste effluent treatment system, if we're using the $1000 per man rem criterion, we are providing marginal reductions in public risk on the order of about $10 million per life saved. If we look at coal plants and providing sulfur removal equipment, scrubbers, on coal plants, depending on whether you're burning high sulfur coal or low sulfur coal, you can be ranging anywhere from $100,000 per life saved to $10 million.

In the area of occupational health and safety, the celebrated case that's before the Supreme Court involving OSHA's benzene regulations has decided to provide a cost-benefit value of about $300 per life saved, and other areas such as coke fume regulation, those regulations have been cited for about 4.5 million.

If you go down the list to environmental protection of about $2.5 million per life saved, into fire protection where the Consumer Product Safety Commission had proposed regulations that were based on a comprehensive cost-benefit analysis that would return a cost-benefit ratio of about a half a million dollars per life saved, and it appears that those regulations will not be imposed because they are perceived to be inflationary and unduly restrictive on the furniture industry.

By comparison, installation of smoke detectors in residences across the nation have been estimated to provide cost-benefit ratios on the order of $50,000 per life saved. But yet, there exists no comprehensive regulatory policy to require they are used. It's more or less left up to the local or regional jurisdictions in those cases.

And you can go into the area of automotive and highway safety where numbers such as $140,000 per life saved have been used explicitly as a cost-benefit criterion for regulations, and various studies have shown that improvements such as airbags could provide cost-benefit values on the order of $320,000 per life saved, but as we all know, those regulations are being delayed or being re-thought because it's perceived that they are too expensive.

And furthermore, and seat belts return a cost-benefit value of about $8000 per life saved. And we all have seat belts in our cars.

MR. WILSON: Is that number allowing for the fact that most of the seat belts aren't used?

MR. O'DONNELL: This number, I believe, is allowing for studies on the use of seat belts and their effectiveness.

And in the area of medical and health programs, various studies have shown that increased use of kidney dialysis treatment units and mobile cardiac units provide cost-benefit ratios on the order of $200,000 or $30,000, and various cancer-screening programs that are voluntary for most purposes, are extremely cost effective in reducing fatality risk; on the order of $10,000 to $80,000 per life saved.

So the number we're proposing of the $100 per man rem is more or less in the median of these things that are certainly not the cost-effective use of funds, but then one questions whether people should be forced to take cancer-screening programs, and you get into the area of voluntary risk reduction versus involuntary. And we feel that a number such as $1 million per life saved, which is reflected in the $100 per man rem, somehow falls in the median range and is probably a reasonable number for a starting point for a safety goal on this issue. And it does, incidentally, fall in line, again, with some of the numbers that have been proposed. The NRC's $1000 per man rem, of course, is well known and it's generally considered to be excessively conservative and that is reflected, I think, by the previous slide. The EPA somewhere along the line proposed a number of the order of $75 per man rem. In Germany from a paper by (?) indicated they are using or considering using numbers such as $100 to $200 of deutschmarks per man rem, which at the current rate of exchange works out to be about $50 to $100. Dr. Rogers from Northeast Utilities in a study he did of occupational exposures and ways of reducing them, suggested a number like $30 per man rem. So the number we're suggesting, $100 per man rem, again is not outside the range of what others have thought reasonable.

MR. SHEWMON: You're mixing here those to employees and those to the general public. Is that right?

MR. O'DONNELL: I don't believe so.

MR. SHEWMON: Would NUREG-0110 apply to -- I've seen it applied to employees.

MR. SHEWMON: No, it's applied to rad waste systems, reducing public risk from normal releases. It does not apply to occupational exposure. I think Dr. Rogers' number is more related to occupational exposures.

The last of our criteria, and it's one that I do not have a number for. It is somewhat controversial within the industry, and many people feel that if you've met the other three criteria that are specified there is no need to establish a limit for core degradation probability. That is, on how often you will have a serious accident involving core degradation.

There's another body of opinion that says that NRC aside or public safety aside, we must ensure ourselves that we cannot repeat a TMI within the next 30 years or whatever, and that from the purely economic point of view, utilities must ensure themselves that the frequency of events such as TMI is very low and therefore, this criterion would probably supplant or replace the other.

Again, there are mixed opinions on this within the Atomic Industrial Forum and within the industry, and I just propose that this has something that is probably ancillary to the other; that if you met the other it's conceivable you wouldn't need this one. On the other hand, if you did establish a criterion such as this, which would have a number of benefits, that is, if you did limit the probability or the frequency of core degradation accidents to the order of one per several decades, that would provide some minimum requirements on accident prevention as opposed to mitigation. Because presumably, if one had the individual and population risk criterion, and that was the only criterion that one had to meet, it's conceivable that one could build a plant without an ECCS provided one had the best containment going and you could, in fact, limit the individual risk to 10^{-5} and the population risk to .1 fatality per year.

It does provide some -- this sort of criterion would provide some minimum requirements for accident prevention; that is, you have to ensure that you do have the ECCS and you do limit the types of accidents that would result in potential exposures to a small value.

It would have the further attribute of reducing the frequency of stress-provoking incidents to the populations in your plants to a small value. And thirdly, it would obviously limit the economic risks associated with accidents. This is not a safety consideration but it is certainly one that is of great interest to the industry and to the utilities.

And fourthly, it would simplify to a great extent the conduction of risk assessment analysis. That is, because it does not involve consequence modeling or site-related assessments of risk.

But there are a number of problems with this, and for that reason I am not prepared to propose or suggest even a number to you; one of which involves the definition of core degradation. It could range from something as simple as the preparation of a number of fuel pins up to complete core meltdown. And until there is some definition of what one means by core degradation, it's very tenuous to suggest the numbers that would go along with that sort of criteria.

This slide attempts to tie things together somewhat for you in terms of the individual risk criterion and the population risk and to give some indication of where one might expect to fall with res ect to the goals that we are suggesting here.

On the ordinate here, -- on this scale is the maximum individual risk in terms of fatalities per year. On this scale is the population risk in terms of fatalities per year per 1000 megawatts. These are the two criteria that we have suggested be included. What I've tried to do here is relate with a very simple model how these things would tie together. And using uniform population distribution model, and also an exponential model for attenuation of individual risk, we can draw lines that represent the risk from, say, a 1000 megawatt unit and a population density of 500 per square mile. And also, for a 1000 megawatt unit sited at a site where the population density was 100. And this would reflect four units at a similar site.

It indicates that if we did draw lines of acceptance of 10^{-5} per year for individual risk, and .1 per year, that it in some way the existing plants that we have on line. If we take WASH-1400 and the level of risk that is reported in that document; that is, about .25 fatalities per year per 1000 megawatt, and the individual risk is somewhat below 10^{-6}, it would give us a point in here. This reflects the fact that if we took the WASH-1400 plant and moved it to a site with 1000 people per square mile as opposed to the Surrey site with about 200, we would come bumping against the .1 fatality per year criterion.

So it indicates -- and also, if we took the WASH-1400 plant and put 12 of those units on a single site, we would then come bumping up against the individual risk criterion.

Again, increasing the number of units at a site increases the individual risk but does not increase the risk per 1000 megawatt. Increasing the population density does do that.

So it indicates that if you have numbers such as these, within the current framework we could probably site plants such as WASH-1400 if one believes the level of risk that is reported, within limits up to high population density sites, wherein one would have to look more closely at the level of risk. And it also indicates that we could consider plants such as Energy Park where you would have a large number of units sited at a low population density site. And these limits would not govern the design or siting of that plant.

It also indicates that if there was an individual plant that was, in fact, 5 times less safe than WASH-1400, we would then be increasing both the individual risk and the population risk, and we would come bumping, again, to the .1 for a situation where a plant was about 5 times less safe than WASH-1400.

What this indicates to me is essentially that by our criteria we are saying that one must get within this boundary without regard to the cost. Once you're within this boundary, one should consider cost-benefit balancing. If one believes that the current generation of plants is approximated by the level of risk that is represented by WASH-1400, we are probably for most plants already within this realm where we should be looking at cost-benefit balancing, with the exception of unusual circumstances such as high population density siting where you may be impinging on the population level, and we should be looking, as our primary consideration, at backfitting and involving new criteria to reduce residual risk further, and primarily the cost-benefit number.

As a further prospective on this, I indicated up here the level of risk that would be associated with a clean 1000-megawatt coal plant, first at a site with a population density of 100 per square mile, and one where the population would be 500 per square mile. And again, these represent the lower end of the scale and would reflect a plant that is burning low-sulfur coal with full scrubbing, and also one in which the maximum offsite concentration of both sulfur dioxide and (?) are one-tenth of the clean air limit.

So this indicates that if we have numbers such as this for nuclear plants, it would be extremely difficult for a coal plant to meet those limits, and if we are really looking at a generic set of safety goals that will apply across the board, numbers such as these may be unduly strengthened for other energy sources such as coal. And that doesn't mean that the level of risk for nuclear power plants is going to shoot up; all it means is that the level at which we apply cost-benefit balancing will be expanded out to another realm. And we are, in fact, in that regime of risk where we should be looking primarily at cost-benefit balancing.

MR. SIESS: I don't quite understand the diagonal lines. What is the variable that moves you along those lines?

MR. O'DONNELL: The variable is -- there is a relationship, linear, between individual risk and population risk.

MR. SIESS: As a function of --

MR. O'DONNELL: As a function -- well, the plant design obviously will drive the individual risk. And the relationship between individual risk and population risk will be described by a constant which is representative of the site condition. And depending on the site, the slope of this line may change one way or the other. It's just meant to provide a simplified model for somehow relating these things. And the number that I've chosen essentially for the slope of this line reflects the relationship of individual and population risk that one would derive from WASH-1400.

There's a number of caveats that I'd like to include in this discussion. Number one is that we have suggested numerical limits here this morning, and it's important to recognize that when you pick numbers you should also pick or identify the level of risk at which you expect to do risk assessment. That is, we're proposing that the numbers, whatever they are, should be -- once they are appropriate for mean value or best estimate, risk estimate, and that if one does more conservative risk estimates these numbers are not appropriate. If you get out and said 90% confidence level and very conservative assessments of risk, numbers like 10^{-5} or .1 fatality per year are not appropriate.

Number two is that whatever initial set of values are picked should be proposed as interim, and for a trial use period of about three years. We just do not have enough experience with using this type of concept and these kinds of approach to say that these numbers are it and these are the final numbers. But it's very important, in our minds, to get on with the task of developing these numbers and starting to use them in risk assessment.

And third is that we should not be totally mesmerized by the numbers in the quantitative goals, and this is particularly important when one comes up against the grey areas or the borderlines when you may be 20% under or over the line, and you cannot, in our opinion, make black and white judgments solely on the basis of risk assessment and numbers. It has to be tempered to some degree with qualitative judgments.

MR. OKRENT: You're out of time for presentation purposes. Dr. Wilson?

MR. WILSON: I have three questions. The first is -- you said you want to take the mean value of the risk; yet, as far as I understand the number from Rasmussen you've taken, and WASH-1400, that's what he thinks is the best value of the risk and not the mean. And roughly speaking, he's got above-normal distribution both of risk and of consequence. And if you calculate mean from that best value, it comes to considerably higher, about two or three or four times higher.

MR. SAUNDERS: This gentleman means median value or maximum.

MR. WILSON: Okay. That I think is an important point because I think, in fact, I would certainly make an important distinction of that point. It's important to be exactly clear on what you're doing. If you were to take the mean value rather than the median, you would also include a procedure for -- you want the procedure to include the uncertainties and giving encouragement to reduce those.

The other thing is we've talked several times about this accident frequency question, rather than the specific individual societal risk. That's really a secondary criterion or divide criterion. Everyone who has mentioned it this morning has said that we want it because we want to separate probability and consequence because we want to concern ourselves with economic concerns. And I think it's very important to make that distinct as a secondary or divide criterion. I think I agree with that.

MR. O'DONNELL: You mean as far as the core probability?

MR. WILSON: Yes. It's not as important --

MR. O'DONNELL: I agree with what you say.

MR. WILSON: It's a procedural thing which I think is right, but it isn't the initial thing you're aiming at.

MR. KERR: In applying your cost-benefit criterion at $100 per man rem, where would you propose to cut off the calculation as far as individual exposure is concerned? Would you have no cutoff? Would you cut off at so many millirems exposure?

MR. O'DONNELL: Well, one can't do cost-benefit on an individual basis; you must do it on a population basis.

MR. KERR: But you have to cut off your exposure at some point, to the individual.

MR. O'DONNELL: I'm not sure what you mean.

MR. KERR: Well, you calculate the exposure over a geographical area of so many individuals. Now, do you calculate out to zero, a millirem or a rem for the individual, or do you cut off at some value in your application of the cost-benefit criterion?

For example, you might argue that one should cut off at background, or an increase that's equal to half background. You didn't mention that and I think it has some influence on the number you get.

MR. O'DONNELL: Yes. I think, for instance, if you use an exponential attenuation model, that after 20 miles or so you don't get significant contributions to total risk, population risk.

MR. KERR: The current Appendix I, for example, requires that you calculate out to a 50 mile radius and take everything you get within that.

MR. O'DONNELL: I wouldn't want to project that. I think you should take the risk assessment out to whatever level or distance at which you get significant contributions to risk.

MR. KERR: Yes, but you have to decide what's significant.

MR. O'DONNELL: I think if you start getting down to below background levels and smaller, the integrated risk begins to make a very small contribution to the total population risk.

MR. KERR: You haven't really decided, then, on a cutoff for individual exposures.

MR. O'DONNELL: No, we have not.

MR. MARK: This is, indeed, just a general question and I don't ask that you answer it all. You referred early in your discussion to all technologies or risk-related activities. And also in the tables there was a variety of things presented. It seems to me that in a discussion such as this, perhaps not in the form considered today but in a form considered for action or alternate discussion, one should select the risks to be put in the scales in some fashion which perhaps would not include such term as "violence". They ought to be risks with either society or some section of society imposes on an individual or a group, consciously and deliberately. And that would hardly put a thing -- it would be questionable whether automobile accidents belong in the comparison.

I'm saying the obvious thing, that this debate over whether something is voluntary or involuntary needs to be avoided.

MR. O'DONNELL: I guess we thought about that and it's a very nebulous quantity to define what is voluntary and what isn't. On one extreme end of the scale one could say well, living near a nuclear power plant is certainly voluntary, and one doesn't have to do that.

MR. MARK: I understand that, but one does probably need to select one's risk comparisons with that kind of a question in mind and not get something in there like violence.

MR. O'DONNELL: I think it's appropriate, if we're going to compare those things, to look at the risk of living near a coal plant, the risk of living near an airport, the risk of living near a hydroelectric dam, and to try to compare things on that basis. Because they are essentially -- they have the same degree of volition involved and the same degree of benefit involved to the individual and to society. We cannot refine things to that extent that we're willing to throw out automobile accidents as a valid basis for comparison because I'm not sure you can throw that out.

It's questionable to what extent driving an automobile in this society is voluntary.

MR. MARK: I didn't want to argue the point.

MR. CATTON: I have a question about your table on cost-benefit ratios. The second item, ECCS, you show roughly $10.00 per man rem. Is that number correct? That doesn't --

MR. O'DONNELL: Yes. That would be if you applied the actual plant as your first engineered safety feature. That is, if you had a plant devoid of safety features and you were going to say well, I'm going to build in safety features to this plant, and if you provided, first of all, an ECCS before you provided a containment or diesel generators. At that level of risk, the ECCS would, in fact, return a cost-benefit ratio of about $100,000 per man rem.

MR. CATTON: Isn't that a bit low for REG GUIDE 1.11?

MR. O'DONNELL: You have to take the context in which it's perceived. That is, if you build a plant without a containment then its level of risk is rather high, and reducing that level of risk is rather easy. That is, if you provided an ECCS to a plant without a containment, your level of risk would drop significantly. It's an incremental, starting from base of the plant with no safety features.

MR. KERR: It's a high return on investment.

MR. CATTON: I understand that.

MR. OKRENT: I think we're going to have to cut off now because we're running half an hour late. We'll take a 10-minute break and if Mr. Temme is here 10 minutes from now we will resume then.

(Short recess.)

DR. OKRENT: Mr. Temme.

MR. TEMME: Thank you, Dr. Okrent, for giving me the opportunity once more to brief you on what is happening in IEEE Working Group 5.4. I would like to start by introducing myself since there are perhaps a number of people here who aren't familiar with what this working group is about.

I am Mark Temme. I am with General Electric Advanced Reactor Systems Department, and also chairing IEEE Working Group SC-5.4. This a working group under the Reliability Subcommittee of the Nuclear Power Engineering Committee of IEEE, focusing on the subject of risk criteria.

Now, I want to mention that I am not going to be giving you any exponents of 10 to deal with this morning. Our working group in its three months of existence has focused on establishing our charter, the parameters of risk criteria, and in general trying to lay the groundwork so we can understand what detailed information we ought to be gathering to complete our objectives.

I think we have made some significant progress toward those goals.

The general topics I am going to cover are what we have agreed upon as the application and scope of the risk criteria which we will be writing in a standard, something about the implementation procedures that relate to the criteria, and finally, a brief view of our near-term plans and perspectives on how we see our activity fitting into the larger scheme of things.

This chart essentially paraphrases the official scope statement which exists now in the IEEE standards files. Essentially we have agreed that as an industry standards committee, our task is to establish rules of practice by which we govern ourselves, and one thing that you might note is that we are not saying we are writing a standard for NRC to use. That simply doesn't really fall within our purview, although, of course, we would welcome NRC's use of the criteria that come out of the standard, and even their adoption of the standard as an official document.

However, we see value in what we are doing regardless of the extent to which the standard finds its way into the safety review and licensing process. And that is silpky ht fhvds us ' b'sis for the intelligent judgments regarding engineering tradeoffs, tradeoffs that deal both with the design of new plants and with the continued operation of existing plants.

One of the first subjects we spent some time on, and much of our effort has been trying to put bounds around the problem we are dealing with, has been what kinds of reactors plants should be governed by the standard and what aspects of their operation.

Our conclusion, particularly influenced by the work of Dr. Slovic and associates, say that it is virtually an intractable problem to develop an acceptable risk criterion by starting at the back end and asking what is acceptable and what is perceived and so forth. We have taken a rather more pragmatic view, looking at what can we do and how can we do it in a responsible manner.

First of all, on the reactor types, we believe we should focus specifically on light water reactors, not because we fail to recognize the existence of other types, but that seems to be what we can do something about now, and as to other reactor types, perhaps we have a little more time to deal with them.

We spent considerable time discussing the issue, and it is an important one, of to what degree should such a standard be retroactive. That is, should it be something that you only look at when you are designing brand new plants, or should it have something to do with plants that are now operating or committed.

We didn't feel that we could responsibly exclude operating in committed plants from such a standard; therefore, our present intent is to write a standard that does include them. However, we think that there may be some wisdom in providing some special provisions for operating plants. We haven't gotten into specific debate over how that will be done or even if it will be done, but what comes to mind are ideas such as different numerical values for the same type of standard, or perhaps grandfathering clauses for certain existing plants.

This is a fairly important issue and we expect to be spending some time discussing it.

Again, after some discussion and debate, we chose to prepare a standard which governs only the risks due to operation of the reactor plant. Some people, particularly some of the consultants to our working group, feel that it is very important to bear in mind other aspects of risk from the nuclear plant cycle.

We don't disagree with that, but we feel those other aspects such as reprocessing and waste storage and so forth, can be treated separately. However, in doing so we are cautioned to keep in mind that there may be some implicit overall risk budget which society will accept or allow to the entire nuclear fuel cycle, and therefore we don't want to use it all up for the reactor plant risk.

We are proposing to address public health and welfare risks, that is, things outside the plant boundary. Occupational risks, of course, can't be ignored, but we think there again is something that can be treated separately and, in fact, to a degree are treated satisfactorily, we believe, by existing occupational safety standards.

Another controversial item was the business of human effects, common cause failures and what we have labeled here as off-site litigation. By the latter we mean such effects as environmental attenuation of radiation, evaluation procedures and so forth, which do affect risk and generally tend to mitigate risk.

We don't feel that we can responsibly prepare a standard or a criterion to be met which leaves out these aspects of risk, and therefore we propose to incorporate them into the standard, recognizing that there are some technical difficulties which follow in that we have to be able to do a risk assessment which includes these effects in order to compare to the standard.

Risks from sabotage and war, again, we specifically exclude from coverage by the standard, essentially because they are fairly controversial and it isn't easy for us to see how they can be managed, at least by the same standard that governs the risk due to operation of the plant.

DR. OKRENT: Could you define sabotage as you are using it?

MR. TEMME: I'm not sure I can precisely. For example, one can probably conjure up acts of sabotage which are simply done by workers in the plant and which initiate an accident which might conceivably fall under the risk of operation of the plant. We haven't come to grips with the fine line between what is sabotage and what is included in normal operational risk. I suspect that is the sort of thing you are getting at in asking the question.

DR. OKRENT: I am trying to understand why it is excluded. It certainly could affect your design if you included it.

MR. TEMME: Well, we are not proposing to exclude it because we think it does not affect the design or that it shouldn't affect the design. We are simply taking the position, at least at this point, that those aspects of the design and the operating procedures for the plant which deal with sabotage need to be treated separately from what we do to comply with this particular standard.

One of the issues that arises, if we choose to incorporate it into the standard, is what is the probability of sabotage. It is very hard for us to see how we could get some appropriate estimates of that. Is that a satisfactory although perhaps not satisfying answer to your question?

DR. OKRENT: Why don't we proceed.

MR. TEMME: All right.

Another item, which I haven't listed here, which we have discussed at some length, and I don't feel we have quite resolved it within our working group, is the question of whether or not a risk that is compared to the standard should include the risk of normal operation of the plant. I think that probably the weight of the consensus at this point is to exclude that as a contributor which we incorporate into the standard; however, it is still a bit controversial within the working group and we need to discuss it some more.

DR. OKRENT: Let me come back to the sabotage question a little bit because I think it is fairly important. Suppose somebody was talking about an L&G facility, with which we are not emotionally involved, and they said we propose to bring this tanker into the heart of New York City but we have an ultrasafe design and it can collide with anything and there is no problem, or whatever, but it happened to be vulnerable to a person carrying some kind of weapon, maybe a high powered rifle, that that could go through the many barriers or so forth.

Should that be excluded from consideration in judging the risk of the imposed importation of L&G?

MR. TEMME: I don't think it should. By the same token, I don't think the risk of the potential of sabotage should or will be excluded from consideration of the risk of nuclear plants. Our only issue here is where in the process of designing and licensing and operating the plant do you consider it.

We feel that there are enough controversial aspects on a technical basis, that is, how do you model the risk due to sabotage, and enough differences in how you would model that risk and analyze it that we ought to separate that aspect of evaluating risk from the aspect of evaluating the risks due to accidents.

That is essentially the argument that we have for excluding it from the standard; not that we don't think it is important. In fact, in discussing this matter we find that we have people in the working group who say, well, we ought to exclude it because it is a very small contribution to the risk. We have other people in the working group who say this is the major contribution to public risk.

That by itself suggests that at least within our working group, we don't really have a good handle on what to do with it.

DR. OKRENT: The reason I am pressing you a little bit on this is you said the standard is written to be used by designers, and --

MR. TEMME: And plant operators.

DR. OKRENT: Now, if instead of excluding it here, you said it is also necessary to develop another standard which will provide guidance to designers with regard to what they should do for sabotage, I might see that it wasn't being left out of this picture. But at the moment it seems to be a hole, and it is, in fact, a difficult design question, but I think the designers are going to have to address it if we are going to do something beyond guards.

MR. TEMME: I don't think we would disagree that a standard or procedures or rules of practice need to be established to deal with it. In that sense, I think we will say that something needs to be done about it. We are not proposing, at least at this point in time, to do something about it as a part of this particular standard.

And some of these other things that I just talked about create similar holes in the sense that the standard isn't going to cover every aspect of the design question.

Now, in beginning to prepare ourselves for coming up with quantitative criteria on risk, the first sort of question we addressed is what kinds of parameters should appear as the top level risk criteria. In general the process by which we are hoping to arrive at a standard is evolving into one in which our consideration of matters such as this, the form of the criteria are largely based on what we refer to as internal considerations, that is, the things that matter to us as representatives of the nuclear industry, the things that motivate us to control risk.

We anticipate that as we get into trying to determine numerical values for these risk measures, that we will focus then more strongly on external considerations such as public perception of risk and so forth. At this point we focused mostly on what seems to be important to us in doing a responsible job of producing an adequately safe plant.

In discussing our list of possible top level criteria, the things that seem to emerge as important to us in choosing among them are to preserve to the greatest extent that we can flexibility to the designer in siting the plant and designing it, and also to the operator in dealing with new information that he has to respond to.

Another objective that was important in our decisionmaking here was the ease of application of the standard. To a degree I view that as an objective that conflicts with the first one. Perhaps the greatest flexibility, for example, would be achieved by expressing the risk in terms of public health effects. That also leads to the most complexity in doing the risk analysis to prove that you have met the criteria.

Our conclusion on top level criteria -- and I emphasize the phrase "top level" here because we have had a lot of discussion about various forms of criteria which we feel may eventually appear in the standard in the form of an allocation from the top level criteria. So there are some things that we haven't ruled out totally.

Our conclusion at this point is that at the top level, the risk criteria that we state can be given in terms of radiological dose rather than in terms of acute and latent fatalities and illnesses and so forth. The reason for this is that we feel that it avoids a certain amount of complexity imposed on the user of the standard. It avoids his need to deal with the controversial aspects of biological effects models and so forth, and he doesn't really lose any flexibility by it in that he is still in a position to take into account remote siting, low population density and those factors which we believe ought to be accounted for.

On the other hand, of course, in order to rate these criteria on radiological dose to public health effects and to compare them to the risks of other form of electrical generation, now our working group has to get involved somewhat with biological dose models, their uncertainties, in order to make our own evaluations. We think we can do that.

We also discussed at some length the idea that contamination of land and loss of the use of land seem like they ought to be important impacts on public welfare, if not health, and therefore they ought to be treated in these criteria, and they didn't seem adequately treated by a measure of dose.

So we are proposing also to have a top level criterion that relates to off-site property contamination. Now, the specifics of both of these have yet to be addressed. In fact, this will be the main topic of our next meeting where we will get into such issues as whether or not to have limits on individual risk, such as limiting the probability that a given individual will receive a certain level of dose, or societal risk, which would be more, perhaps, in the form of the probability distribution of the number of people receiving a given dose, and there are alternatives to both of those.

What kind of dose? Are we able to express this in terms of whole body, or do we have to have several criteria for specific organs? And there are all the questions of expected frequency or probability distributions. All of these aspects of the definition of the criteria we do intend to discuss and arrive at specific decisions on.

DR. KERR: Mr. Tamme, what is the significance of the term "top level"? Is that synonymous with --

MR. TEMME: Specifically what we mean is what is the highest level in terms of relation to public health and welfare risk, what is the highest level at which we have a quantitative criteria in the standards document? Specifically we are aimed at producing a standard which says here is the risk that must be met and here is the risk limit that must be met.

Perhaps it becomes clearer in the next Vu-graph, but there are such measures as core melt probability and so forth which are related to risk, but we view them as being at a lower level in the tier of things.

DR. KERR: Thank you.

MR. TEMME: In fact, to give you a little view of things that we did consider specifically, this is substantially the list of things we have looked at in arriving at our previous conclusions. Health effects I have already talked about. We have considered the risk fee idea as a way of defining risk. While we haven't debated it at great length, there was a general feeling that it wouldn't necessarily be viewed as a responsible way to limit risk, like people outside the industry.

It has the connotation, for example, that if a plant owner has paid his risk fee, then he is no longer liable for anything else that happens. In general we didn't feel that that was something we could adopt. Probably most of the other -- in fact, all of the other items on the list here are in the category of what have been described as hazard statements, released to or from containment, released from site, extensive core damage, and so forth. These are all things that could appear at an intermediate level in criteria.

In every case, the use of one of these items as a top-level criterion limits the flexibility that the user of the standad would have to make the tradeoff decision. Release from containment, for example, does not give any allowance for remote siting or population density and so forth.

We had another criterion that we used here in considering parameters, in that something that you want to assign a probability to must be, at least in principle, something that could be observed. On that basis we decided we would not have a specific criteria on such things as metal trauma. In fact, we also felt the same way about linking cancer fatalities.

There are a couple of other proposals for ways in which we can express top-level roles which we haven't debated specifically: for example, the concept of stating the criterion in terms of comparable risk, simply the risk of a nuclear plant should be no greater than that of a coal plant of the same size.

We will spend a little time discussing that idea also.

The point was made here earlier today that these criteria have virtually no meaning until one defines the analytical procedure by which you show that you have or have not met it, and we feel very strongly in our working group that this is the case. Therefore, for any of this to work, there must also exist a set of rules governing how you use data, where you get data, how you use probability models and so forth that go with the criteria.

Now, it is not within our current charter to write those rules, but there are other standards writing groups that are working in those areas and we propose to defend on their results.

There have been several proposals that risk criteria should be based on limiting the frequencies of important or dominant accident sequences. This could be done either by specifying dominant sequences in the standard or perhaps by giving rules for their identification.

We have discussed this possibility and concluded that we can better preserve the flexibility of the user of the standard to make rational choices by simply addressing the criteria to the integrated result of all sequences, or in other words, more or less as it was done in WASH-1400.

There have been some thoughts presented on criteria which had a point at the lower end below which everything is okay, and some higher point or curve above which nothing is acceptable, and then a region in between which is reserved for cost-benefit or ad hoc decision making. We have debated the wisdom and the need for that kind of criteria, and our present conclusion is that we are aiming toward a simple pass/fail kind of criteria. If you are below it you are okay, and if you are above it, you are not okay.

On discussing the pros and cons and why it might be desirable to have a range in which you do special decision making, the idea of uncertainty in doing the analysis emerged, and it seemed to me, at least, that that was one of the primary motivations for having more complicated kinds of criteria.

We certainly recognized the need to address uncertainty in the analytical results in writing these criteria, and propose to do that. Precisely how, we haven't determined yet, but there will be requirements that the analysis include propagation of uncertainties, and we must then devise some rules for how those uncertainties are used in deciding whether or not the criteria are met.

Now, when I addressed this committee in April, I showed you a schedule which indicated that by about now we would have some tentative numerical goals, and as you can see, we don't. However, I feel that we are pretty much on track with respect to our end point of having a draft standard which we would submit for approval of other standards committees, just prior to the middle of 1981.

We have focused ourselves on trying to understand why we are making the decisions we are making, and debating and resolving and coming to consensus on fundamental issues even at the expense of not being able to conform to a rigid schedule. In fact, we feel that this is probably a more traditional approach to standard writing.

The next things that we will address specifically are writing the detailed specification of our criteria on dose, and then we will collect and review and use the data which we expect to form the basis for coming up with quantitative numbers. These include evaluations of the relation of dose to health effects, looking at the risks of presently operating plants, alternate forms of electrical generation, value impact of risk reduction, and these various things which I referred to before as external considerations.

We have a growing list of consultants to the working group. We have had a lot of interaction with Steve Derby, who worked with Dr. Slovic in developing the approaches to acceptable risk document. We expect that we will be asking Dr. Slovic himself some questions. We have Dr. Smith of the East-West Corporation and Dr. Velosky of EG&G working with us in the area of data, and we expect to identify a few more people as we find the questions that we want to ask.

I would like to close with maybe a note of humility here in addressing the question of how does all of this that I have been talking about fit into the problem or the question of what is an acceptable risk for nuclear plants. This introspective view of our role in the scheme of things puts us rather at the front end.

Again, I want to emphasize that we have taken to heart the massage that has been delivered by Dr. Slovic and others that starting from the other end and asking yourself what should, will or can society accept makes the problem intractable. Therefore, our intent is to do a responsible job of considering cost versus benefits versus safety from the other end, and writing a proposed standard by which we intend to govern our own actions, and allowing ultimately the societal decisionmaking process to react to that.

I might point out that before it gets that far, it has to go through a good deal of review internal to our industry because standards don't get out without that.

I think, if I might make a personal observation, we have a good group of people working on this, and I find that their enthusiasm and their optimism that we can accomplish our objectives have grown considerably in the few months that we have worked together. We feel that this is something that we can do.

Thank you.

DR. OKRENT: Any questions or comments at this time? I will make one comment in passing. It seems to me in the process of adopting standards that it will meet the industry implicitly is taking a position on what it thinks is acceptable risk or accepted risk or tolerable risk or whatever you want to say. I don't think you can avoid the question by saying we are not going to state this explicitly.

MR. TEMME: I'm not sure I follow. I don't think we are trying to avoid the question or to avoid stating explicitly what we think will be acceptable. I don't think we could be rational and just shut that out of our consideration, and we don't propose to.

It is a matter of where you put your emphasis when you begin your problem which perhaps is somewhat involved.

DR. KERR: Associated with that question, it seems to me, is Dr. Okrent's earlier question about the sabotage. If you leave that out, for whatever reason, it seems to me implicitly you are saying it is not important. Is that in line with your view, or is it still an open question?

MR. TEMME: That is not in our view. I guess the question we have to ask ourselves is is that the way it is going to be perceived, because that would bother us.

DR. KERR: I don't see how one could avoid that perception if one is trying to design power plants that have an acceptable -- whatever that means -- risk to the public. If sabotage is an important contributor to public risk, then one ignores it either because one considers it inconsequential or because it is somehow a risk different than other risks and one doesn't have to consider it.

MR. TEMME: I would agree with that if I felt that we were really going to come out with a standard that covered all of the risks of nuclear power.

DR. KERR: I am talking about a nuclear power plant, now, not the whole fuel cycle.

DR. SLOVIC: You are assuming that sabotage is a signficant contributor to risk.

DR. KERR: No, I am saying --

DR. SLOVIC: Then let me remind you that he put that caveat in it.

DR. KERR: No, I said that if one ignores it, one must be assuming it is not a significant contributor. Otherwise, one is not calculating the significant risk of a power plant. I may be completely wrong. I am trying to understand your rationale in excluding it.

MR. TEMME: Our rationale is a fairly practical one. We want to break the problem of determining what risk to design into tractable pieces, and we view this as a separation of two aspects of the problem. Maybe it is not a good separation and could be improperly received. There is a need to examine the relationship between the risks that are initiated by sabotage and the other risks we do propose to cover on this.

What we have not done so far is commit ourselves in this committee to address that, but we aren't saying that it shouldn't be done. In fact, if it begins to appear that what we are producing is not worthwhile because of that omission, then we have to reexamine our position. I think it is debatable at this point whether we should include it or not. Our conclusion is we won't.

DR. OKRENT: Thank you.

The next speaker is Mr. Bernero of the NRC.

MR. BERNERO: Good morning. I am Bob Bernero of the NRC staff. I have some cats and dogs for the Subcommittee.

First of all, I know a subject of interest to the Subcommittee is flood risk, at least to Dr. Okrent, and we have a recent memorandum, an internal memorandum in my staff, about the flow code effort, which is a flood prediction code.

I am sharing the memo with you. I haven't had too much time with the staff to discuss it, but it seems to indicate that the flow code would calculate that three times 10 to the minus 4 exceedence probability for exceeding the probable maximum flood on the Susquehannah River at a band of, I think it is almost a decade up and down for 95 percent confidence to a 5 percent confidence, by one technique; and it is roughly an order of magnitude below that by another technique.

I don't know whether this aggravates your concern or subdues it.

DR. OKRENT: What does below mean?

MR. BERNERO: It is a lower probability of exceeding the probable maximum flood of 1.6 million cubic feet per second. So I distribute those. There are 15 or so copies of that memorandum.

The budget status, which I think you are interested in, is as of this time, as you know, we have two budget levels, one we call the PPPG, Policy, Planning, Something and Guidance, and then we have a requested level. At this time the user offices, who are to endorse at least 85 percent of our budget, have endorsed well above the probabilistic analysis section there, that Systems and Reliability Analysis Decision Unit. They have endorsed well above the PPPG level, and the Office of Research includes the remainder of that budget in the 15 percent discretionary allowance.

So in the sense that the systems and reliability analysis budget for TAS is concerned, it is sufficiently endorsed by user offices to be approved, but whether that means we get all the money remains to be seen because, as you know, the funding for FY 81 and 82 is still somewhat suspect.

DR. OKRENT: Would you repeat what you said?

(General)

MR. BERNERO: It is probably obscure if you haven't followed the thing.

DR. OKRENT: Is the endorsement at the lower level or the upper level of FY 82f?

MR. BERNERO: I better repeat it so that the audience isn't totally bewildered. In the NRC for the Office of Research and Programs there is a requirement that no more than 15 percent of the research budget may be discretionary, that is, not requested or approved by offices such as Nuclear Reactor Regulation or Nuclear Material Safety and Safeguards.

Consequently, our budget is shared with them, shown to them, and they endorse it or do not endorse it, in pieces. We have a budget that hinges on how much money the FY 82 budget will be. We have a level the Commission has proposed to us for guidance purposes, and we call that, because the Commission's guidance document is called that, the PPPG level.

There is also the proposal by the Office of Research that even more money be alotted for research, and that is a higher number, not surprisingly. Within that research budget is a decision unit or element called Systems and Reliability Analysis, which is substantially what all of us think of when you say PRA or risk assessment or something like that.

That element of the research budget has been fully endorsed by the user offices. In addition, they have endorsed some of the increment between that PPPG level and the full requested level. And the remainder not endorsed by user offices has been endorsed by the director of research, Bob Budnitz, which thereby sanctifies it within his 15 percent discretionary allowance.

So, in effect I have full approval for the higher budget figure, but that doesn't mean I will have the higher budget figure. We are facing constraints not only in dollars but in manpower.

DR. OKRENT: Just one more detail question. In the information that I have, it wasn't clear to me whether IREP and that kind of thing is any more in SARA, Systems and Reliability Analysis.

MR. BERNARO: Yes, IREP is in there. You don't have detailed notes on it, on what it is. I will send you those separately. But it is in there.

DR. OKRENT: Is there a new sub-element?

MR. BERNARO: No, it is in Systems Analysis.

DR. OKRENT: It is.

MR. BERNARO: Yes. It is in a sub-element of Systems and Reliability Analysis which is called Systems Analysis. It is Sub-element number 3, I believe. I am not clear on the number.

DR. OKRENT: It wasn't covered in this memorandum on --

MR. BERNARD: Perhaps it would be best if I went over that with you afterwards.

DR. OKRENT: Thank you.

MR. BERNARD: So, with that, we are to meet again with the full committee on July 8th, is it, or 9th, sometime next week, and I am not sure how much more information is developing back at the ranch, but there is quite a bit of concern. The manpower constraints is turning into something that is quite difficult.

I would like to talk a little bit about our work in development of acceptable risk criteria. We are not at a stage where we have answers yet. We have got a body of work going on that should bring some significant fruit out by the end of this fiscal year, by October. I see a problem here and now that I think needs pursuit, and I would like to raise that problem here in this forum in view of the audience as well as the subcommittees present.

The work of the NRC in research on acceptable risk is including a number of elements. We have at Brookhaven National Lab an effort doing a literature survey or a literature review of all of the people who have used risk analysis, quantitative analysis of one sort or another, so that we have a convenient repository of that information for evaluation and comparison.

In addition, in what I would call the philosophy of using risk assessment or quantitative standards or goals, we have two efforts going on. In Dr. Slovic's work which we have contracted for, he has already produced a substantial work. It was February of this year, I think, that the draft was published and circulated, both here and to certain people abroad for comment.

We have gotten some comments back on that document, and he is working on the final, and I believe it is August or so, something like that.

VOICE: The end of this month.

MR. BERNERO: At the end of this month, the end of July, we should have this ready. This will be a discussion of approaches, I believe. He could characterize it to you far better than I could, but it is a discussion of approaches to the use of acceptable risk criteria.

In parallel with his effort, we have Brookhaven National Laboratory working with us on a somewhat more regimented approach, and it includes the concept of if you can't agree on acceptable risk, at least try to agree on what is unacceptable. Along with that, Brookhaven is working to develop quantitative criteria, the associated data and methodology that might be used to demonstrate that you satisfy those criteria or you don't.

In addition, since the question comes up in this context, they are rebaselining WASH-1400. They are taking the human error out so that one might be able to look at accident sequences, the whole risk picture, and say something, in effect, about how much is the design's fault and how much is the human.

Now, that is a pretty delicate thing because, really, when you go into these things, the human fault can contribute 50 percent or more of the risk. In a way, one might make the analogy to try to sort out how reliable should an automobile be to be separated from how reliable the driver should be, or how drunk or how sober.

DR. OKRENT: Excuse me. When do you expect there to be some kind of output from the activities you just described by BNL?

MR. BERNERO: All of these should have reports out in each of these areas in this fiscal year, by October. We expect to have a report. There should be f[illegible] BNL reports: one on literature review, one on the accep[illegible] risk decisionmaking process, a separate one on t[illegible] data and methodology, and a separate one on rebaselining the WASH-1400. You should get all of those by October of this year.

We are very much interested in using the probabilistic approach in licensing, not necessarily to replace. We have to be cautious. The idea of quantitative licensing i a long way off, quantitative in the sense of replacing deterministic criteria. But this could be so valuable as a supplement. If it is going to be a supplement and deliver these values so that you have a systematic evaluation of things and you know where to look, where to apply priorities, this has to mean something to both the industry and the regulators.

If you are going to use risk assessment, it has to be a two-edged sword. It has to be considered as something that can point the need for a safety improvement and also point the unnecessariness of an apparent safety improvement. A ratchet wrench should have a reversal switch on it.

Now, one of the difficulties we are foreseeing, and much of this has been covered by some of the earlier speakers, is the societal aversion to high consequence events is generally apparent, but how much should that be brought into our consideration in acceptable risk criteria for nuclear.

We do have the problem, and it has been discussed many time, that the nuclear accident has the capability of killing immediately or later through the latent cancer fatality. It is clearly wrong or inaccurate to sum up latent cancers and immediate deaths; they are not the same. But by the same token, one shouldn't just wave off the latent cancers as negligible.

So some rational method is needed to compare those. The special societal perception of nuclear risk, I sometimes think we feed it as much as deal with it. In showing compliance with probabilistic criteria, we have many very real problems, not the least of which is the difficulty of predicting human failure.

In this area we have a lot of separate work that you are aware of in the human failure analysis, but in there lies one of the places where we can fall into fruitless argument to a very great extent, because in a way, the human error contribution is a natural place to fight the battle if one is inclined to drive the probability of something happening up or down, because it is so influential on the outcome.

There will be bound to be a lot of controversy about the human errors used in risk assessments.

Now, we have previously recommended -- or suggested, I should say -- a probability of core damage. Many of the standards proposed go to the actual risk of the nuclear power plant, the risk of a fatality, the risk of a man/rem exposure as an index of bodily damage.

That requires you to go all the way through the probability of events, the models of containment failures, the models of release, the models of radioactive transport, the models of emergency response by the public and by officials moving them in or out or sheltering them. It is an extremely complicated and long and tortuous calculation. It picks up all sorts of baggage along the way and a lot more uncertainty, and it may not be the most practical way to go. It may not be the most practical standard to start with.

That is one of the reasons we are suggesting the possibility of going only as far as the probability of severe core damage. Someone said earlier you can't define severe core damage. There is a lot less uncertainty defining that than there is demonstrating and carrying along the uncertainties in that demonstration of actual radiological exposures from a class 9 sort of accident.

DR. OKRENT: Excuse me. Are you suggesting that this be the only criterion?

MR. BERNERO: No, something like an initial one.

DR. OKRENT: What do you mean by an initial one?

MR. BERNERO: You know, ultimately if there are criteria for fatality, for trade-offs and such, those are good, and they would include due treatment of both the probability and the consequence models appropriate to the calculation of whether you satisfy that standard or not.

But now, where we stand with difficulties in treating uncertainties, with difficulties in agreeing on methodology, in difficulties in reaching even the standards we would try to meet, it may be appropriate to look for a temporary expedient at the core damage level.

DR. OKRENT: I would discourage you from that.

MR. BERNERO: Well, let me go to my final slide because you have heard this before about the warning range. In fact, I think one of the earlier speakers even mentioned it. The probability of core damage for operation and design, we suggested, might be that one times 10 to the minus 3 per year for a single plant is a possible threshold of change or correction, with a warning range a magnitude below it. And then if one went in and removed the human element, perhaps just design only, taking the human error out, that a core damage probability of one times 10 to the minus 5, with warning range a magnitude below that is a criterion that might be used.

DR. SHEWMON: Pardon me. Have you defined core damage?

MR. BERNERO: No, we haven't. What we have done, it is like WASH-1400. WASH-1400 used the shorthand version "core melt" for "substantial core damage." It had no way to analytically disti guish the failure of the system. In simple terms, if e design of the system needs one high pressure injection pump to cool the core adequately according to some ECCS analysis, and one low pressure injection pump, let's say, just arbitrarily, just one of each, at a certain condition have to be operating, if one of them fails, the analytical system can't analyze whether that would, indeed, lead to substantial core damage; butit presumes it.

DR. SHEWMON: So if you are substantially overhearing a significant fraction of the core --

MR. BERNERO: No, I would be more inclined to define it as if you are substantially not in satisfaction of the theoretical cooling requirements. There are theoretical cooling requirements in the plant design and in its safety analysis --

DR. SHEWMON: I suspect we could work with the first law of thermodynamics and prove that those two are close enough to equivalent.

MR. BERNERO: Well, there is probably a fairly conservative bias in the theoretical cooling requirements. That is the thing that the risk assessment just doesn't have the time or the resources to sort out.

DR. KERR: This discussion reminds me a little bit of a story my wife told me. When she was taking ninth grade algebra, she had a colleague in the class who passed in an exam and said "I can't work these problems, but here are some I can work."

(General laughter.)

MR. BERNERO: There's a lot of truth to that.

I would like to point out, in case the Subcommittee hasn't sat down and thought about some of the things they have heard, that there are an awful lot of quantitative risk assessments or probabilistic risk analyses going on right now, and I am not even sure that this list is complete. The industry, to my knowledge, the owner, is doing Big Rock Point. The owner, on NRR request, is doing the Limerick Plant.

Zion and Indian Point have joined together. They had a preliminary risk assessment done on the two plants, and they are now -- wait a minute. I'm not sure the preliminary assessment covered more than Indian Point. I have to say that. Did it cover Zion as well? I only read the part on Indian Point.

In addition, they have contracted with Picker, Low and Garrick and others, and they are doing a much larger effort as a follow-on to that, a much more complete and substantial risk assessment. Ed Zebroski mentioned that NSAC is working on Oconee now. A fairly large effort on Oconee has just gotten under way.

NRC has near completion the Crystal River evaluation done under IREP in a program we call RSSMAP, Reactor Safety Study Methodology Application Program. We are just completing the reports on Sequoyah, Oconee, Calvert Cliffs and Grand Gulf. Those will be published shortly.

The IREP continuation, which we are just getting under way tentatively right now -- we originally had Zion and Indian Point in it, and at their request, because of this new large effort, they are shifting out of it, but we do expect to go ahead with Arkansas Unit 1. That is the B&W plant at that site.

Millstone 1, which is a boiling water reactor of fairly early vintage, Browns Ferry 2, and Calvert Cliffs we hope to do further in IREP.

Lastly, and I believe I am supposed to brief you on i. tomorrow, we just did a very short-term miniature risk assessment of Indian Point II and III for the Commission as part of the Task Force.

If I look at all this, it is obvious that in 1980 and 1981 we are going to have before us a lot of quantitative risk assessments while we are discussing the philosophy of dealing with probabilistic risk assessment and its standards or criteria of acceptability.

I like to use the expression that when the regulators get up in the morning, and Harold Denton, in particular, as the principla legal regulatory of reactors, he puts on his trousers and accepts the continued operation of three dozen reactors with the left leg and three dozen reactors with the left leg.

Now, if he has before him a risk assessment, as he does today, warts and all -- Crystal River has gaps, omissions, challengeable assumptions, and it is in peer review and people are criticizing and commenting on it, but there it stands on the table in front of him and his people, and it says there is one serious release sequence that has a probability of one times 10 to the minus 4 per reactor year.

And then over in a basket there is a whole bunch of things, a probability of about 10 to the minus 5 per year, but they all add up to smaller releases with the combined probability of about two times 10 to the minus 4 per year, and all of those sequences for the small releases seem to be heavily laced with human error, human difficulty, human confusion or misoperation.

What is he supposed to do with that? What is the owner supposed to do with it? While we are talking for the next two or three years, that thing sits on the table and it gets refined day by day, month by month, and there is a very serious question of whether we really do need an interim criterion or at least a policy dealing with such things.

Now, Saul Levine just gave a talk in Stockholm, Sweden, and I think it derives from the draft message that one of the earlier speakers spoke to. I brought two copies. I am sorry I didn't have 15 copies for the Subcommittee, but I at least had two. I brought them along.

Basically, if you go into WASH-1400 -- this is the last figure of Saul's talk -- but if you go into WASH-1400, you will find this figure again, all of the miscellaneous risks, and he constructs in there a logic which says if the rest of the risks of everyday life are bounded at the lower side by something like air crashes on your head where you are on the ground, it would not be unreasonable to give a safety role ten times lower for the nuclear power contribution to everday life risk.

That is consistent with the philosophy that the new risk shouldn't lurk above the background of everyday risk for people. And then this was the WASH-1400 early fatality curve for 100 nuclear power plants on the presumption that Surry and Peach Bottom are indeed prototypical of 100 nuclear power plants, and there is serious question about that.

Using a thesis proposal, Levine combined early and latent fatalities against this curve here representing 100 nuclear power plants. That gives one the ultimate way to compare -- and Saul, I know, has suggested to me that he thinks that is the idea way to compare plants, draw the risk curve, draw these risk profiles for immediate fatalities, for latent fatalities, for early injuries, for property damage, for whatever, to draw them and compare them.

Well, if you do that you are getting caught up with that complexity of modeling and who did the analysis and which version of the code did he use and whose meteoroligical model, and so forth. Things can bog down. Comparing risk curves coming from two different sources can be extremely difficult, and what we are getting in that list I showed you before is a lot of risk assessments, and in many cases, from a lot of different sources. We don't have standardized approaches.

Well, that leads me to suggest the consideration of something, and this is so tentative that it is written on flammable celophane with washable disappearing ink, not in your handout, and this is not completely inconsistent with the suggested core damage probability.

Perhaps interim criteria like this, not formally drawn, not a rulemaking, but at least what to do until the doctor comes, until there is a better and more rigorous way to judge, that if you see a reasonable point estimate of core damage having a probabilty in excess of 10 to the minus 2 per year, it might indeed be appropriate to say fix that damn thing in days or hours; in effect, shut the plant down.

For this range, 10 to the minus 3, to 10 to the minus 2 probability, fix it, but there are months; fix it at the next shutdown, that sort of thing, a reasonably prompt repair or correction. In many cases, remember, some of these things are so heavily laced with human factor contributions it is changing procedures or it is making fairly minor changes to a plant to correct some risk contributors. Event B out of the Surry Plant in WASH-1400 is a classic example. The dominant risk contributor could be suppressed mightily with a simple testing regimen.

Then, in the range of 10 to the minus 4, to 10 to the minus 3, there are years to deal with it. That one in Crystal River I mentioned. If one takes that at face value, one times 10 to the minus 4 probability, a blackout kind of a core melt, you could take your time and really give that careful consideration on should there be another D.C. bus in there, should there be another auxiliary feedwater pump or some independent dedicated shutdown heat removal system.

One could deliberate responsibly on that. This whole range here, to either fix in years or possibly not fix, is reasonable rather long-range deliberation, and 10 to the minus 5 probably is in the range for good, that you don't need to do anything with it.

Now, I am not sure that we can say this is the criteria. I put this up. This is highly personal. I really think that we have to have a way, a sensible way, not too conservative and not too optimistic or tolerant, to deal with the information we are coming up with. The only other alternatives, there are really two.

We can have all these risk assessments and try to compare them with these elegant risk curves, and any of you who have ever looked at a CCDF curve drawn by one party and compared it to another know that they don't even use the same size graph paper. It really is a difficult job. The only other alternative is to keep doing that and put them on the wall like works of art.

To a very great extent we did that with WASH-1400. WASH-1400 was done, and then for five years -- well, four years after the final was published -- all we did was argue about that report, and we did precious little with the messages in it. And there was a grievous penalty on March 28th of 1979 because of that very thing.

So I think it behooves us to think in terms of a very short-range goal, how to deal with these things here, and this is one possible approach. But I gather you would discourage us from going after core melt probability.

DR. OKRENT: Well, you are talking here about something which is related to but I would say is different from the more general subject of quantitative risk criteria. This is only intended to provide the regulatory staff some basis for thinking about a need to take action and so forth.

MR. BENERO: It is a preliminary risk criterion, really.

DR. OKRENT: Well, let me explore a little bit. If it is this simple -- well, I need first to have some definitions. This is the probability of core damage from all sources or from any single source? As show here.

MR. BERNERO: No. This is the probability of core damage, and this gets into a very sticky area. This is the probability of core damage from internal or external events, I would think.

DR. OKRENT: Is it the sum or is it per source of --

MR. BERNERO: It would be the sum. Now, one can deal with individual risk contributors, and that example from Crystal River is a good one. There is a single dominant accident sequence. I don't know what the seismic risk contribution in Crystal River is. No one has looked at it, to my knowledge. But I can deal with that one accident sequence in this range. I can get a perspective on the system design in operation, sub-sum, and deal with it with some perspective.

DR. OKRENT: Let me go on with just a brief look at this. I am reminded of a movie I saw some time ago. It seems to me it had to do with 1984, which is now getting nearby, and some kind of a revolution that was to make all people equal. But then as society evolved, some people became more equal.

I think some accidents are more equal than others. So it is conceivable to me that there might be one kind of an accident where you could say 10 to the minus 4, fix in years, sounds plausible if it involves core damage. There can be another accident, maybe conbined with the site, where you would say I don't have years.

So I have made the point, I think. This is still an incomplete portion of the decisionmaking process I would make, even for interim criteria.

DR. SHEWMON: Now that you are convinced, maybe you can explain it to me. Are you saying that this thing you are concerned about, whereas you weren't concerned about the other one, is an accident that would have much higher consequences, or what?

DR. OKRENT: Sure, because there could be one event where if it occurs, you expect the containment has a high chance of mitigating the situation, and furthermore, it may be in a rather remote site. Another event also leading to core damage could have a probability of a half or some big number of leading to a large release, and furthermore, there may not be such a good site.

DR. SHEWMON: A minute ago he said, if I understood him correctly, that this is only an undercooled core, and he leaves out of this thing possible routes from there to severe damage, like what can you do and that sort of thing. Or are you now getting into that?

MR. BERNERO: No. Let me use the Crystal River sequence as an example because I think they display Dr. Okrent's concerns. In the use of such an arbitrary scale, the responsible use of it, there is one sequence at one times 10 to the minus 4 that is a severe release. It is a Class PWR-II release. That is that sort of a blackout sequence.

There is a sum of sequences that adds up to about two times 10 to the minus 4, each of which I call low grade core melts. They are lower class releases, much, much lower consequences. So that if someone is looking for the perspective of how to judge that particular analysis of that particular plant, that risk assessment carries right through to a projected release, although it didn't do any elegant calculations of it. It was using some limited judgment.

It is a remote site. The Crystal River site, in case you are not familiar with the demography, is quite remote. It is a remote site. It has one serious release right at this line, and then a collection of cats and dogs that adds up to be about right here of far less serious releases, and that has to be considered in making any decision on what to do.

DR. SIESS: Bob, in limiting this criterion to core damage, does that mean that you are satisfied that there are not accidents that do not involve core damage that cannot have greater consequences? Spent pool fuel drops?

MR. BERNERO: The spent pool fuel accidents, I think, are not in the same class of risk reach as core damage accidents.

DR. SIESS: I am talking about not risk, but consequences.

MR. BERNERO: That is what I mean. The potential just doesn't seem to be there. In our studies we are not pursuing them at all.

DR. SIESS: Not just the classical accident, but let's say dropping a shield block into an open reactor or dropping a cast on top of all the spent fuel. What are the consequences of those?

MR. BERNERO: We haven't done any new work that I know of since WASH-1400 in my group.

DR. SIESS: I have seen some analyses recently that show fairly high consequences, but I don't know what source they are based on.

MR. BERNERO: I know there is one argument afoot these days, I think on the Salem case, that you coiuld start a zirconium fire in the spent fuel pool and the thing would go off like a flare, and you would get a great big dispersion term or something. It is in one of the spent fuel pool expansion cases. I am not familiar with the particulars on it.

DR. SIESS: But what you are proposing is based on the assumption, at least, that core damage does lead to the greatest damage.

MR. BERNERO: Yes, that that is where the risk is. Now, I would emphasize that in our perspective, our view, core damage is not synonymous with substantial off-site risk. It probably varies a great deal from plant to plant, but in WASH-1400 and in later studies, there is evidence that says that in some cases the majority of "core melts" or severe core damages won't cause substantial off-site releases, and indeed, one can have a fair number of curies released without killing people.

So there is some kind of a multiplier that is probably plant specific that says given the probability of severe core damage, the probability of substantial release is something lower and the probability of serious bodily harm to people off-site is lower still. But that is where one gets into the elegance of modeling that makes it more difficult.

Right now I think that is something that has to be done subjectively in some crude framework like this as an interim measure. I think we ought to face that problem.

DR. OKRENT: Yes. Will you identify yourself, please?

MR. THOMAS: I am Garry Thomas from NSAC. I think you have just made the point I was hoping to make, the fact that you use an interim criteria based on core damage alone. Should it not be based, from a regulatory viewpoint, on consequences to the public? Core damage by itself can be stopped. In the middle of an accident it can be stopped and result in essentially no consequences to the public. That was one point.

The second point is somewhat related. You have talked about negative impact of the operators, that is, taking out operator errors. Have you talked about or are you considering in your studies the positive impact of the operators where they can, in fact, aid in stopping an accident sequence progression.

MR. BERNERO: To answer the first, I agree with you that core damage isn't a risk except to the owner of the plant who has to pay to fix it up. It is a risk insofar as it can hurt the people off site, do substantial bodily harm. In an ultimate sense, you have to consider the actual reach, the fraction that does cause the off-site harm.

As far as the positive value of humans, we are trying to do that, but again, the systems are imperfect. The modeling is imperfect for dealing with that.

DR. OKRENT: I think in fact there are sort of two different things on the table at the moment. One is the general subject of quantitative risk criteria, and the other is Crystal River.

MR. BERNERO: As an example.

DR. OKRENT: As an example, but as an actual example, not a hypothetical one. With regard to the latter, I would like to request, both on behalf of myself and the Committee, that we manage to get copies of whatever it is that Mr. Denton is able to look at when he is putting on his trousers.

(General laughgter.)

DR. OKRENT: I have been pressing the staff to tell us what they were learning from that study, and up to now with not very great success. It seems to me that it would be of value, and perhaps, in fact, you might almost say require that the staff and the ACRS discuss that specific study and what should be done as a specific item.

For example, if I go back to what the staff said in, let's say, considering what to do with regard to environmental impact statements, the estimate of an accident having the kind of release you have just discussed was about four orders of magnitude lower probability than what you have said. In other words, 10 to the minus 8 per year were the kind of numbers that were used in 1972 or so, and even in WASH-1400, the estimate of a PWR-II is, I suppose, roughly two orders of magnitude.

MR. BERNERO: Well, 1-1/2 or so.

DR. OKRENT: All right. I said roughly, but in that ballpark. So this is something separately to think about, and also I don't know what the uncertainties are here. I assume there is a range up and down. This is probably a best estimate number you are giving us rather than a 90 percent confidence number.

So I think, in fact, you should try to get the documentation to the Committee as soon as possible and try to schedule at a early but effective time a basis for discussing the report and its implications. This kind of thing would likely come up, but again, I have indicated in my own thinking you have to somehow factor in whatever the other factors in addition to the core damage probability are in the decisionmaking.1

I don't think we should try to arrive at an answer on the interim criteria via this particular subcommittee meeting or try to discuss this at a great length.

DR. KERR: I think we should commend the try.

DR. OKRENT: Oh, indeed. I think if you don't try to put something specific on the table to talk about, it is hard to make progress. So I agree with what Dr. Kerr said. I think Whipple and I some years ago tried to provoke discussion by putting proposals on the table in that light.

I will take one or two comments and then we are going to adjourn for lunch.

MR. NOYES: Larry Noyes, Philadelphia Electric Company. I would reiterate your comments about the (inaudible words). I am really concerned about the fact that perhaps you find this sort of plan of the uncertainty of (inaudible word) modeling with perhaps an even greater uncertainty of subjective determination of what the consequences are.

DR. OKRENT: Dr. Wilson.

DR. WILSON: I would reiterate my point I made earlier, that although I think you almost certainly have to do this, I think any time it gets written down as a proposal, it should get written down in a larger framework or the larger framework would get lost.

DR. OKRENT: There is a hand at the back. Will you please identify yourself?

MR. WEAVER: Bill Weaver from Babcox and Wilcox. I would like to ask if the PAS has endorsed or reviewed the IREP on Crystal River. We at B&W have reviewed it pretty conclusion and come to the conclusion it is a very poor analysis, and that recommendations on changes should not be based on that particular analysis.

MR. BERNERO: We are in the midst of a process we call peer review where we have had extensive comment from the owner, from B&W, from other parties within the NRC staff, and we have had extensive meetings on it. We have j st concluded those and now we are compiling what all the comments are.

DR. OKRENT: It sounds to me like it is a semi-public document, in view of the people --

MR. BERNERO: Well, there have certainly been a lot of bootleg copies of it floating around.

DR. OKRENT: As I have indicated to you before, I find it curious and also quite disturbing the flow of information to the ACRS is sometimes what it is.

DR. KERR: We could try the Freedom of Information Act.

(General laughter.)

DR. OKRENT: Indeed, I was going to ask Mr. Quittschreiber to find out what the procedures are whereby an ACRS member, an individual, can use the Freedom of Information Act, just as an example.

DR. SIESS: I think perhaps we are not considered to be peers.

(General laughter.)

DR. OKRENT: I think it will be awkward if I have to use the Freedom of Information Act to get information that I think should have been sent to the ACRS. This is an old point. I think things are improving over what they were in recent years, but you mentioned the earlier four reports, for example, that sit around in draft form in the PAS for years, not available to the ACRS.

VOICE: One.

DR. OKRENT: Well, again, I think it is something to continue to reflect on. I will leave that as a request for Mr. Quittschreiber. As an example I may just try to do that.

Well, I think what we might do is instead of taking the next speaker and having a delayed lunch, we will pretend we are almost on schedule and will reconvene at 1:30, assuming the next speaker, the chairman and the recorder of this session are here, and the rest of you can take longer if you wish.

(Whereupon, at 12:30 p.m., the meeting was recessed, to reconvene at 1:30 p.m. the same day.)

AFTERNOON SESSION

(1:30 p.m.)

DR. OKRENT: The next speaker is Dr. Joksimovic.

DR. JOKSIMOVIC: Good afternoon. I would like to thank Dr. Okrent and the members of the Subcommittee and the staff for inviting me to present our views on this important subject. I am Voyen (phonetic) Joksimovic, Manager of Safety and Reliability at General Atomic. I manage our ACGR light water reactor, waste management and synthetic fuel safety and reliability program.

I brought two of my close associates with me, W. J. Haughton and F.S. Zombec, who have participated heavily in preparation of the statement that you have before you. I have also taken the liberty of appointing Joe Haughton to be my Vu-graph operator to save a few minutes of my precious time. With your permission.

With that, I will start out with the first Vu-graph. I will try in the first statement to reflect the views of designers, since I have to deal with them on a daily basis. What we have observed in this country and other affluent societies is that the regulatory objectives have often been stated as those that have to achieve the highest standard, lowest achievable release of radioactivity or carcinogens, and from the standpoint of designers and from the standpoint of general design criteria, the kind of st.ff that I have to deal with, I maintain that they are neither reasonable nor adequate.

When these objectives are translated into practice, then we end up with things like releases as low as reasonably achievable, we end up with things like general design criteria plus regulatory guides. I maintain, and it is obviously my opinion, highly subjective, that both have experienced serious shortcomings. They are ambiguous and they tend to make designers concentrate on on meeting requirements rather than stimulatinag innovation and legitimate tradeoffs.

I also maintain that the safety standards achieved are largely dependent on individuals assigned to particular plants.

So what is the remedy to all this? I hope it is in the introduction of quantitative safety goals. I have a reason to believe that they will provide consistent safety criteria and thus firm framework for designers. They will, hopefully, avoid ratcheting of plant requirements. They will leave designers with freedom to create effective design solutions to real problems they face, like plant availability.

Additionally, they will protect public and investment more effectively. They will enable designers to get on with the job of designing future reactors, and they may contribute towards better public understanding of nuclear safety.

Now, with that I would like to state that at General Atomic we have been rather vocal in voicing our concerns about the licensing process, and we have been advocating now for a number of years the use of probabilistic risk assessment techniques, in safety assessment, safety R&D recommendations, and the licensing decisions.

On the screen I have a summary of some of the activities that we have been involved in, like a statement before Professor Lewis' (?) committee in December of 1977. We have also offered a statement before the Udall committee in February of 1979. We presented a paper at the AIF workshop in May of 1979.

Our president, Dr. Agnew, has written a letter to Kemeny in September of 1979. Bill Haughton and I have presented a paper at the Knoxville conference dealing specifically with the safety goals, and we have a so-called licensing topic report on selection of design basis accidents using PRA as a rational, which is currently being printed.

Now, we have stated our views in the Knoxville paper as a kind of progress report. We have an ongoing program in this field which the Department of Energy funded, and some key elements in our approach are kind of evolving. I have a list of them which tend to summarize the kind of notions that we have created so far.

The first one I have on the list is professional judgment. The reason why we use that is that we recognize it is up to the society as a whole to decide what the safety goals should be. The society hasn't done that, so we believe we can assist in this process to the society that we serve, simply so we can get on with the job and design future reactors.

We term this professional judgment, and there will be a figure I explain subsequently which gives you a rough idea what I mean by that.

In our approach we do take into account competitive risk studies as a very valuable background. However, we do emphasize the knowledge of probabilistic risk assessment studies as WASH-1400, such as our own AIPA studies, such as the Deutsche Risikostudie, and numerous other studies that are currently in progress.

We draw significant knowledge from it which enlargens our horizons and enables us to put in better perspective various comparative risk studies. What we also do is we believe that some key elements of the existing regulatory framework ought to be retained. One concept that in our thinking should be retained is the concept of design basis accidents.

However, we believe that the PRA should focus that process towards a more rational and systematic selection of events. We have a report on that subject that I will also spend a few more minutes later on explaining what that concept is.

So basically we are talking about using a design basis accident type of region, and we are also talking about adding to regions, one which we call safety margin and the other which we have called safety research. I will explain in a minute the definition of those.

We are also in our work trying to reflect the perception of the public. I recently have been reminded that Dr. Kissinger in his book said that perceptions are more important than reality. We have tried to account for that by specifying the concept of no identifiable public injury. In this context, the identifiable public injury means changes to human body tissue or fluid caused by radiation which can be identified by medical examination.

The other thing we have converged on is the use of limit lines in a similar fashion to the work performed in Great Britain by Red Farmer, Kinchin and Peter Congendon. In our considerations we have recognized that in a democratic society, we can talk about a spectrum of politics which can range from what we call an emphasized risk policy to a balanced risk policy.

We have recognized that for the public safety goals to be meaningful, that we have to talk about acceptable rather than nonacceptable ones, and in that context we think that one has to talk in terms of acceptable individual risk, acceptable societal or collective risk, and acceptable public property risk.

We do endorse the subject of reliability goals and investment risk goals, and we are working on this, but this particular presentation is completely focused on public safety goals.

Now, here is the diagram which we have reproduced from the Knoxville paper, which makes an attempt to specify on one diagram what are the major steps in development and implementation of public safety goals. The first block over here deals with all the necessary steps one has to go through before one can come up with what the goals are.

To the right we have listed some key implementation activities, and what I am talking about over there is there is a need for all of us to speak, by and large, the same language. It doesn't mean that there is only one language which is reccgnized. We can take the type of situation like the United Nations where we could have four or five languages, or in other words, we may have four or five acceptable. I don't believe we can have 27, 39, whatever.

I think there is a need to go through a streamlining process so that we can communicate and we can effectively implement the safety goals once they are set. Having gone through these type of activities, there are a variety of applications which we are recommending, but I do not intend to spend any time on that since it is outside the scope of my presentation today.

This portion over here emphasizes the knowledge of probabilistic risk assessment studies, and the portion above emphasizes the knowledge of comparative risk studies. There is a block over here which is public perception of the nuclear risk. It is tough for us to handle that; however, we try.

Here is a diagram which we have termed the map of quantitative safety regions. The primary objective of this diagram is to relate accident frequencies versus design rigor. We do have an actuarial data base from operation of nuclear power plants in the United States and elsewhere; howeaver, we have to go beyond that data base to predict the kind of accidents of lower frequency. PRA is the only tool useful to accomplish that, so PRA and that type of notion is used to search and try to find the type of events which may influence safety, and hence those events can become design basis accidents. By that we mean that they are subjected to rigorous scrutiny.

The power plant designers can then take the first step in mitigating against the consequences of such accidents which might otherwise be unacceptable: for instance, being here rather than here, by designing equipment and procedures for operators so that the consequences can be kept suitably low and/or the likelihood of the accident can be reduced.

This entails a cost to the plant, and hence one can perceive that these accidents pass through some kind of toll gate. The resulting reductions in impact of potential accidents also reduces investment risks. Hence, we are proposing that the accidents which are likely to happen pass through this toll gate or design basis process, which is to say that any accident with a frequency such it will be likely to happen in the United States over the total currently projected operating history and the design basis process.

Even events which have only a 50 percent chance of happening are included in this region. Based on this percent and taking a U.S. program of 166 reactors having a typical lifetime of 40 years each, the mean frequency at the lower boundary of the region is found to be 10 to the minus 4 per reactor year, as shown in the figure.

And hence that is our definition of design basis region. We realize we cannot stop there, that there is a need to draw further down, and hence we have defined the safety margin region to provide a proven margin against some events whose probability of not happening in the current U.S. program is between 50 and 90 percent.

The lower mean frequency limits of this region of 10 to the minus 5 per reactor year is compatible with a 90 percent assurance that an accident sequence below this region will not occur. Events predicted to lie in this region would not be expected to occur; hence, there should be no blanket requirement to automatically design for them.

A suitable margin against the unlikely chance that they may happen can be obtained by some form of design capability, which may be inherent or designed into the plant, or maybe brought to the plant if a potential severe accident threatens to occur.

Design capability requirements also entail a cost to the plant, and hence the notion of a second toll gate is appropriate.

Now, there are some rare events which can be predicted below the safety margin region, that is, having a mean frequency below 10 to the minus 5 per reactor year. We call that a safety research region. A further step down in the frequency to 10 to the minus 6 per reactor year is equivalent to 99 percent assurance that an accident sequence below this frequency will not occur.

These sequences should not impact public safety. However, the significant members of this group which may have high consequences are candidates for safety research studies and experiments. The results of these studies would be subjected to peer review to determine conclusively which region these frequencies were in and whether the consequences were low enough.

In these studies we are going to take several years to achieve meaningful results. The plant designs might need to include features to mitigate such events until the assurance of their insignificance was confirmed. Additionally, many people believe that high consequence events between 10 to the minus 5 and ten to the minus 6 are unacceptable.

Hence, on an interim basis it could be suggested the design capability requirements also be considered in this region.

Now, the research region does extend below 10 to the minus 6 per reactor year; however, it will be devoted to determining the correct frequency for extremely rare events. Sequences having accident frequencies below 10 to the minus 6 have no impact on public safety and therefore would not require any additional design considerations.

Having defined this, I am now going to deal with specifics. We have reached the level when we have the first draft of our quantatati for individual risk, societal risk, and public property damage. This is a diagram that we have composed for the individual risk, which is a line drawn through three points.

At the point of 10 to the minus 4, this characterizes our design basis region and a consequence of no identifiable public injury. We have interpreted this to mean something like protective action guide in a situation of 5 rems. We also learned from talking to various radiobiologists that if an individual is exposed to that kind of low radiation, that it cannot be detected in his bloodstream, and hence the point meets our intent of no identifiable public injury.

We extended the line in the spirit of our attempt to retain as large fraction of the framework as possible. Then we had plotted another point at 10 to the minus 6, which is the point of no early fatality. We took the threshold to be of one percent, which is translated into a dose of approximately 200 rems.

Having drawn this limit line, it happens that it passes at 10 to the minus 5 checkpoint. To give you the perspective of where some other lines are, we took the liberty of interpreting Ed Zebroski's points -- I think he left, but anyway I will still go ahead and tell you where we think they are.

Here is his first point, core melt, and I think here is the second point, containment failure. Also, here is a point from the Farmer diagram, and we have plotted two other points. One is the dose exposed, the maximum an individual was exposed in the TMI accident, and the second one is the TMI EAD (?) dose of one rem.

When we drew the line, we established there is a slope of minus 1.2, which said that basically this is our risk aversion factor for individual risk goal.

Now, the subject of societal risk goal is much more complex. As I have said, in democratic societies there could be a spectrum of policies ranging from a policy that emphasizes some risks irrespective of their contribution to total risk. We call that emphasized risk policy. On the other hand, we could have policy which is based upon a balanced risk, where an attempt is made to take into account what are the comparative risks, and then the nuclear risk is put in that type of perspective.

Hence, having decided that one can have a range of public policy limit line, then we faced the problem of finding where the points should be. Consistent with our approach of no identifiable public injury, we have plotted a point of 10 to the minus 4.

With regard to the balanced risk policy limit line, we felt that in a rational society, that we could talk about a small fraction of latent cancers in a population surrounding a particular nuclear power plant, and if we can choose a number which is a small fraction of the incidence rate that we should be acceptable to a balanced risk policy type of society. The number that we chose was in between one-tenth and one-hundredth of the spontaneous cancer rate, which was down to something like 100 latent cancer fatalities.

Having plotted that point, we have concluded that others in the society are very risk averse to this type of possibility where a large number of people can be exposed to a disease like cancer, and hence we felt that that should be reflected in a slope which is greater than minus 1.2, which was individual risk, and we chose a slope of minus 1.5, which is comparative to the kind of notions that Red Farmer introduced about 14 years ago.

So we simply drew the line through this point at a slope of minus 1.5. On the other hand, we felt that we should also come up with a rationale for a very restrictive type of society which emphasizes some risks like nuclear, and we have emphasized, using our professional judgment, where the point would be, and we have converged on a number like .1 latent cancers. We drew the line through that point, the same slope of minus 1.5.

Now, the last thing that we did is to come up with a proposed quantitative safety goal for the public property risk. We have given the least consideration to this in our curren work, so this can be viewed as a first cut, our first attempt to do that. I will give just a few lines of thinking behind this curve.

We have been hired to be consultants to the insurance companies, and we have learned on that contract how they think. What we have observed is that they are concerned about losses in the range of $10 million to $300 million dollars. So we took those types of points into consideration.

We have also observed that they are not interested in frequencies below 10 to the minus 5. That is the kind of cutoff frequency for their thinking. So we have plotted 300 million at 10 to the minus 5 as our point, and we drew our line to a slope of unity because we decided that those people are very rational in their business life, and hence there is no need to have any particular risk aversion.

Hence, this line is basically, as I said, our first cut, and it encompasses the notions from dealings with insurance company type of people. In doing so, we have taken into account absence of the Price-Anderson Act, and we have also taken into account this kind of risk category they are willing to put on anything, which is $300 million.

Thank you very much.

DR. OKRENT: I am a little bit intrigued at your comment that they are not risk averse. The fact that they have a cutoff on what they are willing to insure suggests that they are quite risk averse, in a certain sense.

DR. JOKSIMOVIC: I think that is true. However, I think when you try to get some kind of a slope, they are perfectly willing to accept a linear one, as long as you come to some kind of risk ceiling beyond which they are not willing to risk.

DR. OKRENT: We have a few minutes for discussion. Are there any questions or comments?

(There was no response.

DR. OKRENT: I guess not. Thank you.

Let's continue. Next Dr. Griesmeyer will discuss the current stages of development in work he is doing as an ACRS fellow.

DR. GRIESMEYER: We are working on a risk management framework, trying to come up with ideas as to how do you manage the risk. In part of that framework there is going to be some quantitative decision, quantitative risk acceptance criteria or non-acceptance criteria, depending on how you define them.

In setting up this risk management scheme, you have to realize that basically you are talking about a number of decisions on various interacting levels. You have to decide whether you need the electric power. You have to decide what sort of technology you are going to use. You have to select the site. You also have to decide whether the risk is acceptable or not.

At each one of these levels you are going to have to look at economic impacts, sociopolitical impacts, environmental impacts and health and safety impacts. So what we have is a set of interacting decisions and a set also of interacting impacts for each level of the decision.

So public health and safety risk acceptance criteria is not going to be unique because it depends upon economic impacts, sociopolitical impacts and environmental impacts, as well. So what we are trying to do is come up with some sort of rational approach that looks at risk acceptance, public health and safety and environmental impact in a way that is compatible with the rest of the decision.

If the risks and health and safety impacts and environmental impacts, the two alternatives are roughly comparable, depending on whatever weighting factor you use, then the economic and sociopolitical impacts will control the decision.

So not only are the risk acceptance and public health and safety impacts and environmental impacts not unique, but they are also not necessarily controlling in the decision. So the risk management framework is going to have to reflect this.

So we are trying to come up with some basic premises that we can use to approach this. The first premise that we used was that the public health and safety should be treated in a consistent manner. You can define consistent however you might, but at least you should attempt to do that.

This would naturally lead to the idea of considering the risks and impacts of all the alternatives, the alternative of technologies or the alternative of not generating the power in the first place. This also leads to an idea that if you are going to look at risk perceptions, you are also going to have to understand how you are going to do that, because soliciting risk perceptions often determins the perception, your method of soliciting. Perceptions are also very changeable over time and they might not be useful or suitable for long-term management of risk.

So the next thing is that because of equity considerations, no individual should be unduly burdened by risk. This seems to be a natural thing to look at, so it leads to a natural separation of individual risks and societal impacts, or aggregates of societal impacts.

Another problem is that the satisfaction of one group's issues should not have a large detrimental on the desires of society as a whole. Basically what we are getting at here is there has to be some way of making the decision, some way of deciding when you are done wit this decision, when you are going to go on and either not build the plant or build it. There has to be some method of closure.

Then the aggregate impact of a particular facility should not stress the resilience of society, or threaten the resilience of society. So to some extent you do have to include risk aversion because a large catastrophe does threaten society in a way different than a lot of small individual impacts.

VOICE: Why do you use resilience as the word?

DR. GRIESMEYER: Oh, because it sounds good.

(General laughter.)

DR. GRIESMEYER: No, it basically talks about the ability of society to repair itself. Like when you go to Las Vegas, if your expected loss over a long period of time is 3 percent, then you may be able to handle it. But if you don't have a big enough bankroll and you get messed up in one of the fluctuations, then you have lost although your expected loss isn't too bad. That is why you have to consider resilience.

So basically we have come up with a management framework that has three features. It has the safety profile, which is just the probabilistic risk assessment, both site and plant specific. In this assessment you are going to have to include the uncertainties, propagate them in an explicit way so that the risks you come out with are going to be essentially risk distribution.

And then, of course, there are going to be uncertainties in this. The uncertainties are going to be as to data, modeling uncertainties, and a particularly troublesome type of uncertainty, and that is what have you left out? There isn't really a good way of finding that. So the safety profile might have difficulty in covering that particular aspect.

The safety profile should be able to identify important accident sequences, systems and so on, and they should be able to be ranked according to safety impacts. There should be some sort of certification of the practitioners and methodology so that there is some standardization of these risk profiles.

The risk profile should be used as a design tool. It should be used as a licensing tool, and it also should be used as a tool to modify procedures and operations in a way that reflects the safety impacts.

The problem with all these risk assessments is that there is really a simple way of coming up with a real deterministic answer of what the risk is or what the cost is going to be. You are never going to be able to come up with a deterministic answer. You are not even going to be able to come up with people to agree what the risk is or on the evaluation of the risk.

So there has to be some way of closing this argument: what is the estimated risk. So we have tentatively proposed a science court. Basically the idea behind this is to come up with some means with statutory authority to say this is the level of risk. Of course, NRC now is going to -- the level of risk that you are going to use in the decision process.

Of course, it can be updated and it has to be updated as you go, but there has to be some way of closing this part of the argument and getting on to deciding whether it is acceptable or not. And also you have to review the safety profiles. So this is the intent behind this idea of a science court. Many people have suggested it, and many people have also said that it can't possibly work, but that is the intention behind it.

The next thing is a quantitative set of decision rules. These decision rules are compromised between a judgmental and an analytical approach. There is going to be a need for some judgment because there isn't going to be a final answer, but there also has to be some analysis. So these are going to be some sort of pragmatic set as a compromise between analysis and judgmental procedures.

These decision rules, now, are going to be necessary conditions. That are, they are going to be more like non-acceptance criteria. The reason for this is that they are only part of a much larger decision process: whether you are going to need the power, which technology to use, which site to use, and then whether, once you have decided all these, whether it is safe enough.

So these really are non-acceptance criteria in that sense, in the sense that they are conditions for consideration as one of the alternative means to supply the power. Also, there are uncertainties involved. The uncertainties are data and modeling uncertainties, and these to some extent you can specify in your distribution in an explicit way, but there is also the problem of omissions, such as design errors, things that you are not going to necessarily catch with your analysis.

To that extent, they also can't be acceptance criteria. They can only be part of the decision process because of the uncertainties. Then reflecting this, we are also going to have upper risk limits. These are the non-acceptance limits. There is going to be a discretionary range in which you have to have special consideration, and then there is going to be the goal level of risk. This is lowest. You may be roughly an order of magnitude higher than the upper risk limit.

This discretionary range allows for the consideration of uncertainties, and it also allows for the consideration of competing risks. So that is the reason for the discretionary range. The upper risk limits, then, are these qualifying limits.

Something to go along with this is these risk limits should not be too restrictive because if they are too restrictive, then you make it so it is not possible for them to be met, and this technology is no longer a viable alternative although it may have less social, political, economic or safety impacts than an alternative.

So you have to be careful in making sure that your risk limits are not too stringent.

In the decisions rules we have come up with certain categories: the hazard state limits, individual limits, societal impacts, the public health and safety and property and resource damage, and societal impact reduction. The hazard states, as has been discussed earlier today, is something some people think we should have and some people think we should not have.

Basically it is quite clear that you can have a large societal impact even if you don't produce a large amount of radioactivity; financial loss to the utility, which is reflected in utility bills; and also there is a large trauma to the population. So there is some reason for using hazard states as at least part of your criteria. Not only that. If you structure your analysis so that you pay attention to these hazard states, you might be able to more adequately find precursors of much more important events. So the hazard states are basically the states of increasing severity.

The next thing would be individual risk limits or considerations, and these are basically in the interest of equity. Then there are societal, public health and safety impacts. These have to be some sort of aggregated impact of society, a weighing of costs and benefits. And then you also have to consider property and resources.

Resources is a particularly hard problem because you cannot modify as many of the benefits that you get out of various resources, such as important aquifers or a large amount of farmland that is near a reactor. If it is contaminated, it might be what we would consider a catastrophe and it might be very difficult to monetize this in some of the conventional ways.

And then there also has to be societal impact reduction, and basically, since society only has a limited amount of resources for risk reduction, this should be done in some sort of cost-effective way. Of course, in defining this cost effectiveness you have to evaluate what the change in impact is and you have to somehow or other monetize it or say what the limits should be, so there again it is difficult. So these are the basic decision rule categories.

t8
flw t7
er
77
bfml

Looking at hazard states, the first state is significant core damage. Now, what I have done here to come up with an operational definition about the requirements, they do have to be workable and they do have to be (unintelligible) the fine that by which you mean.

All of the hazard states and all the risk limits that we do are really proxies for something. If we can cover most bases, then probably the ones that you misses -- or at least some of the ones you missed -- will be handled. So, we set an operational definition here.

You can reduce the problem -- the limit to the rationales that you think that it should not happen either of two limits here. The probability of core damage in the life of the plant is less than 1/10th; or less than 1/30th of the lifetime. These can be translated to expect a reactor a 35 to 40 year lifetime to be the frequency, 3×10^{-3} per reactor year; or it could be used -- this would stress -- the most stringent requirement on probability of core damage would stress prevention a little bit more than the less one.

So, this is it. You should spend a lot more time in the profession of that, at least as much as you can. A way of handling the uncertainty now is the same that all of these should be met to the expected value, the mean value. There are some uncertainties. The risk has a distribution.

There will be one here. It is the expected value of that distribution. (Inaudible) to include uncertainties. A possible way of doing that would be saying that the level of risk is measured by the expected value, or some fraction of higher confidence bounds.

So, that is what we have done here. These numbers, of course, are meant to be examples of those two hard and fast numbers. So, we see here -- we say that the risk level has to be (unintelligible) per reactor year for core damage, or 5×10^{-3} for core damage at 90 percent confidence bound.

If the distribution was very wide, this will control; now this will control. So, this is just a possible means of including the uncertainties of the risk estimate. It does not include the uncertainties involved in the possible emissions, such as design errors. Design errors -- it is very difficult to discover design errors and it is very hard to include them in the published safety analysis.

So, really aren't handled by this, or this possible approach for looking at uncertainties. Large scale core melt is the next hazard state. Again, it has an operational definition. By large scale, I mean more than 10 percent of the reactor core.

The rationale here is that it should not be very powerful in the reactor programs. At the least, it should be less than 1/2 of the whole reactor program, considering 5000 to 7000 reactor years. This turns out to be the frequency of core melt

bfm3

should be (inaudible) reactor year.

Again, there is 95 percent (inaudible). So, this would be the purge numbers before core melt about this size. Now, by having a very stringent requirement on ability of core damage, we are trying to say that we should try and stress prevention.

By having this conditional release, large scale and controlled release 10 percent greater the iodine retort, having it less than 100 that you had a core melt, large scale core melt defined in this way. When I say, given a large scale core melt, I mean given the average large scale core melt that you would have found in your analysis -- some sort of operational definition so that you know what you mean.

This would be less than 100 per melt, less than a chance of frequency of 1 in 100 per melt. Again, a factor for the lighter confidence bound. This would be basically a possible approach for dealing with hazard states.

Now, with the individual risk limit, I came up with two possible proposals here. The first proposal is that the probabilit of death due to reactors -- somebody living near a reactor site should be less than 1 in 1000. So, the best thing is that his probability of death is very small due to this reactor.

This would work out to be a probability of death to be 2×10^{-5} per site year. The individual has to be balanced by death benefits or should be small compared to background. So,

bfm4

that's the basic question that we come up with in this type of criteria. Then, again, this criteria is melf-mem (ph) on the face of it by the hazard state that you satisfy the hazard state limits.

However --

DR. KERR: What is the probability of death due to a reactor during the lifetime of a reactor?

DR. GRIESMEYER: During the lifetime of a reactor, right. That will (inaudible) per year, about like this. Again, by specifying a confidence limit a little bit higher, you are taking into account some of the uncertainties in the estimates of death risk, which are not taking into account the uncertainties that it do -- there do emissions.

So, we have a conditional probability. This is the one that will probably control, because if you met the hazard state with this, you have already met these limits. So, the controlling one, probably this conditional ability of death, given a core melt.

This requires that there is some sort of mitigation, some sort of problem -- defenses against release. Also, it required a look at the effectiveness of your evacuation plan.

DR. SHEWMON: That's any core melt?

DR. GRIESMEYER: Any core melt defined in terms of that average before.

DR. SHEWMON: You have at least three definitions or

bfm5

two definitions of core melt. There is only one --

DR. GRIESMEYER: There has only been one large scale core melt -- large scale core melt greater than 10 percent of the core. There is a number of ways you can do this. That is just a probabilistic average.

So, this is just death due to the reactor. Then we have centurated (ph) of the probability of early death. Say that should be a factor of five less. So, this is just any death due to the reactor. This is death -- an early death due to the reactor.

This is one possible way of handling the risk limits. Now, I'm putting on a lot of numbers here. They are basically the upgraded structure, one or two of the numbers would control the limits. You would have to worry about the rest of them. These numbers are basically to illustrate.

Now, this proposal looks at another possible proposal for individual lists. It says that you should still look at early and delayed death, but you do it in a little different way.

Now, applying "R" as the individual chance of death per site year, per reactor site year. "RM" is the individual chance of death given the meltdown. "LF" fraction of age deaths, creating life expectancy before the onset of fatal symptoms. We find a measure of insult here as the possibility of death times the fraction of the life lost before -- or the portion of life lost after the onset of symptoms.

bfm6

So, it is weighted to loss of life expectancy. It only goes to the onset of symptoms, so that the period that that symptom is not included. You can avoid some of the problems of degradation, how do you measure that.

DR. OKRENT: Excuse me. You said period of death, do you mean period of life?

DR. GRIESMEYER: Period of life.

(Laughter.)

So, again, we have the same sort of expected value, then 95 percent confidence limit that is a little bit higher. Again, it is just a way of approach to handle the uncertainties in the estimate.

This is the same thing, only it is conditional on core melt. It is the same sort of criteria. All of these conditional core melt criteria are -- they are conditional on the probability of core melt factor of roughly 100 higher than of the hazard state that would require you to meet, attempt to meet.

It is saying that there can possibly be some emissions, some larger uncertainties than you thought in your core melt probabilities. Okay.

Societal impacts is the next catergory. For various reasons, we haven't used a frequency consequence limit line of basically because everytime I saw one, I never know what it meant.

DR. SHEWMON: Whereas, this is all clear the first time.

(Laughter.)

DR. GRIESMEYER: Okay. So, in societal impacts, we are trying to come up with a measure of the societal impacts. We want to put limits on it.

Now, the societal impact is going to be some function of a bunch of sequences that have frequencies in some sort of consequence. This function here may be complex, it may be very simple. If it is the frequency times the consequence, then it's just the expected consequence. That would be the measure of social costs. Then you would want to put a limit on it.

For example, the social cost function would be one with risk aversion, frequency times consequence to some power. In the previous discussion, that was frequency times the number of deaths to the 1.5 power. It's that sort of thing that we mean by this cost function.

DR. SAUNDERS: What are you summing over there with that expression?

DR. GRIESMEYER: Over events that -- any event that has consequence that is measureable. It is the somewhat expected consequence of this.

DR. SAUNDERS: Is societal cost a functional cost of the frequency in which it occurs?

DR. GRIESMEYER: The frequency times the consequence.

DR. SAUNDERS: That's what you have down there, but up there you go up --

bfm8

DR. GRIESMEYER: Up here I am saying that it is very difficult to decide what this should be. The first impression -- the first guess that everybody does is the frequency times consequence. That is expected consequence.

DR. SAUNDERS: So, "F" is the function of the accident.

DR. GRIESMEYER: Of the accident.

DR. SAUNDERS: "M" is a function of the accident.

DR. GRIESMEYER: Of the accident.

DR. SAUNDERS: Your sum over the percent of possible accidents.

DR. GRIESMEYER: Possible accidents. Possible accident sequences.

DR. SAUNDERS: And the cost?

DR. GRIESMEYER: The expected cost per year.

DR. SAUNDERS: Oh, it's expectation.

DR. GRIESMEYER: Yes, it is the expected cost measured by this function. Okay.

DR. SAUNDERS: As long as it's an expectation with foresight.

DR. GRIESMEYER: Yes, it is an expectation. Normally, when you do a risk assessment, this is distributed and this is distributed. This is just to illustrate the cost.

So, looking at health and safety impacts, we have done basically the same sort of thing again. We have said that both health and safety, deaths due to the reactor, measured by some of -- by some cost function. The units of this would be equivalent deaths per year.

If the cost function now is just frequency times consequence, then this is just an expected consequence per year. If it is frequency times consequence to some power, some weighted one, risk adverse social costs, okay, whatever it is shouldn't react further.

This is -- should be 10 -- ten per year. The reason why I picked this number is it is roughly equal to the lower bound on expected cost per year for a large coal plant. So, basically, in order to qualify now for consideration of an alternative to meet the need for power, this should not be any more risky and have much more social cost than the alternatives.

This is meant to be an upper bound. It is only meant to be qualifying a goal which is what the factor of ten lower than this may be dependent upon. There is this discretionary range below this. You also have competing risks and other design criteria in this area.

DR. LOWRANCE: Just for a point of clarification, the ten deaths per year comes from -- it's a comparison to deaths in operation of a coal plant?

DR. GRIESMEYER: Yes, an operation of a coal plant.

DR. WILSON: Not the mining part?

DR. GRIESMEYER: Yes. Public health effect due to the physical stuff coming out. This is the lower bound. That is why we feel comfortable that this is now; that this is quite a bit higher than those proposed earlier today.

The idea here is that these are just qualifying limits, and the upper limit. The goal is roughly the order of demand tube (?) below this.

These are not the only requirements that the reactor has to satisfy. Okay. Again, we have a conditional one here. Given that you have a core melt of -- a probability of 10^{-2} years, you would still meet this criteria in the expected number of deaths per reactor per melt with 1000.

That's where this number comes from. Again, we are not weighted to this number, but this type of thinking might possibly allow us to include -- or somehow or other account for things that have been left out in the analysis.

Again, we separated early deaths from all deaths due to the reactor. We, essentially this is a factor of five as before.

The next one is property and resource damage. This is mainly one of the most important, also one of the most difficult. You can monetize your losses. They did so in most of the reactor safety studies that have come out. The problem is that they didn't monetize everything. This is going to be the equivalent across all technologies, some of roughly the same as you require for a coal plant.

You really come up with some problems in that many of

bfm11

the damages due to a coal plant are environmental damages, degradation of a forest, or killing all the fish in a stream.

It is very difficult to monetize this. So, the things that you do monetize, possibly a way of limiting it would be to say the value of the energy generated per year -- that the expected monetary losses should be less than some fraction of the value of the energy generate per year.

I really wouldn't have good way of peeking what this number should be, but I do think that we do need to have a limit on the expected cost monetized -- everything you could monetize. You also have to look at the things that you can't monetize.

There should be some sort of limit; and it possibly could take this form. Again, you would meet this to the expected value of -- you would have to meet this in maybe the 95 percent confidence zone. You would have to meet this slightly less (inaudible).

For some resources, it is difficult to monetize. Their loss would be a very large catastrophe. Some resources such as very fertile farm land, an important (unintelligible), it may be that you should require special consideration when it is possible to really damage this, say beyond some level of effect like a potential of damaging ten percent of the resources.

Or something on this sort of -- you should have special consideration on the licensing. I don't know what to do as far as a quantitative safety goal, but it does seem clear that you

bfm12

are going to have to have special consideration for at least some of the large resources, some of your important resources.

It may be that the only way you can consider these is through siting restrictions. It does seem to be important that you at least take these into account. The next thing would be some sort of societal impact reduction.

Now societal impact reductions as low as the can be reasonably achievable time criteria, you know, the marginal cost limits. The range, we saw earlier today ranges from $10,000 per life saved to roughly $5 million. In some special cases, $300 million and a billion per life saved.

So, we take this relatively highlimit, because for the -- this is the change in early death per life of the plant. So, this is the total number of the total change in early death due to your improvement in the life of the plant. We say that this should be $5 million per equivalent death.

Again, if you don't have any risk aversion, this is per expected death. If you have risk aversion, it weights more heavily the large scale accidents, then it drops down to a million dollars per death, or equivalent death.

VOICE: Early death?

DR. GRIESMEYER: This is early death, $5 million for early -- equivalent early death. A million dollars for death. Again, these are tentative numbers.

This last one, I don't know how I can justify it except

bfm13

that this would be the change in expected loss. It is saying that you should spend twice as much to avoid the loss.

This would be a form of insurance, especially for the large losses; it may be very much better to avoid them rather than to have to absorb them. So, it may be justifiable to have something of this form, again, very tentatively.

So, for particular safety improvement, you would calculate the change in economic loss, expected economic loss per year or per the life of the plant; change in early death; and a change in total deaths.

It should be less than -- if the cost is less than this, you have to require the improvement. Notice that the early deaths get counted twice so it is $6 million for early death. This is tentative for illustrative purposes.

Basically, what we're going to do is come up with a way of risk management in a way that is compatible with the rest of the decision process. Also, to develop some of the structure of this process. The numbers that we've given are basically for illustration purposes, a place to start talking about adjusting either way or limiting. Thank you.

DR. OKRENT: Okay. Questions? Could you give your name, please?

MR. SUDNET: Bill Sudnet from the Nuclear Safety Analysis Center. I'm intrigued by a lot of the concepts that you put up over there. They raise a question to me as to mode

bfm14

of implementation. This scheme appears to be one that goes beyond the normal context of reactor regulations. I wonder if you have given any thought to a way in which a society like ours might go about implementing something like this?

DR. GRIESMEYER: I think it would be quite difficult to implement, I think. It would be very difficult to get some sort of standard method of doing the safety analysis. It is very difficult for people who decide: "Okay, this is what we're going to use with these decision rules."

So, I think it would be quite difficult, but I think that what you could do is you could do a safety analysis and apply these rules on the site, and do it a couple of times to find out what sort of things are possible.

What implications does it really have, more of an interation than "Let's try it."

There are certainly going to have to be some more methodology development as far as, say, earthquake analysis and fire analysis. These are both in the developmental stages, and usually they do a fairly cursory analysis. This would probably not be good enough for this sort of thing, expecially because if he has such a large uncertainty. There is a penalty for uncertainty.

So, I think it would be difficult, but I think it would be useful because we do have to efficiently allocate our resources. We do have to look at the decision problem, to the risk problem

bfm15

in a way that is compatible with the whole issue problem.

This seems like a way of starting to flesh out an approach to that.

MR. SUDNET: My question was really directed more towards the administrative than the technical. I understand that it is a serious technical problem, but probably even more serious to get the governmental machinery to operate in a fashion like this.

If the machinery can't operate like this, then the technical end of the problem is for naught.

DR. KERR: Do you think that this is any more complicated and detailed, for example, the operation of a typical public service condition in a state where rates are made? This is almost simple compared to some of the administrative functions in the system.

MR. SUDNET: I understand that those are less than optimum in times --

DR. KERR: They exist. In some senses, they are working. Is it your view that this would be more complicated than that?

MR. SUDNET: I really don't know that.

DR.OKRENT: Dr. Lave?

DR. LAVE: Well, you reminded me of the old story about the opera singer, a person with psychic powers, and the ACRS fellow who are standed on a desert island with lots of cases of b16

beans. Nobody has a can opener. The opera singer, and the cans don't open. The person with the psychic power thinks about the cans and the cans don't open. The person from the ACRS says, assume we had an electric can opener powered by a nuclear reactor --

(Laughter.)

The difficulties I have with this don't revolve around either the technical or administrative complexity, they really revolve around the kind of assumptions that you have to make. I think that is my problem with whether doing this as a shadow exercise is going to turn up anything.

Let me be rather specific. You say -- quote correctly that you need to have some way of closing off these issues. So, deus ex machina, we have a science court that is going to resolve all of the issues.

The problem is there is no democracy that I know about that resolves these issues. That is, you always have some people in the crowd that kick and scream long after the decision has been made. Indeed, the decisions only get named in a very slow fashion. Science courts don't help very much.

Then, when I go over to your basic premises, I'm in love with every one of your basic premises, except that I can't keep from help remark that it is scientists who keep on looking for coherent policies and democracies don't produce them.

Democracies produce higglety-pigglety policies that are

bfm17

all generalizations of special cases about whose brother-in-law and where they lived, and whether there happened to be a lake there, or whether it rained that day. So, I agree that in my ideal society, my platonic society, I want to have public health and safety treated in a consistent fashion. I agree.

I don't think there is any way that there is any policy in the United States that either now or in the future is going to treat these things in a consistent fashion.

Your third assumption is that no group can hold a society hostage. Gee, that's my society, too, except that I am a member of the ACLU. I think its minority rights ought to be protected. I don't know quite how to square the two. That's okay. I'll believe it for right now.

So, as a logical exercise, if you are a philosopher developing this, then I think it would be interesting to see how it consequences of the system. If you're not a philosopher developing this, then I don't quite see what it has to do with any systems that would remotely have to do with, or what the exploration of it is going to have to do with bubble current.

I'm sorry. That sounds very anti-intellectual.

VOICE: (Inaudible.)

DR. OKRENT: Excuse me, in the first place, if the member of the audience speaks, let him always get his name so the reporter knows who he is. Let's let Dr. Griesmeyer comment if he wishes, at this point.

bfm18

DR. GRIESMEYER: Basically, I realize that this is quite theoretical, but I think you are going to have to start with something. I think that the proposal that we heard by the Nuclear Regulatory Commission representatives of having this limit on core damage, so at the other end of the spectrum. What we're going to try to do is come someplace in the middle.

You're going to have to sharpen both ends, I guess. You are going to have to make proposals like these, and proposals like those in order to get anywhere.

DR. OKRENT: I would like to comment, if I can on one part of this. I am somewhat of a co-worker on this.

DR. KERR: You're going to blow my theory, which was that he was over 30, and he was under 30.

(Laughter.)

DR. OKRENT: I'm not sure if he's under 30 or not.

(Laughter.)

It's always seemed to me that one of the most difficult aspects of trying to propose the use of quantitative risk criteria was the problem, and assessing the risks.

In other words, that I think you might develop a recipe that everybody follows, and they all got the same number. To me, that wouldn't change the fact that they are very uncertain. So, that's unreal to me. I think there continue to be a large uncertainty in what these are actually.

I don't think it is going to be practical to, as Dr.

Griesmeyer said, has some deterministic way of finding out which is the right analysis. I really don't think that's practical. In fact, when we first started doing this, we spoke to Harold Greene, who you know, to see what his thoughts in the area was. Did this destroy the whole approach or not.

At least as I interpret it, his perspective, there needed to be some kind of a legally acceptable and desgnated process whereby judgment was made. Given this, then one might have a qay of living within this framework of undertainty.

So, I guess my own feeling is that a way of specifying a legally designated process is to have "a science court." There might be a specific group appointed by the President or whatever that do this. The thought was not to be the NRC itself who is doing this, but some other group who made just this judgment.

Then, the NRC would apply the decision rules which it imposed from the Federal Register and so on. In the absence of something like this, I am afraid the discussion would go on forever and you'll have all kinds of judges trying to review these matters or something.

So, it is not clear to me that it is impossible to do. I think there can be admonitions to such a science court by the Congress, if it were to set this up, to the effect that, well, if you need to err, err on the side of prudence and something of this sort. In other words, however -- the Congress has put general wording in the EPA and various groups saying that this b 20

they way we want you to go on a certain area.

If this won't work, I need to know what will, because absent a way of handling this large uncertainty and large variation that I expect from people who are trying to do their best job and have no conscious bias, even unconscious bias, there is going to be, I think, a large uncertainty.

DR. LAVE: Let me just respond briefly. I don't -- since you're going to make me talk for 20 minutes in a little while, I don't want to persue that now. I think that the objection to the science court is not that it is not -- be something that a court of appeals would not uphold.

Then, you had all sorts of nuclear reactors that were started that veriuos courts upheld, but they never got finished anyway. That is that one of the problems in a democratic society is that if you have a lot of unrest, it will bubble to the surface. If you have a nice legal system, it will bubble to the surface in the legal system. There will be the challenges.

If there is not a legal system, someone will sit in. That is, you will have passive resistance, the way we've got civil rights. If that doesn't work, you'll have somebody blowing up something, depending on how they feel about it.

So, simply having the legal processes all picked right doesn't help if you have to satisfy what is the underlying public feeling. I guess the more general point I wanted to make is that the discussion is basically of a discussion about scientific criteria and scientists following their way through, really has not brought in a public perception, which is where I thought we started on this anyway.

That is, if you talk about how safe is safe, or what is acceptable risk, then we keep on saying that you have to find out what is acceptable risk; finding out who accepts it, or what they are willing to accept.

DR. GRIESMEYER: I think that you can deal with that somewhat. It is going to sound like a sidestep, but the acceptance for the technology is based on socio-political and economical criteria. Maybe as much, maybe more the risk.

The risk is another part of the decision, it interacts with it. It is not -- I think that if you handle the risk in the way that we propose, it is compatible with the rest of the decision process. I think that the rest of this decision process, our main objections to the nuclear reactors have nothing to do with the safety of the light water reactors. They have something to do with, maybe, waste disposal, something to do with proliferation.

That's not what we're dealing with right here. I think that it might be possible in a method such as this, is the safety aspects of the LWR and not really say anything about the overall acceptance of the technology.

I think that, at least the thing that we're doing right here is, is this safe to society which, in some sense but bf 2

not totally, that is not the same question as it is acceptable to society.

DR. KERR: The question with which I certainly am grappling is not acceptable risk and not -- it's a problem that we have. We have to have something that has general acceptance in the middle of society. Risk is an important element, but it may not be the only important -- also, I guess I have a good bit os skepticism about science courts because I think the questions that have been posed for them that this is one of, deal both with policy and with interactions for consideration.

Scientists have difficulty enough being rational and they deal only with technical questions. When they start getting into policy questions, they are the most irrational group that I have ever heard of anywhere.

I just don't -- I don't think a science board is liable to be any more capable of dealing with these questions than a court of law.

DR. OKRENT: Excuse me. Let me comment. I think there is a misunderstanding involved. Theintent is that the technical part of the question of what the risk is is what the science court gives a best judgment on, if you want to put it that way.

The question of what constitutes acceptable risk which will -- and so forth -- what monetary values you should put on; these are the decision rules that the NRC would propose in the Federal Register, and if the Congress didn't like them, they would

bfm23

tell them so in all kinds of hearings.

They could even pass a law. So, there are really two different parts to this. The science court is not really involved in, well, the societal impact judgment as to acceptability or non-acceptibility.

They're intended to look at the available information and say, "This is our best judgment of the -- whatever it is -- the 90 percent confidence."

DR. KERR: They (inaudible) regulatory case, if you have to have a science court, you have technical questions on which there is a wide divergence of opinion. So, you're going -- you have a controversy to start. In the fact of a science court does not mean that you make the controversy any less. You perhaps give to some of the people the responsibility for deciding between --

DR. OKRENT: I absolutely agree, we have a continuing controversy on abortion. We have a court set up to arrive at a decision which is the law of the land until there is a change in that decision, if there is. They have a certain framework in which they arrive at a decision for the law of the country.

DR. KERR: But is a science court going to be equipped to deal with that issue, then, than the courts that we have? It is not (inaudible) but it is.

DR. OKRENT: I would be reluctant to ask lawyers with no trackable background to arrive at a judgment that is the

bfm24

probability of an accident leading to core melt, where reactor C is larger than 10^{-3} or 10^{-4}.

So, I think it is hard enough for people to understand in great detail what is going on.

Let's see, Dr. Wilson?

DR. WILSON: I wondered if -- there is a discrepancy here. I personally don't like the science court idea, either, but that -- once we've got the procedures of pointing to someone like the NRC to do just that sort of question.

It would provide you truth from the right Commissioners. If you don't, it might to the right science court anyway. (?)

(Laughter.)

So, the question is -- the one illuminating -- one thing that occurs to (inaudible) about the question on the valuance of this acceptibility. That is, if everybody in this room could be thinking about those things, typically those who have been quantitatively thinking about some of those questions, and were to get the brilliant discovery someone would make, or some technique that would reduce the risk; has it indexed the probability of core melt, the individual risk; the societal risk that all reduce back to intent for a nuclear reactor affecting nothing else about the system, would this affect the acceptability at all?

In a large number of people, it may make no difference. I think that is some sense, I think that one has to distinguish

bfm25

here what one is trying to do -- what I believe one is trying to do by this.

I'm trying to do -- the only think that a person is really (inaudible) for the moment is to set criteria. We satisfy this limited part of the task, bearing in mind that there are some people who, for completely different reasons, think reactors are for the birds; although there is still work for the devil.

These are completely relevant to the problem. Have you any way of avoiding the question of doing just that part of it and leaving the other part temporally to one side and let that be business in another matter, whatever that matter may be.

DR. OKRENT: Dr. Lowrance?

DR. LOWRANCE: I have two comments on the science court I can't put myself -- (inaudible) -- scientific pluralism in general. Having read an undue share of National Academy of Sciences and Technology Assessment and other such reports, it bothers me to think of legitimizing any one group of scientists too much.

(Laughter.)

I do think once they are placed under the public domain, it is very hard to call for revisions again and again. I think any party looking into reactor safety over the last 15 years, let's say, would have wanted to change its mind at varousl points along the way, and would still want to.

It has never been clear to me just how a science court b1-26

would do that.

My second comment is any time a scientific tribunal of some sort looks at an issue, it must deal with the question of the definition of the problem. We had a perfect example of where a science court set up would be hard-pressed with the sabotage issue that you yourself pursued this morning.

If somebody instructed the science court, the regulators let's say, some of the science court the question of assessment of risk of a plant.

Say, "Oh, by the way, don't look at sabotage, or don't worry about earthquakes or whatever."

The panel would find itself, or anyone, get back into the social realm and say, "Look, don't draw the boundaries around it for us."

The consequences go in all sorts of directions, both on technical and social sides to argue just in setting up the question. That brings in all kinds of subjects, if I may, and considerations.

end t8

t9 flws.

DR. OKRENT: I think with my previous statement we need a method, I think, of providing a legal decision. Now, you could say PASLB does this, or the Commissioners or something. It isn't clear to me right now that either of those two do it in a current framework. They might in some future one.

I must say I am not trying to defend the idea of a science court. I would be happy to have a better working proposal. I do think there is a way of dealing with the question of a court living too much on precedent. I mean, I think we have seen in the NRC and the AEC before, and even in the ACRS that precedent is important on future thinking. It is also true, of course, in the Supreme Court, and you may say that is good or it is bad.

I mean, you are saying about scientists, I think, what lawyers probably say about lawyers. But you can build in ways of having a certain amount of change in the court's constitution every number of years, and if the people are selected properly, in fact, they won't be bound by what is decided before. They will have to be convinced.

They are really a mechanism. The question is, what do you have that is better? I am not saying anything is good, but my question would be, what is better?

DR. SHEWMON: The question is whether you are going to abolish the NRC to do it better. I think that is more practical point. I mean, the Commission. But that is probably not why we are here today.

DR. OKRENT: Let's see. There were some hands at the back. Would you give your name, please? The nearest hand first.

MR. THOMAS: Jerry Thomas, MSAC. I am going to change the subject away for a minute from the science court. I would like to congratulate you as the first of seven speakers to quantify core damage and put it as a step-wise progression. I think that is a very important part of risk assessment. In fact, approach it in a step-wise process without increasing danger to the public. You are the first of seven to do it.

Everyone else has stood up and said, core damage is bad, with no definition of what core damage is.

DR. OKRENT: Yes?

MR. O'DONNELL: Ken O'Donnell, EBASCO.

I have found remarkable similarities between your presentation and mine, and therefore I am highly approving of your presentation.

I just wondered if among the four criteria you have, whether there is a hierarchy among them. That is, do you have to meet all four of them or is one more important than another?

DR. GRIESMEYER: I think when we look at the numbers, one will be more important than another, and one will control, but I think that you have to have them all applied simultaneously, and this is just a batch of constraints.

MR. O'DONNELL: I guess in particular the last one, which would be the cost benefit criteria.

DR. GRIESMEYER. Well, the cost benefit criteria, if you haven't met the other ones, the cost benefit criteria won't be good enough.

MR. O'DONNELL: But if you do meet the others, and you don't meet the cost benefit --

DR. GRIESMEYER: You still have to do it.

MR. O'DONNELL: -- you still have to do it.

DR. GRIESMEYER: You still have to do it. That is how we envision it.

DR. OKRENT: Dr. Slovic?

DR. SLOVIC: I would like to comment on two things that have been running through the various presentations, and I think they are exemplified in your presentation, that concern me. There seems to be a distinction between an approach which sort of stays within the technical realm, and that is, say, what I saw as the IEEE approach, where you say, now, supposing as technical experts we try to see what would be reasonable to us, a reasonable rate of growth, a reasonable criteria, and so forth, as opposed to a broader approach which brings public acceptability into the picture.

We all know that is very difficult to do, but yet some of the speakers today have tried to acknowledge that, and others have said, let's take a more limited view.

I think it is important to keep that distinction in mind when considering these proposals.

Now, if we think about the broader perspective, where we are concerned about public reaction, then something, I think, important becomes apparent from a technical standpoint, and that has to do with this risk diversion notion, which you very clearly highlighted as relevant to your criteria, and also previous speakers have noted it as well.

I am uneasy with the notion of viewing risk diversion as kind of a coefficient that you can tack on to sort of vary the expected value, or modify expected value. I think that the reaction, the impact, the social impact, the social cost of a significant event, accident, what have you, is not a ry meaningful or very neat function of the magnitude of that accident.

You yourself note that in your paper, where you say early on, you say, "As evidenced by the Three Mile Island accident, abnormal events at nuclear power plants can be quite costly even without radiologically induced health effects."

I think that is very significant, but then I see the later part of your paper is not really taking proper cognizance of that. What I am getting at is that the impact of an accident is often a function of what that accident signals about the probabilities of such accidents, the degree to which the technology is under control, and so forth.

So, an accident that has relatively small health effects could have a very big impact by its secondary, third order effects, and so forth, if we shut down the industry for a significant period of time, which has immense costs of a great variety.

So, I really am concerned about the ability to use that simple functional relationship between magnitude and cost, and I wonder if one buys that technical problem, acknowledges that, what that means for criteria such as these.

DR. CRISEMEYER: Well, I will admit that I was looking for the ring under the light, and I could come up with an example in this way. I am not saying that the cost function, whatever it is, is simple. I am just saying that it will take a while to develop something. Maybe this form is not -- because it is such a complex function, and it varies with time, it varies with a lot of different things, but you do have to come up with something that is operational, and I don't know how -- the reason why -- one of the main reasons why I put in that I thought it would be important to have hazard states was because of the signal value of the precursors to a large accident.

That was one of the main reasons why the hazard states were in this, because of the signal value of seemingly small accidents, or small by some criteria. It is a very difficult problem, and I really don't know how to approach it.

DR. SHEWMON: It occurs to me, just in listening to some of the other presentations as well, that one implication of this is that maybe we ought to examine more carefully the notion of this probability of core degradation sort of thing.

That is, it may be that the consequences of any significant accident are so great that we want to limit the overall, the aggregate probability of such an accident to some very small number, say, within the next 20 or 30 years, so then if we say we want this cumulative probability to be below such and such a point, then we can work backwards to what that implies from the design standpoint, given that the response to any kind of accident at all, like at TMI, would be so great that it really would swamp the kinds of costs you get from ten fatalities, 100 fatalities, so many latent cancers.

That is really not the point. It is the massive social response in the event of another TMI-like accident.

DR. GRIESMEYER: Right, and that is why I had the core damage thing, and that is why I suggest that it be once in 30 lifetimes. That is the idea, is to make it so that that doesn't happen very often, and criteria that have been given earlier today for that same thing would be once in 30 years for the whole reactor program.

So, I think that you do have to take that into account in some way, and we are always groping for a way of doing it.

DR. OKRENT: If I could comment there, you could try to pick a very low number, like one in a million per reactor year, and if then that led to, let's say, damage to the core similar to what occurred at TMI, and then the probability would certainly be very small in the next 30 years that this would occur if the 10^{-6} had been achieved.

On the other hand, by picking the number, it doesn't make it so, and it doesn't affect the design of the existing plants.

It is not completely clear to me that in developing an approach either like this one or one of the others we have heard today, that one should give what I will call excessive emphasis to the thing you expressed concern about, that another non-damaging accident could have huge ramifi::ations.

Society doesn't have to react that way, and in fact I think we see that there are countries which are going ahead with nuclear power programs more aggressively, in fact, since TMI, not less. I think it i because they see a clearer economic need and so forth, pol.tical need, a variety of needs.

So, while not one to minimize the fact that there are differences of opinion in the country, I think there still is reason to try to see what is a possible structure for the country going about its business.

DR. WILSON: Is part of your response to the whole thing there, is that partially conditioned by the fact that your belief that a mass response to TMI, which is, I share, has been rather greater than was in fact warranted by the size of the accident, and therefore one's belief that if another TMI were to happen moderately soon, there would be another massive response which would be greater than would be warranted by the size of the accident. This is the sort of thing Paul has brought up, that we ought to somehow explicitly take that into account, that this particular feeling might disappear with time in the United States, as it doesn't seem to appear, for example, in France, and hardly in the Soviet Union, and that perhaps to put a criterion in concrete right now which represented something

9

which might -- a present thing which might be transitory -- might be an error, whereas a calculated health effect, if the calculation is right, would not be transitory.

DR. GRIESMEYER: Well, the limit that you put on that health effect might be transitory.

DR. WILSON: But if you calculate it, and calculate it right, that is a number which is eternal, so to speak. However, assessment of a public perception, that may be very real, and very important, might be a thing that changes with time, and it has to be looked at rather closely.

DR. GRIESMEYER: Public perceptions by themselves are not adequate, because they do vary in time. It won't work for long-term management.

DR. WILSON: I think that is somewhat of what you were -- I am trying to sort of pull that out of what you were saying as what was implicit in what you were saying. I wasn't quite sure whether it is what you meant.

DR. KERR: I think that is what he meant.

(General laughter.)

DR. OKRENT: In fact, you stated fairly directly what I meant, but let me comment in two ways on that. I have at times certainly published things which indicated the personal opinion that there was a need for society at least to think about optimum use of its resources, and to think about risks in a broader context than only the directly observable public health and safety.

In that context, though, I have also indicated I think it is not unreasonable that we try to make nuclear reactors safer than previous or existing technologies, and I think the general society would want the new coal plants to be safer if they knew how to do it, let's say, in a practical way, and so forth.

So, I think, as I say, that nuclear plants have a more stringent target is not a bad idea, although if it turned out to rule out their use, you would have to go back and see, what does this mean to the economy and so forth. So, I don't say that in an unqualified way.

I do think you have to keep in mind what people now feel, but I do think also that what people now feel changes. If I look outside the area of nuclear reactors, I can think of a couple of areas around the world where ther have been really drastic changes. I will give one from the U. S. first.

I think the question of civil rights, the public attitude toward civil rights in this country is very much different t: ay than it was when I was a boy, okay, or even more recently. If one insisted that one had to take the public perception into view in molding policy, one would have stayed the way it was, and it would not have changed, because there was a certain perception which is different

11

from today's.

If I want to take a more drastic characterization of a change in public perception, I think I could argue that in the mid-thirties the public perception of what Hitler was doing in Germany was not all that bad by the majority of the people. There is a very different perception today.

So, these are different than the nuclear one, but I think they are not irrelevant to the question of the nuclear one, and while one is saying public perception is not the only factor, it is not necessarily the dominant factor in trying to decide what a country should do, I would myself prefer to stay with some of the principles that Dr. Griesmeyer showed, like, you know, no undue ill effect to any individual.

That is a tenable premise under a variety of situations.

DR. WILSON: But you do, I think, believe that in trying to arrive at an acceptable risk -- public perception has a fairly important role to play.

DR. OKRENT: I do agree.

DR. SHEWMON: You would also acknowledge, I am sure, that scientists' perceptions change over time as well as they learn more about the situation.

DR. SHEWMON: Only insofar as they are part of the public.

DR. OKRENT: I hope so. I will put it even more strongly.

DR. SHEWMON: I would like to ask a question in that line, and I guess -- I don't know the consultants particularly well, but I get the impression that you have worried about public risk at least as much as the rest. Have you also worried about rates of change?

One of the things that struck me this morning from O'Donnell's talk was the fact that we now kill 2.5 people by electrocution each year with what a power plant puts out. I don't know what you can trust about what is written in the press, but then if you people bring up every so often articles about what was printed back at the turn of the century when AC power was coming in, or a little bit earlier, and how horrible that was going to be.

One of my own hopes is that indeed reactors are safer than most things, and maybe I will live long enough to see the public acknowledge that, but aside from my biases, have you ever tried in your studies to do things on how fast these things do change?

DR. SLOVIC: I haven't studied that, and I am not aware of a lot of work in that area. I think it is a fascinating and important problem. Obviously, the public has changed drastically in its reactions to sorts of technologies over time, but just what will happen with the

13

nuclear case is not so clear, because it has certain special characteristics. In fact, there are very low probabilities involved, so you really can't get the kind of experience that would lead you to kind of accommodate to the -- The fact it is radiation and not some other sort of impact may also be important.

So, it is a little difficult to project what the rate of change will be, and I don't know of any really good research on that problem.

DR. MARK: Paul, I think you might want to take into account on that same question that it does change. Electricity, which you give as a nice example, where there was horror and consternation when -- proposed to go electric, and you see no sign of that now. It is accepted as a must. There would be worries about fire if they went back to gaslights today.

On automobiles they changed. They were supposed to be dangerous. The public attitude has changed, but the fact is that cars are probably just as dangerous as people thought.

(General laughter.)

DR. OKRENT: Are there comments now on what Dr. Griesmeyer was presenting? I recognize that the consultants are going to have time to present a coherent set of comments.

(No response.)

DR. OKRENT: Well, let's see. I guess we could take a ten-minute break, and then begin with Dr. Kastenberg and Dr. Johnson. That is going to take about an hour instead of the 30 minutes shown on the agenda, so we will end up being 30 minutes behind when they are done.

All right. We will reconvene, then, at 3:30.

(Whereupon, a brief recess was taken.)

DR. OKRENT: Let's begin the presentation by Kastenberg and Johnson. We may have to interrupt this after the first half in order to catch Dr. Slovic's comments, because he has to leave by 10 to 5:00.

DR. JOHNSON: My name is Dave Johnson. Bill Kastenberg and I have looked a little bit at some of the implications, the technical implications of Mike's criteria that he has proposed, and we hope by this to feed back to him for further development of the criteria.

The presentation comes in three pieces. First, I would like to make a comment or two on individual risk in nuclear power. Bill will comment on societal risk as interpreted by Mike's criteria, and in particular what social costs are implied when one considers risk diversion in a simple model as well as uncertainties.

Finally, I would like to apply the criteria to a 1,000 megawatt coal plant, and to compare accepted technology, and what this criteria will tell us about that

15

application.

Several people have referred to a calculation of the risk to an individual from a nuclear power plant, simply outlined on this chart, and basically what was done was the probability of an individual receiving a lethal dose for a particular release cateogry in WASH 1400 was calculated, and that was combined with the WASH 1400 release probability, and these -- the product of these terms was summed to estimate the individual risk

The bottom line is not really claiming that this is the risk to an individual per se, but it is comparing to Mike's criteria, which -- this is by an order of magnitude less than his expected value criteria.

The question then becomes what about uncertainties, and one had to compare the uncertainties in this calculation with, say, his 90 percent, and if you believe that the uncertainties come, say, through the frequency of the releases, then -- say they are off by a factor of 10, they are still well within the 90 percent confidence limits.

The more interesting question becomes then an estimate of the individual risk of latent effects. We tried to estimate that by a simple manner. We tried to estimate this in a rather simple manner. Basically, we assumed a frequency release distribution similar to what Dr.

16

Joksimovic described, in which the frequency of the release was varied inversely to the dose that an individual received.

We didn't feel that one could limit the risk calculation, a latent risk calculation to an individual, only toward accident scenarios, but would rather have to consider all releases.

We modeled the frequency of release versus the dose to an individual near a power plant, and as you would expect, the alpha here is a power describing the frequency of release to the dose. We used an inverse power. We knew the cutoff iose for low doses.

And as you would expect, this -- the calculated individual risk for latent death, delayed death is rather sensitive to the parameters --

DR. QUITTSCHREIBER: Just a second, Dave. Is there a Dr. Collins here?

(No response.)

DR. JOHNSON: We are not trying to sell this as a model of reality, but the bottom line is that risk calculations have more latent risks to an individual near a power plant, and from what we could tell from our simple model, perhaps just ignoring non-accident cases, this is not quite correct. It is conceivable that a simple model such as this would yield latent risks which were approaching limits that Mike has talked about.

17

Are there any questions on this before I --

DR. WILSON: I don't understand the latent risk. Do you mean, latent risks in -- releases, or --

DR. JOHNSON: Risk due to -- risk of developing cancer, say.

DR. WILSON: Due to an accident?

DR. JOHNSON: No, due to non-accident releases.

DR. OKRENT: Routine releases.

DR. JOHNSON: Routine releases, with a very simple model of the frequency of release.

DR. OKRENT: What are you assuming is the millirem per year from routine releases?

DR. JOHNSON: It is a distribution, quite similar to what Dr. Joksimovic put up as his release criteria. Again, I am not claiming this is modeling reality, but it is something that most risk calculations have simply ignored before, and assuming a simple model, be it realistic or not. Whereas the risk of acute death is well below the limits Mike has talked about, the latent risks that a simple model would give you would be approaching the limits Mike talked about.

DR. KASTENBERG: Do these releases conform to the tech specs, for example, for plant operators?

You used the term "normal operation," and I am trying to understand what one means by normal operation.

DR. JOHNSON: Well, I shouldn't have used that term. I should say non-accident releases. Any releases that are not covered by WASH 1400 category releases.

There is certainly a large number of releases that WASH 1400 does not consider. Normal releases --

DR. KASTENBERG: Okay, but these releases might, for example, be beyond Part 20?

DR. JOHNSON: Yes.

Bill has some comments now on the societal -- the implications of the societal risk criteria.

DR. KASTENBERG: Let's see. What I will try to show you are some calculations that we did for societal risk, and as part of the framework that Mike presented before, he mentioned that one thing that you would have to produce for a given reactor is a safety profile, and not having the manpower, funds, computer, and so on to produce safety profiles, we decided to use the so-called safety profiles of WASH 1400, which most of you are familiar with, which I just summarize on this first vu-graph.

Basically, we worked with the complementary cumulative distribution functions for early fatalities and for latent fatalities. When we started, we were also going to look at genetic effects and property damage, and then on the next vu-graph we were going to look at land contaminated and relocation area, with the ultimate result looking at the

19

resource question, but we didn't get that far.

Also, I want to point out, I will also use for illustrative purposes a series of complementary cumulative distribution functions which were calculated during a study at Sandia Laboratories, which has become somewhat controversial, but we will just take them as examples, as hypothetical reactors.

Basically, they are the surry reactor placed on different sites and then scaled to the reactor that is on that actual site. But for us, it is a hypothetical reactor, because it is not the reactor that is on that site. We just look at them as some curves that we are going to play with a little bit, and they would be representative of a number of reactors.

I want to point out one other thing which is interesting in some of the numerical results which I will show you. If you just look at Curve Number 3, which represents a particular reactor, and Curve Number 4, there is an interesting feature. The two curves cross each other, so that one curve has a high frequency of low consequence events and the other one has a high frequency of high consequence events.

The relative ranking in terms of societal risk changes depending upon how you characterize the societal risk.

DR. MARK: These are -- with population distributions?

DR. KASTENBERG: Yes, these have the actual population distributions on those sites. Again, I did it for early fatalities for a number of sites, and there is a curve there, a set of curves for latent fatalities.

One thing that should be pointed out is that on latent fatalities, you can get a little confused looking at WASH 1400, in that in the original curves for latent fatalities, they had latent cancer fatalities for a year, and in the Sandia work it is totally for cancer fatalities, that is, 30 years, the total during the 30 years following an accident.

So, the curve shifts by an order of magnitude when you take in the total latent fatalities.

So, these are the so-called safety profiles that we used to generate some numerical results just to see how Mike's societal criteria work out, and the expressions that we used to generate the societal risk, basically, using these complementary cumulative distribution functions, you can write the risk or the expected value of the risk as just -- over the region of interest, of the consequence times the derivative of the complementary cumulative distribution function, and for those of you who like to say it is the area under the CCBF itself, if you integrate that by parts,

21

you will get that also.

Then, also, to look at risk aversion, or as Mike had called it in a previous note, the expected social cost, you just take some higher moment of this integral, so it is just X to the alpha, and at the time I guess Mike had proposed a value of X equals one and a half -- alpha equals one and a half. When alpha is one, it is just the expected value of the risk. When alpha is greater than one, then you have some built-in risk aversion in that event.

Then, we will compare some of the integrals both for the nuclear phase and for the coal phase via one of the criteria that Mike showed before, that the expected social cost is a function of alpha, has to be less than or equal to ten deaths per year, and then again, in one of Mike's earlier drafts, he and Dave Okrent had proposed a value of $2 million per death averted, and I noticed in today's presentation he has now split it to $5 million for early deaths and $1 million for latent deaths, and maybe if you take a weighted average you will get two times 10^6 dollars per death.

So, again, I think Mike made the point, and I think we will make the point that these are not hard and fixed numbers, but I just used those as an example of how this kind of framework would work with a given alara principle.

22

The other thing you can do is play around with this cost as the cost of the improvement, delta EC of alpha is the change in deaths per year that this improvement will get you over Y years. Mike had the symbol L before. That has to be less than two times 10^{6} per death averted. If you play with this a little bit, you could also ask, what is the maximum that you should spend, and if you just rearrange it a little bit, C alpha would be the most that you should spend.

These are the kinds of things we calculated both for the nuclear case and for the coal case. I will show you how some of these numbers work out. Just running through the WASH 1400 curves, we checked them against what was in the WASH 1400, I calculated the expected value for early fatalities to the average curve and get 4.4 times 10^{-5} deaths per year. For the latent, total over 30 years is 2.7 10^{-2}, which is the same as you find in the executive summary of WASH 1400 within the numerical accuracy of how I read the numbers off the curve.

Genetic effects I totalled over 30 years. Property damage, the expected value is 5.5 10^{3} dollars per year. Relocation area and decontamination area, very small numbers if you take an expected value, and that is somewhat deceiving, because if you recall those curves for decontamination area and relocation area are very flat all the way out in area, and then all of a sudden drop very drastically.

It is hard to know whether it is the model that they used, that they only integrate out to 50 miles, or whether there is some physical thing that causes the curve to drop off suddenly.

These are the base line numbers that we worked with, and that you get from looking at the WASH 1400 curves. And then, just repeating the same calculation for the curves for those different reactor sites, just to see what kind of variation that you might get, again, looking at early deaths, you get quite a spread in the expected value of the risk.

At the site at San Onofre, 10^{-8} deaths per year, all the way up to Zion, Indian Point, 10^{-4}, for early deaths per year, well below Mike's criterion of less than ten deaths per year, and on the latents, those curves are hard to read, to try to do a little analysis, and I can only do some of them, and the ones that I could get the values for, Zion and Indian Point, the numbers are on the order of five times 10^{-2} deaths per year.

That is total deaths per expected value, and I guess what this tells you is, if you take a straight expected value as the societal risk, they are well below Mike's criterion of ten.

24

DR. OKRENT: If you use the reactor safety study numbers.

DR. KASTENBERG: Yes, if you use the reactor safety study numbers, and you assume that the methodology is correct, and for these hypothetical reactors, they would pass Mike's criteria basically.

One other thing --

DR. WILSON: Any one of them, including Indian Point and Zion?

DR. KASTENBERG: Exactly. Well, it is hypothetical, because it is not the Indian Point reactor that is being analyzed here. It is the surry reactor on the Indian Point site.

DR. WILSON: Sure, I understand. It is the site w w r s n h re

DR. KASTENBERG: Right.

DR. SIESS: Does this include an alpha yet?

DR. KASTENBERG: No alpha yet. This is straight expected value.

The other thing is, as everyone knows, but it comes up again, that the latent effects dominate the early effects of these kinds of calculations, as you might expect.

Then, I went and did a series of calculations to see what the effect of the alpha is in trying to incorporate this diversion via this artifact of X to the alpha, and just

25

to see what happens for early deaths per year as a function of alpha, again, the first column is the straight expected value. If you go to an alpha of one and a half, things -- the expected value goes up or expected social cost goes up, and if you go to alpha equals two, it still goes up again, but even with alpha equals two for early deaths, you still have yourself below Mike's criterion of ten deaths per year, and an alpha equal two, basically, I would interpret as saying, one death is one death per year, one at a time, ten deaths is 100 deaths per year, one at a time, and so on and so forth.

Now, you can do the same exercise for latent effects, and I only -- again, I was only able to pick three of those curves because they are bunched together, and here, you see the result is very, very sensitive to the risk diversion factor. Even with an alpha of one and a half, you start to get close to Mike's criterion of ten deaths per year, and if you go to alpha equals two, you get 53, 400, 149. You are way, way over Mike's criterion.

So, if society were to adopt a criterion which said, the expected social costs cannot be greater than ten deaths per year, and you had to have an aversion factor of alpha equals two, none of these reactors would pass, based on that criterion.

So, the results are very sensitive to what you

26

might want to assume for risk aversion.

Okay. Now, the next thing I tried to do was to apply the alara principle to see how that would work out numerically, and what I did here was to pose a problem, and the problem is to take for the case of latent deaths the first set of curves done at Sandia for latent deaths which had the Indian Point and Zion sites, and I don't remember which one was which, and I just posed the problem, suppose you wanted to make an improvement on the reactor represented by Curve Number 1 so that it was equal in risk to the reactor safety study curve, and that it would cost you first $1 million to make that improvement, and then $10 million to make that improvement.

Is it cost effective at $2 million per life saved as a function of risk averted, and what you see basically is without risk aversion at $1 million for that improvement, without risk aversion, it is cost effective on my $2 million per death averted or life saved. At $10 million it is not cost effective, and for the second reactor, in either case, it is not cost effective without risk aversion.

As soon as you introduce a little bit of risk aversion into the calculation --

DR. SIESS: I don't see on that first that one is and one not.

DR. KASTENBERG: It isn't because of the number of deaths that you are saving in going from Curve 1 to the reactor safety study and you save so many deaths per year --

DR. SIESS: What do you compare to get the answer you got up there on the table? How do I decide from the table whether it is or is not?

DR. KASTENBERG: The criterion is that if you can spend less than for $2 million or less per life saved, then you should spend it, and there are different numbers of lives saved in reactor Number 1 going to RSS than reactor Number 2 going to RSS.

DR. SIESS: I was looking for two under Reactor Number 1.

DR. KASTENBERG: Oh.

DR. SIESS: What is the $1 million and $10 million again?

DR. KASTENBERG: I just posed that suppose the improvement would cost you $1 million. Is it worth it? And then I said, suppose the same improvement were to cost you $10 million. Should you do it?

DR. SHEWMON: Each one saves so many lives per year?

DR. KASTENBERG: I have to go back to the --

DR. GRIESMEYER: Well, it is equivalent deaths, isn't it?

DR. SHEWMON: What are the numbers written into alpha for one? I don't understand that.

DR. KASTENBERG: It is a factor of ten. This is $1.1 million per life saved, and here it is $11.5 million.

DR. GRIESMEYER: And it is equivalenet death, too, because when you use -- because you have weighted the large accident sum on alpha equals one and a half and alpha equals two.

DR. KASTENBERG: We are just talking about the first column. As soo. as you move to some risk aversion, it is always cost effective.

DR. SHEWMON: In the cases you said where one was and one wasn't --

DR. KASTENBERG: I was comparing these two numbers, and then I was looking at these two numbers. I wasn't looking at these two.

DR. SHEWMON: The Curve 1 versus Curve 2, not the one versus two.

DR. KASTENBERG: Right.

DR. OKRENT: I am still missing something. When you spend $1 million in Case 1, what is it you accomplish with that $1 million?

DR. KASTENBERG: I am making the expected value of Curve 1 have the same expected value of WASH 1400.

DR. OKRENT: All right, so you are moving, in effect, from a surry light reactor as at Indian Point or

Zion --

DR. KASTENBERG: To surry on the composite site.

DR. OKRENT: -- to surry at the composite site.

DR. KASTENBERG: To surry at the composite site.

DR. OKRENT: All right.

DR. KASTENBERG: I am making it as good as if -- and I just -- because I don't know what it would cost to do that, so I picked $1 million and $2 million just for illustration.

DR. OKRENT: All right.

DR. KASTENBERG: Then, the second point is when you introduce a little bit of risk aversion into the calculation, it is almost always -- well, in all of these cases it is cost effective. You will always meet the alara criterion, and it would tell you with risk aversion built in you should make the improvement.

E TP9

DR. KASTENBERG: Then the second point is when you introduce a little bit of risk aversion into the calculation, it's almost always -- well, in all of these cases it's cost effective. You will always meet the ALARA(?) criterion, and it would tell you with risk aversion built it, you should make the improvement.

And on the next vugraph I wanted to see what's the most you should spend, and you can see what happens. The number gets tremendously large as a function of risk aversion. In other words, if you look at straight expected value for curve one to make it equivalent to Surry on a composite site, the most you should ever spend is $1.7 million, and curve two, $.9 million.

As soon as you put in risk aversion, you go up quite a bit, two orders of magnitude of dollars. And if you go to this alpha equals two, you're spending more than the reactor is worth. I should say you're allowed to spend more than the reactor's worth to make the improvement.

And the illustration here, I think, is that this kind of a characterization of risk aversion causes the ALARA-type considerations to be very, very, very sensitive. Or to put it another way, if people are so risk averse to this, you're going to have to be spending lots and lots and lots of money, if you force people to spend $2 million for Life-Save, if you can accomplish it.

DR. WILSON: But this is for the latent cancers, not for the --

DR. KASTENBERG: This is latent, yes. I based all this on latents.

DR. WILSON: Which is quite a real distinction.

DR. KASTENBERG: Yes, right.

Now, one might argue that people are not as averse to latent effects, latent deaths as they are to early deaths, and I could reproduce the calculations for early deaths as well, I suppose; but this seems to be more interesting.

Then I had been discussing this with Dr. Okrent when I completed the calculation, and he suggested that rather than trying it as a function of aversion, why not try it as a function of uncertainty. And I picked up on that and thought well, it may be that people are not averse to the fact that there are a lot of deaths, but another school of thought would be people are averse because you're very uncertain as to how you characterize the safety profile.

And if you look at it, you can do the same kind of calculation as a function of uncertainty, and that's what I tried to do here. That is, with zero uncertainty, again assuming that you knew exactly that RSS was correct and you knew exactly that the calculation for Zion and Surry and Zion/Indian Point were correct, should you spend $1 million to make the improvement, or should you spend $10 million.

And again, it's the same sort of thing I showed before, same numbers with zero uncertainty, here you should spend the money, here you shouldn't, and here you shouldn't. Then you say what if I'm uncertain by a factor of 10, both in RSS and in the Surry/Indian Point calculation, should I make the expenditure of a million or $10 million. And you see what starts to happen is as you get to plus 10 uncertainty, in both cases for curve number one you make the improvement. For curve number two, if it were a million dollars, yes, you'd make the improvement; if it were $10 million, no, you wouldn't.

And then you can go to an extreme and say suppose I were off by a factor of 100 in both calculations, then in all cases you would make the improvements. And basically what you're doing is paying for your uncertainty.

DR. WILSON: I thought you were in fact uncertain in both calculations. Are they varying together when you do that or are they varying separately? I mean, are you saying one is uncertain relative to the other, or that both are multiplied by 10?

DR. KASTENBERG: They're both multiplied by a factor of ten and then both multiplied by a factor of 100.

DR. WILSON: Okay. But they're not uncertain in respect to each other.

DR. KASTENBERG: No, no. And again, you could do that as another variance on this --

DR. WILSON: Surely.

DR. KASTENBERG: -- To see again how sensitive these calculations are to that.

And then last but not least, I asked the same question as I asked in the case of the risk aversion calculation, and that is, what's the maximum that you should spend given no uncertainty, a factor of 10 uncertainty, and a factor of 100 uncertainty. And again, the amount that you should spend starts to get very, very large as a function of uncertainty, and just scales by factors of 10 for each factor of 10 in uncertainty.

So a factor of 100 uncertainty, again you're spending close to what the plant is worth to make the -- you're allowed to spend what the plant is worth to make the improvement if you insist on this ALARA principle and then insist on such a large uncertainty.

Okay. So that, I guess, sort of gives you an idea of, hopefully as you look through the vugraphs at your leisure, gives you an idea of what Mike's and Dr. Okrent's framework means in terms of trying to do some calculations with respect to numbers.

DR. OKRENT: Excuse me. Before you go on --

DR. KASTENBERG: Yes.

DR. OKRENT: I'm not sure I recall just how you did the uncertainty calculation in Table 4.7.

DR. KASTENBERG: I took the expected value for the RSS, latent deaths per year, I had obtained originally, multiplied it first by a factor of 10, then by a factor of 100, and then I did the same with the Surry at Indian Point and Surry at Zion, took them with no -- assuming they were correct and assuming they were off by a factor of 10 and then assuming they were off by a factor of 100.

DR. OKRENT: All right.

DR. KASTENBERG: Then I took the difference in the deaths per year, multiplied it by the 30-year plant life, divided 1 million by that number, 10 million by that number, and then compared that to $2 million.

DR. OKRENT: Fine.

DR. KASTENBERG: Then David and I were fortunate in finding an article in the literature having to do with a similar problem for coal, and we went through the same analysis, and David will run through and tell you what we came up with there.

DR. OKRENT: David, before you start that, I wonder again if you could help me recall in that Table 3.3 you showed, what was the definition again of alpha, nu and --

DR. JOHNSON: Okay. Alpha was a power to which the frequency density distribution varied as the dose to an individual varied, the inverse power of alpha. Alpha was one and it varied inversely; after two it varied with the square of the dose, one over the square of the dose.

Nu, to make this function well-behaved with small doses, nu was a co-axis of which the frequency distribution became a constant.

DR. OKRENT: Okay. And Arsebel(?)?

DR. JOHNSON: Arsebel is simply the integral of this function with the action of the relationship, the cancer function.

DR. OKRENT: All right. Well, the last question, and then what's phi over nu to the alpha minus one?

DR. JOHNSON: That's a measure -- there is obviously a large variety of arsebels on that chart. To help me decide which one of those are completely out of hand, that last column is a measure of the frequency of releases which lead to the small doses. If you get numbers like ten to the minus seven, that tells me that the model is not realistic. If I get numbers that correspond to doses that are in the millirem range of ten to the minus two, that tells me that it's not completely out of hand.

DR. OKRENT: Okay. Well, that helps. I'm not sure I know then -- the statement that the Q risk is ten to the minus five per year a boundary condition?

DR. JOHNSON: That's an assumption. It's based to get the normalizing factor. I had to assume a --

DR. OKRENT: All right. Fine.

DR. JOHNSON: But then the bottom line there is the latent risk is an appreciable portion of the acute risk in some of the cases.

DR. OKRENT: There was a hand toward the back. Yes.

SPEAKER: I would like to pose a question. In carrying say an alpha of one versus an alpha of two, a tenfold difference in the alpha one case and a hundredfold difference in the alpha two as far as the (inaudible), I'd just like to know is it reasonable to assume then that your, I think it was $2 million in debt over a willingness to pay, is it reasonable to assume that that value will remain constant for this type of equipment, or is some of this slightly fallacious?

DR. KASTENBERG: Well, the criterion is formulated in terms of effective social cost, however you arrive at that, and that's how it was used in this case.

Now, it may be that you want to change it with risk aversion, but it's built in in this case. I suppose what you could do is not choose an alpha at all and say, you know, if there were no risk aversion, it's $2 million per death averted or life saved, and with risk aversion I'll make it $10 million per life saved.

SPEAKER: Okay. (Inaudible) on the basis for the choice of alpha. I was thinking in terms of an alpha that was chosen on the basis of perhaps a perception as opposed to actual dollar cost.

DR. KASTENBERG: There's a companion paper that Okrent, Griesmeyer and Simpson wrote where they look at some natural disasters and other things to try and show what the range of alphas are that people see in technologies and in natural phenomena. You might want to look at that.

SPEAKER: I can see someone perceiving that the accident is 100 times worse rather than actually being 10 times worse, ... still not be willing to pay 100 times more than (inaudible).

DR. KASTENBERG: Well, actually --

SPEAKER: I understand your answer.

DR. GRIESMEYER: Oh, yeah. It's just a social cost function however you make it.

DR. OKRENT: If I can make one comment, there are some papers in the literature where people have derived or proposed alphas. There's one I can think of by Slezen and Ferrer -- I think they're the authors -- where they just looked at available statistics for accidents, if I recall correctly, beginning at 10 fatalities and going to 1,000 or something, and plotting them. And they observed that they fell off at a certain rate, and from this, if you wished to, you could derive the inference that society is risk averse according to (inaudible) power law.

DR. KERR: Whose power law?

DR. OKRENT: Well, it would be the power law that fit this data, and you can get alphas like two or even a little larger depending upon what data you fit.

I think part of the intent --

DR. SHEWMON: Are you saying that because that's the way the accidents are distributed, society must be welded that way, is that it?

DR. OKRENT: No, no. I'm not saying that. I'm saying it has been suggested that this --

DR. SHEWMON: By some people you choose to quote then.

(Laughter.)

DR. OKRENT: Well said. If you would let me finish, I'll complete the thought.

It has been suggested that one might use, for example, a square law, and I have seen suggestions for larger numbers than a square law. Of course, if alpha is one, there is no risk aversion. The report referred to earlier by Griesmeyer, Simpson and myself was one where we tried to show that if you used an alpha of two or three and took a published risk assessment like the one done for Canby Island, and calculated the expected social cost of this, because of the fact there were large numbers of fatalities for accidents going between one and 10 or 20,000, you computed expected social costs that were staggering, if you used an alpha of two or three. And in fact, the intent is to show that in fact society does not really design its facilities and operate its facilities as if it's concerned about these events with such a power factor, okay.

Now, what Dr. Kastenberg did was just to look at what would be the inference of trying to use an alpha of whatever it was, one and a half and two, if you took the numbers from RSS, to give one again a feeling for that set of numbers, what did it do?

I guess earlier Chris Whipple, who is still, I think, on the same wavelength as I am, said he didn't think one should use much risk aversion. That's what he said in an earlier report. And we had this in mind then. These studies are intended to illustrate the problems you get into if you try to use these things, and in fact, if you applied them to other things than nuclear, they come out even more staggering, let me put it that way, in their impact.

Dr. Whipple.

DR. WHIPPLE: I might add, Dave, that in the paper you and I did together, we looked at the relationship historically on large accidents and did not replicate the data you mentioned in that other report. We found the historical slope was quite close to one.

DR. OKRENT: If you bring in bigger events, that's right, it flattens out. It doesn't keep falling off.

DR. WHIPPLE: It depends on what time period you take your data from, but we did not find exponents of greater than two at all for any time period.

DR. SHEWMON: What sort of a case can be made for the fact that alpha is less than one? It seems to me that there may well be a saturation effect here, or it's not at all obvious to me that a DC-10 killing 300 people in an afternoon is worse than 30 accidents that kill 10 spread over a year or something. Is there any evidence that there is?

DR. SLOVIC: The basic function of almost all psychological perceptions is exactly that sort of function with a coefficient of less than one; that is, the difference between 2,000 deaths and, you know, 2,050 is smaller than the difference between, you know, 10 and 60. That 50 death difference disappears as you go up on things.

So, in fact, people probably have internally multiple functions with alpha less than one and alpha greater than one simultaneously. And there is a certain aversion to a large accident, but on the other hand, we also have this perceptual response which is just the opposite.

DR. SIEGEL: If my memory is correct, I think Brenula(?) used a power less than one to dispel the so-called St. Petersburg paradox which is quite antiquated.

DR. OKRENT: Well, there is a philosophical argument that Professor Bergstrom, who is an economist, gave which would be to the effect that for large events, you would use an alpha less than one, and it could be that in general the people who were grieving for one another would be less.

DR. SHEWMON: Using your earlier argument, we could say that since there have been wars quite regularly through history and they've killed an awful lot of people, we must not be too concerned about killing people. But then let's get on with this.

DR. OKRENT: Okay. I think it should be understood that the calculation that Dr. Kastenberg did was not intended to indicate that one should use an alpha of one and a half or two, but why one might not use it.

DR. JOHNSON: About two years ago Griffiths at Brookhaven published a risk methodology of looking at coal power plants, and Bill and I decided it would be an interesting exercise to reassess their findings in view of Mike's criteria to see what these criteria would imply.

We're not saying that coal and nuclear should be made equal at risk, but simply we thought it would be an instructive exercise to apply these criteria to well-accepted technology.

The assumptions the Brookhaven book used were that the health effects were based on a population exposed to a sulfate concentration, and he developed a model somewhat like CRAC's that (inaudible) sulfate concentrations over a region. The health effects were correlated to sulfur emissions and didn't take into account specifically the other bad actors possibly in the coal, so that these were taken into account through a correlation.

A linear damage component was assumed with the subjective distribution of what the slope of this linear function should be. The health effects were only calculated out to 80 kilometers, and the description of the plant is giving at least 5,000 megawatt electric, 3 percent sulfur coal, 75 percent capacity (inaudible), etcetera.

And I guess the only other important assumption is the same meteorological data we used with all four hypothetical sites. The sites chosen were all in the Pittsburgh area shown on this diagram. The sites were numbered one, two, three and four, and there is an 80 kilometer radius.

Not terribly surprising, site four will turn out to be the best site in regard to health effects, and it's the only one that does not include the Pittsburgh area.

The results are shown in this diagram which is a cumulative distribution of the health effects for the four hypothetical sites, and what we did was simply treat that as a type of distribution one gets when one does a risk study of a reactor and see what the values (inaudible).

These are the results of our calculations. First of all, the populations within the sites are given in the first line. The societal risks are simply the deaths per year read off the previous diagram with the 90 percent confidence value.

Now, we don't really know enough to speak about the maximum exposed individual, but by simply dividing the societal cost, the societal risk by the population, we can get an average individual risk. And simply looking at these as compared to Mike's criteria, you're mighty close, if not over, in either case, so with the expected value and the 90 percent confidence value.

Now, the control co. s were estimated -- this is not going to show up very well -- the Brookhaven group looked at the literature and tried to estimate what are the control costs for a whole plant. What you can't see here actually is plotted the fraction of sulfur removed versus control costs in millions of dollars per year. And it's a curve that you would expect costs nothing at zero sulfur removed, and rises exponentially essentially if you remove 100 percent of the sulfur.

It's a gray band that's plotted here.. That takes into account the various estimates of control costs.

So what we asked was what would it cost to improve plants one, two, and three so that they are just as (inaudible) plant number four.

I should have pointed out that plant four had an expected social cost of 10 deaths per year, which is Mike's (inaudible) criteria.

To improve plant one to make it equivalent to plant four we'd have to remove about 58 percent of the sulfur at $15 million per year. Another way of looking at that is approximately $10 per death averted. And if we look at this final column of dollars per death averted, we see that any of these strategies would be allowed under Mike's criteria. But also we would note that the costs would be on the order of tens of millions of dollars per year.

To return to plant number four which has an expected social cost of 10 deaths per year, we still have to apply the ALARA principle to that plant, 10 deaths. If you can improve on that social cost with less than $2 million per death averted, then you should go ahead and do it. And if you use expected values, if we saved an additional 6 people --

DR. KASTENBERG: No. Four.

DR. JOHNSON: An additional four people.

DR. KASTENBERG: You start with 10.

DR. JOHNSON: You start with 10 and you save an additional four people, it would mean 40 percent of the sulfur was removed and would cost $2 million. And likewise, for an additional five people at site number four, it is also right on the borderline of the $2 million per death averted.

At the upper 90 percent confidence bound, that's the value we choose, therefore; in other words, five deaths per year. At the upper 90 percent confidence bound, that's the appropriate place to make this comparison. It would go all the way down to six deaths a year averted at a cost of $52 million per year, in which case at this confidence level there is still one death per year social cost.

So once again, depending on the confidence level you choose to work at, the social cost in dollars per year that these criteria imply when they're applied to an accepted technology are quite (inaudible).

This final graph shows the strategy of improving plants one, two, and three to the improved site number four; that is, the site which now has an expected cost of five deaths per year. And once again, the numbers just get incredibly larger, many tens of million dollars per year.

DR. LAVE: Excuse me. Why is it that it gets cheaper to save lives at the 90 percent upper confidence bound than it does at --

DR. JOHNSON: They're saving more lives per given amount of dollars.

So I think though the basic lesson learned from all these studies is if we look at the implications of these criteria, it can be quite costly any way you look at it. So we have to be extremely careful how we choose values of alpha for the risk aversion, how many dollars we allow to be spent per death averted. It makes a staggering difference in the total of costs to the society.

DR. KASTENBERG: I wanted to just make one comment. This curve that we used for all the calculations is the cumulative probability in the sense that you're 90 percent confident that you have 86, or 150, or 114, and 27 deaths or less. So it is that cumulative probability distribution.

DR. JOHNSON: It's slightly different than a reactor case.

DR. KASTENBERG: Right. And when we say expected value, we actually integrated the curve and took the first moment of the curve; so when we say expected value, it's not the 50 percent confidence limit, but it is the actual expected value of that curve.

DR. LAVE: Excuse me. One more question, another thing I don't understand about Table 5.4, your (inaudible) is 1.3, and you're getting them to the point where they will cause five deaths per year expected value, correct?

DR. JOHNSON: Yes.

DR. LAVE: And yet you are removing less sulfur than you're removing from plant four to get down to five deaths per year, and I don't understand that.

DR. JOHNSON: It's a fraction of sulfur removed.

DR. LAVE: I understand but, for example --

DR. JOHNSON: They are still on -- they are different sites.

DR. LAVE: Let's try again. Table 5.3, in order to get down to five deaths per year, we're removing 50 percent of the sulfur from plant number four.

DR. JOHNSON: Right. You're releasing more sulfur in that case than in any of the other cases. See, the expected social cost is still five deaths per year in all those cases, but you're exposing more people on the Table 5 quota.

DR. SHEWMON: These are all completely equivalent plants with the same fuel?

DR. JOHNSON: They're hypothetical.

DR. SHEWMON: Okay.

DR. JOHNSON: The only thing changing is the site.

DR. KERR: This is only e .y deaths.

DR. JOHNSON: No. This is a steady state case. It's just correlation with sulfates.

DR. OKRENT: Well, I think what I'm going to propose is we go on to let Dr. Slovic give his comments now since he has a deadline, and then we can come back to a discussion of the previous talks if it's wished.

Paul.

DR. SLOVIC: I really don't have any coherent set of comments.

DR. OKRENT: Dr. Slovic will give us his incoherent comments.

(Laughter.)

DR. SLOVIC: I think I'm even more confused now than I was this morning, although I found the route to that confusion and the path to that conclusion quite interesting.

Let me just reiterate a couple of the points I made earlier. It seems to me that there's an important distinction that we should consider a little more carefully between going through these exercises for our benefit as scientists, say, or as designers, as technical people, and going through them as political exercises where we're really going to face up to the political implications and bring the public in on these discussions and these criteria.

I think it makes a big difference, and whereas I can see some hope for the types of analyses that are being presented today as technical exercises, I'm more skeptical with regard to their success in the public arena.

With regard to public perceptions, again let me note that I think that the model of social cost that incorporates an alpha parameter is not the right model because I don't think social costs work that way. I think that the social costs are a function of what the impact of a particular accident will have on society, and that impact is not just immediate and latent deaths and genetic defects and property damage; I think that's only one small component of the social costs.

An important component again in the political world is what will happen to the industry, for example, in the event of a certain type of accident and what then will happen as a function of what happens to the industry, you know, and if you have cascading second and third order effects, which can be immense really and which may have health effects greater than the health effects that we're usually talking about here.

So I think that needs to be -- the kind of model that would take that into account needs to be thought about. It's really a very different kind of model than has been proposed here.

DR. KERR: But, Dr. Slovic, it seems to me that the points you're making now don't necessarily have to do with whether one uses an aversion model or not; it seems to me what you're saying is that the number of parameters that are being considered as consequences are not sufficiently great, that it's more than health effects that one should consider but other consequences as well rather than how one -- I think what Dr. Griesmeyer was doing was saying given a consequence, here's the way one can incorporate risk aversion. It seems to me you're saying one has not treated all the consequences that you consider to be important.

Am I missing the point?

DR. SLOVIC: I agree with you. I'm saying we have to consider secondary, tertiary consequences in the model; and I think it's hard to do that with an alpha parameter, because the alpha parameter, at least as it's used in some of the examples here, is conditioned on the magnitude of the accident. And I don't think magnitude is really the parameter that's important here. You can have a very small accident which will have immense consequences, and somehow we have to get that into the model. That's really all I'm saying.

With regard to the science court notion, we can debate that I think quite extensively. Let me just comment that I would worry that a court type system which might be adversarial in nature and highlight some of the conflicts among scientists might actually be problematic inasmuch as it would leave the public more uncertain and more confused than it might have been before.

I don't have an alternative to suggest, but I am a little bit skeptical about certain forms of the science court.

I've noted that many -- that oftentimes in developing an approach towards these quantitative criteria, one takes recourse to looking at risk statistics -- for example, the automobile statistics or the violent death statistics that we saw earlier today. And I think here it's important to consider the possibility that nuclear power is really very unique in its characteristics. This point was touched upon earlier. Dr. Mark asked a question about it. I think it's worth emphasizing that at least in the psychological studies that we have done, nuclear power virtually stands alone in the characteristics that it has vis-a-vis things like not only voluntariness but the potential, perceived potential for catastrophe and other adversive characteristics.

So I think we have to be careful in drawing guidance from things like automobiles and other hazards which really have very different kinds of distributions, time effects and so forth involved.

DR. KERR: I agree with you except it seems to me that indeed one can draw a considerable amount of guidance, and the guidance is people don't make decisions based on risk numbers.

DR. SLOVIC: Yes.

Those are really the only comments I wish to make at this time.

DR. GRIESMEYER: It may be that adequate protection of the public is possible without acceptance by the public. And acceptance by the public again is another question that may have to be answered in another forum than what we're trying to do right here.

DR. SHEWMON: Well, but if you're trying to just protect the public, then it's not clear why you want to be 10 times as safe for this as for something else, whereas I thought part of it is if we do penance this way, why (inaudible) their hearts or something.

(Laughter.)

DR. KERR: To me it's interesting to observe in reading some of the history of the development of steamboats or railroads that there was concern about risks but there was also concern about these radicals who were moving people at 15 and 20 miles an hour, which were almost sinful speeds. There was an aversion to risk, but there was also an aversion to what was looked at as a real social change and a change in attitude that probably was not meant to occur in nature.

DR. WILSON: An interesting thing about railroads, in the first passenger railroad, the opening, we killed a member of Parliament, and that may show that sort of level of standard we put on our members of value of life of the members of Parliament. It didn't stop. The proposition didn't arise for about ten years yet. I mean, it really wasn't the same thing, and we haven't had an instance of nuclear power killing anybody. I mean, there's a real distinction that the technologies have. But they did start in a pretty miserable way as far as the safety was concerned.

DR. KERR: That history is one of the things that led to one British comment I saw about Three Mile Island which said that you should kill people rather than frighten them.

(Laughter.)

DR. OKRENT: Dr. Catton, did you have a comment?

DR. CATTON: I was only going to mention gasoline as the fluid of the devil.

(Laughter.)

DR. SHEWMON: I thought it was alcohol.

DR. CATTON: Well, that, too.

DR. OKRENT: Dr. Kastenberg.

DR. KASTENBERG: I just wanted to make a comment, something we didn't point out in going along with Paul's aversion to have a risk aversion factor in alpha, one thing that's interesting if that if you change the ALARA number, we found in the coal study the choice of whether to make the improvement or not, it's very sensitive to that. And in the Brookhaven paper they made a statement, which we have not tracked down, but they made a statement that the new EPA standard is $600,000 per death averted; and if you use the $600,000 rather than $2 million, many of the cases where you would have made the improvement at the $2 million, at the $600,000 you would not. It's very sensitive to that, and the same thing with the uncertainty; it's very sensitive to that. You don't necessarily have to have a risk aversion factor to cause you to make improvements. There are other things that could force you to make improvements -- the ALARA number that you use, the uncertainty you use, and so on.

DR. WILSON: Where does that $600,000 come from?

DR. KASTENBERG: It was in the Brookhaven paper. As I say, I haven't been able to track it down, but the Brookhaven paper claimed that the new EPA standard could be interpreted as $600,000 per life saved.

DR. WILSON: Oh, on --

DR. KASTENBERG: Sulfur removal.

DR. WILSON: On sulfur removal, okay.

DR. KASTENBERG: Yes.

DR. OKRENT: Paul, I wonder if I could press you in the following way. If we assume, perhaps naively, that there are or will be pressures to develop some kind of plausible proposal at least to test out in some way and that one may not have the luxury of doing ten years of research in this area to develop a proposal, what would be your best first draft proposal for risk aversion, for arriving at what Mike Griesmeyer called closure as to what is the legal definition of the calculated risk from a plant, and that's where the science court came in.

In others words, you know, suppose the President called you in and said I want you to come up with some kind of proposal we can start thinking about. You can research it afterwards.

(Laughter.)

DR. OKRENT: Or better yet, if you give me a proposal, I'll give you a research contract.

(Laughter.)

DR. SLOVIC: You ask tough questions. I don't know. I have a different background, and that is, my orientation is more that of decision analysis; and I'm not a decision analyst but neither am I an engineer, so I probably shouldn't be even answering your question.

I guess I would want to look more closely at what the decisions were that were so crucial. I would want to segment it out by the types of decisions that were under consideration and the alternatives that were relevant for each decision, and then try to build a fairly complicated structure of all the relevant factors and parameters. Just like when you ask, you know, what's the probability of a certain type of accident at a reactor, you go through a complex modeling job of structuring how you would even construct an answer to that. I think you have to do the same kind of modeling to construct an answer to the kinds of questions that you're asking, and that modeling gets involved both with various sorts of possible events and consequences and probabilities and costs. And, you know, that would be the approach.

Now, to some extent what we've heard today is, you know -- does that, although it seems to be operating at a more general level. You know, it's kind of an answer for more situations. And what I really don't know the answer to is the utility of working at this more general level and trying to get numbers that have utility for a wide variety of situations as opposed to being more specific and focusing on narrower classes of decisions and getting numbers that have less general applicability.

end tp 10

DR. OKRENT: Let's see, I gather you were discussing the risk version portion of my two-pronged question.

DR. SLOVIC: Well, I was trying to answer the question of what best approach to propose to someone who needs an answer. I can see I confused you.

DR. OKRENT: No. If you want me to give you a specific, well-defined problem on which you could provide a proposed approach, I could give you one that the NRC is grappling with now. If they have a reactor at a site which is on -- -- number ten times more populated on the average than their average site, and if one just then assumes that all reactors are equal in this regard, they are probably even emitting radioactivity, whatever that is, how should they approach this question?

DR. SLOVIC: What is the question specifically?

DR. OKRENT: In other words --

DR. SLOVIC: Whether the site is acceptable?

DR. OKRENT: Yes, that is one question, and that if they decide it is acceptable on what basis would they do it. If they decided they needed to change the reactor to match the site, on what basis would they do it, and so forth?

DR. SLOVIC: Well, do you want an answer right now?

DR. OKRENT: No. I tried to make the question specific to see if that would help.

DR. SLOVIC: Yes. I don't feel that I am really the best person to answer that. I don't know.

DR. GRIESMEYER: Well, there is a tradeoff between completeness of the decision process and the cost of delaying decisions. And so if you make a set of decision problems that includes the cost of delay of the decision, and then you can complete a long decision analysis, do it; where time of decision and the resources that fit to make the decision becomes one of the aspects of the decision -- -- and in that sense then you are going to figure out a way of cutting off some of your completeness and trade it off for time, because time is money and delay. And if you can structure a decision analysis in that way, I think that would be the sort of thing we are trying to do, is there is definitely a tradeoff between the resources we can devote for this particular decision if we have got 75 similar decisions to make.

DR. SLOVIC: Well, there is always a tradeoff between time and effort and so forth, but I don't think that is the problem here. I think the problem is knowing what you would do. In many cases the decision is important enough that you could spend, you know you have the resources and you can devote the time. The question, I think, is, okay, given you have the time and resources, what sort of modeling would you do and how would you do it. And I think again it is sort of a state of the art, you know, with regard to the modeling and the analysis that is the problem right now.

I think once we make some progress in that direction then maybe some general principles could be developed which could then be applied more efficiently.

DR. OKRENT: Any other questions anyone wants to put to Dr. Slovic?

DR. WILSON: I am not quite sure when you say is it efficient. I mean to understand what model to use. I am not -- most of what we have discussed today has been basically what might be called the engineering model of how to calculate actual risk in engineering systems. As far as I understand you, you are not using "model" in that sense, are you, but the model of how to think it for this? Is that right?

DR. SLOVIC: I am thinking of an analysis that would enable you to choose among alternative sites, alternative designs. You know, that being the problem, then how would you structure all the relevant factors? What are the various consequences and what are their probabilities and so forth? How would you amalgamate the diverse consequences, this sort of thing?

DR. SIESS: Could you model the present

decisionmaking process for licensing nuclear power plants, essentially up through the -- let's say that the initial decision of the ASLB is the decision we are referring to. Can you model all the steps up to that? Or could you?

DR. SLOVIC: I don't know enough about the details of the process. One often can model the decision process of an individual, an expert, a position, some other sort of expert. Decision processes th t are very complicated and seem to be almost unarticulable.

DR. SIESS: That describes the one I am talk'ng about.

(Laughter.)

The present process involves the technical engineering decisions, environmental decisions. It involves lawyers quite heavily, both rational and irrational intervenor actions. I don't know how it can get much more complicated.

DR. SHEWMON: Let me change the subject. What do you mean by the word "modeling"? Do you mean "try to predict"? Do you mean "putting a vent tree down"; instead of -- -- the pads they go down -- yes-no decisions being made?

DR. SLOVIC: I guess I mean imposing a structure on the analysis that will lead you to some answer or some decision. But I would consider a fault tree or vent tree a model that is used to come up with an estimate of risk. And I think similarly if you want to make a decision one can think about how you would structure all of the relevant items of information in order to help you make that decision, analyze the problem in proper depth.

DR. OKRENT: I think we are at ten to five. We had better let Dr. Slovic catch his plane. Are there any points anyone wants to raise with Drs. Johnson or Kastenberg on what they presented?

I have a question which is raised by the discussion where you gave the average risk, Dr. Johnson. I think you said that you didn't -- of estimating the risk of the most affected individuals by a single plant.

DR. JOHNSON: Right.

DR. OKRENT: Have you seen any such estimates that have been made?

DR. JOHNSON: No.

DR. SAUNDERS: I would like to add, it seems that in the plan for the site that you chose you had circles drawn around them as if there were CPE at each site. It seems to me that they should then function as a -- -- and god knows what.

DR. JOHNSON: One of the assumptions in the model was that a single set of meteorology data was used at each site.

DR. SAUNDERS: Oh, I see. So it didn't -- -- with the circles --

DR. JOHNSON: All the circle basically did was select a different target population.

DR. SAUNDERS: So you have the same dispersion at each site?

DR. JOHNSON: Right, I understand.

DR. WILSON: So there is probably about one-third of the total effect in this case.

AUDIENCE: Dave, were you just asking for the risk to the individual at the coal plant or the nuclear plant?

DR. OKRENT: The coal plant.

Yes?

AUDIENCE: Dr. Okrent, on that point, in the paper that myself and a colleague published last November in Nuclear Safety, we did attempt to estimate the maximum risk to an individual from a coal plant. And we looked at the maximum offsite concentrations that are emitted under the Clean Air Act. And using a linear model, we estimated the risk at about two -- -- minus four to the maximum exposed individual.

DR. OKRENT: That is per year.

AUDIENCE: Per year, yes.

DR. OKRENT: And can someone tell me, do they generally meet these limits on the average over the year?

DR. SAUNDERS: Yes.

DR. OKRENT: Okay, so it is not -- okay, thank you. Are there any other points that arise out of this presentation?

If not, then why don't we, I guess, go from my right to my left if that is an orderly approach, however it is received -- Dr. Lave.

DR. LAVE: (inaudible) of the Emperor Justinian reads that sodomy caused earthquakes. I didn't worry very much about that in Pittsburgh, but I understand the people in San Francisco worry about it a lot.

(Laughter.)

I think that it is that kind of thing which really characterizes most of what we have been doing today, which is that you have to get the scientific facts right before there is very much else that you can do, or at least the perception of scientific facts right.

There is a very nice paper that I read which helped a lot with trying to answer David Okrent's question, the one he has been effecting on what I have done, about what the actual decision, what rules should we have. The paper is one by John Jackson and Howard Conrichter at Penn, who are complaining bitterly about what the NRC has done to date and who talk about the process by which scientists, and particularly engineers, attempt to convert value questions into scientific criteria issues, which are more or less drawn arbitrarily out of the air without recognizing that value questions are value questions and don't get dealt with drawing some numbers arbitrarily out of the air.

So that I would describe most of what we have been doing today as not coming to grips at all with why it is that nuclear reactors have not met with the greatest of public approval and success, or at least why it is that there are significant elements in the population that oppose them.

Let me develop that before I try and give you some suggestion for what to do about it.

I mentioned earlier that consistency is something that logicians try to impose on us and that we all try to impose on our graduate students, although every time somebody actually talks about the scientific method in practice as distinct from what it is we try and teach high school students, it doesn't seem to work that way. But in particular, consistency just doesn't go at all when we are talking about politics; that is, that policies get made without any particular regard for consistency as a virtue, and I think that we cannot -- it is not necessarily fruitful to compare public decisions in one area with public decisions in another and say, see, they are inconsistent, as if that was a great source of discovery of something or other.

But one of the things that I noted today is that virtually everybody, except perhaps David Okrent, has assumed exclusively that nuclear power ought to be safer than any other technology. I really wonder about that. Why is it that we have that kind of a notion. I think that many of us, or at least they seemed to have expressed it, and I will leave myself out of that since I haven't said it yet, many of us believe that there are value differences that are inherent in nuclear power somehow in comparison with other kinds of technology. And when this group comes out with statements of that sort, it is just not difficult to see that the general public is also going to have some value conflicts, indeed probably more general value conflicts than we have here.

One of the points that struck me in the midst of all of this discussion about ten to the minus four and ten to the minus eight is the issue of verifiability. That is, how in the world does one know that something is ten to the minus eight, and this got mentioned I think at the first meeting that I attended of the subcommittee, where I think that one of the major issues in the public's mind is not whether ten to the minus seven or ten to the minus five is s'ed dnoufg aut re'lly how it is you knov thas is ir sem sn the minus five.

I think there is a fairly general interpretation, that Three Mile Island was an indication that those probabilities are not as stated but in fact they are much larger, and not a factor of two larger, but perhaps ten to the four larger than had been said before.

And so the issue of verifiability strikes very large, and I must say that as a practitioner of all this I am highly skeptical of what those numbers are, as to whether they are right within a factor of two or factor of ten or factor of a thousand, particularly when one is trying to introduce a fact such as sabotage, terrorism in there.

And so verifiability seems to me to be this very large area we haven't heard very much about, and I guess I would have thought that before one can go public with any proposal that you had better have something quite concrete to say about the extent of verifiability.

One minor point in all this is there seems to be this great reason to compare risk with background levels, assuming that the implication is that if we are going to build reactors we ought to install them and say either Miami or Wilkesbury where we have this enormous population of 65 and older of whom are at daily risk not ten to the minus three, but more, a greater risk than -- more like ten to the minus one.

Apparently if we were to put a reactor there, then these people would be much more accepting of the risk. Or alternatively, we could put the reactors in Denver because they get so much more radiation anyway. I must say I think that those are entirely spurious, that the comparison with backgrounds may have some psychological effects when you are presenting a number to the public, but I can't believe that they have more than that.

Let me just try and be a little bit more concrete. That is, I think that today has some fruitful, there were some fruitful effects of today, not because we heard different people arguing for ten to the minus five or ten to the minus six, but because at the time we were actually going to make some progress somebody has to come up with some concrete proposals which are fleshed out so that we know what the proposal is, what the implications are, whether they can be administered and so on. And I think that there was some of that today, where we actually began to get some fairly concrete proposals and began to look at something of the implications.

The problem is that those proposals are being made to this distinguished subcommittee, which is in no position to resolve these value conflicts that I started off with. Whatever your distinguished backgrounds in the past, nobody elected you to resolve their value conflicts. I don't imagine many of you have had very much experience with all that, and at least on the basis of some of my dealings with your chairman I would not think that you were terribly skilled --

(Laughter.)

-- in resolving these kinds of conflicts.

DR. OKRENT: We value independence, whether we agree with it or not.

There is a single institution in our society that, and this is for the whole of the United States, actually is duly constituted to resolve value conflicts, and that is the United States Congress. It is the only body that is elected by all of society specifically with the goal of looking at these conflicts between individuals and resolving them somehow.

And the only time we are ever going to get any definitive resolution of these value conflicts is from the Congress.

On the other hand, the Congress is bent on getting reelected each time, and they know better than to muck with something that is going to get them into trouble. And so time after time they have tried as hard as they can to not resolve value conflicts that they don't have to resolve. And one of the favorite techniques has been to create an independent regulatory commission, take problems that they can't resolve, shove them off onto the commission, and then when the commission makes a judgment say gee, that was a dumb thing to do.

So although my first preference would truly be to get Congress to resolve these value conflicts, it is probably not going to happen very soon and thus we are still left with the question that David wants me to address; namely, what in the world can this subcommittee recommend to the NRC.

There I think that I don't have anything very powerful to recommend just because you are the boss. But I think that if I have something to recommend, it is that the NRC ought to begin to act as if it were in fact the body engaged in resolving value conflicts. That is, it ought to use these time-honored practices that the Congress uses; namely, to publish a set of proposals in the Federal Register, to put out a notice that you are going to hold hearings on these, to then start holding some hearings on these various proposals, have various groups from society step forward and have their say or have a very skilled chairman who manages to let everybody believe that their comments are highly valued and will be taken into account in the final decision even though assuming you can't please everybody, and that after holding all these hearings probably for at least a year and two and getting everybody exhausted in the course of all this, compiling tens of thousands of pages of testimony, then the NRC can come out with some kind of judgment, at which point the courts and the Congress will look at this record and say, gee, if they were that careful it must have been a good job.

And in the course of it -- let me not be cynical -- I think that they would learn something from the hearings, that in the course of developing these views their proposals would be clarified, the value conflicts would be clarified, there would be some ability to trade off, to give a group something rather in return for getting something. I think that their proposals in the end had some chance of working, as distinct from this one that says well, any idiot can see that if your chance of death is only, is ten to the minus two or greater per year, that therefore it must be that an increased risk of ten to the minus six is a negligible one and you ought to accept it.

The fact is that people in our society don't accept that for whatever reason. So that I think that if the Congress isn't going to resolve the value conflicts, which I would hope it would but don't think it will, then the NRC is going to have to resolve them, and the only way that I know of offhand for doing that is to get into the whole tedious business of holding hearings.

And that is all that I have to say.

DR. KERR: (inaudible) decisions?

DR. LAVE: Well, in the best tradition of the Supreme Court, all of their decisions are five to four opinions, and then wind up changing ten years later when the court gets changed around, I don't see why these should be different.

But in particular, when you have value conflicts which cut right across the sensitive points of society, then damn it, you ought to have a three to two decisions, because they are that close.

And I think the public by and large recognizes that when you have a five to four decision from the Supreme Court that there is a lot of agonizing that went into that, and even if it didn't cut your way, then you say, well, gee whiz, as soon as we get one more justice on there we will bring up the case and get it reversed.

And sure, that is hope springs eternal. Democracy has gone on from one year to the next. They don't settle an issue for all time.

That is part of the way in which one placates people on value conflicts of that sort, that it isn't set in stone, that there will be opportunity to argue it again if there is a new commissioner or there is some new evidence of some sort.

Now that doesn't mean that every time you get a new Supreme Court justice that you settle all the old cases again, but at least hope springs eternal.

DR. OKRENT: Well, Lester, as you know, the ACRS recommended a little over a year ago now, I guess, that the NRC try to develop such criteria and in fact propose them to the Congress, that the Congress could say whether they thought these seemed okay, or at least they didn't say they were terrible or wh tever.

I have to assume that if I were a commissioner and was trying to decide should we tackle this thing and should we have hearings, as you just proposed, I would like to think that there are going to be some plausible approaches to discuss. In other words, we are not going to be proposing to hold hearings on something, and the whole thing is going to turn out -- people say, you fools. I mean this is just the wrong thing to deal with.

So I think there still is a need to be able to have something that at least can serve as one of the points of discussion. In fact, I have reason to think that this is one of the kinds of thought processes that the commissioners would undergo, and I think they are interested in seeing whether one or more plausible ones can't be provided as a starting point.

I think in fact the very -- especially if one of these, or more, could be provided I think that they would then go into the very step you have described. I think that there would be a good probability that they would.

DR. LAVE: Yes, I think that it is essential that out of this process that you try to get not a consensus on some acceptable risk but that instead you try and numerate a broader way of approaches as fully fleshed out as possible, which would under my proposal be published in the Federal Register and serve as a basis for comment, along with what other ones people want to propose.

But for the other part let me just put to you again that if you were the Congress you would not let one of your independent regulatory agencies force you to take on a task like deciding what acceptable risk is.

I think that the position the Congress is going to be in is after the NRC has made its decision then the Congress will test the wind and say it is okay or you are a bunch of damn fools, and that is the role they enjoy, they feel comfortable with and they have had lots of experience with.

(Laughter.)

DR. OKRENT: They may in fact, as I tried to indicate earlier, choose not to publish it but also not to say it is unacceptable, in other words, and allow it to proceed unless it gets into trouble.

DR. KERR: What they might do is sit back, as you say, and not bless it and then wait for the first accident to occur and then say what a bunch damn fools you are.

DR. OKRENT: Bernero?

DR. BERNERO: Bob Bernero, NRC. I would recommend if you haven't seen it already, the Commission has gone in the general direction Dr. Lave suggests on the Indian Point case, that order the Commission put out on May 30th, I think it was, something like that, which sets up a hearing board to gather extensive evidence of a rather murky mixture on whether Indian Point is too populous, too risky or whatever, and then the Commission would make a judgment on that in the end.

But I think it suffers from what Dr. Okrent was pointing out. It doesn't really have a cohesion to it in its present form. It is just a mixture of all sorts of things.

DR. LAVE: Let me add one caution, since the Commission is going to get into this, and that is you ought to take a look at the kind of hearings that OSHA or EPA has held and know that you had better get some very tough people in there to hold those hearings, because, first of all, they will last a long time; secondly, people will not be polite to one another, and so on. But then when you are dealing with fundamental value conflicts, that is all of what goes on. And we have some fundamental value conflicts here. But I think it is better to get these settled in advance insofar as you can than it is to pretend that they don't exist.

DR. OKRENT: Dr. Shewmon?

DR. SHEWMON: Doesn't anybody else do things as rationally as we are trying to do? What I have in mind is now that these numbers of ten to the minus six and ten to the minus four and ten to the minus eight, and separate is the question they are unverifiable, and we can discuss that over a beer some night and exchange stories. But let's come back to the basic question of ten to the minus six is a nice -- -- and maybe it is the thing that always comes to a technical person's mind or some of them, but is it done in this country by these other agencies you are talking about or do people just back out, well, I think it would be $600,000 if OSHA did it that way or EPA.

DR. LAVE: There have been a number of attempts to go out to this. There are these -- -- called economists around who get in and start advising agencies to do damn fool things, and agencies always knew in the past that you never face squarely up to any issue that you could manage to skirt around, and we have had in Congress that have gotten high up into various agencies so that, for example, the National Highway Transportation Safety Agency at one point published a memo that said that for analytic purposes a human life would be worth $200,000 in all of their decisions. And they stuck with that for, I think, 36 hours before the roof fell on their heads.

But as a matter of fact, the Federal Aeronautics Administration has a value of life of $300,000, which it uses in every one of its decisions. In every look at aviation safety they explicitly look at how many lives would be saved on a probabilistic basis. They value all those lives at $300,000 apiece, they compare them to costs, and they decide whether to go ahead or not to go ahead, or better still, whether to apply it to this airport or to that one. And they decide this airport should get it and that airport shouldn't.

So there are two agencies. When you start taking a look at the health area rather than accidents, since accidents are fairly clear-cut, somebody dies or they don't die, and you know the accident took place. But when somebody dies of a lung cancer, you don't know whether that was due to their cigarette smoking or whether that was due to have been irradiated.

When you look at the health area, then agencies like the FDA just loves to have the delaying clause around where they can just lean their elbows on it at any time the going gets tough.

But there are a couple of areas, the Toxic Substances Control Act created a nightmare for EPA, since under that act they are explicitly required to balance benefits and costs or risks and benefits somehow, and EPA has found it virtually impossible to implement TSCA because of this provision.

The FDA must explicitly balance benefits and risks when it decides to license a new drug. And mostly what the FDA does is to decide not to act for right now, but to ask for more evidence.

So these kind of decisions are being made explicitly by agencies. They try and develop criteria for them. They find it extraordinarily difficult and uncomfortable to make a judgment, and in general they try to gather more evidence to avoid making judgments. In the end when they make a judgment they face up to these.

And then almost invariably they wind up getting sued. So there isn't any easy road, but there are other agencies that are doing it.

DR. OKRENT: Yes?

AUDIENCE: I think there is one point where Congress does (inaudible) they actually go through the budget every year, and they are implicitly making some sort of decisions in value conflicts.

DR. LAVE: No, that is right. There is a way in which the Congress guides -- well, first of all, Congress every year in every one of its judgments resolve value conflicts. Every time they go through the budget process they decide how much money for Headstart and how much money for Defense, how much money for the Nuclear Regulatory Commission. So they do settle all these things every year. It is just that in the midst of that they had just as soon not take on any that they don't have to, like some sort of universal decision on what the safety of nuclear power ought to be.

There are a number of inferences that one can make about what it is that Congress does or doesn't do. The problem is that Congress moves around faster than the common law so that it is very difficult to draw any general thread out of congressional decisions because the Congress changes every two years, and with congressional change you get vast differences in the kinds of value resolutions that they come to on all this.

Nor would I particularly advise that the NRC spend a lot of time studying recent congressional decisions, since I don't think there has been enough consistency there so that very much would be learned. I think that they would be much better off by going directly to the people if they can't get the Congress to address it specifically.

DR. OKRENT: Thank you, Lester. Dr. Wilson?

DR. WILSON: Now the question of what I can add to this. I circulated a statement of some other thoughts beforehand and which have now changed of course. One of the problems is I open my mouth so often I can't remember what I said the last time.

(Laughter.)

So in somewhat of a short length of time I have changed. But the general framework, it seems to me, of what I think is, are the following, that it is important if you are trying to discuss (inaudible) to put it in a logical framework. Otherwise, (inaudible) what set of decisions are there that people want to ask with which this risk of the nuclear power (inaudible).

Should I list here that nuclear power stations are dangerous or should I not? Should I (inaudible).

So the whole set of decision questions which involve different constituencies, and the question is who decides them.

Now that is a whole difficult question which we haven't understood (inaudible) but you must realize that that is the framework in which the whole thing we are talki..g about must fit.

So here we get a jump in our logic, and I said there were several constituencies. There are probably 200 million of them in this country. (inaudible)

So then these are the groups that one tends to come up with -- the occupational risk, the occupation, the people engaged (inaudible), the individual of which this seems to be the site boundary, and the general people -- I made a slight distinction there because I think it is worth making, because a lot of people have said one in a billion risk per year (inaudible), but we know that it won't get down that low (inaudible), and then the society, it might be 200 a year, the whole technology, or 10 per year per power station.

Then comes the question, are those the exact numbers. It seems to me we can talk about and discuss. Now one thing that we have to mention -- -- uncertainties in all of these, and of course that is where we end up -- verifiability, and I think that is where the problem, some of the problems come, and I think that (inaudible).

We had to understand the uncertainties, which include of course the fact that we excluded sabotage, which may or may not be important, most important. The fact of the uncertainty in the figures (inaudible).

. So this gets you to how you implement something like this in the face of uncertainty, bearing in mind you have got to have a calculation procedure which is definite for calculating risks.

I would think from what I have listened to the difficulty of verification is somewhat the same point as the uncertainty. It is applied to different features, and uncertainty is established mathematically.

(inaudible)

So we have to go on, I think, and leave them out. And one of the problems is if you leave them out, and mind you people are all the time leaving them out, you get yourself into trouble.

Now Rasmussen's approach to it, and the median was to allow normal distribution, and I know that has gotten into trouble with a lot of people who perceive this -- Richard Garwin, for example (inaudible). And what do you mean by this, that when you have got a medium risk and so on?

Well, I think that is moderately, well, leaving a rather normal distribution. That is about normal distribution, I think (inaudible).

Well, I don't know what that means, but let's suppose it means that the sigma is one and the two-sigma -- -- we wanted the probability rather than the consequence, and of course since you multiply the two together you have to hold them together.

And so I would include this (inaudible) a mathematical procedure for coping with that. As far as I know, no one has done, and it could (inaudible).

Now the societal risk question is that it is simpler, and I don't have to mention that here. We know, for example, 80,000 people are going to get cigarettes, lung cancer from cigarette smoking every year, and that is fairly well known. However, if we were just calculating that from some model, which had an uncertainty in it, we could just say, well, this uncertainty, the sigma of (inaudible). It is all important to me whether that 80,000 is really 80,000 or whether it really is 100,000, which means society has to cope with a different number of -- -- different number of hospitals, it has to -- -- all the different things of society would be different on that basis.

So on that basis I don't know what one does, but it might be -- -- to go to a 99 percent -- -- or two-sigma limit, or something like that -- -- and I think they have to be handled differently, and that is very difficult.

Once one does that I think one copies with some of the societal problems of large accidents, because -- -- the uncertainty in the risk of very large accidents is very much greater than (inaudible).

So the last thing I want to talk about is relatively important.

ape 12
RC ACRS
/1/
ield/
atfield

So the last thing I want -- thing I should have mentioned before, I'm not sure that the -- (WORD UNINTELLIGIBLE) principle is much more fundamental and important in the framework than a (WORD UNINTELLIGIBLE), and that's because (WORDS UNINTELLIGIBLE) should be thought about as (WORDS UNINTELLIGIBLE). We should, in some sense, be starting with that and deriving all the other things from it, if we knew how to do that logic.

And so it's important not only at the end, after you have done those in minutes, it's because (WORDS UNINTELLIGIBLE) if we can (WORDS UNINTELLIGIBLE) far out from the (WORD UNINTELLIGIBLE) principle and industry gets a little more guts than they have at the moment, and we see, in fact, a hundred off in these limits from what (WORD UNINTELLIGIBLE) would tell you, then I calculate that out and say, well, I'll start suing NRC or something for their giving too flat a criteria.

And so that's a judgment in which, in some sense, we are judging these criteria which (WORDS UNINTELLIGIBLE) against that (WORD UNINTELLIGIBLE) principle.

And now I would, the other thing I'd emphasize, I think, today, I think it's important if we put these things to get the framework right, and in that framework it's the core melt, et cetera, criteria for the hazard indices are the large or secondary criteria. That's not to mean to say that they aren't very important. They are very important -- I mentioned them in my first talk (WORDS UNINTELLIGIBLE). They are most

JO-2

important, it seems to me, for three reasons. Number one is because they're objective, more objective than some of the others. Number two, they are what -- they do address the public perception of the particular accidents of the moment, the (WORDS UNINTELLIGIBLE) Three Mile Island. Number three, when you're talking about this kind of consequence, you want to break up the different elements of an event tree as much as possible which you're not quite sure one or the other. And so you want to both make sure that core melt has a low probability and after a core melt you have low consequence, because even if you go by sabotage, if you have it properly quantified (WORDS UNINTELLIGIBLE) core melt. And so that using the consequences of core melt, it helps in -- in sabotage. So that breaks up that into little bits.

But one must be aware that it is a secondary criteria, and then, of course, it would change, it would be (WORDS UNINTELLIGIBLE) that may change a little more than the others.

Now, having said this, I think one now has to play the whole record backwards, because it gets to the point that Lester Lave was talking about, that what we've got to end up with is something people are going to accept in the end and which finally goes to (WORD UNINTELLIGIBLE) the Congress and the people. And so although one comes up with these sorts of engineered things, one then has to go back, are one -- is one, in fact, reaching (WORDS UNINTELLIGIBLE) these first few things, O-3

are you answering those questions. And those aren't questions (WORDS UNINTELLIGIBLE) can answer. We can only make guesses of what the things are. (WORDS UNINTELLIGIBLE) would, presumably, address. And I don't see any (WORDS UNINTELLIGIBLE) forward of that kind of change. But maybe -- you've got to make sure that the (WORDS UNINTELLIGIBLE) we are starting at the beginning and keep addressing that method but each time.

Those are my feelings, Dr. Shewmon.

DR. SHEWMON: Would you put back the part where you have societal risk (WORDS UNINTELLIGIBLE) and explain what you said again.

DR. WILSON: Okay. What I pointed out here, as I say, (WORDS UNINTELLIGIBLE) risk, suppose we take the risk calculated by the WASH-1400 methodology, and it's, what, .02 per year (WORD UNINTELLIGIBLE) per reactor year, it comes out to be. Now, there's uncertainty in that, where the -- if I understand that (WORDS UNINTELLIGIBLE) puts sigma one equals one and sigma two equals one, approximately. As Rasmussen (?) said, it might be five times bigger or might be five times smaller. Then you could put them equal to two and say (WORDS UNINTELLIGIBLE). Now, it turns out to be, turns out smaller, that's the thing to be thought about.

Now, when I say that, that would mean I've got -- or should be combining those in some sense. And what I'm saying is, one wants to somehow go to (WORDS UNINTELLIGIBLE) kes a full probabilistic distribution, say it's twice the standard deviation of (WORD UNINTELLIGIBLE) normal distribution. (WORDS UNINTELLIGIBLE).

DR. SHEWMON: And you'll take the main thrust from sigma, and not say it equals two sigma.

DR. WILSON: Well, we'll say if we want to protect things, let's go up to the -- we'd go up, up -- because there is, having said that, because there isn't a station anywhere in the (WORDS UNINTELLIGIBLE) NRC and I'll make them, I'll lean over backwards, I'm being a little safe. That's what I would put in there. I wouldn't take the median value or the mean value. I would take something more -- more (WORD UNINTELLIGIBLE) more conservative, more pessimistic than that.

And that does two things for you. It gives you a (WORD UNINTELLIGIBLE) -- well, one thing, you gain incentive for everybody who is (WORDS UNINTELLIGIBLE) process of adducing risk and the process of reducing the understanding of risk; if you've got that uncertainty in there in that basis, you've got both uncertainties (WORDS UNINTELLIGIBLE).

DR. SHEWMON: Thank you.

DR. WILSON: It's -- it's fairly consistent. And I must apologize, by the way, to Dr. Griesmeyer and Dr. Kastenberg for treating them as a completely degenerate pair of ACRS fellows (WORDS UNINTELLIGIBLE).

(Laughter)

JO-5

(WORDS UNINTELLIGIBLE) I referred to Dr. Griesmeyer -- to Dr. Kastenberg when it was Dr. Griesmeyer had done something and vice versa.

DR. OKRENT: Any questions for Dr. Wilson?

Dr. Lowrance.

DR. LOWRANCE: My comments are very brief. There must be some point of diminishing return here.

I generally think that this same quantitative modeling approach to probabilistic risk assessment ought to be encouraged. I think the kind of work the fellows is doing, are doing is very, very useful. I do believe that at some point these proposals need to get firmed up a little bit and exposed a little more broadly to public comment, both for smooth tactics as well as for real refinement. And I don't think that's a bad consideration.

I'd feel a bit better about it if they backed off on the Science Court proposal. I think that's what we have a Nuclear Regulatory Commission for. We set up regulatory authorities so we can knock them down when we disagree with them. And putting one more set of scientists in there, it seems to me, doesn't help a whole lot.

I would hope, then, that we'd take some of these general approaches and try to apply them to a particular real situation and say how do you get from the very broad kinds of things that, for instance, the fellows have been doing, how do we get from those, which really in a sense are overall summed performance characteristics, to design characteristics and operation characteristics for a real power plant in a real place.

I think also one might go back -- and perhaps this has been done -- but go back and look at some accidents, such as the Three Mile Island sequence, and ask how those are accommodated within some of the models that are being considered. And so there seems to be a little loose talk about Three Mile Island now as though it were not a bad accident at all, I mean, seeing no one was killed, or maybe one person in the long run, or something of that sort, when, in fact, it seems to me, it was a damn near miss. I don't think that's a point of just hysterical perception. I really think some things went far out of design, so to speak, and there again it's that domain between design and some gross performance that I -- that needs to be examined a little more carefully.

And perhaps this has all been done internally, but I am not as aware -- I'm not aware of that myself.

If the -- if some of these ideas get exposed in hearings over the next few years, they could be conducted, I think those hearings could well be conducted, by a good administrative law judge. That's the kind -- I agree with several comments that you need somebody tough to run those hearings, and I think the NRC could do itself a favor by getting an experienced, scientifically trained administrative law judge to run

O-7

the hearings.

DR. OKRENT: Tough and beyond.

DR. LOWRANCE: Well.

(Laughter)

Just briefly.

DR. OKRENT: Could I ask two questions, or at least explore two points.

You suggested it might be useful to expose, at least, let's say, the kind of thing that the ACS Health has been working on through a broader segment of the section. And I -- you all are supposed to be the first cut at that broader segment, of course.

(Laughter)

DR. LOWRANCE: Well, we have done a little bit of that.

DR. OKRENT: Yes. Now, what would you suggest as a possible mechanism or possible mechanisms for doing that? Do you have any favorite ideas there?

DR. LOWRANCE: No. I in general favor structured hearings as inviting certain kinds of people in the first round, certain, let's say, experts, and notions of quantitative analysis in general. I'll think about it. Perhaps privately, or later in a letter, I can suggest a format. I think just opening up and saying, "Here are some ideas. Does anybody want to come to Washington and comment?" would just be -- wouldn't be very productive. Other agencies have tried that. The FTA does that routinely with its proposals in the Federal Register and gets lots and lots and lots of feedback and doesn't quite know what to do with it all. So more structure than that I think would be good. But I don't have a specific proposal.

DR. OKRENT: All right. Now, getting back to the science court, I'd like to explore that a little bit.

My understanding of the current Nuclear Regulatory Commission is that the members of the Commission itself are not selected on the basis that they should be able to review probabilistic analyses and judge which is valid or what the real answers is, or so forth. And I don't anticipate that that is going to be an important criterion for a Commissioner. So if the Commissioners are the ones who are going to arrive at the judgment on what is the most probable level of risk for a specific reactor, it seems to me that they can't arrive at the judgment from within: they have to take someone's advice.

So you could say, "Well, all right, the Commissioners can set up a group, the NRC staff perhaps, who do it for them." Well, in a sense, you could say that the NRC staff are acting in that capacity now. However, they end up, as I observe it, being parties in developing what I'll call safety positions, and sometimes they'll come into an analysis with a position in mind and I don't know whether the analysis is prejudiced; I have my own suspicions -- okay? -- that, you know, perhaps subconsciously. But as I have indicated, many times, I think, in this business, it's easy to get a range of answers quite honestly, not even knowingly. In fact, if you didn't remember what you did last year, you might very well get a different one this year, you know, about the benefit of new data and so forth.

So it's for that reason, at least, I'm a little skeptical that the NRC staff are the ones who should act as the judge as this is the -- and also they end up, as you well know, in a protagonist position sometimes with the industry, sometimes with intervenors, and so forth. So it would seem to me if it's -- if there is to be some kind of a group that provides a -- a judgment that this is our best estimate of what the actual risk is, there would have to be another body within the NRC if it worked within the NRC. At least, that's my own feeling. Now, maybe that's the wrong route to go, but if so I need to know what the other route is to closure of the question of this is the level of risk that we're going to use as the working level until we know something else. It's in the absence of any mechanisms that I can see that exist that, in fact, I suggested to Griesemeyer we have to propose something to meet this need and it needs to have some kind of legally acceptable status, it seems to me.

So, you know, maybe the term "science court" is the wrong one to use, because the -- it has a connotation in people's minds. So we need to invent a new word or term or something. But I need to see another way of doing it if not via what I will now call a kind of a science court but not the science court that was discussed in Science.

DR. KASTENBERG: Could I suggest thinking about your question later on and not just answering it now. If someone looked at the hearings that led to your final acceptance criteria for (WORD UNINTELLIGIBLE) for cooling system -- I don't know that this is the answer, I'm saying that this is representative of a very difficult problem that WASH handled by hearings and a decision then was made by the Commission. That might convince you to use that approach, it might convince you that it is an approach not to use, but at least it's -- that was not a (WORD UNINTELLIGIBLE).

DR. OKRENT: No. Well, let me comment on that. In the first place, that wasn't intended to be a generic solution. I don't think the Commissioners could have lived through doing this reactor after reactor after reactor. I don't believe the Commissioners themselves tried to become technical experts. In fact, I know very well that after the staff had reached its position and testified and everybody else testified, the Commission got two or more very well thought of individuals from the technical community to be consultants directly to the Commission; and they relied strongly on the advice from these consultants. Which I think is (WORD UNINTELLIGIBLE). I can't remember, but one or more of the Commissioners then was a lawyer, none of them were experts in low-heat ECCS.

DR. LOWRANCE: This is a common affair in lots of government agencies, probably 30 or 40 of them. And I don't see anything wrong with the scientific assessment going on within the agency, and I don't see any real advantages in setting up another bodies that's somehow a little bit more independent.

There may be. But independent experts always have their biases, too. And it's -- I'm not sure you solve any special problems by turning to an independent group, whether it's the National Academy of Sciences or some other such body.

DR. WILSON: The English hearing on windscale was handled in a interesting way. They had -- it was before a judge who -- and they had two technical assessors who were highly competent people, one, Sir Edward Pochin and I can't remember the name of the other, a distinguished chemist, and both of which were to advise him on highly technical -- on the technical questions and the final -- and would write part of the final decision. And I think the net result of that report was actually technically -- it had a lot of technical validity, soundness to it. And there's no question that he, in fact, asked a lot of -- a lot of technical questions were asked of British Nuclear Field Services which the intervenors didn't get around to asking. They weren't trained enough. So there was, in fact, a lot of facts brought out in that hearing.

And so that wasn't a -- I mean, that wasn't a science court, but it was a group of assessors brought out with a judge and it was -- it was part of a general public hearing process.

DR. LOWRANCE: And any -- any regulatory group can turn to other agencies or consultants for attention to special issues. So the NRC, even if it were doing most of the assessment in-house, could still, as it always has done, turn to outsiders for substudies of parts of the puzzle.

DR. OKRENT: I think in that case, in fact, the -- wasn't the inspectorate, or whatever term to the hearing at windscale, the government set it up.

DR. WILSON: It was set up by -- well, that is standard in the United Kingdom -- set up by the government and report to the government and to the House of Commons.

DR. OKRENT: That's a kind, I would say that's a kind of science court.

DR. WILSON: That is a kind of science court.

DR. SIESS: I need some clarification. My understanding, it's your concept that the science court would, essentially, certify the safety critical, or whatever you report to it as, for each reactor (WORDS UNINTELLIGIBLE), so that would have to be for each reactor, as an application was tendered of operating reactors as the case may be.

DR. OKRENT: Are you -- I would -- now -- I'm --

DR. SIESS: The point is, the rest of the question, how long do you think it would take for each, each review? Are you thinking of something that would be done in a few days, or a few months?

DR. OKRENT: All right, let me comment. When Dr. Griesmeyer gave his presentation, he indicated two things that the science court might do. One was to, I'll use the word, set the level of risk which is to be applied to this plant. In other words, they would take the information, say, "We think this is our best judgment on the level of risk of this plant. And now go by the decision rule that you and the Congress and everybody have worked out," or whoever it is. Then he said they might also, if you will, certify the risk profile for the plant.

These are two separate things, and the science court might not necessarily do the second of these, because that would be a more detailed kind of a function, and not having to say, "Yes, we've gone through each one of these event trees," and so forth and so on.

DR. SIESS: So the first you would have the staff, presumably, argue one risk level, the applicant argue another risk level, the intervenors perhaps still another, and the court would come out with its decision as to what, in its opinion, it thought was correct.

DR. OKRENT: Yeah.

DR. SIESS: And that would be applied against the decision base.

DR. OKRENT: And I must say, one reason why I favor this not being within the NRC is, I suppose you'd say, to improve the public confidence in the process. Indeed, I think that's a not unimportant factor, and I think that it would work to some extent.

Dr. Bernero.

DR. BERNERO: I wonder if you would comment, the typical atomic safety and licensing board seems to fit a crude description of a science court: it's a lawyer and two scientists of one sort or another -- perhaps not of the stature you think of. I wonder why wouldn't that fit your description.

DR. OKRENT: Well, I agree with the last comment. I don't think it has the necessary stature. And it would have to be a hearing of a completely different sort than the way it has been drawn. Right now, as you well know, from -- what? -- let's say third parties like Harold Green have had to say about the way the hearings are run, there's not a lot of satisfaction.

But it -- I don't -- it might be, as we just heard, in the U.K., they set up a group not unlike that, with some legal and some technical competence, and very highly qualified, for what needed to be done.

All right, Dr. Whipple.

DR. WHIPPLE: I think windscale was a public inquiry had to be put. And I think what David's is referring to is some kind of risk council.

JO-15

DR. OKRENT: I'm sorry, risk council?

DR. WHIPPLE: Yes. I think what (WORDS UNINTELLIGIBLE) something like a risk council.

DR. OKRENT: I'm not sure how you're using the term.

DR. WHIPPLE: Well, I think you want to distinguish it from a public inquiry, don't you?

DR. SIESS: It would be a public inquiry but of limited scope. The scope would be limited to the risk assessment.

DR. OKRENT: That's right. So I think (WORDS UNINTELLIGIBLE) --

DR. WHIPPLE: I'd call that a risk council, which, you know, I mean, it's divorced from the normal type of public inquiry process, which can be, you know, public inquiry processes you set up (WORDS UNINTELLIGIBLE) by the calendar. And I don't think the same thing can really be said about the (WORDS UNINTELLIGIBLE). I think you're looking for some sort of body which can be called a risk council or whatever. A body of experts.

DR. OKRENT: Dr. Kastenberg.

DR. KASTENBERG: Just a question. In the example that was cited before, where the staff would come in and tell what it thought the acceptable risk was, and the intervenor might come in, and the applicant --

DR. OKRENT: Excuse me, not what the acceptable risk was -- what that level is.

DR. KASTENBERG: Right, the level of risk.

Would the role of the Commissioners change in this kind of a structure? What would their role be?

DR. OKRENT: Well, if it proceeded the way we have proposed, the Commissioners would be very active in trying to decide on what decision rules the Commission was going to -- like, they would be ultimately the ones who decided on the decision rules, unless the Congress overrode them or the President appointed new Commissioners to change the decision rules or something, but as an independent Commission they would act for society, in effect, in adopting -- only -- I think it's up -- they're the ones who have the authority now to adopt such rules.

SPEAKER: Dr. Okrent, if this body who were to pass on the risk is outside the NRC, from whence would they derive their statutory authority or power?

DR. OKRENT It would take an act of Congress.

SPEAKER: So this would be, essentially, a substitute for Congress. Congress would delegate the authority to this group.

DR. OKRENT: This -- for a specific function, namely, to arrive at a judgment on the level of risk. And again, the Congress would do this in the framework of certain decision rules that were going to be used (WORDS UNINTELLIGIBLE).

DR. WILSON: The problem with that -- that, in

England, it's worked because the system of public inquiry set up by the, requested by the House of Commons and set up, and set up in the way the windscale inquiry was done, is done for all sorts of things -- national security, under Lord Demming (?), and it's this sort of one thing or another, and there's a large history for that.

Whereas you're talking about setting up something rather similar in terms of an inquiry for a particular purpose, and there is no -- there's not precedent for this in our -- in the system of government here. And to what -- to the extent that these things have been gone through, for OTA and such, they haven't been particularly -- they don't -- they haven't had that stature and they haven't been that good. And so I am -- I worry about it being the first of its kind, so to speak. It will probably -- it may or may not be a great success, and it could be -- it could be just in deep trouble. And that's what --

MR. HANRAHAN: David Hanrahan (?), Nuclear Regulatory Commission.

I think as I understand what you have said as it has evolved in the hearing, it strikes me your last description of it described the Nuclear Regulatory Commission: established by law by Congress, appointed by the President with congressional advise and consent, to make those judgments. And somehow now to, because we're somehow unhappy with the present operation, to appoint a further board of scientists, who would now, somehow, 0-18

by some magic (WORDS UNINTELLIGIBLE) which they haven't in the field of nuclear power yet, smacks of hubris to me.

DR. SIESS: Well, the present Commission does not have the expertise in risk analysis. The object here would be to have --

MR. HANRAHAN: But the present -- to make the present and all previous Commissions and all future Commissions are assigned a role under the law to make these determinations of what's safe, that a license can be issued and that they -- that the nuclear facility can be operated without undue risk to the public health and safety in the common defense and security, and they make that judgment.

DR. SIESS: But this is only one step toward that judgment, one-half of one step toward that.

MR. HANRAHAN: And I go back to the process that Dr. Okrent outlined, where he felt uncomfortable with the staff because it may have come with some biases or positions. I would suggest to you that any scientist will come with the baggage of their experience in previous positions.

Then he suggested the Commission will need perhaps some advice and guidance to overwhelm -- overcome this staff bias in guidance. I would remind you that the Commission has a statutory organization which is supposed to provide them with just that sort of advice.

DR. LAVE: I think the other problem is that you

JO-19

suggest the Commissioners don't have the technical expertise. I guess that it would be fun to try and get some agreement in the technical community as to who had that ability.

(Laughter)

I bet you the set, if you did a Venn diagram the set, would be empty.

DR. SAUNDERS: A bunch of disconnected nodes.

DR. OKRENT: Well, I can see we don't have unanimity --

(Laughter)

-- on so simple a concept. Dr. Saunders.

DR. SAUNDERS: Well, may I be excused from going to the podium? I -- my comments will be so short as to not make it worth my while.

I want to address myself to the question why is it that a significant element of our population is so opposed to nuclear power. Since I have very little congress with that portion of the population which is below average, or that proportion which has mental instability, I shall confine my remarks as representative of the remaining set.

The first reason that I find why they object to nuclear power is because of the lack of discontinuance. If you -- and because of the implications -- if you use nuclear power you are locked into caring and nurturing the waste for a quarter of a million years. This presupposes a social stability, as well as a geological one, which is beyond our comprehension.

The second reason is, essentially, the unverifiability of ten to the minus sixth. May I say I have some difficulty with this myself. I know there are elements of science where predictions of the order of ten to the minus six have been made and have been, subsequently been verified. However, when it comes to, say, the Rasmussen study, you find ultimately that inside the calculations are measures of personal opinion and basion analysis which are not verifiable in the usual sense that scientific statements are verifiable.

I believe that a probabilistic analysis should essentially be a generalization of what a deterministic analysis would be if you had infinite resources and infinite time. It lumps together in the probabilistic manner distributions of quantities, the physical behavior of metal, or weather behavior, which anybody can verify on the average, but there is never an interjection of personal opinion. To fold together the opinion of a hundred experts and to say that's closer to the truth than you started with, I think is (WORDS UNINTELLIGIBLE) certainly unverifiable.

Thirdly, amongst many elements of our population, it is the opinion that some of us are not only unkempt, absent-minded, but -- and also are simply prostituted by the master from whom they receive a paycheck. And I don't know what we can do about that; I must confess somehow I am tempted at times to serve in such capacity.

Now, lastly, to comment on the fact that what an accident signals is far more important than the magnitude I think is a very pertinent comment. To point out simply that the DC-10 accident in Chicago had a lot higher consequences to the McDonnell Douglas Corporation than did an accident of almost the same magnitude which was in times past due to somebody carrying three sticks of dynamite on board and blowing himself up. And in the same sense, I think, Browns Ferry did not have near the consequences for the nuclear industry as did Three Mile Island. We certainly did learn a lesson there: we just can't agree on what the lesson is. And today Dr. Kerr tells me that it is killed people don't scare. I think that's -- that's saying something very close to the truth.

Now, today we've heard -- now, about the presentation today, if I may have a minute -- today we've heard two approaches. One says lump everything into the probability of core melt and do an analysis, I think, which -- on which everyone can agree, and don't worry about the details. Another approach which says break things apart into conditional probabilities, do the analysis as closely as you can in each one, and then synthesize them together in order to make a probabilistic analysis of the consequences to the community.

If I had a vote, I would vote for the latter. I think that's a much better way to do things.

Now, lastly, Edward Everett (?) said that he could

JO-22

talk for ten minutes -- it would take him a month to prepare; but to talk for two hours, he would be ready at the drop of a hat. And so I close now someplace in the middle.

DR. SHEWMON: Would you restate your last point about the kinds of probability calculations you felt more at home with.

DR. SAUNDERS: Insofar as there is no interjection of engineering judgment, so-called engineering judgment, into the disposition of parameters which are known to nature but not to the statistician, I much prefer an analysis which separates the consequences, conditional consequences, given a core melt, the probability of such an event happening.

DR. SHEWMON: To? To what?

DR. SAUNDERS: To public health and safety.

DR. SHEWMON: You said you'd rather go through the calculation given a core melt, what are the probabilities of something, than what other approach?

DR. SAUNDERS: To just lumping everything together into saying core damage.

DR. CATTON: Split the speculative part from the part you can do well, reasonably well.

I think that was right.

DR. SAUNDERS: Yes, that's right.

Is that sufficient? I'm for doing -- I think that analysis is so cheap relative to everything else that even given

JO-23

the errors that are sometimes made it's always wise to do that. You can always lump it together.

DR. SHEWMON: Mm hm. Okay.

DR. OKRENT: Any comments or questions?

DR. WILSON: I'm not quite clear, of that, what you -- you say you want to end up, so to speak, with the individual risk of health to an individual, to people, and the total effect on society, you want to end with that, and not stop anywhere at this core damage thing, is that what you said?

DR. SAUNDERS: Well, I mean to say that the ACRS should have the option of looking at the assumptions that are made on each of these separate points in the analysis. They should all be broken out so we can get those if we want to look at them.

DR. WILSON: Oh, good. Oh, yes. Absolutely.

DR. OKRENT: Dr. Apostolakis.

Speak up. You're far away.

DR. APOSTOLAKIS: I'm quite confused, I must admit.

(Laughter)

DR. SHEWMON: Congratulations. You stuck with it so long.

DR. APOSTOLAKIS: I don't understand what this means (WORDS UNINTELLIGIBLE) another is not. Does it mean that the (WORDS UNINTELLIGIBLE) are smaller? Because it would seem to me (WORDS UNINTELLIGIBLE). I don't see why the analysis (WORDS

JO-24

UNINTELLIGIBLE) is better than the analysis for core melt. There are (WORDS UNINTELLIGIBLE) too. I don't think all of us understand the same thing by (WORDS UNINTELLIGIBLE). And finally what you (WORDS UNINTELLIGIBLE). I mean, engineers always (WORDS UNINTELLIGIBLE) probabilistic risk analysis tries to do is to quantify (WORDS UNINTELLIGIBLE). And to penalize them because they are trying to do that I would think would be unfair.

So I don't understand the words (WORD UNINTELLIGIBLE) analysis and (WORDS UNINTELLIGIBLE).

DR. SAUNDERS: Well, something is verifiable, in principle, when another investigator, given the same set of data, or given the same phenomenon to be investigated, can take physical measurements as to the fracture resistance of the metal, as to the -- as to its various strength, as to the weather that's extant, and arrive at the same conclusions.

You can never do that whenever the opinion that's expressed is a subjective opinion of the person who's doing the analysis. That is unverifiable in principle, as far as I am concerned. I cannot do an introspection -- I can't do introspection and obtain the same answers that you do.

DR. SHEWMON: Let me give a different example of something that bothers me a fair amount. And that was the recent decision we went through on ATWS, where you had the equally competent statisticians, let's say, looking at the same data and coming to the numbers which were several orders of magnitude difference on the probability of all the control rods not going in when they were supposed to. And the best you can say is that honest men differ.

And I would say that the probability of the control rods not going in is, for me, unverifiable within a few orders of magnitude. And that bothers me a lot when we start making policy on those things.

And I guess my biggest hangup on the exercise we're going through here is, are we just going to end up with another set of things where one set of experts come in and say, "Gee, it's ten to the minus seven per year," and the staff comes down and says, "Gee, the best we can guarantee is, it's only ten to the minus three" and there you are? To me that's unverifiable. I don't know what it is to Dr. Saunders.

(Pause)

DR. KASTENBERG: Is it fair to ask Dr. Apostolakis if he is now unconfused?

(Laughter)

DR. SAUNDERS: Not a fair -- that's not a fair question, no.

DR. APOSTOLAKIS: You're making, then, (WORDS UNINTELLIGIBLE). Most of them are (WORDS UNINTELLIGIBLE) whether its' ten to the minus six (WORDS UNINTELLIGIBLE).

DR. SHEWMON: Yeah. I likened that one time to

JO-26

proving that there's not any pink elephants in my front yard because if I look out with a certain frequency I haven't seen any and therefore that's the probability. And that is the argument the staff came in with. Now, they took -- it hadn't happened yet, so we can only guarantee it's so low. And other people came in other ways.

But, you know, ten to the minus three wasn't going up and that's why we're prescribing another set of fictions on ATWS. One man's perception.

DR. OKRENT: Let's see. Dr. Catton, did you have some comments?

DR. CATTON: I'd like to add a couple of comments.

First, I agree with Dr. Saunders. And I'll comment more on that.

(Pause for microphone adjustment)

Okay, I'd like to make a comment about professional judgment.

I think professional judgment depends on whether you're buying or selling. And that tends to change the answers significantly.

With respect to a science court, I think the idea of a science court is good, but I really don't know whether or not it'll work, or even how one might implement it.

One thing I didn't understand about the presentations was this business about latent deaths and early deaths. To me

JO-27

it seems backwards.

I think the costs associated with the stress and neurosis of a latent death would be far greater than the costs of an early death.

And just to give you an idea about neurosis, when there's no way to agree on things, the impact of the negative group, or intervenors, is very costly, and for a small 100-kilowatt reactor over a two- or three-month period it costs $30,000. And I don't know if you can ratio that number up to a big reactor, but I imagine you could get pretty close.

As I see it, there are answers to two separate questions that are needed. I think first: what is the frequency and consequences of an event? And second: what is acceptable? I'm not sure the second will ever be answered.

When one factors in the 15 to 20 percent neurotic element that was cited by Zebroski, to me it's really more a religious than a scientific question, and I think it's better left to Congress.

As to the first part, when one uses probabilistic approaches, one has to incorporate physical phenomena into the analysis even when the phenomena is not understood. I don't belong really in the same group as this one, being an engineer. To do this gives me a very uncomfortable feeling. And I'll give you an example.

The steam spike calculated by the March Code to result

JO-28

from that TLN B-prime accident is based on assumed interactions between molten fuel and water. And, you know, you guess at a radius of a particle, you guess at a heat transfer coefficient, and you get an answer. It turns out the steam spike governs design of mitigating devices, which are then to be assessed in a probabilistic way. I find that rather unsatisfactory.

END TAPE 12

ape 13

In my view, one has to limit probabilistic methods to applications where the phenomena are understood or be willing to accept a great deal of conservative and also a great deal of negative reaction on the part of the lay public. I don't really consider myself lay public, yet I hold this negative feeling.

On the other hand, there are examples where these methods have been used to direct research towards phenomena that need understanding in order to improve or decrease the uncertainties. And this was demonstrated very nicely by Ray DeSalvo.

Just to reiterate, I feel very uncomfortable when I'm faced with markoff chains and comments about basion (?) methods being used to assure me I am safe. Insurance is more comforting.

I think a great deal more effort needs to be spent on increasing one's confidence if probabilistic methods are to be acceptable to the lay technical public. And further, I think any such efforts would be wasted if they're directed towards the 20 percent element.

By the way, I don't really mean to call them neurotic. The group that I'm faced with is not neurotic. I really don't understand what their problem is. They seem to be reasonably intelligent people except when the question is nuclear energy.

DR. OKRENT: Before we break up, in the first place, it's not adjournment time yet, half a minute, and if I'm careless I'll finish early.

I think it's --

DR. KERR: We won't tell anybody.

DR. OKRENT: -- important to see whether the subcommittee members have any recommendations for either specific things that they think it would be useful to do or specific kinds of things they would like to see if the consultants would do or ideas on how we should proceed in the next couple of months.

DR. SHEWMON: What do you see as the output of this? A letter? Presumably.

DR. OKRENT: Well --

DR. SHEWMON: Someday? Or what?

DR. OKRENT: We were asked, by either one or more than one Commissioner, I can't remember, to try to provide suggestions if we could. In June, I think it was, the esteemed member on my right suggested that we try to get something of a preliminary nature that could be a starting point, we'd put it as the first of some iterations back and forth, that we might be able to transmit in that context as a way of reaching, you know, and end of a first stage and to go on.

So there might be a kind of a report, let's say, modified, let's say, from what we heard Dr. Griesmeyer discuss, perhaps with examp'es, or whatever, which might be proposed for the committee to send along for thought. Something like this. That's one possibility.

If that isn't a viable possibility, we should ask ourselves what should be the alternate one, do we have to provide multiple proposals in the first go-round or so forth.

Dr. Mark.

DR. MARK: Well, these are very trifling remarks, I'm afraid.

I think it has enough bad connotaions that it would be worthwhile accepting essentially higher discretion to avoid the use of the term "science court." It's not -- it was fine between you and Harold Greene. I think it's not certain that it's fine here today.

It's certainly a really open area. It's a very unfortunate term and could be, I think, if I heard you, and maybe I didn't understand enough, that your function could be covered by the use of such terms as a panel, or whatever, of assessors who are going to assess the evidence and say what they think the probabilities are.

I would think that for a short-range communication with

JO-31

the Commission, I wouldn't necessarily, in fact, one may even avoid suggesting that this would be a permanent new board, but it would be an ad hoc board for the purpose of trying to give comments on today's, either today's situation or the situation which is within reach today, if one chooses the best technology that is in use, or something of that kind.

And although the Commission is charged with the responsibility, I certainly agree with the feeling that it's probably better if they should be urged to establish such a board. I don't want to use the ASLAB term, and yet it would have that relationship to them as something they can do, and if they put the right people on it, they have the legal authority, and they can then decide if they like their comments and these are the numbers we will advance to the next level where acceptability has to be assessed. This first step wouldn't have acceptability as its charter at all, but rather assessment of the situation, which can be used for discussion after the Commission has certified that they think these are the things in a wider audience, such as the public or the Congress.

Really I went further than I think I know enough to go. But the science court should probably not get frozen into the discussion.

DR. OKRENT: Other comments? Chet?

DR. SIESS: It seems to me that the work that's been done by the fellows and similar work that's been proposed by

JO-32

industry representatives goes a long way toward trying to express safety in some kind of quantitative terms. These come out to be probabilistic -- that's inevitable, since we're not talking about zero risk and there isn't zero risk.

How you relate this to acceptable levels of safety I don't know. The point has been made that we have different constituencies. If Congress is our constituency we go about it one way, or somebody goes about it one way. If the public or individual members of the public are the constituency it is clearly a question as to whether they understand the measures of risk that have been proposed or if they do understand them whether they agree with their perception of risk, because if they don't agree with their perception of risk they're not going to believe it anyway.

So I think we've got to propose some measures of risk -- individual, societal, et cetera -- that have some promise. Now, how to relate those to the problem of what's acceptable, I think the proposal of the hearing, it may be a two-year, another ECCS hearing, has some merit. It may not work. It will define the issues. It might tell you whether something else would work; I don't know.

But I think the first step has been made. I get a certain amount of comfort from the fact that the approach by the Atomic Industrial Forum and the approach as developed by the fellows, and some of the other things I've heard, are very close to each other.

But what's acceptable, and acceptable to whom, I don't know.

And as far as a science court, the only argument I could make for that is that without one the Commission is going to ask the ACRS to decide these things, and that's not only the wrong people but we don't want the job.

(Laughter)

And I do think that something like the ASLB, that is, a hearing with expertise in the risk assessment field, just like they appoint boards with antitrust expertise or at one time environmental, a board with risk assessment expertise would satisfy regulatory and legislative mandates. The Commission can ask for help in an area where they lack the expertise.

Now, whether that -- I don't think that would have the prestige of the science court that Dave envisages, because it's being appointed by the Commission and the Commission itself, just like the Commission staff, has no credibility.

DR. OKRENT: In which circles?

DR. SIESS: Well, in the circles -- if they had credibility we wouldn't be asking the questions we're asking now.

DR. OKRENT: Dr. Kerr.

DR. KERR: I think we have gone far enough and collected enough information that is usable, that we -- I mean, the subcommittee and the ACRS and the fellows and staff -- need to put together a discussion paper in which we describe some of the alternatives. And I don't think we describe just one thing but some of the alternatives that have been discussed, try to discuss their advantages and disadvantages, bring this to the full committee for some airing of the committee, and then see if we can't agree to send something to the Commission, not with the idea that this is a finished product, but with the idea that here are some things that we have explored, here are the problems we see with them, here are the advantages, the disadvantages, and get a reaction: does the Commission think that this is the direction in which they'd like to see some additional activity take place; does it make sense.

It seems to me this also has an advantage. I believe that although Congress as a whole might not undertake this, from what I've read, there is interest on the part of some congressmen and some congressional staffers in applying this sort of thing to other areas as well. And I think it would be well to get as much input into what we have done so far as we can before we take perhaps further (WORD UNINTELLIGIBLE) steps.

DR. OKRENT: Any other comments? Any --

DR. WILSON: One last thought. If you -- one way of getting wider information is to get this particular proposal of the -- of Dr. Griesmeyer in some journal where many more people can look at it and criticize it.

DR. OKRENT: We may try to do that.

JO-35

DR. KERR: May I make one last comment, that apropos of engineering judgment, one of the esteemed members of the ACRS once said that the difficulty with engineering judgment is that it requires both engineering and judgment.

(Laughter)

DR. OKRENT: Well, I think we have reached the end of the agenda. So, instead of asking for further comments, I'll thank you all for your active participation.

END TAPE 13

(Thereupon, at 6:25 p.m., the meeting was adjourned.)

CERTIFICATE OF REPORTER

This is to certify that the attached proceedings

before: NUCLEAR REGULATORY COMMISSION

In the matter of: ACRS - Subcommittee on Reliability & Probabilistic Assessment

Name of Proceeding:

Docket No.:

Place: Inglewood, California

Date: July 1, 1980

were held as herein appears, and that this is the ORIGINAL transcript thereof for the files of the Department.

Reporter

PUT IN TRANSCRIPT

SAFETY GOALS -- ATTRIBUTES AND A POSSIBLE FORMULATION

Edwin L. Zebroski
Nuclear Safety Analysis Center
Palo Alto, Ca.

VISUAL AIDS FOR PRESENTATION TO

SUBCOMMITTEE OF THE

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

July 1, 1980

Los Angeles, Ca.

SAFETY GOAL - I

- "RISK ENVELOPE" ESTIMATE BY WASH-1400 USEFUL TECHNIQUE BUT NOT OF ITSELF A WORKABLE TOOL FOR DESIGN, OPERATION, & REGULATION.

- ABSENT A PRACTICAL SAFETY GOAL. THERE IS TENDENCY OF ALL REGULATION TO STRIVE FOR NEAR-ZERO RISK FROM ANY DEFINED HAZARD.

- MEMBERS OF BIO-ETHICS COMMUNITY (DNA, SACCHARIN, EXTREME LIFE SUPPORT MEASURES, ABORTION CRITERIA, ETC.). NOTE THAT EXTREME REDUCTIONS IN A SPECIFIED RISK OFTEN INCREASE OTHER, LESS WELL-STUDIED RISKS.

- PRESENT LEGISLATION PROVIDES NO GUIDE FOR REGULATION TO AVOID EXCESSIVE INCREASED IN ALTERNATE RISKS OF HUMAN MISERY AND DEATH (E.G., DEPRIVATION, SOCIAL CHAOS, INFLATION, POSSIBLE CONTRIBUTING FACTOR FOR WARS) FROM DILATORY EXPLOITATION OF DOMESTIC ENERGY CAPABILITIES.

- ONE MEASURE OF PENALTY TO SOCIETY; NEARLY ONE TRILLION DOLLARS ADDED FUEL BILL IN THIS CENTURY DUE TO DELAYS, CANCELLATIONS, OR NON-COMMITTMENTS OF NUCLEAR UNITS.

ELZ:cic
6/12/80

SAFETY GOAL - II Attributes required

- Requires definitions of practical methods for design & operating decisions
- Must provide an objective basis for regulator-utility analysis and agreement on what is "safe enough"
- Must be clearly a "non-zero" risk goal and methodology
- Must be describable in terms which are understandable and acceptable by reasonably informed (and emotionally stable) laymen
- Must provide for full use of best-available data and decision processes
- Should make use of relative risks of main alternate sources of electricity, including social costs of shortages, interruptions, and sharp increases in costs

ELZ:cic
6/12/80

SAFETY GOAL - III   One possible formulation of Safety Goal

1. Reactor design and operation to insure that expected time to core-damaging accidents is not less than 30 years.*

2. Reactor and containment system design and operation to maintain assurance of not less than 99.9% probability of termination of the accident without radiation release leading to a total dose of 1 REM to any member of the public.

3. Use relative risk assessment methods (similar to conventional engineering trade-off studies) to establish need for, or adequacy of, design or operating improvements which establish that criteria (1) and (2) above are met, using existing operating experience as reference base.

4. Use statistically rigorous formulation with defined confidence levels and permissible error bounds, where needed, and include cummulative effects of actual total population of operating reactors.

5. Nuclear risk to be maintained at no more than one-third of the total risk of the two largest alternate sources.

6. Emergency plans to provide 99% assurance of total population dose less than 5000 man-rem even if containment failure were to occur after a core-damaging accident which was not safely terminated.

7. Improvements to reduce risks to one-tenth or less of main practical alternate sources to be sought, but to be implemented only if shown to be cost-effective and with no measurable effect on cost or timely availability of energy.

* For total population of civilian reactors in operation in the U.S.

ELZ:cic
6/12/80

ALTERNATE STRATEGIES FOR MEASURED IMPROVEMENTS IN REACTOR SAFETY AND RELIABILITY

- ASSUME THAT MANY ELEMENTS OF DESIGN, OPERATION, MANAGEMENT AND REGULATION ARE GENERALLY INADEQUATE; IMPROVEMENTS REQUIRE MAJOR CHANGES:
  - PROMULGATE CONTINUING CHANGES IN DESIGN
  - PROMULGATE CONTINUING CHANGES IN PROCEDURES, TRAINING, AND MANAGEMENT REQUIREMENTS
  - INCREASE THE EXHAUSTIVNESS AND PUNITIVENESS OF REGULATORY ACTIONS

- OR -

- NOTE THAT NEARLY ALL ELEMENTS OF DESIGN, OPERATION, AND REGULATION HAVE FUNCTIONED ADEQUATELY FOR NEARLY 500 REACTOR-YEARS OF OPERATION; IMPROVEMENTS REQUIRE MAINLY THE PREVENTION, OR DETECTION AND TIMELY REMEDY, OF THE OCCASIONAL LAPSES FROM GOOD PRACTICE:
  - INCREASE RIGOR OF TRAINING AND EDUCATION TO AVOID LAPSES IN OPERATION AND MAINTENANCE
  - MINIMIZE CHANGES IN DESIGN OR PROCEDURE UNLESS BENEFIT IS PROVEN BY DEFINITIVE ANALYSIS
  - INCREASE THOROUGHNESS OF RECORDING, ANALYSIS, AND FEEDBACK OF OPERATING EXPERIENCE
  - INCREASE USE OF POOLED RESOURCES TO ATTACK GENERIC ASPECTS OF PRACTICAL AND REGULATORY PROBLEMS TO CONSERVE UTILITY RESOURCES FOR PLANT-SPECIFIC ITEMS
  - CONDUCT EVALUATION, INSPECTION, AND REGULATION IN A PROBLEM-SOLVING AND TUTORIAL MODE, RATHER THAN PUNITIVE MODE, UNLESS THERE IS EVIDENCE OF WILLFUL NEGLIGENCE OR MALFEASANCE

ELZ/6/25/80

WORK IN PROGRESS RELATED TO PRACTICABILITY OF SAFETY GOALS

- DERIVATION AND TESTING OF THE SENSITIVITY OF A GIVEN EVENT TREE PROBABILITY TO:
  - THE NUMBER OF LINKED EVENTS IN THE TREE
  - THE TIME-VARIATION IN THE STATISTICAL DISTRIBUTION FUNCTION OF COMPONENT FAILURE PROBABILITIES
- DERIVATION OF THE EXPECTED EFFECT ON STATISTICAL UNCERTAINTIES OF INCREASED EXPOSURE TIME VERSUS INCREASED POTENTIAL BENEFITS OF OPERATING EXPERIENCE.
- ASSESSMENT OF THE CUMULATIVE EFFECT ON REDUCTION OF PROBABILITY OF UNANTICIPATED BRANCHES OF EVENT TREES, CONSIDERING THE EFFECTS OF:
  - INCREASED RIGOR OF RECORDING, ANALYSIS, AND FEEDBACK OF OPERATING EXPERIENCE
  - INCREASED DEFINITION AND TRAINING IN PROCEDURES AND SIMULATOR TESTING
  - PERIODICALLY IMPROVED REANALYSIS OF DESIGN BASIS USING PLANT-SPECIFIC EVENT TREE ANALYSIS
  - THE CUMULATIVE EFFECTS OF DISCIPLINED REGULATION, EVALUATION, AND INSPECTION.

EHS/6/25/80

ACCEPTABLE RISK

- o LITERATURE REVIEW
- o PHILOSOPHY
  - SLOVIC ET AL
  - BNL ON UNACCEPTABLE RISK
- o DATA & METHODOLOGY FOR TEST
- o REBASELINING WASH-1400

BERNERO

ADVANTAGES OF A PROBABILISTIC APPROACH IN LICENSING

- ASSESSING CONTRIBUTIONS TO SAFETY SYSTEM UNAVAILABILITY FROM VARIOUS CAUSES

- SYSTEMATIC EVALUATION OF A BROAD SPECTRUM OF POTENTIAL ACCIDENT SEQUENCES

- DETAILED ANALYSIS OF ACCIDENT CLASSES 3-9. EVALUATION OF REALISTIC, PHYSICALLY CONSISTENT POTENTIAL ACCIDENT SEQUENCES

- ASSESSMENT OF CONSEQUENCES TO THE POPULATION AROUND THE REACTOR SITE

- UNIFIED FRAMEWORK FOR SAFETY ANALYSES

GROUNDRULES FOR PROBABILISTIC CRITERIA

- SIMPLE RATIONALE BEHIND THE CRITERIA

  UNDERSTANDABLE AND ACCEPTABLE TO BOTH TECHNICAL AND LAYPERSONS

- CRITERIA SHOULD CONSTRAIN RISK OF NUCLEAR POWER TO BE AT LEAST AS LOW AS OTHER COMPETING POWER GENERATING TECHNOLOGIES

- CRITERIA SHOULD NOT INCREASE RISK FROM NUCLEAR POWER ABOVE ITS CURRENT LEVEL

- FULFILLMENT OF CRITERIA SHOULD BE EASILY DEMONSTRABLE

- CRITERIA SHOULD INCORPORATE SOCIETY'S AVERSION TO LARGE CONSEQUENCES

SOME DIFFICULTIES OF PROBABILISTIC SAFETY CRITERIA

- SHOULD RISK CRITERIA INCORPORATE SOCIETAL AVERSION TO HIGH CONSEQUENCE EVENTS?

- WHAT RELATIVE WEIGHT SHOULD CRITERIA IMPLY IN TREATING EARLY DEATHS vs. LATENT CANCER DEATHS?

- SHOULD RISK CRITERIA INCORPORATE SPECIAL SOCIETAL PERCEPTION OF NUCLEAR RISK?

- SHOWING COMPLIANCE TO PROBABILISTIC CRITERIA - AREAS OF CONCERN

  BEST ESTIMATE vs. EXPECTED VALUE

  DEFINING UNCERTAINTIES

  DEGREE OF CONFIDENCE ASSOCIATED WITH RARE EVENTS

  LARGE VARIATION IN ESTIMATES OF SOME EVENTS ASSESSED BY EXPERTS

  DIFFICULTY IN PREDICTING HUMAN FAILURE

CURRENT ANALYSES

BY INDUSTRY

- BIG ROCK POINT
- LIMERICK
- ZION/INDIAN POINT
- OCONEE

BY NRC

- CRYSTAL RIVER
- SEQUOYAH
- OCONEE
- CALVERT CLIFFS
- GRAND GULF
- ARKANSAS 1
- MILLSTONE 1
- BROWNS FERRY 2
- INDIAN POINT 2/3

ONE APPROACH FOR CORE DAMAGE PROBABILITY CRITERIA

P_{CD} = CORE DAMAGE PROBABILITY (PER RY)

- CALCULATE POINT VALUE FOR P_{CD}

- IF P_{CD} ABOVE UNACCEPTABILITY LEVEL

  THEN IMPROVEMENTS REQUIRED

- IF P_{CD} WITHIN WARNING RANGE THEN

  CASE BY CASE EVALUATION

- WARNING RANGE ADDED TO ACCOUNT

  FOR CALCULATIONAL UNCERTAINTIES

CRITERIA FOR P_{CD}: OPERATION AND DESIGN

- COVERAGE: SPECIFICALLY DEFINED INITIATING EVENTS, SYSTEM FAILURES, COMPONENT FAILURES TESTING CONTRIBUTIONS, HUMAN ERRORS
- UNACCEPTABLE LEVEL:

  $P_{CD} > 1 \times 10^{-3}$
- WARNING RANGE:

  $1 \times 10^{-4} \leq P_{CD} \leq 1 \times 10^{-3}$

CRITERIA FOR P_{CD}: DESIGN ONLY

- COVERAGE: SPECIFICALLY DEFINED INITIATING EVENTS, SYSTEM FAILURES, COMPONENT FAILURES
- UNACCEPTABLE LEVEL:

  $P_{CD} > 1 \times 10^{-5}$
- WARNING RANGE:

  $1 \times 10^{-6} \leq P_{CD} \leq 1 \times 10^{-5}$

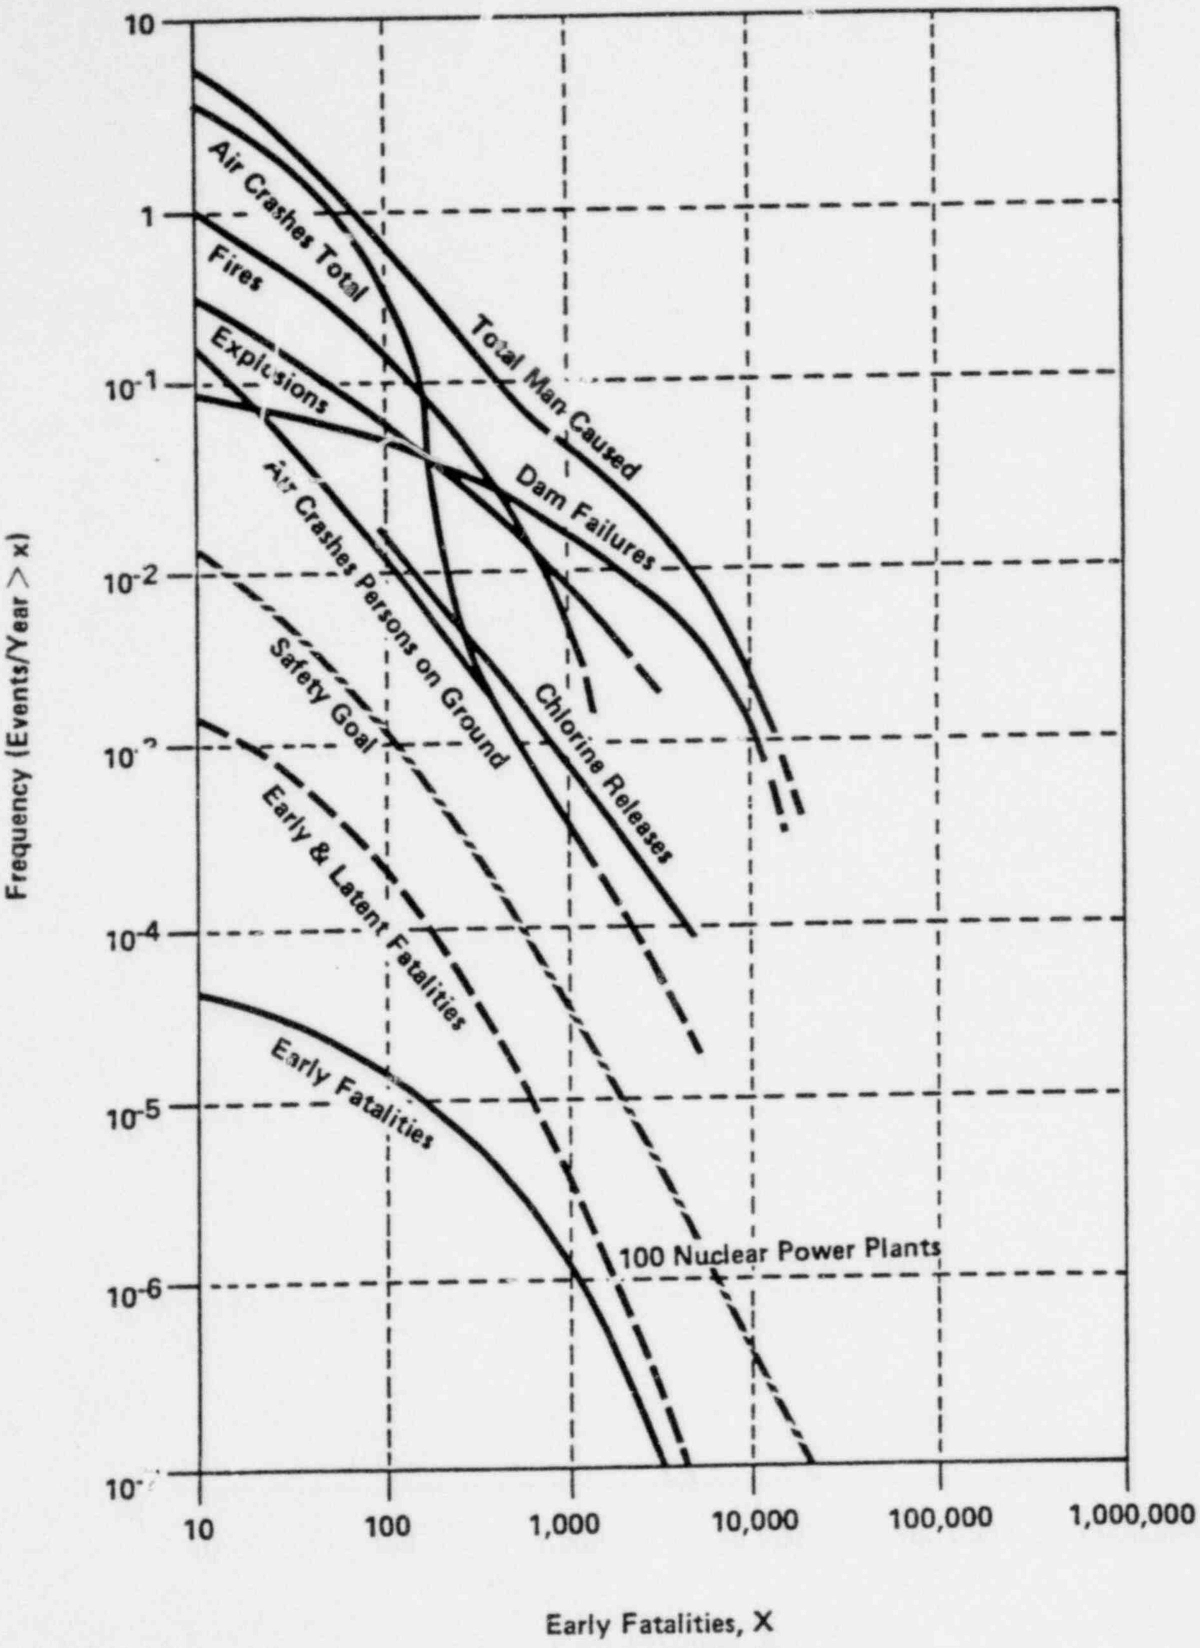

10
1
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
10^{-7}
Frequency (Events/Year > x)
Air Crashes Total
Fires
Explosions
Total Man Caused
Dam Failures
Air Crashes Persons on Ground
Chlorine Releases
Safety Goal
Early & Latent Fatalities
Early Fatalities
100 Nuclear Power Plants
10
100
1,000
10,000
100,000
1,000,000
Early Fatalities, X

FIGURE 4

TEMME

IEEE WORKING GROUP

NPEC - SC 5.4

RISK CRITERIA

PRESENTATION TO

ACRS RELIABILITY AND PROBABILISTIC ASSESSMENT

SUBCOMMITTEE

JULY 1, 1980

M. I. TEMME
CHAIRMAN, W.G. 5.4

IEEE NPEC SC-5.4

CONTENTS

- APPLICATION AND SCOPE OF RISK CRITERIA
- IMPLEMENTATION PROCEDURES
- PLANS AND PERSPECTIVE

IEEE NPEC SC-5.4

PROPOSED APPLICATION OF RISK CRITERIA STANDARD

- THE STANDARD IS WRITTEN TO BE USED BY DESIGNERS AND OPERATORS OF NUCLEAR POWER PLANTS FOR MAKING TECHNICAL DECISIONS RELATIVE TO THE DESIGN AND OPERATION OF THE PLANTS

- PERCEIVED VALUE OF THE STANDARD IS THAT IT WILL PROVIDE A CONSISTENT BASIS FOR MAKING INTELLIGENT ENGINEERING TRADE-OFF DECISIONS

M. I. TEMME
JULY 1, 1980

IEEE NPEC SC-5.ᵁ

SOURCES OF RISK COVERED

- LIGHT WATER REACTORS
  - OPERATING PLANTS
  - COMMITTED PLANTS
  - NEW PLANTS
- REACTOR PLANT
  - RISKS FROM OTHER PARTS OF THE NUCLEAR FUEL CYCLE NOT TO BE GOVERNED BY THE CRITERIA
- PUBLIC HEALTH AND WELFARE
  - OCCUPATIONAL RISKS NOT GOVERNED BY THE STANDARD
- RISK CONTRIBUTIONS OF HUMAN EFFECTS, COMMON CAUSE FAILURES AND OFF-SITE MITIGATION ARE TO BE INCLUDED
- RISKS FROM SABOTAGE AND WAR EXCLUDED

M. I. TEMME
JULY 1, 1980

IEEE NPEC SC-5

PARAMETERS OF TOP-LEVEL RISK CRITERIA

- IMPORTANT CONSIDERATIONS:
  - FLEXIBILITY IN SITING AND DESIGN
  - EASE OF APPLICATION
- CRITERION WILL BE STATED IN TERMS OF RADIOLOGICAL DOSE LIMITS RATHER THAN HEALTH EFFECTS LIMITS
- ADDITIONAL CRITERION RELATED TO OFF-SITE PROPERTY CONTAMINATION WILL BE INCLUDED

M. I. TEMME
JULY 1, 1980

IEEE NPEC SC-5

PARAMETERS CONSIDERED AND REJECTED

- HEALTH EFFECTS (ACUTE, LATENT FATALITIES, ILLNESS)
- RISK FEE
- HAZARD STATE LIMITS
- RELEASE TO CONTAINMENT
- RELEASE FROM CONTAINMENT
- RELEASE FROM SITE
- EXTENSIVE CORE DAMAGE

M. I. TEMME
JULY 1, 1980

IEEE NPEC SC-5.4

IMPLEMENTATION PROCEDURES

- THE MEANING OF NUMERICAL CRITERIA IS DETERMINED BY THE ANALYTICAL PROCEDURES FOR THEIR IMPLEMENTATION
- WG 5.4 WILL DEPEND ON OTHER STANDARDS WRITING GROUPS THAT SPECIFY PROCEDURAL DETAILS OF RISK ANALYSIS
- THE PROBABILISTIC CRITERIA WILL RELATE TO INTEGRATED RESULT OF ALL IMPORTANT ACCIDENT SEQUENCES, RATHER THAN INDIVIDUAL SEQUENCES
- THE CRITERIA WILL APPLY TO "PASS" - "FAIL" DECISIONS

M. I. TEMME
JULY 1, 1980

IEEE NPEC SC-5.4

NEAR TERM PLANS

- PHILOSOPHY
  - WG 5.4 IS GIVING FIRST PRIORITY TO A CONSISTENT DECISION-MAKING PROCESS, RATHER THAN CONFORMANCE TO A RIGID SCHEDULE
- INITIAL TARGET - DRAFT STANDARD ~ MID 1981 STILL APPEARS FEASIBLE
- NEAR TERM WORK
  - DETAILED SPECIFICATION OF CRITERIA ON DOSE
  - REVIEW OF THE DATA AND BASES FOR QUANTIFICATION - E.G.,
    - RELATION OF DOSE TO HEALTH EFFECTS
    - RISKS OF PRESENTLY OPERATING PLANTS
    - RISKS OF ALTERNATE FORMS OF ELECTRICAL ENERGY GENERATION
    - VALUE-IMPACT OF RISK REDUCTION

M. I. TEMME
JULY 1, 1980

IEEE NPEC SC-5.4

PERSPECTIVE

- STANDARDS ORGANIZATION

↕

- INDUSTRIAL PRACTICE

↕

- SOCIETAL EVALUATION OF ACCEPTABILITY

M. I. Temme
July 1, 1980

JOKSIMOVIC

REGULATORY OBJECTIVES

- HIGHEST ACHIEVABLE STANDARD
- LOWEST ACHIEVABLE RELEASE OF RADIOACTIVITY OR CARCINOGENS
- UNTENABLE, NEITHER REASONABLE NOR ADEQUATE AS PLANT DESIGN CRITERIA

REGULATORY PRACTICE

- RELEASES AS LOW AS REASONABLY ACHIEVABLE
- GENERAL DESIGN CRITERIA PLUS REGULATORY GUIDES
- BOTH EXPERIENCED SERIOUS SHORTCOMINGS - AMBIGUOUS, TEND TO MAKE DESIGNERS CONCENTRATE ON MEETING REQUIREMENTS RATHER THAN STIMULATE INNOVATION AND LEGITIMATE TRADEOFFS, SAFETY STANDARDS ACHIEVED DEPENDENT ON INDIVIDUALS ASSIGNED TO PARTICULAR PLANTS

QUANTITATIVE SAFETY GOALS

- WILL PROVIDE CONSISTENT SAFETY CRITERIA AND THUS FIRM FRAMEWORK FOR DESIGNERS
- AVOID RATCHETING OF PLANT REQUIREMENTS
- LEAVE DESIGNERS WITH FREEDOM TO CREATE EFFECTIVE DESIGN SOLUTIONS TO REAL PROBLEMS THEY FACE
- WILL PROTECT PUBLIC AND INVESTMENT MORE EFFECTIVELY
- WILL ENABLE DESIGNERS TO GET ON WITH THE JOB OF DESIGNING FUTURE REACTORS
- WILL CONTRIBUTE TOWARDS PUBLIC UNDERSTANDING OF NUCLEAR SAFETY

GA BACKGROUND

STATEMENT BEFORE NRC RARG GROUP	-	DEC. 1977
STATEMENT BEFORE UDALL COMMITTEE	-	FEB. 1979
AIF WORKSHOP PAPER	-	MAY 1979
AGNEW TO KEMENY LETTER	-	SEP. 1979
KNOXVILLE CONFERENCE PAPER	-	APR. 1980
LTR ON SELECTION OF DESIGN BASIS ACCIDENTS	-	JUN. 1980

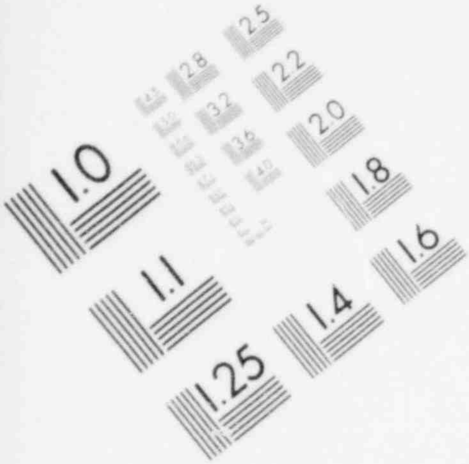

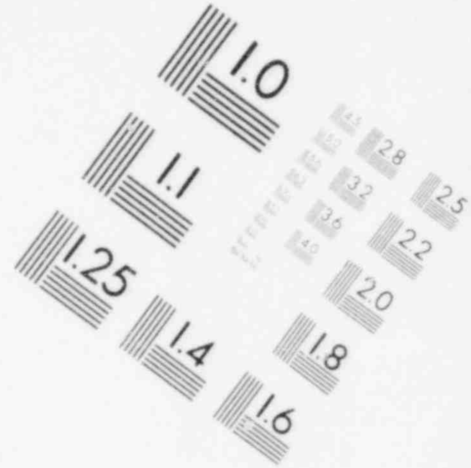

IMAGE EVALUATION
TEST TARGET (MT-3)

6″

MICROCOPY RESOLUTION TEST CHART

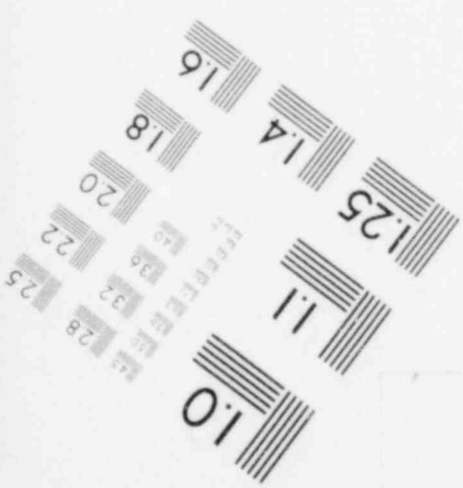

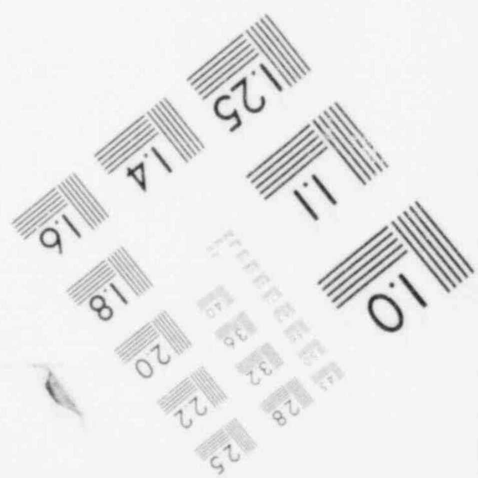

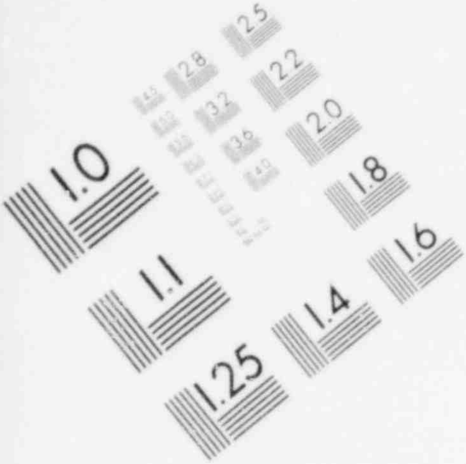

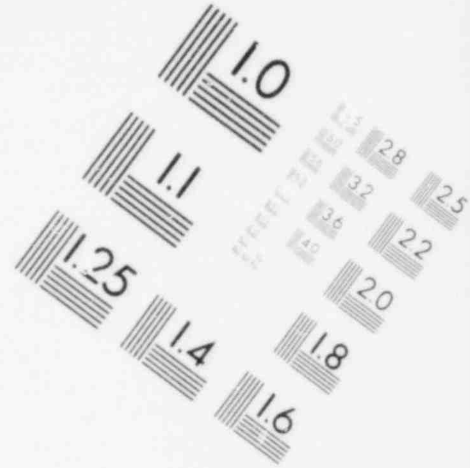

IMAGE EVALUATION
TEST TARGET (MT-3)

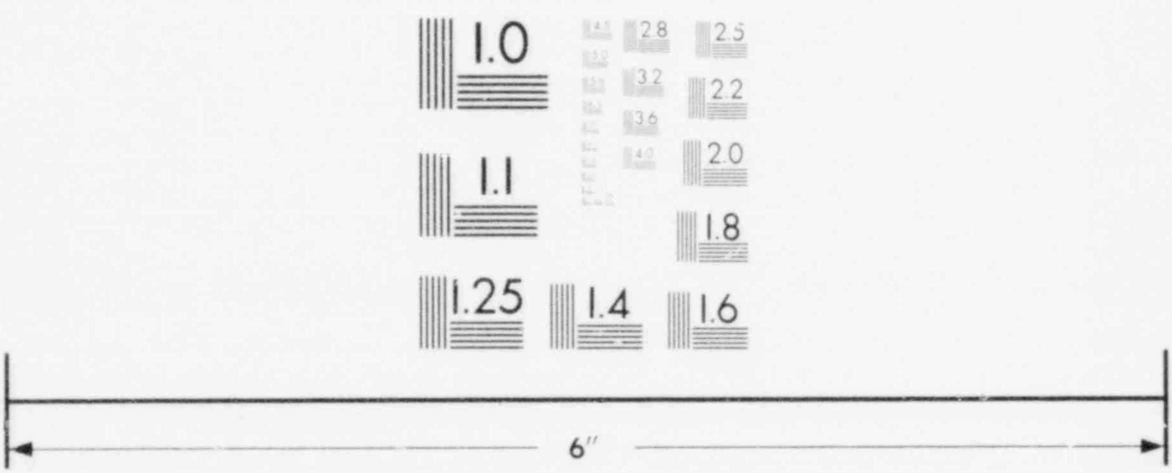

MICROCOPY RESOLUTION TEST CHART

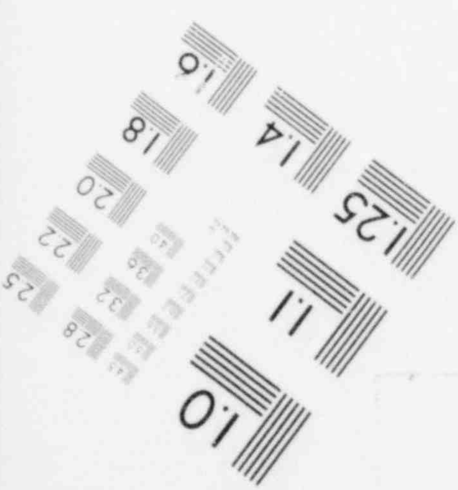

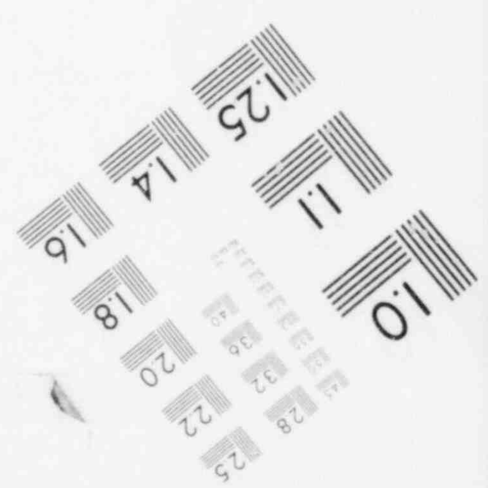

KEY ELEMENTS (CONCEPTS) OF GA APPROACH TO QUANTITATIVE SAFETY GOALS

- PROFESSIONAL JUDGMENT
- RISK BUDGET VIA COMPARATIVE RISK STUDIES
- EMPHASIS THROUGH KNOWLEDGE OF PRA BUT RETAINING SOME ASPECTS OF DETERMINISTIC REGULATORY FRAMEWORK
- RETENTION OF DBA BUT PRA FOCUSED
- ADDITION OF SAFETY MARGIN AND RESEARCH REGIONS
- NO IDENTIFIABLE PUBLIC INJURY
- USE OF LIMIT LINES
- RANGE FROM EMPHASIZED TO BALANCED RISK POLICY
- ACCEPTABLE INDIVIDUAL, SOCIETAL, AND PUBLIC PROPERTY RISKS

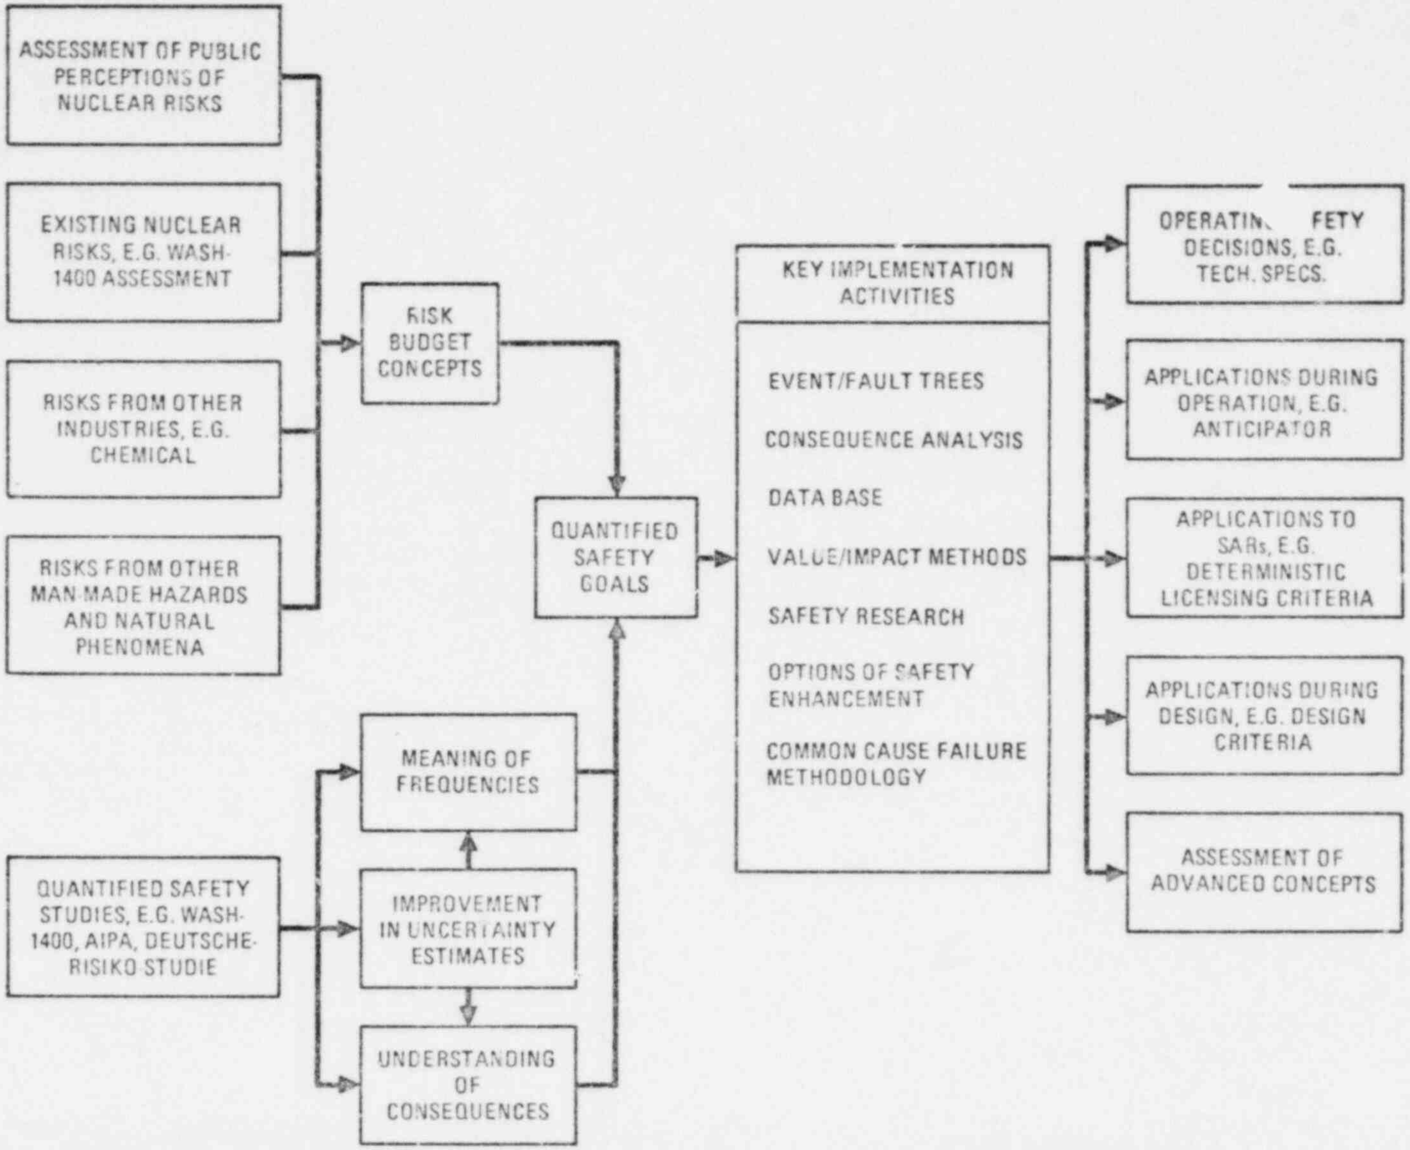

Figure 1. Major steps in development and implementation of quantified safety goals

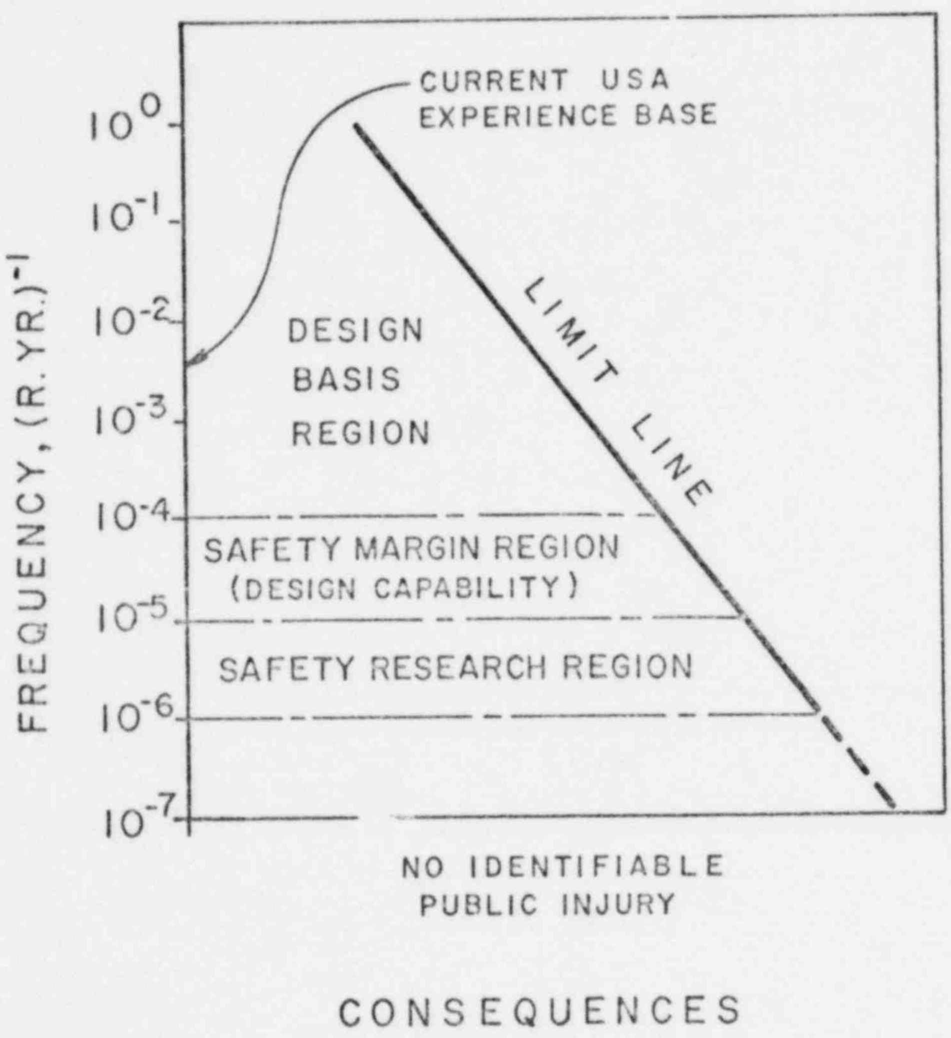

FIGURE 2

MAP OF QUANTITATIVE SAFETY REGIONS

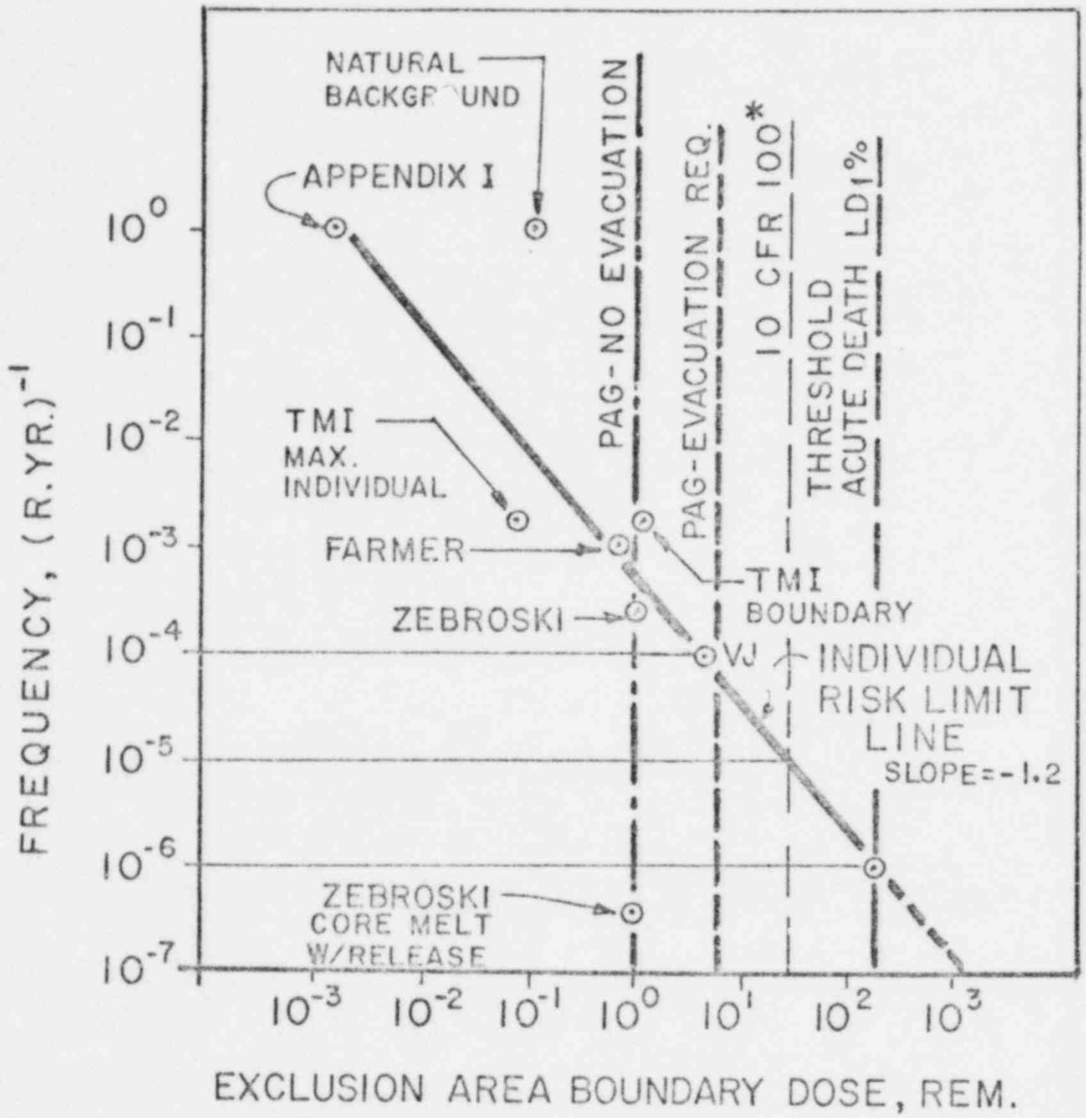

* ASSUMING REALISTIC EVALUATION OF DOSES

FIGURE 3

QUANTITATIVE SAFETY GOAL - INDIVIDUAL RISK TO MAXIMUM EXPOSED INDIVIDUAL MEMBER OF THE PUBLIC

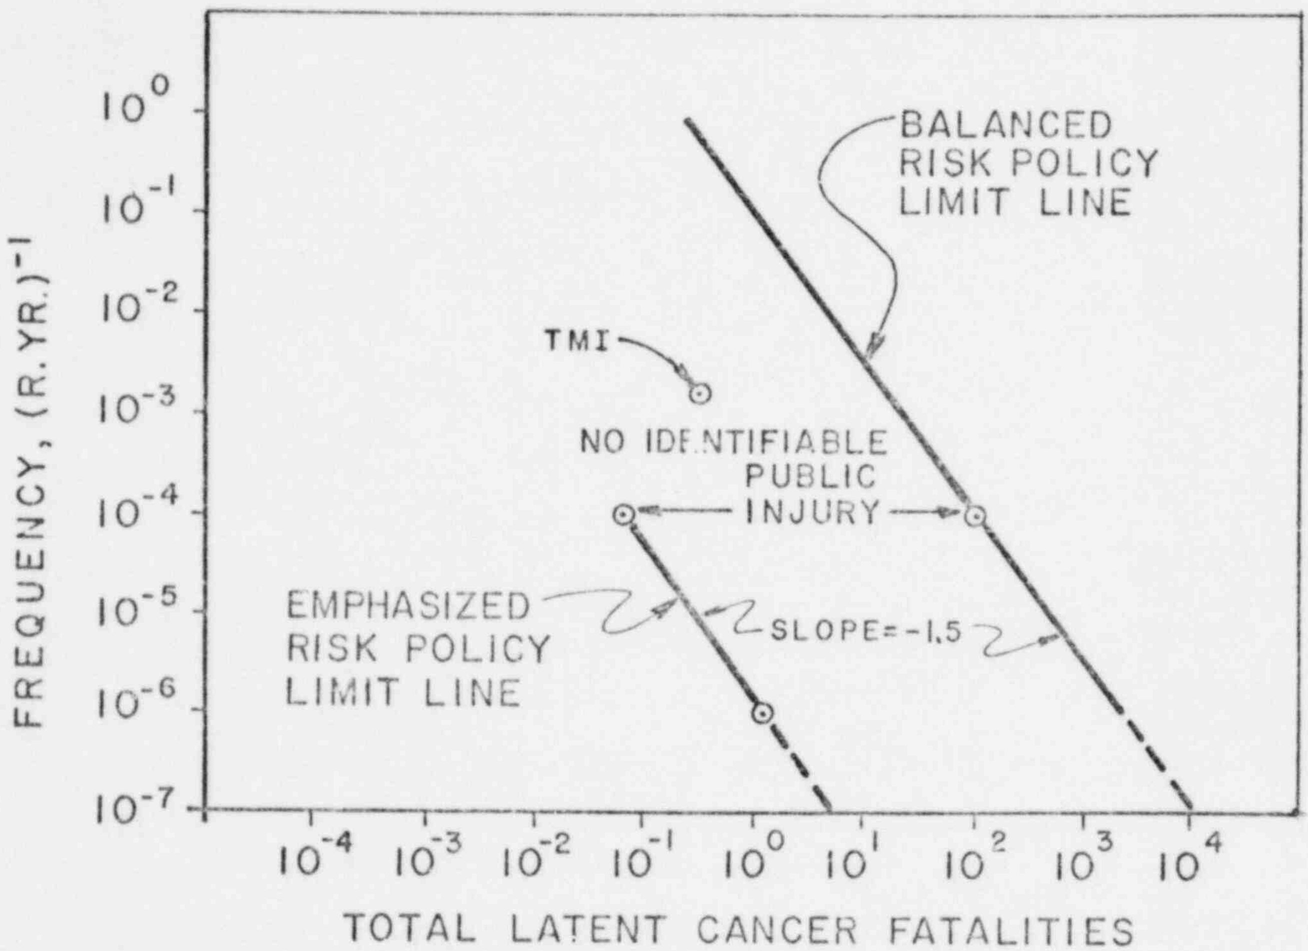

FIGURE 4

QUANTITATIVE SAFETY GOAL - SOCIETAL RISK

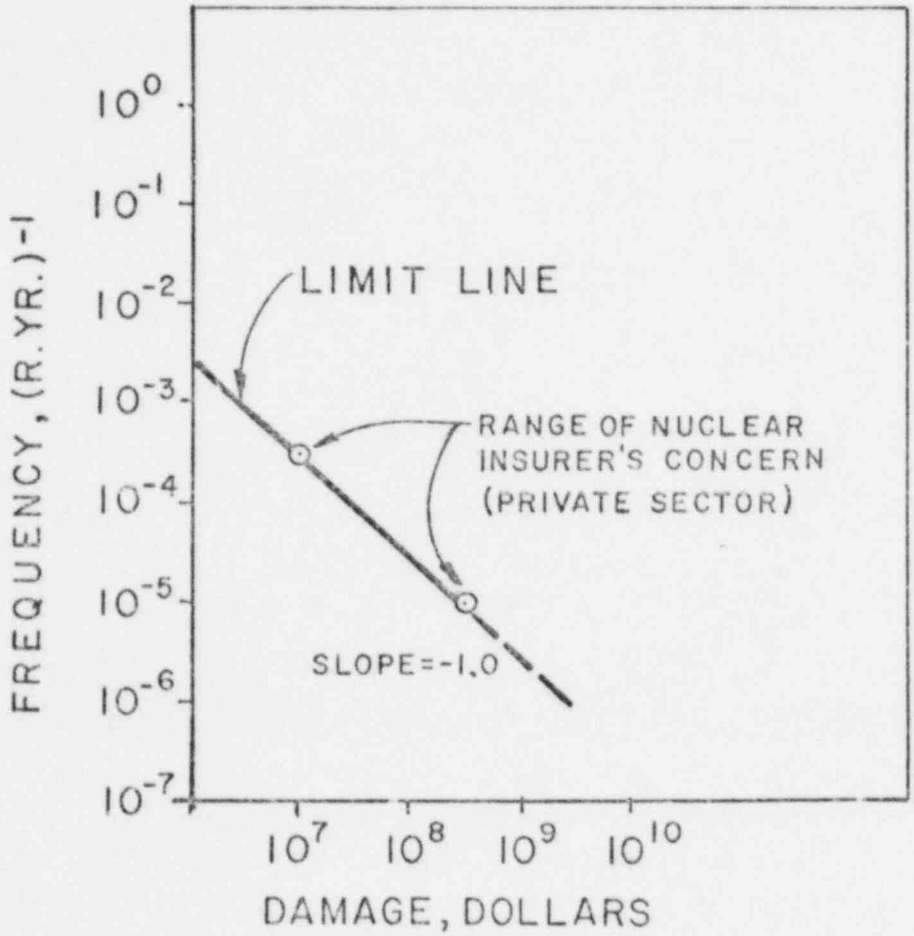

FIGURE 5
QUANTITATIVE SAFETY GOAL-
PUBLIC PROPERTY RISK

GRIESMEYER

Decision Levels

Need for Electrical Generating Capacity

Technological Alternatives

Site Selection

Risk Acceptance

Considerations

Economic Impacts

Socio-political Impacts

Environmental Impacts

Public Health and Safety Impacts

Basic Premises

Public health and safety should be treated in a consistent manner.

No individual should be unduly burdened by risk.

Satisfaction of one groups wishes should not have large detrimental effects on the desires of society as a whole.

The aggregate impacts of a particular facility should not unduly threaten the resilience of society.

Features

Safety Profile

Science Court

Review of safety profiles.

Legally binding closure of issues concerning estimated levels of risk.

Decision Rules

Necessary Conditions – not sufficient.

Upper risk limits, discretionary range and goal.

Decision Rule Categories

Hazard States

Individual Risks

Societal Impacts

Public Health and Safety

Property and Resource Damage

Societal Impact Reduction

Hazard States

State	Probability Limit		Frequency Upper Limit: Expected	Frequency Upper Limit: 90% confidence
Significant Core Damage ($\geq 10\%$ Noble gas inventory leaking into primary coolant)	A	$P_c \leq \frac{1}{10}$ in reactor lifetime	$f_c \leq 3 \times 10^{-3}/r\text{-}y$	$f_c \leq 1 \times 10^{-2}/r\text{-}y$
	B	$P_c \leq \frac{1}{30}$ in reactor lifetime	$f_c \leq 1 \times 10^{-3}/r\text{-}y$	$f_c \leq 5 \times 10^{-3}/r\text{-}y$
Large Scale Core Melt ($\geq 10\%$)		$P_m \leq \frac{1}{2}$ in U.S reactor program	$f_m \leq 10^{-4}/r\text{-}y$	$f_m \leq 5 \times 10^{-4}/r\text{-}y$
Large Scale Uncontrolled Release ($\geq 10\%$ Iodine inventory)		$P_R \leq \frac{1}{100}$, given a large scale melt	$f_R \leq .01/\text{melt}$	$f_R \leq .05/\text{melt}$

Individual Limits – Proposal A

Probability Limits	Upper Frequency Limit		Conditional Probability Limit given core melt	
	Expected	90% Confidence	Expected	9 % confidence
Probability of death due to a reactor $P_d \leq .001$	$f_d \leq 2 \times 10^{-5}$ per site-year	$f_d \leq 1 \times 10^{-4}$ per site-year	$P_{d/m} \leq 2 \times 10^{-3}$	$P_{d/m} \leq 1 \times 10^{-2}$
Probability of early death due to a reactor $P_{ed} \leq .0002$	$f_{ed} \leq 5 \times 10^{-6}$ per site-year	$f_{ed} \leq 2 \times 10^{-5}$ per site-year	$P_{ed/m} \leq 5 \times 10^{-4}$	$P_{ed/m} \leq 2 \times 10^{-3}$

Individual Risk Limits – Proposal B

Definitions	Measure of Insult	Upper Limit	
		Expected	90% confidence
R = Individual chance of death per site-year R_m = Individual chance of death given a meltdown L_f = Fraction of age adjusted remaining life expectancy before the onset of fatal symptoms.	$D_I = R(1-L_f)$	$D_I \leq 10^{-5}$ per site-year	$D_I \leq 5 \times 10^{-5}$ per site-year
	$D_{I_m} = R_m(1-L_f)$	$D_{I_m} \leq 10^{-3}$/melt	$D_{I_m} \leq 5 \times 10^{-3}$/melt

Social I act Limits

Social Cost Functions

$$E = \sum G(f, m)$$

Social Cost Limits

$$E \leq L$$

example

$$G(f, m) = f m^{\alpha}$$

Public Health and Safety

Deaths due to reactor

$$E_d = \Sigma G_d(f, m_d) \qquad \left(= \Sigma f m_d^{\alpha}\right)$$

$$E_d \leq 10/\text{reacto-year}$$

Deaths given a core melt

$$E_{d|m} \leq 1000/\text{melt}$$

Early Deaths due to reactor

$$E_{ed} = \Sigma G_{ed}(f, m_{ed})$$

$$E_{ed} \leq 2/\text{reactor-year}$$

Early Deaths given a core melt

$$E_{ed|m} \leq 200/\text{melt}$$

Property and Resources

Monetized Losses

$E_e \leq F \cdot$ (value of energy generated/year)

Special Resources (farm land, important aquifer, etc.)
Beyond some level of effect (say potential loss of 10% of local resource) special consideration is required in licensing

Societal Impact Reduction — ALARA

Marginal Cost Limits

$\$5 \times 10^6 / (\Delta E_{ed} \cdot \ell)$

$\$1 \times 10^6 / (\Delta E_d \cdot \ell)$

$2 / (\Delta E_e \, \ell)$

Practicality defined as

$$C \leq \ell \left[2 \Delta E_e + \$5 \times 10^6 \Delta E_{ed} + \$1 \times 10^6 \Delta E_d \right]$$

O'DONNELL

GENERAL CONSIDERATIONS ON USE OF PRA IN REGULATORY PROCESS

- PRA SHOULD SUPPORT DETERMINISTIC REQUIREMENTS
- USE PRA AS BASIS FOR GENERIC REQUIREMENTS, NOT AS LICENSING CONDITION FOR INDIVIDUAL PLANTS
- ADOPT PRA AS BASIS FOR JUSTIFYING CHANGES (ADDITIONS OR DEPARTURES) IN DETERMINISTIC REQUIREMENTS
- PRA SHOULD BE DONE AS REALISTICALLY AS POSSIBLE, WITH DEGREE OF UNCERTAINTY OR CONSERVATISM EXPLICITLY STATED
- ESTABLISH QUANTITATIVE SAFETY GOALS FOR PRA-BASED DECISION MAKING
- DEVELOP COMMON PRA METHODOLOGY

AREAS REQUIRING DEVELOPMENT OF COMMON PRA METHODOLOGY

- EVENT/FAULT TREE ANALYSIS LEVEL OF DETAIL
- COMPONENT FAILURE DATA BASE
- COMMON CAUSE FAILURE
- SYSTEM INTERACTION
- CONSEQUENCE MODELLING
- VALUE-IMPACT METHODOLOGY
- HUMAN FACTORS
- QUANTIFICATION OF CONSERVATISM AND UNCERTAINTY

MOST USEFUL PRA APPLICATIONS IN REGULATORY PROCESS

- VALUE-IMPACT ANALYSIS FOR NEW REQUIREMENTS
- PRIORITIZATION AND RESOLUTION OF GENERIC UNRESOLVED SAFETY ISSUES
- GENERIC RULE MAKING PROCEEDINGS (SITING, EMERGENCY PLANNING, CORE DEGRADATION)
- ESTABLISHING PRIORITIES FOR SAFETY RESEARCH
- DETERMINING NEED FOR BACKFITTING (SYSTEMATIC EVALUATION PROGRAM)
- DETERMINING NEED FOR PLANT SHUTDOWN ORDERS
- TECHNICAL SPECIFICATIONS (LCO AND TESTING REQUIREMENTS)
- EVALUATING OPERATING EXPERIENCE

ESTABLISHING QUANTITATIVE SAFETY GOALS

I — ESTABLISH BASIC PRINCIPLES

II — DETERMINE WHAT SHOULD BE QUANTIFIED

III — DEVELOP NUMERICAL VALUES

IV — ESTABLISH MEANS OF APPLICATION

BASIC PRINCIPLES FOR SAFETY GOALS

- GENERICALLY APPLICABLE TO ALL TECHNOLOGIES OR RISK RELATED ACTIVITIES
- ACCEPTABLE SOCIETAL RISK SHOULD REFLECT SOCIETAL BENEFITS
- NO INDIVIDUAL SHOULD BEAR INORDINATE SHARE OF RISK
- PROMOTE OPTIMUM ALLOCATION OF RESOURCES IN REDUCING RISK

ELEMENTS TO BE ADDRESSED IN QUANTITATIVE SAFETY GOALS

- INDIVIDUAL HEALTH EFFECTS
- POPULATION HEALTH EFFECTS
- COST-BENEFIT RATIO
- CORE DEGRADATION PROBABILITY

INDIVIDUAL HEALTH EFFECTS CRITERION

THE INCREMENTAL RISK OF ADVERSE HEALTH EFFECTS TO THE MAXIMALLY EXPOSED INDIVIDUAL IN THE VICINITY OF A NUCLEAR PLANT SITE SHOULD NOT RESULT IN A SIGNIFICANT INCREASE IN ANNUAL MORTALITY RISK OR IN SIGNIFICANT SHORTENING OF EXPECTED STATISTICAL LIFE SPAN.

<u>SUGGESTED GOAL</u> 10^{-5}/YR INDIVIDUAL MORTALITY RISK (MEAN VALUE)

<u>BASES</u>

- REPRESENTS A SMALL FRACTION OF EXISTING BACKGROUND RISK

  ~ 0.1% OF TOTAL MORTALITY RISK
  ~ 1% OF ACCIDENT MORTALITY RISK

- COMPARISON WITH OTHER RISK CONTRIBUTION

	AVERAGE INDIVIDUAL* RISK
MOTOR VEHICLES	2×10^{-4}
VIOLENCE	2×10^{-4}
FIRES	1×10^{-4}
AIR TRAVEL	1×10^{-5}
FALLING OBJECTS	1×10^{-5}
ELECTROCUTION	6×10^{-6}
RADIATION INDUCED CANCERS	1×10^{-5}

* IT IS EXPECTED THAT THE MAXIMUM RISK TO AN INDIVIDUAL FROM ANY OF THESE SOURCES IS CONSIDERABLY HIGHER THAN THE AVERAGE RISK

SOME PROPOSED NUMERICAL VALUES FOR INDIVIDUAL RISK CRITERION

NRC – RES	10^{-5}/YR UNACCEPTABLE
	10^{-6} – 10^{-5}/YR WARNING RANGE (CASE BY CASE EVALUATION)
WILSON	10^{-5}/YR NEAR SITE
	10^{-6}/YR NEXT TOWNSHIP
OKRENT	2 X 10^{-4}/YR ESSENTIAL ACTIVITY
	2 X 10^{-5}/YR BENEFICIAL ACTIVITY
	2 X 10^{-6}/YR PERIPHERAL ACTIVITY
	ASSESS RISK AT 90% C.L.
CORKERTON ET AL (CEGB)	10^{-5}/YR PUBLIC
	10^{-4}/YR WORKER
WASH 1400	8 X 10^{-7}/YR
GERMAN RISK STUDY	1 X 10^{-6}/YR
AIF	10^{-5}/YR

POPULATION HEALTH EFFECTS CRITERION

THE INCREMENTAL CUMULATIVE RISK OF ADVERSE HEALTH EFFECTS TO THE EXPOSED POPULATION PER 1000 MW(e) OF NUCLEAR POWER CAPACITY, CONSIDERING THE PROBABILITY AND CONSEQUENCES OF EVENTS INTEGRATED OVER THE SPECTRUM OF POTENTIAL ACCIDENTS, SHOULD BE NO MORE THAN A SMALL FRACTION OF THE AVERAGE BACKGROUND INCIDENCE OF HEALTH EFFECTS.

SUGGESTED GOAL 0.1 FATALITY/YR – 1000 MW(e) (MEAN VALUE)

BASES

- REPRESENTS NEGLIGIBLE INCREMENT OF EXISTING BACKGROUND RISK
  ~ .001% OF TOTAL MORTALITY RISK*
  ~ .005% OF TOTAL CANCER RISK*
- COMPARISON WITH OTHER RISK CONTRIBUTIONS

	NUMBER OF FATALITIES/YR	PERCENT OF TOTAL RISK
MOTOR VEHICLES	50,000	2.5
VIOLENCE	40,000	2
FIRES	7,500	0.4
AIR TRAVEL	1,800	0.09
FALLING OBJECTS	1,200	0.06
ELECTROCUTION	1,100	0.05
RADIATION INDUCED CANCER	2,000	0.1

*FOR 200,000 MW(e) TOTAL NUCLEAR CAPACITY

SOME PROPOSED NUMERICAL VALUES FOR POPULATION RISK CRITERION (Implied From Risk Curves)

LEVINE	0.2 FATALITIES/YR
WASH 1400	0.02 FATALITIES/YR
GERMAN RISK STUDY	0.4 FATALITIES/YR
AIF	0.1 FATALITIES/YR

COST-BENEFIT CRITERION

THE BENEFIT, IN TERMS OF POPULATION RISK REDUCTION, AFFORDED BY A CHANGE IN PLANT DESIGN OR OPERATING PROCEDURE SHOULD BE COMPARABLE TO THAT WHICH IS GENERALLY ACHIEVABLE THROUGH ALTERNATE INVESTMENTS OF THE COST OF THE CHANGE IN OTHER AREAS OF PUBLIC RISK REDUCTION.

SUGGESTED GOAL $100/MAN-REM

BASES

- EQUIVALENT TO $1 MILLION/LIFE SAVED
- COMPARABLE TO MEDIAN COST-BENEFIT RATIOS FOR OTHER HEALTH & SAFETY PROTECTIVE MEASURES

COST-BENEFIT RATIOS FOR VARIOUS HEALTH AND SAFETY PROTECTIVE MEASURES

	COST-BENEFIT RATIO ($ MILLIONS/LIFE SAVED)
NUCLEAR POWER PLANT DESIGN FEATURES	
RADWASTE EFFLUENT TREATMENT SYSTEMS	10
ECCS	0.1
CONTAINMENT	5
D-G SETS	1
HYDROGEN RECOMBINERS	3000
COAL POWER PLANT DESIGN FEATURES	
HIGH SULFUR COAL WITH SO_2 SCRUBBERS (85% REMOVAL)	0.1-1.4
LOW SULFUR COAL WITH SO_2 SCRUBBERS (85% REMOVAL)	0.7-10
OCCUPATIONAL HEALTH AND SAFETY	
OSHA BENZENE REGULATIONS	300
OSHA COKE FUME REGULATIONS	4.5
ENVIRONMENTAL PROTECTION	
EPA VINYL CHLORIDE REGULATIONS	4
PROPOSED EPA DRINKING WATER REGULATIONS	2.5
FIRE PROTECTION	
PROPOSED CPSC UPHOLSTERED FURNITURE FLAMMABILITY STANDARDS	0.5
SMOKE DETECTORS	0.05-0.08
AUTOMOTIVE AND HIGHWAY SAFETY	
HIGHWAY SAFETY PROGRAMS	0.14
AUTO SAFETY IMPROVEMENTS (1966-1970)	0.13
AIR BAGS	0.32
SEAT BELTS	0.08
MEDICAL AND HEALTH PROGRAMS	
KIDNEY DIALYSIS TREATMENT UNITS	0.2
MOBILE CARDIAC EMERGENCY TREATMENT UNITS	0.03
CANCER SCREENING PROGRAMS	0.01–0.08

SOME PROPOSED NUMERICAL VALUES FOR COST-BENEFIT CRITERION

NRC (RG 1.110)	\$1000/MAN-REM
EPA	\$75/MAN-REM
WEST GERMANY	DM 100 – 200/MAN-REM (\$50 – 100/MAN–REM)
ROGERS	\$30/MAN-REM
AIF	\$100/MAN-REM

CORE DEGRADATION PROBABILITY CRITERION

A LIMIT SHOULD BE ESTABLISHED FOR THE PROBABILITY OF ACCIDENTS INVOLVING SERIOUS CORE DEGRADATION SUCH THAT, GIVEN THE EXPECTED POPULATION OF REACTORS, THE RECURRENCE INTERVAL FOR ACCIDENTS AS SERIOUS AS THREE MILE ISLAND IS ON THE ORDER OF ONE PER SEVERAL DECADES.

DISCUSSION

- ESTABLISHES MINIMUM REQUIREMENTS FOR ACCIDENT PREVENTION – PREVENTS UNDUE EMPHASIS ON MITIGATION

- REDUCES FREQUENCY OF STRESS PROVOKING EVENTS FOR POPULATIONS NEAR PLANTS

- LIMITS ECONOMIC RISKS OF ACCIDENTS

RELATIONSHIP AMONG INDIVIDUAL RISK, POPULATION RISK, POPULATION DENSITY, AND STATION CAPACITY

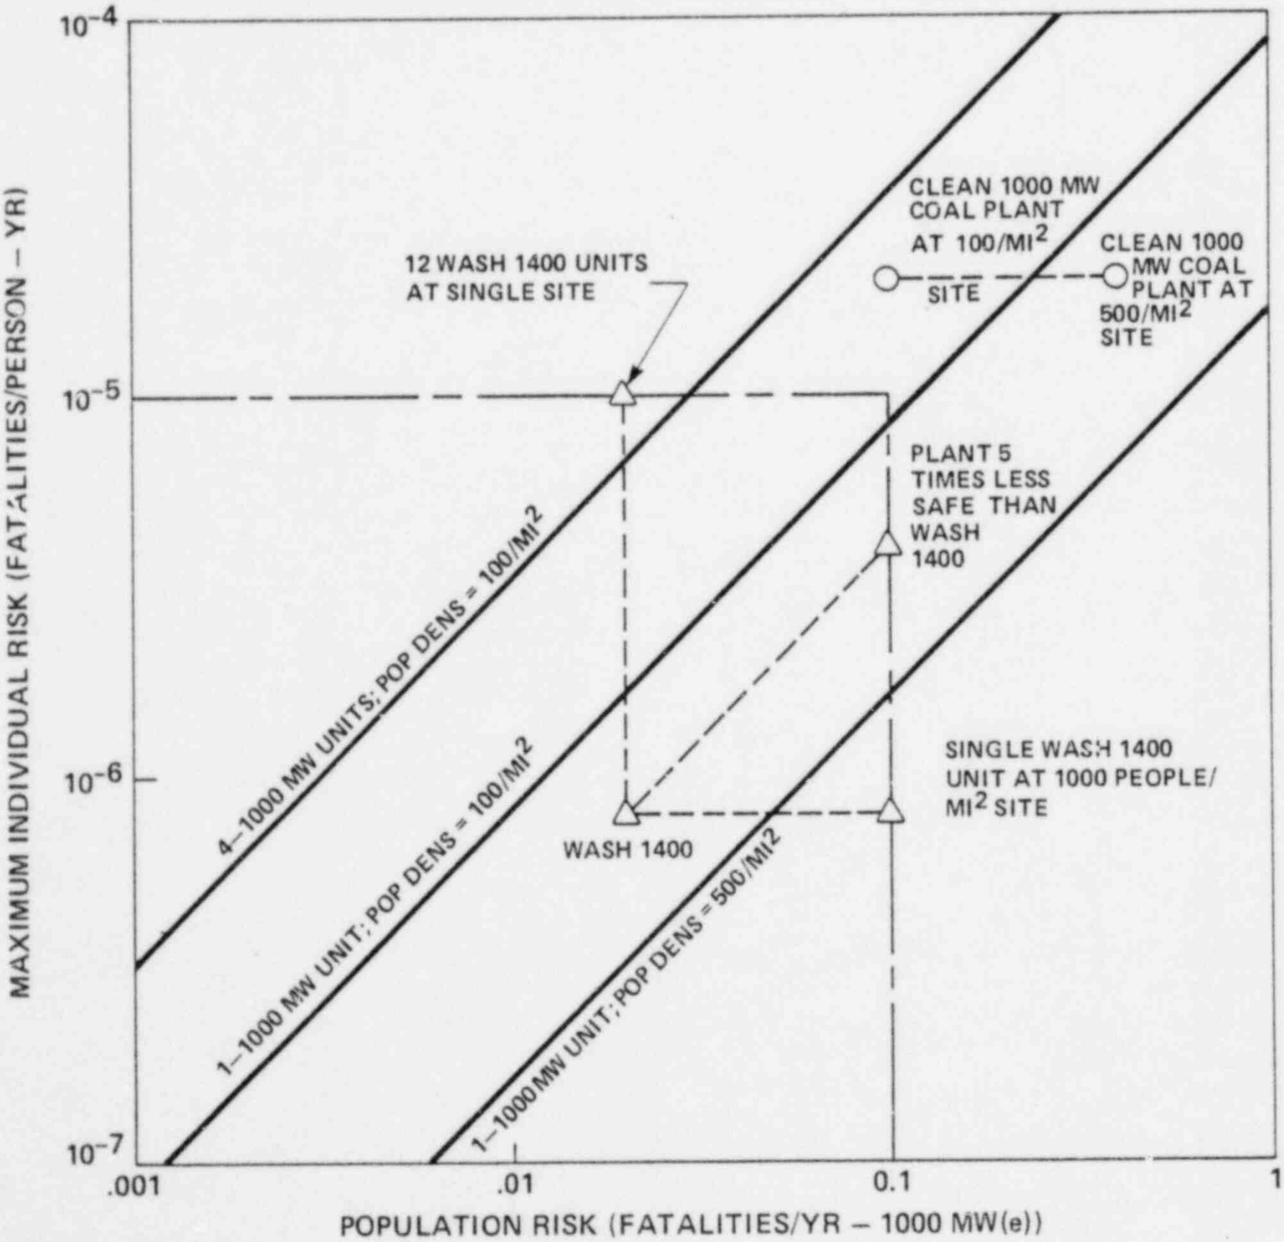

CAVEATS FOR QUANTITATIVE SAFETY GOALS

- SUGGESTED NUMERICAL VALUES SHOULD BE USED WITH MEAN VALUE (50% C.L.) ESTIMATES OF RISK. HIGHER VALUES APPROPRIATE FOR MORE CONSERVATIVE (HIGHER C.L.) ESTIMATES OF RISK
- INITIAL SET OF VALUES SHOULD BE INTERIM FOR TRIAL USE PERIOD OF 3 YEARS
- QUALITATIVE JUDGEMENT MUST SUPPLEMENT QUANTITATIVE GOALS – PARTICULARY IMPORTANT IN BORDERLINE CASES

O'DONNELL

REGULATORY OBJECTIVES

- HIGHEST ACHIEVABLE STANDARD
- LOWEST ACHIEVABLE RELEASE OF RADIOACTIVITY OR CARCINOGENS

∴ UNTENABLE, NEITHER REASONABLE NOR ADEQUATE AS PLANT DESIGN CRITERIA

REGULATORY PRACTICE

- RELEASES AS LOW AS REASONABLY ACHIEVABLE
- GENERAL DESIGN CRITERIA PLUS REGULATORY GUIDES
- BOTH EXPERIENCED SERIOUS SHORTCOMINGS - AMBIGUOUS, TEND TO MAKE DESIGNERS CONCENTRATE ON MEETING REQUIREMENTS RATHER THAN STIMULATE INNOVATION AND LEGITIMATE TRADEOFFS, SAFETY STANDARDS ACHIEVED DEPENDENT ON INDIVIDUALS ASSIGNED TO PARTICULAR PLANTS

QUANTITATIVE SAFETY GOALS

- WILL PROVIDE CONSISTENT SAFETY CRITERIA AND THUS FIRM FRAMEWORK FOR DESIGNERS
- AVOID RATCHETING OF PLANT REQUIREMENTS
- LEAVE DESIGNERS WITH FREEDOM TO CREATE EFFECTIVE DESIGN SOLUTIONS TO REAL PROBLEMS THEY FACE
- WILL PROTECT PUBLIC AND INVESTMENT MORE EFFECTIVELY
- WILL ENABLE DESIGNERS TO GET ON WITH THE JOB OF DESIGNING FUTURE REACTORS
- WILL CONTRIBUTE TOWARDS PUBLIC UNDERSTANDING OF NUCLEAR SAFETY

GA BACKGROUND

STATEMENT BEFORE NRC RARG GROUP	-	DEC. 1977
STATEMENT BEFORE UDALL COMMITTEE	-	FEB. 1979
AIF WORKSHOP PAPER	-	MAY 1979
AGNEW TO KEMENY LETTER	-	SEP. 1979
KNOXVILLE CONFERENCE PAPER	-	APR. 1980
LTR ON SELECTION OF DESIGN BASIS ACCIDENTS	-	JUN. 1980

KEY ELEMENTS (CONCEPTS) OF GA APPROACH TO QUANTITATIVE SAFETY GOALS

- PROFESSIONAL JUDGMENT
- RISK BUDGET VIA COMPARATIVE RISK STUDIES
- EMPHASIS THROUGH KNOWLEDGE OF PRA BUT RETAINING SOME ASPECTS OF DETERMINISTIC REGULATORY FRAMEWORK
- RETENTION OF DBA BUT PRA FOCUSED
- ADDITION OF SAFETY MARGIN AND RESEARCH REGIONS
- NO IDENTIFIABLE PUBLIC INJURY
- USE OF LIMIT LINES
- RANGE FROM EMPHASIZED TO BALANCED RISK POLICY
- ACCEPTABLE INDIVIDUAL, SOCIETAL, AND PUBLIC PROPERTY RISKS

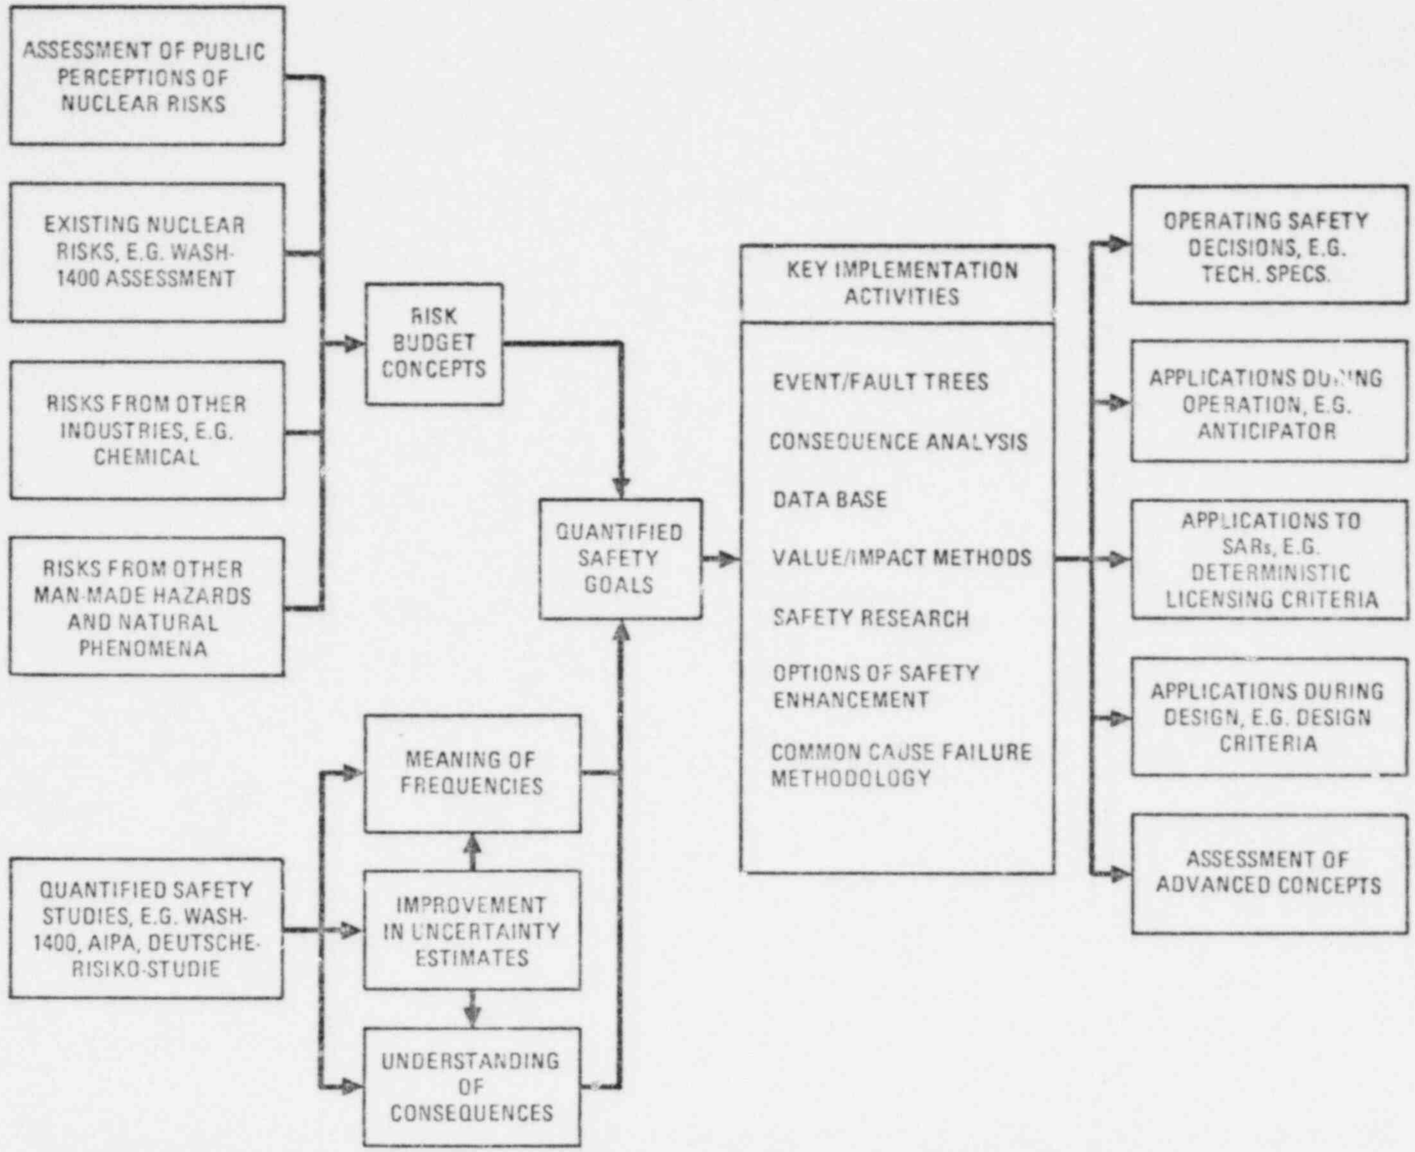

Figure 1. Major steps in development and implementation of quantified safety goals

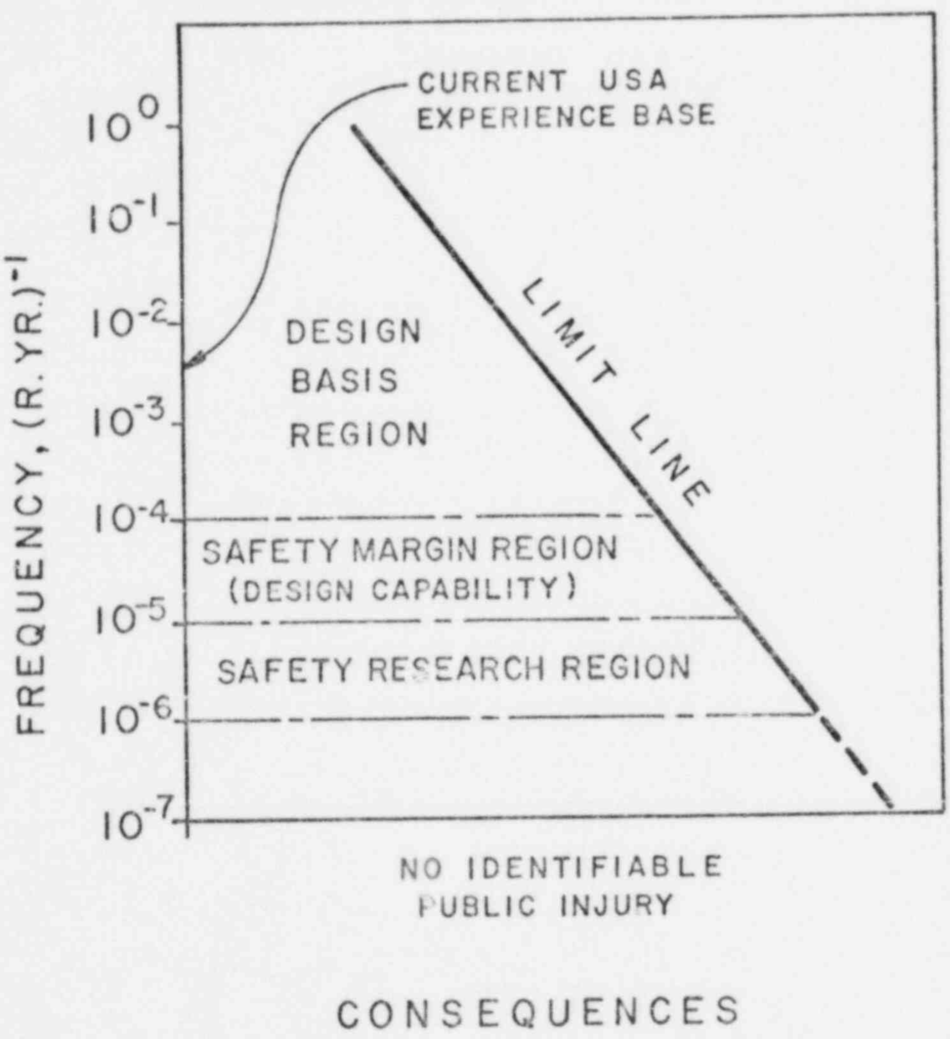

FIGURE 2
MAP OF QUANTITATIVE SAFETY REGIONS

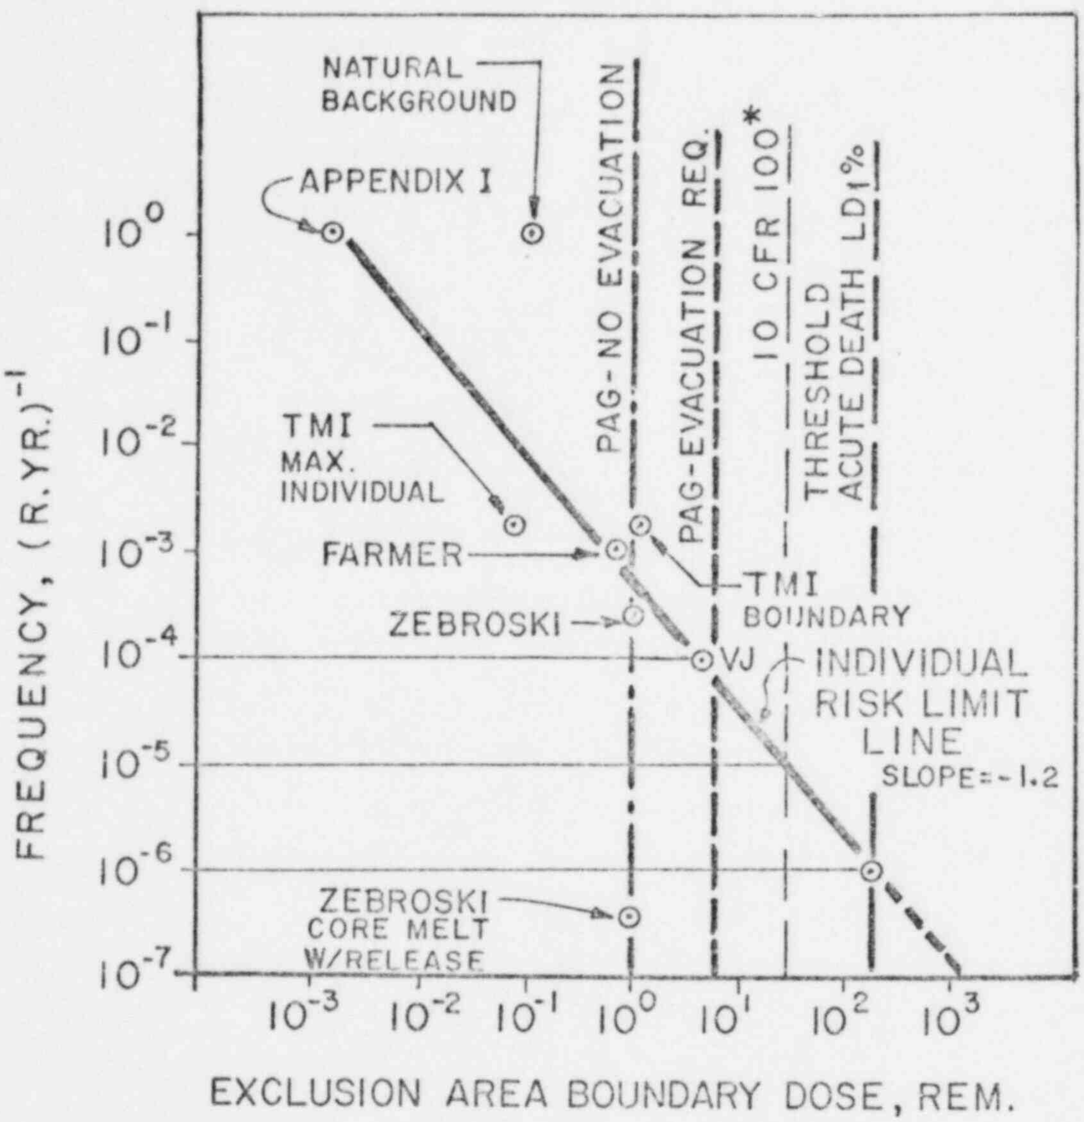

* ASSUMING REALISTIC EVALUATION OF DOSES

FIGURE 3
QUANTITATIVE SAFETY GOAL - INDIVIDUAL RISK TO MAXIMUM EXPOSED INDIVIDUAL MEMBER OF THE PUBLIC

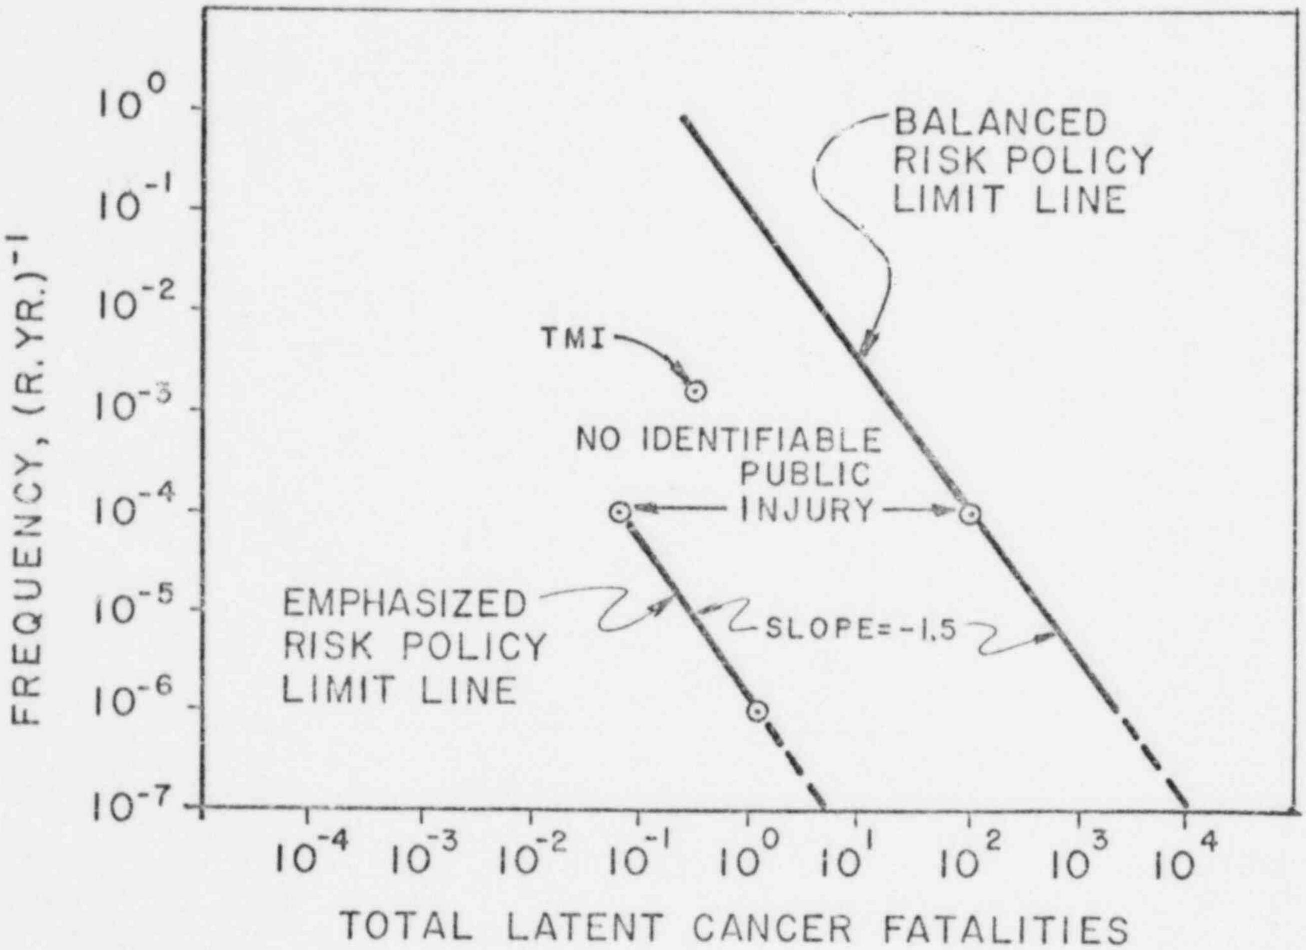

FIGURE 4

QUANTITATIVE SAFETY GOAL - SOCIETAL RISK

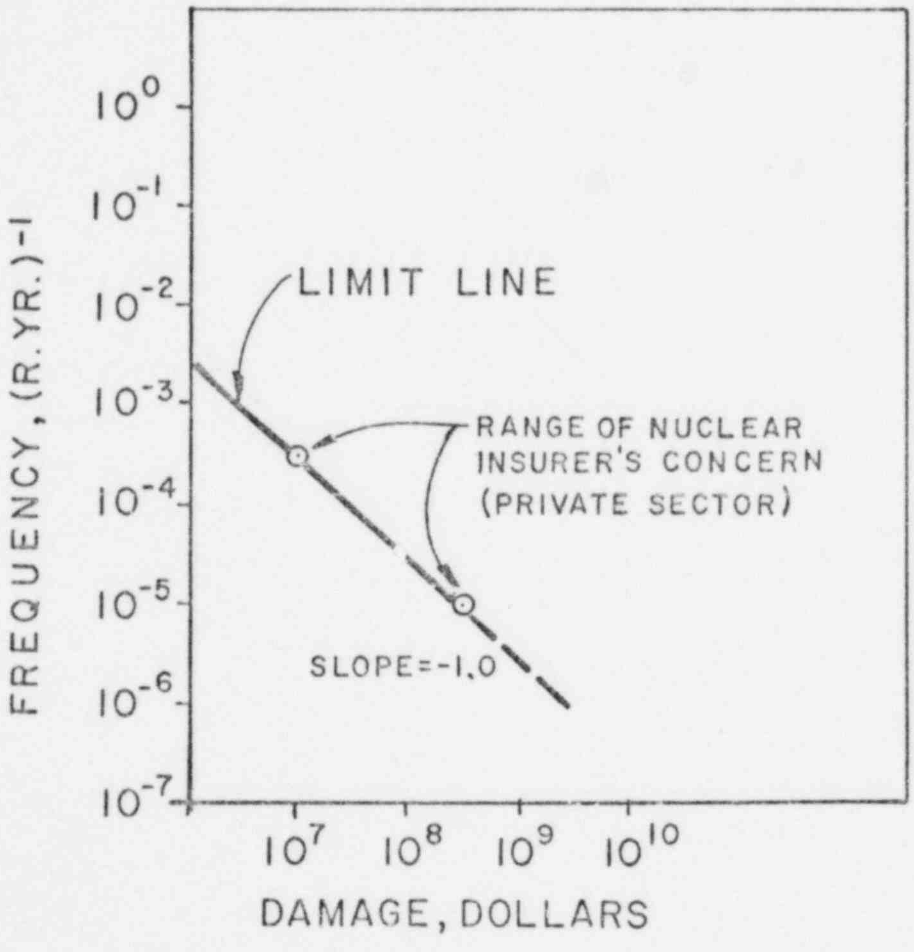

FIGURE 5

QUANTITATIVE SAFETY GOAL-
PUBLIC PROPERTY RISK

JOHNSON

June 17, 1980

NOTE TO: William E. Vesely, Head
Methodology Applications Section
Probabilistic Analysis Staff, PAS

FROM: James W. Johnson
Methodology Applications Section
Probabilistic Analysis Staff, PAS

SUBJECT: FLOE CODE SUMMARY REPORT

Attached please find a brief FLOE code summary report. I have included an example application of the code to annual peak discharge data from the Susquehanna River. These data have also been forwarded to EG&G so that their code may be exercised with actual discharge data.

James W. Johnson
Methodology Applications Section
Probabilistic Analysis Staff
Office of Nuclear Regulatory Research

Attachment: FLOE Code
Summary Report

FLOE CODE SUMMARY REPORT

INTRODUCTION

The Flood Level Occurrence Evaluation (FLOE) Code is being developed to provide point and interval estimates of the probability that large flood discharges are exceeded. The code utilizes both expert assessments and historical flood occurrence data to obtain estimates of flood exceedence probabilities and their associated uncertainties. Postulated relationships between regulated and non-regulated flows, as well as the associated uncertainties are also incorporated in the code.

This report provides a brief description of the methodology used in the FLOE code along with a status report and future plans. Finally, as an example application, the code is applied to discharge data from the Susquehanna River at Harrisburg, Pa. This example is for illustrative purposes only.

METHODOLOGY

In the FLOE code, we confine our attention to the maxima of the annual hydrograph peaks (i.e., the largest annual flood flow), and estimate the probability that the annual maxima exceeds a specified value. Previous approaches to this problem have preceeded along the following two lines:

(a) It is assumed that the largest annual flood peaks have the third asymptotic distribution of the maxima (of a sequence of independent identically distributed random variables (Gumbel's Law)) and the tail probability of the asymptotic distribution is then estimated.

(b) the hydrograph is considered to be the realization of a stochastic process, and the probabilistic properties of the maxima of the stocastic process then analyzed.

The approach taken in the FLOE code is analogous to (a) above, in the sense that a distribution is assumed for the largest annual flood peaks, but the estimation of the tail properties is undertaken within a Bayesian framework.

For notational purposes, we let the random values of the unregulated hydrograph peaks be denoted by $\dot{X}_u$, and the regulated hydrograph peaks by $\dot{X}_R$. Also let $X_u = \log_{10}\dot{X}_u$ and $X_R = \log_{10}\dot{X}_R$.
For convenience, $X \sim G(\alpha,\beta)$ will denote the fact that a random variable X has a gamma distribution with a shape parameter α and a scale parameter β.

Benson (1968) has given a comprehensive summary of an investigation conducted by a federal interagency group on the distribution of the random variable $\dot{X}_u$. The group has recommended that the log-Pearson Type III distribution be used for the $\dot{X}_u$. A random variable X is said to have a log-Pearson Type III distribution if its density function is given by:

$$h(X;\alpha,\beta) = \frac{1}{\Gamma(\alpha)|\beta \ln 10|}\left(\frac{\ln X}{\beta \ln 10}\right)^{\alpha-1} e^{-\ln X/\beta \ln 10}\ \frac{1}{X}$$

where $1<x<\infty$ if $\beta>0$ and $0<X<1$ if $\beta<0$.
Equivalently this means that $\log_{10}X \sim G(\alpha,\beta)$ if $\beta>0$ and $\log_{10}(\frac{1}{X}) \sim G(\alpha,-\beta)$ if $\beta<0$.

Even though there does not appear to be any theoretical explanations which support the use of this distribution, we have chosen to conform initially with the recommendation of the interagency group. Our approach can and will be amended to account for other forms of the distribution of X_u.

In order to estimate the probability that the annual maxima, X_u, exceeds a specified value, say X^*, we shall have to obtain an estimate of:

$$\Pr(\dot{X}_u \geq X^*) = \Pr(\log_{10}\dot{X}_u \geq \log_{10}X^*)$$
$$= \Pr(X_u \geq X)$$

where $X = \log_{10}X^*$. If $\beta>0$ then $X_u \sim G(\alpha,\beta)$ and

$$\Pr(\dot{X}_u \geq X^*) = \int_X^\infty \frac{e^{-\beta s}(\beta s)^{\alpha-1}\beta ds}{\Gamma(\alpha)}$$

$$\underline{\underline{def}}\ \overline{G}(X^*,\alpha,\beta)$$

If $\beta<0$, then $-X_u \sim G(\alpha,-\beta)$ and

$$\Pr(X_u \geq X^*) = \int_0^{-X} \frac{e^{-(-\beta s)}(\beta s)^{\alpha-1}\beta ds}{(\alpha)}$$

$$\underline{\underline{def}}\ \overline{G}(-X;\alpha,-\beta)$$

Thus our parameter of interest is $\overline{G}(X^*;\alpha,\beta)$

If $\beta>0$, and $\overline{G}(-\chi^*;\alpha,-\beta)$ if $\beta<0$, where both $\overline{G}(X^*;\alpha,\beta)$ and $\overline{G}(-X^*,\alpha,-\beta)$ lie between zero and one. We assume that it is known a priori whether $\beta>0$ or $\beta<0$. So for the sake of discussion we assume that $\beta>0$. The case

$\beta<0$ is treated in a similar manner. Since we have adopted a Bayesian point of view, we shall have to endow a prior distribution on $\overline{H}(X;\alpha,\beta)$ for two distinct choices of X, and then using the unregulated and regulated data obtain it posterior distribution. To obtain the posterior distribution we hypothesize a simple relationship between regulated and unregulated discharges.

FLOE CODE STATUS

The first version of the FLOE code is operational and has been exercised on some "test data". This version of the code permits the estimation of exceedence probabilities for both unregulated and regulated flows. The code has three basic options. The first option permits the specification of a sequence of percentile points at distinct discharges. The second option allows only the specification of δ and $(1-\delta)$ percent bounds at distinct discharges. At δ percent region of α,β points is determined and equal probability assigned to each of the α,β points. The final option permits the acutal specification of the joint prior distribution of α and β. All three options are now operational at EG&G.

The three versions of the code can each be exercised with unregulated data only or with unregulated and regulated discharges. If only unregulated data is used, the exceedence probabilities will be calculated in terms of regulated discharges. In estimating exceedance probabilities for regulated flows, it is assumed that a simple linear relationship exists between regulated and unregulated flows (specified by a parameter θ). The parameter specified by this relationship is assumed to be uniformly distributed with the hydrologist specifying the δ and $1-\delta$ percent bounds. The user must also specify the discharge levels at which exceedance probabilities are to be computed. The output from the code provides the exceedance probability distributions at each discharge point specified. In addition to the probability density function, the δ and $1-\delta$ percentiles, the mean and median of the exceedance probability is also provided.

The FLOE code also has the following plotting capabilities:

a. Plots of the average, median and the 5th and 95th percentiles of exceedance probability as a function of discharge.

b. Plots of both complementary cumulative distribution functions and density functions of exceedance at specified discharge levels.

c. Plots of the prior and posterior densities of the parameters involved in the Bayesian analysis. In particular, the joint probability mass function of the gamma distribution parameters α and β are shown and a separate plot of the density of the transformation parameter θ.

PLANS

Several routines in the FLOE code are being investigated in an attempt to provide more efficient algorithms. These routines include the algorithms translating hydrologist judgement into a region in the α,β plane for the prior distribution. Log Scale plots of the reciprocal of the inverse of the gamma, $G(\alpha,1)$, distribution function as a function of α have been generated, in order to obtain transparencies which can be used to show the exact shape of the regions in the α,β plane satisfying a wide variety of possible hydrologist specifications.

Other plans for the FLOE code include extensive sensitivity studies for the estimation of exceedance probabilities for unregulated flows. These sensitivity studies will include the effect of different prior bounds, the effect of different prior distributions, the effect of different discharge levels at which prior information is provided, the effect of mesh size in the α,β region. We will also develop automated procedures to determine the "appropriateness" of the log-Pearson type III distribution as the sampling distribution. We will also investigate the numerical stability of the exceedance results, considering the use of importance sampling and approximations for the prior distribution to make the code execute more efficiently.

EXAMPLE

To illustrate the FLOE code methodology an example evaluation for the Susquehanna River was performed. Although actual Susquehanna discharge data was used, this analysis is for illustrative purposes only. Before the methodology can be realistically applied, a comprehensive sensitivity study must be performed. In addition, no assessments of exceedance probabilities at the selected discharges was elicited from experts. Instead, we have based our posterior exceedance estimates on some assumed prior bounds.

An 87 year record, for the period 1891 through 1977 was used in this analysis. The actual annual peak Susquehanna discharges are provided in Table 1. All annual peak discharge data are assumed to be unregulated flows.

Consequently, exceedance probability estimates are obtained only for unregulated discharges. We assumed, in the first example calculation, that prior information in terms of bounds on exceedance probabilities was available at discharges of 2000,000 cfs and 350,000 cfs. These bounds were provided as the 5th and 95th percentiles. As 200,000 cfs, the assumed 5th and 95th percentiles were IE-1 and 7E-1 respectively. At 350,000 cfs, the assumed 5th and 95th percentiles were IE-2 and IE-1 respectively.

The 5th and 95th posterior exceedance probabilities are plotted in Figure 1 for discharges ranging from 1,700,000 cfs to 2,000,000 cfs. The median exceedance probability estimates range from 1.3E-5 at 1,700,000 cfs to 4.5E-6 at 2,000,000 cfs. As examination of some of the detailed printput from this example case suggest that the code is "fairly sensitive" to the grid size for the parameters (α and β) of the Log-Pearson type III distribution. Additional studies are planned in this area.

Previous analyses of discharge data from the Susquehanna River have suggested that the Log-Pearson type III distribution may not be an appropriate sampling distribution. In particular, it has been suggested that there are two underlying causal mechanisms associated with the process. Graphical analysis of the data suggest that if this is the case, the extreme annual discharges (assumed to result from one causal mechanism) is the dominant contributor to exceedance probability at the large discharges. So to provide an alternate example analysis we identified seven extreme annual discharges and assumed a Log-Pearson type III sampling distribution. The 5th and 95th posterior exceedance probabilities are plotted in Figure 2 for discharges ranging from 1,7000,000 to 2,000,000 cfs. The median exceedance probability estimates range from 2.6E-4 to 1.4E-4.

As a second illustration of the FLOE code, the above examples were repeated assuming log-normal percentiles on the exceedance probabilities at discharges of 200,000 and 350,000 cfs. Figure 3 provides the 5th and 95th percentile estimates on exceedance probabilities assuming the 87 year record. Exceedance probability estimates, assuming only the seven extreme estimates, are provided in Figure 4.

SUSQUEHANNA DISCHARGE DATA

Year	Discharge (cfs)
1891	408,000
1892	270,000
1893	324,000
1894	613,000
1895	230,000
1896	265,000
1897	180,000
1898	315,000
1899	228,000
1900	238,000
1901	249,000
1902	449,000
1903	276,000
1904	298,000
1905	306,000
1906	210,000
1907	247,000
1908	297,000
1909	297,000
1910	332,000
1911	178,000
1912	249,000
1913	402,000
1914	358,000
1915	286,000
1916	379,000
1917	155,000
1918	288,000
1919	294,000
1920	423,000
1921	178,000
1922	278,000
1923	261,000
1924	324,000
1925	379,000
1926	166,000
1927	323,500
1928	252,400
1929	235,000
1930	177,000
1931	153,000
1932	245,000
1933	269,000
1934	141,000

Year	Discharge (cfs)
1935	242,000
1936	740,000
1937	231,000
1938	178,000
1939	210,000
1940	418,000
1941	244,000
1942	290,000
1943	412,000
1944	212,000
1945	252,000
1946	494,000
1947	214,000
1948	308,000
1949	220,000
1950	300,000
1951	416,000
1952	324,000
1953	216,000
1954	242,000
1955	177,000
1956	338,000
1957	250,000
1958	281,000
1969	230,000
1960	382,000
1961	392,000
1962	270,000
1963	249,000
1964	484,000
1965	136,000
1966	265,000
1967	182,000
1968	202,000
1969	139,000
1970	343,000
1971	224,000
1972	1,020,000
1973	209,000
1974	205,000
1975	529,000
1976	239,000
1977	254,000

FIGURE 1 POSTERIOR 5th AND 95th PERCENTILES
PRIOR INFORMATION:PRIOR BOUNDS
DATA:EXTREME PEAK DISCHARGES

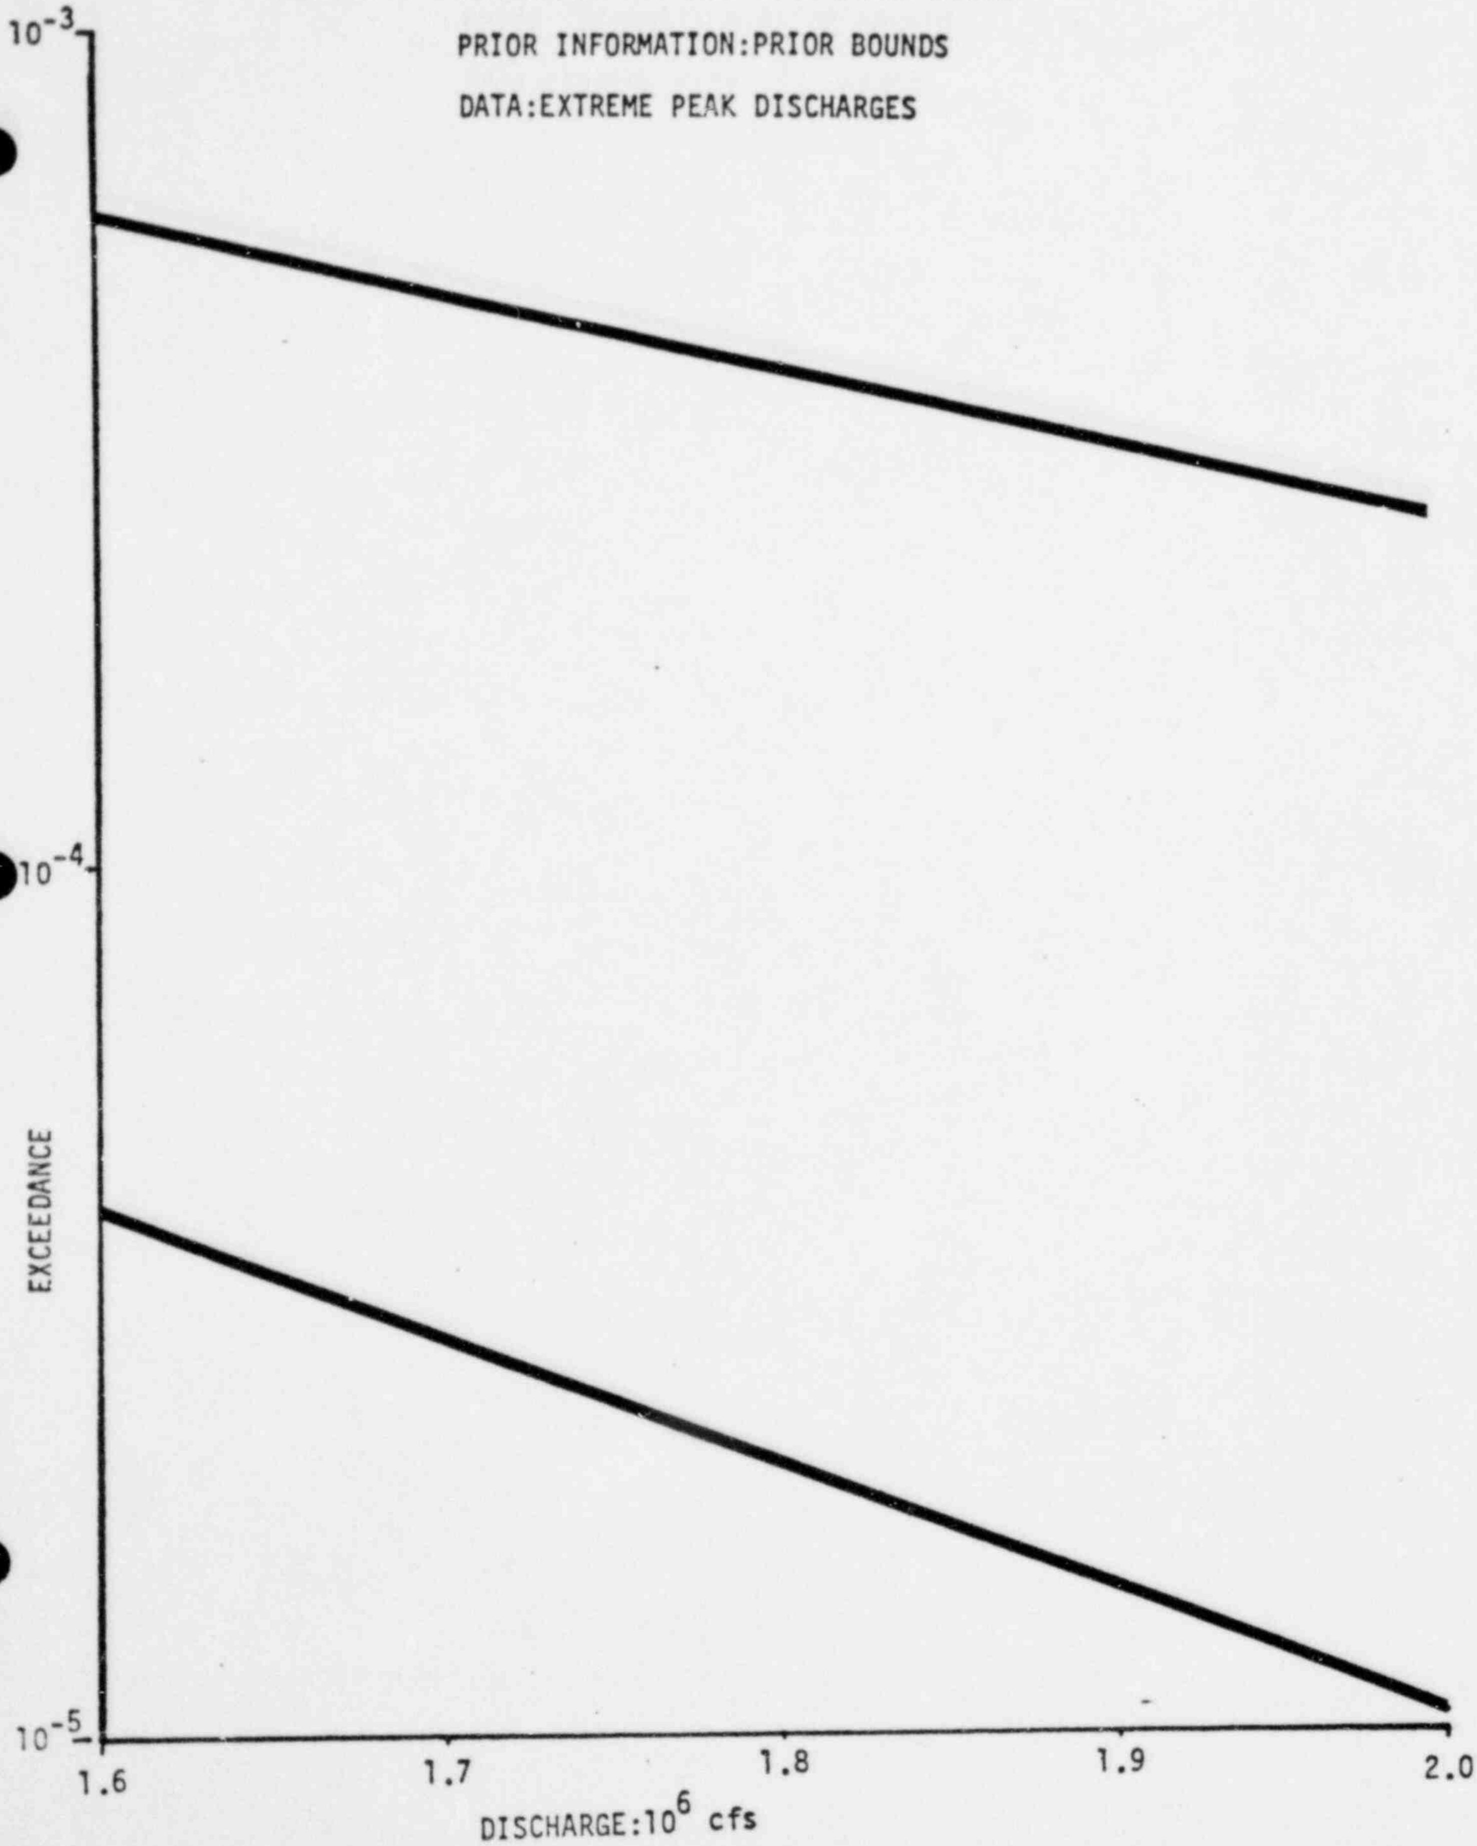

FIGURE 2 POSTERIOR 5th AND 95th PERCENTILES

PRIOR INFORMATION:PRIOR BOUNDS

DATA:87 YEAR RECORD

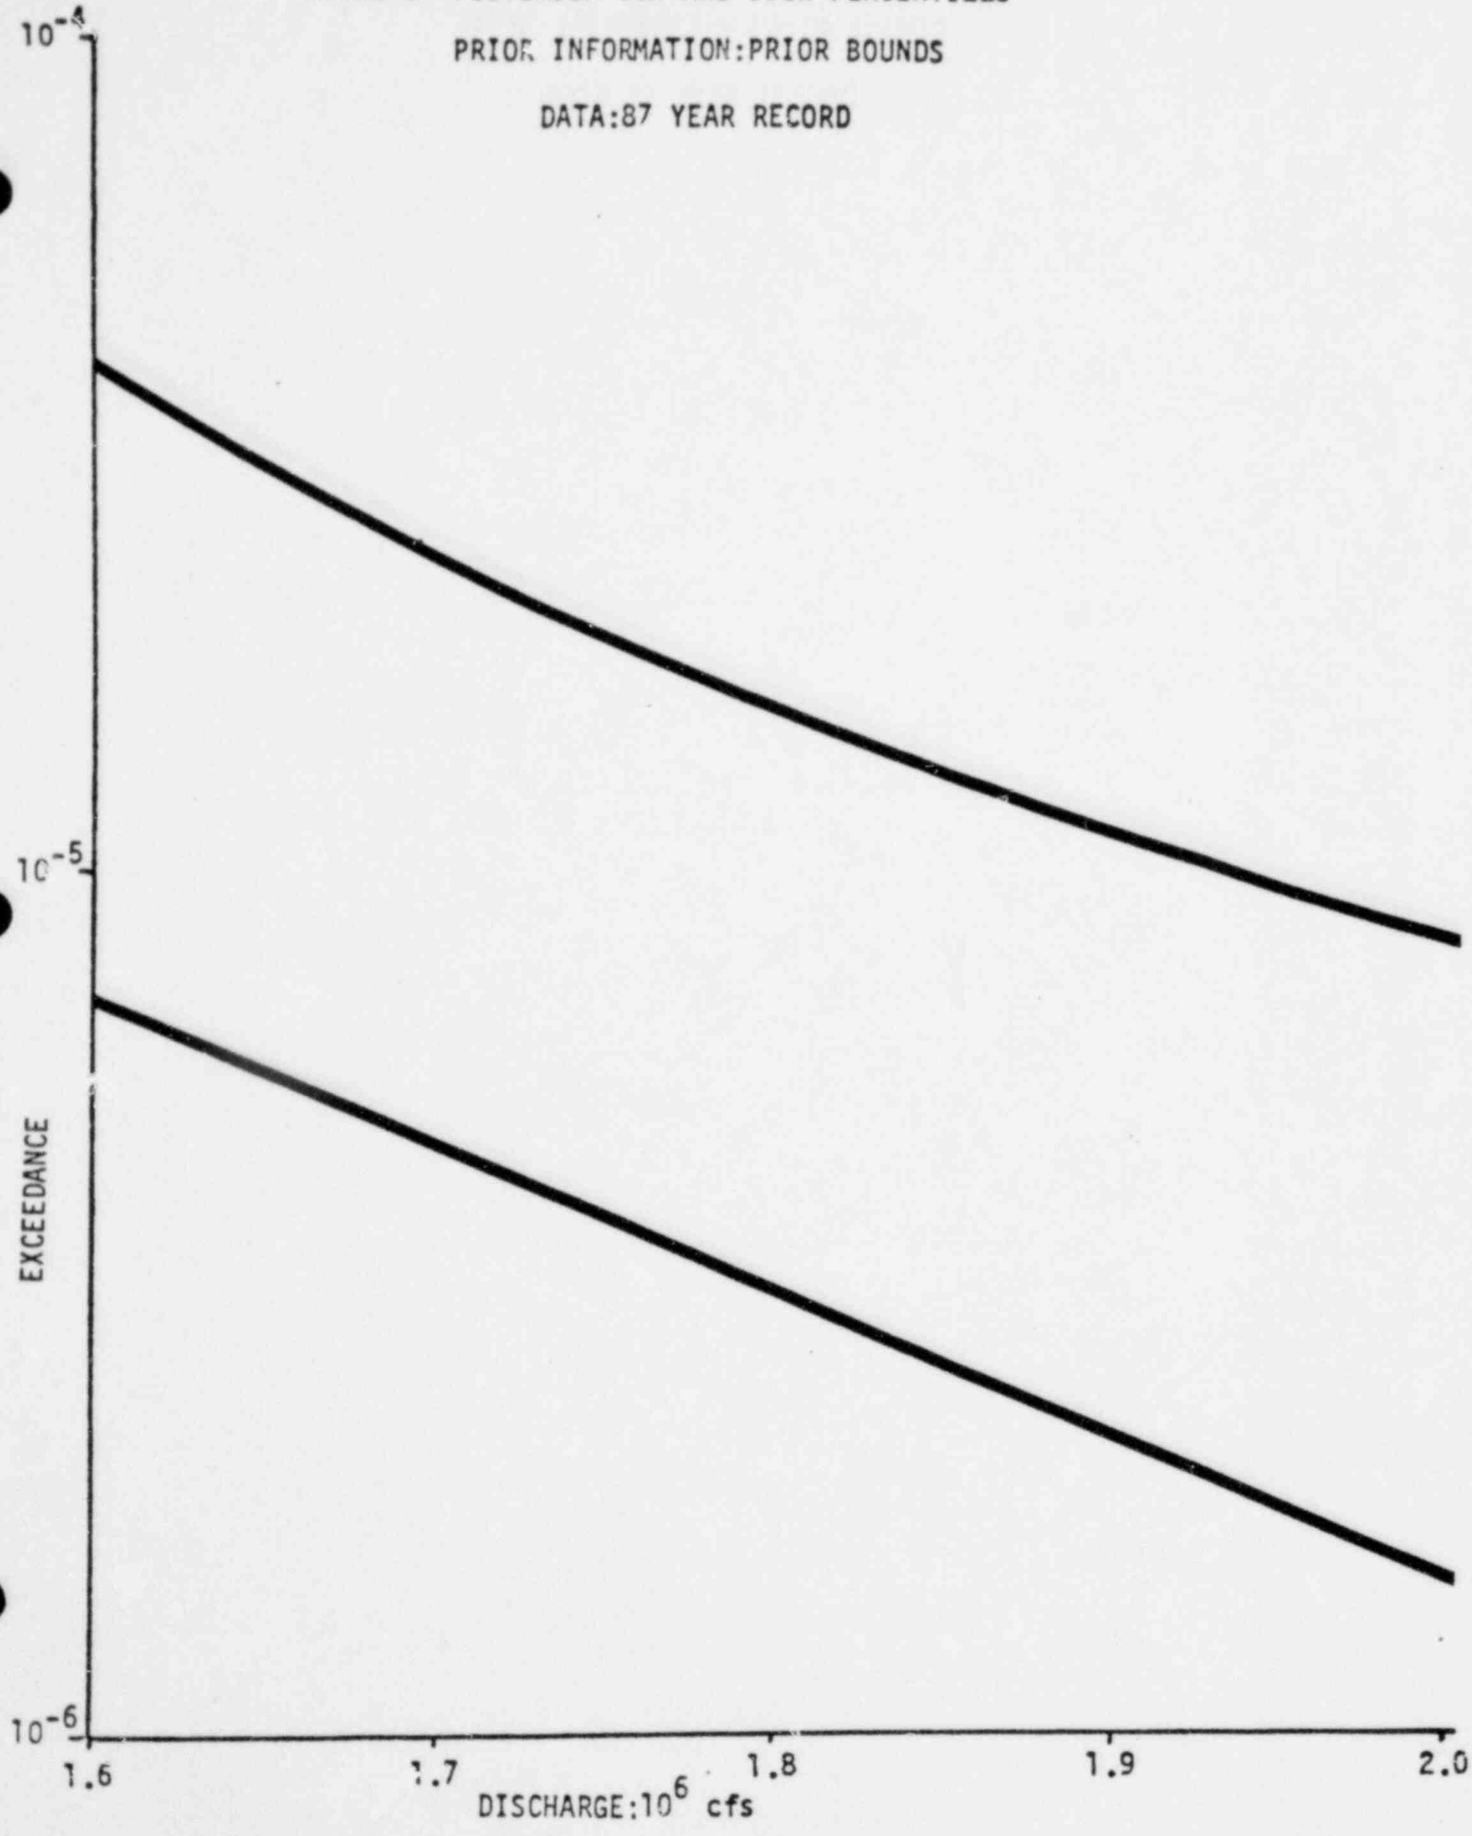

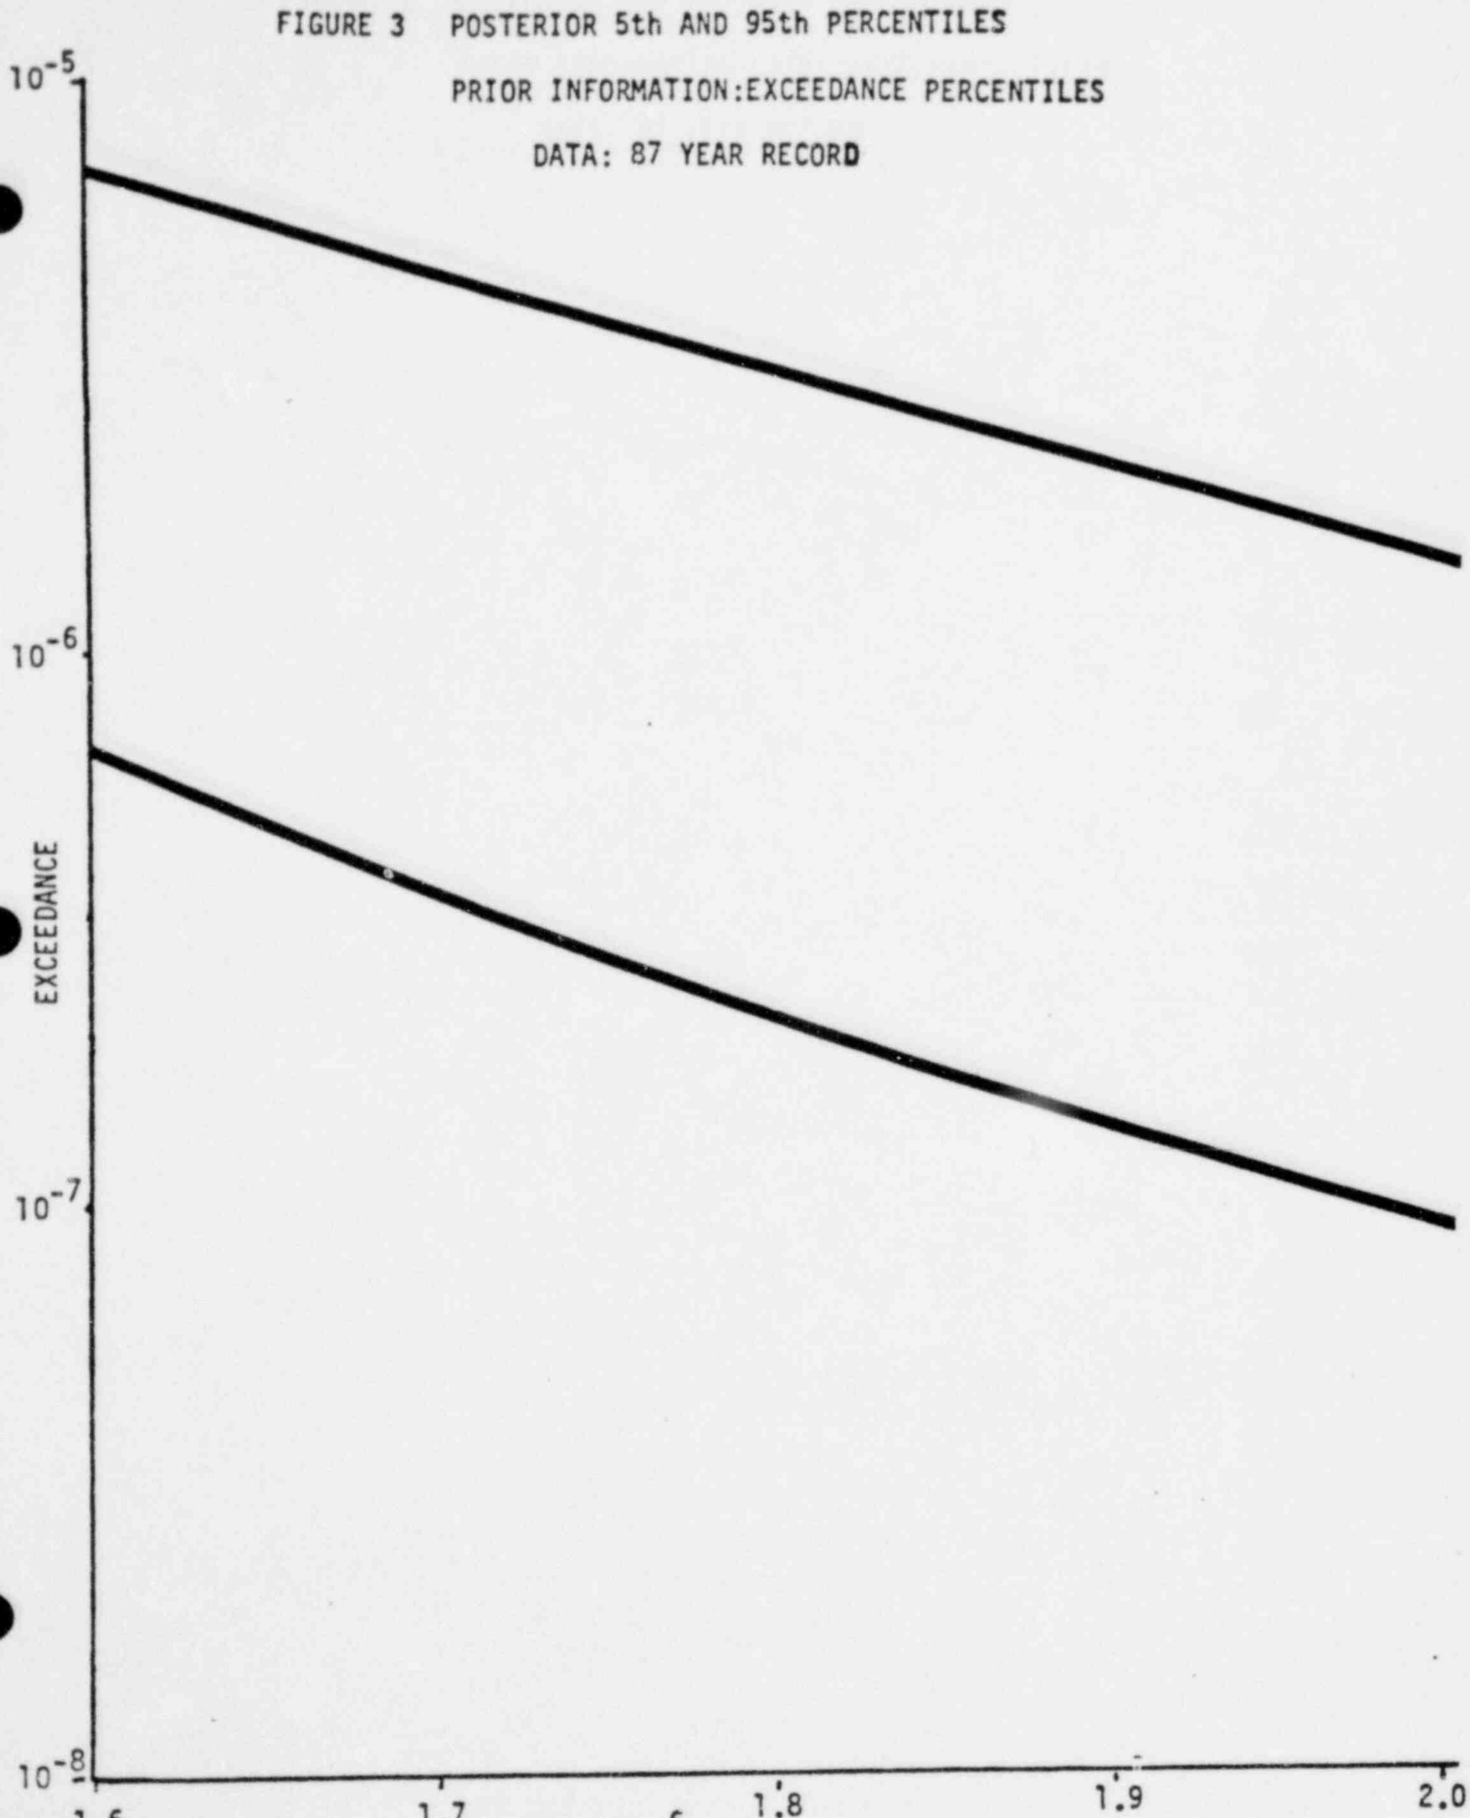
FIGURE 3 POSTERIOR 5th AND 95th PERCENTILES
PRIOR INFORMATION:EXCEEDANCE PERCENTILES
DATA: 87 YEAR RECORD
10^{-5}
10^{-6}
10^{-7}
10^{-8}
EXCEEDANCE
1.6
1.7
1.8
1.9
2.0
DISCHARGE:10^6 cfs

FIGURE 4 POSTERIOR 5th AND 95th PERCENTILES

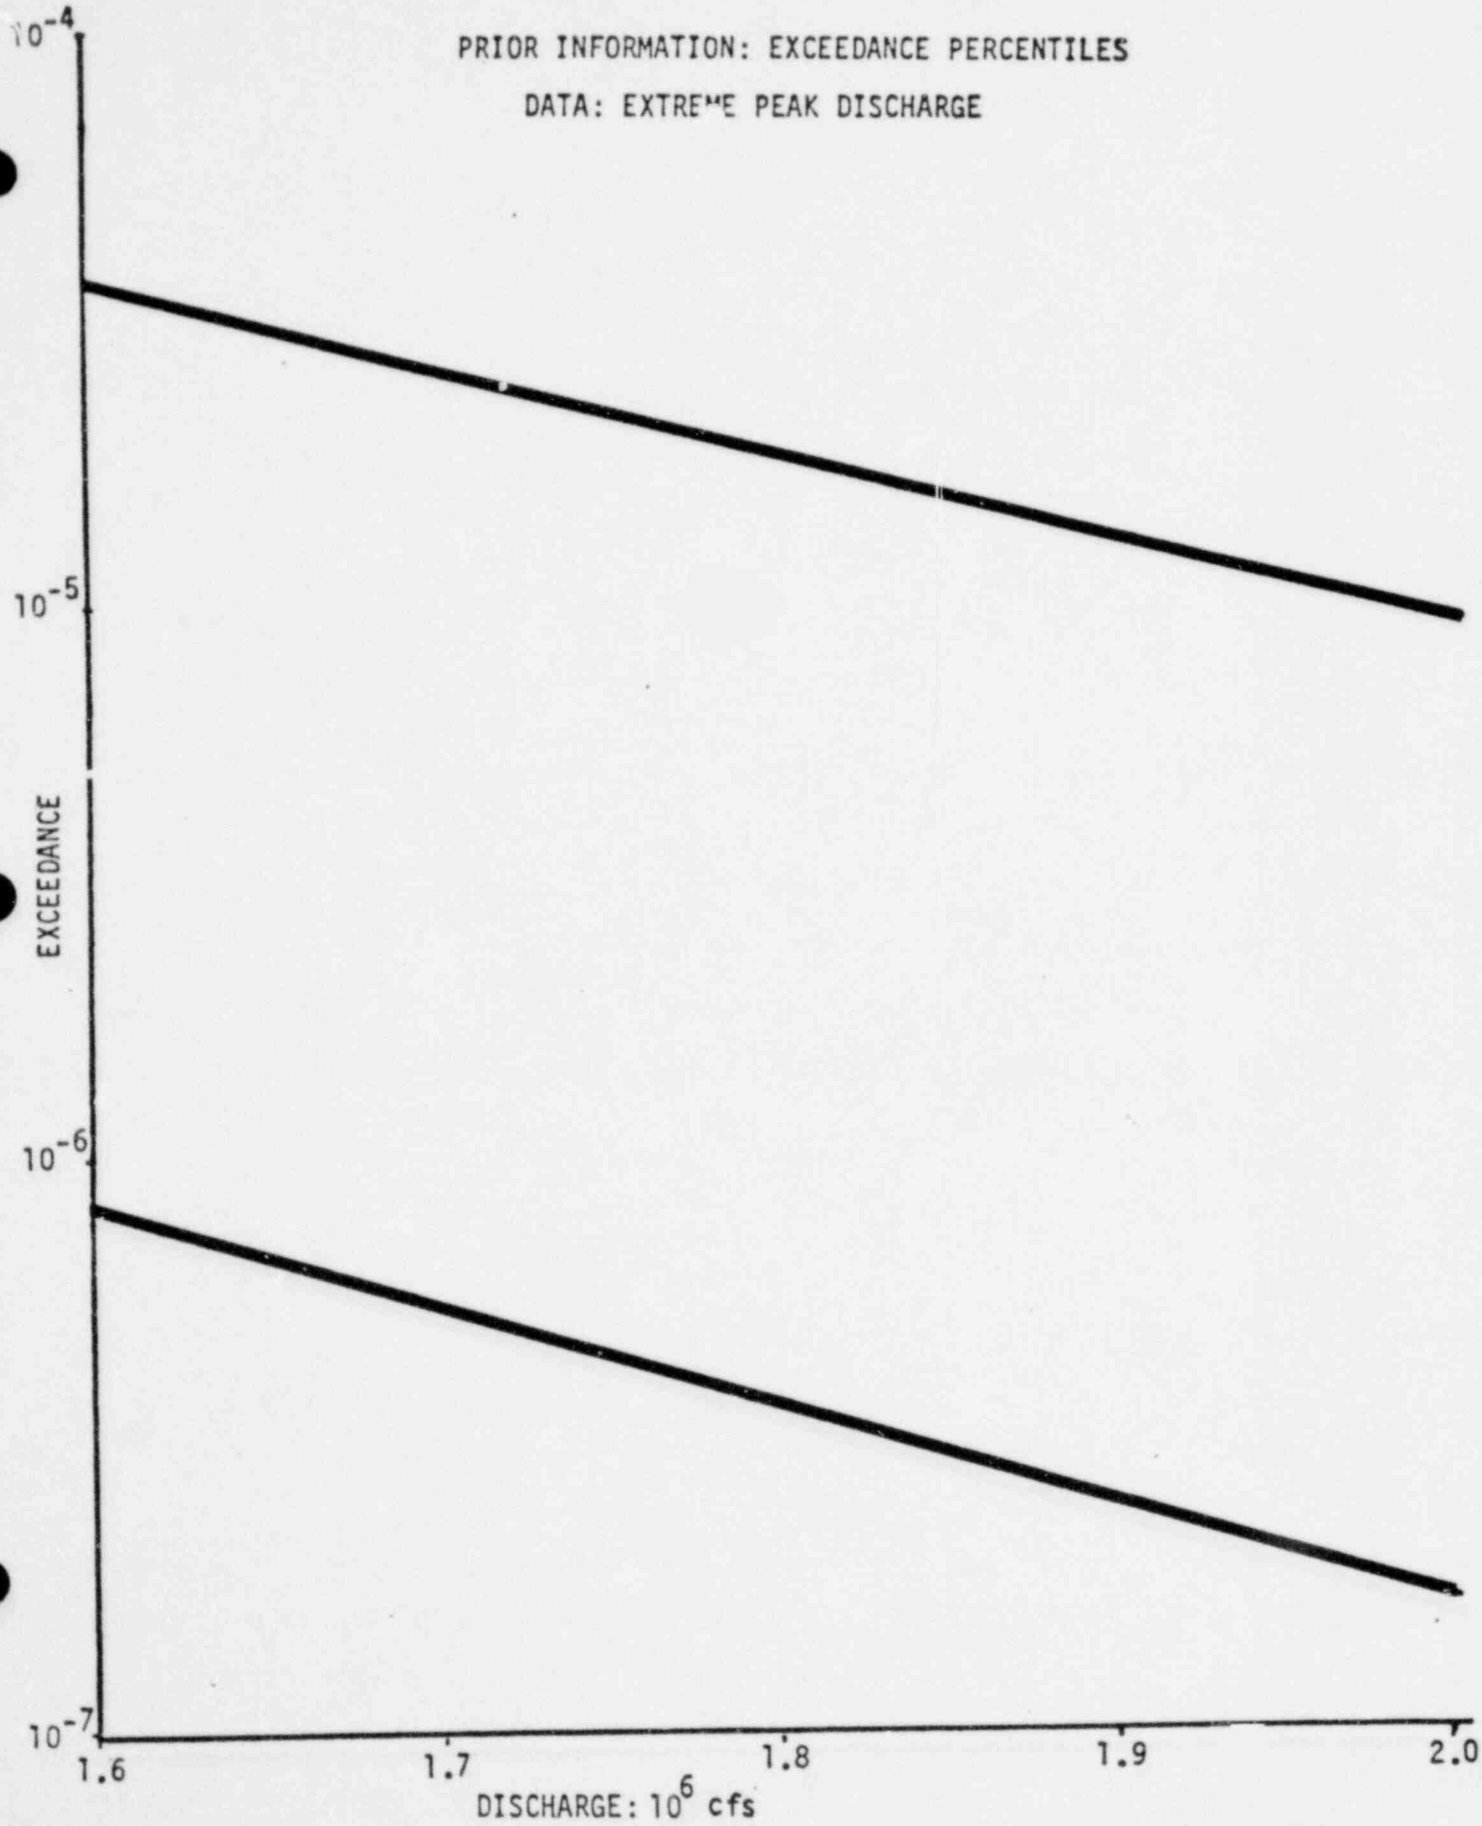

REFERENCE

Manuel A. Benson, "Uniform Flood-Frequency Estimating Methods for Federal Agencies," Water Resources Research, Vol. 4., No. 5., 891-908, (1968).

~~GRIESMEYER~~
KASTENBERG
~~JOHNSON~~

TABLE 3.1

Case 1: Risk of Immediate Death to Individual Originally Located 0.5 to 1 Mile from Reactor

Release Category*	Probability (P_i) Dose ≥ 200 Rem	Frequency (F_L) of Release (yr^{-1})*	P_iF_i
1A	0.028	$4x10^{-7}$	$1.1x10^{-8}$
1B	0.018	$5x10^{-7}$	$9.0x10^{-9}$
2	0.022	$8x10^{-6}$	$1.7x10^{-7}$
3	0.029	$4x10^{-6}$	$1.2x10^{-7}$
4	0.031	$5x10^{-7}$	$1.6x10^{-8}$
5	0.028	$7x10^{-7}$	$2.0x10^{-8}$
6	0	$6x10^{-6}$	0
7	0	$4x10^{-5}$	0
8	0	$4x10^{-5}$	0
9	0	$4x10^{-4}$	0

$$R = \sum_i P_iF_i = 3.5 \times 10^{-7} \; yr^{-1}$$

* from RSS [2]

TABLE 3.2

Case II: Risk of Immediate Death to Individual Originally Located 1 to 1 1/2 Mile from Reactor

Release Category*	Probability (P_i) Dose ≥ 200 Rem	Frequency (F_L) of Release (yr^{-1})*	P_iF_i
1A	0.029	$4x10^{-7}$	$1.1x10^{-8}$
1B	0.013	$5x10^{-7}$	$6.7x10^{-9}$
2	0.015	$8x10^{-6}$	$1.2x10^{-7}$
3	0.032	$4x10^{-6}$	$1.3x10^{-7}$
4	0.029	$5x10^{-7}$	$1.4x10^{-8}$
5	0.026	$7x10^{-7}$	$1.8x10^{-8}$
6	0	$6x10^{-6}$	0
7	0	$4x10^{-5}$	0
8	0	$4x10^{-5}$	0
9	0	$4x10^{-4}$	0

$$R = \sum_i P_iF_i = 3.0x10^{-7}\ yr^{-1}$$

* from RSS [2]

Table 3.3

Latent risk to individual given that the acute risk is 10^{-5}/yr. The assumed cut-off dose for early effects is 200 Rem.

α	ν (Rem)	R^{ℓ} yr^{-1}	$\frac{\varphi}{\nu^{\alpha-1}}$ yr^{-1}
1.99	1	1.2×10^{-6}	2×10^{-3}
1.99	10	6.0×10^{-7}	2×10^{-4}
1.99	0.005	3.8×10^{-5}	0.4
1.01	1	2.0×10^{-9}	1×10^{-7}
1.01	10	1.9×10^{-9}	1×10^{-7}
1.01	0.005	2.0×10^{-9}	1×10^{-7}
1.5	1	1.9×10^{-7}	7×10^{-5}
1.5	10	1.6×10^{-7}	2×10^{-5}
1.5	0.005	3×10^{-7}	1×10^{-3}

for an assumed cut-off dose for early effects of 350 Rem

α	ν (Rem)	R^{ℓ} yr^{-1}	$\varphi/\nu^{\alpha-1}$ yr^{-1}
1.99	1	2.3×10^{-6}	3×10^{-3}
1.99	.005	6.6×10^{-6}	0.6

$$\mathcal{R} \equiv -\int_{x_{min}}^{x_{max}} x \frac{\partial N(X \geq x)}{\partial x} dx$$

$$E_c(\alpha) = -\int_{x_{min}}^{x_{max}} x^{\alpha} \frac{\partial N(X \geq x)}{\partial x} dx$$

$$E_c(\alpha) \leq 10 \frac{\text{deaths}}{\text{yr.}}$$

$$\frac{\text{Cost } (\$)}{\Delta E_c(\alpha) \cdot Y} \leq \frac{\$ 2 \times 10^6}{\text{death averted}}$$

$$C \leq \$ 2 \times 10^6 \, \Delta E_c(\alpha) \cdot Y \equiv C_\alpha$$

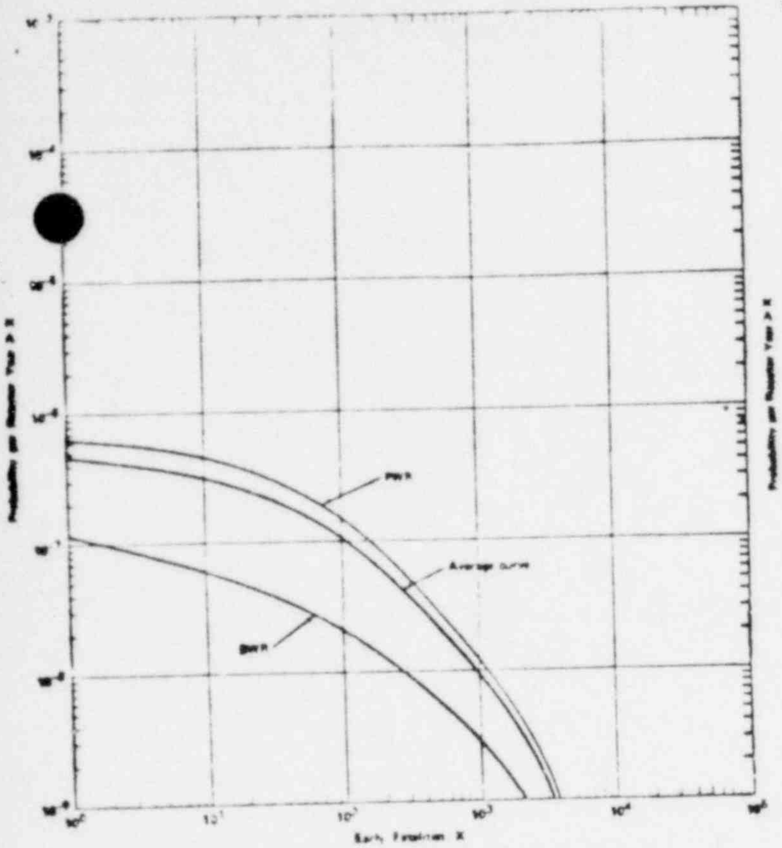

FIGURE 4.1 Probability Distribution for Early Fatalities per Reactor Year

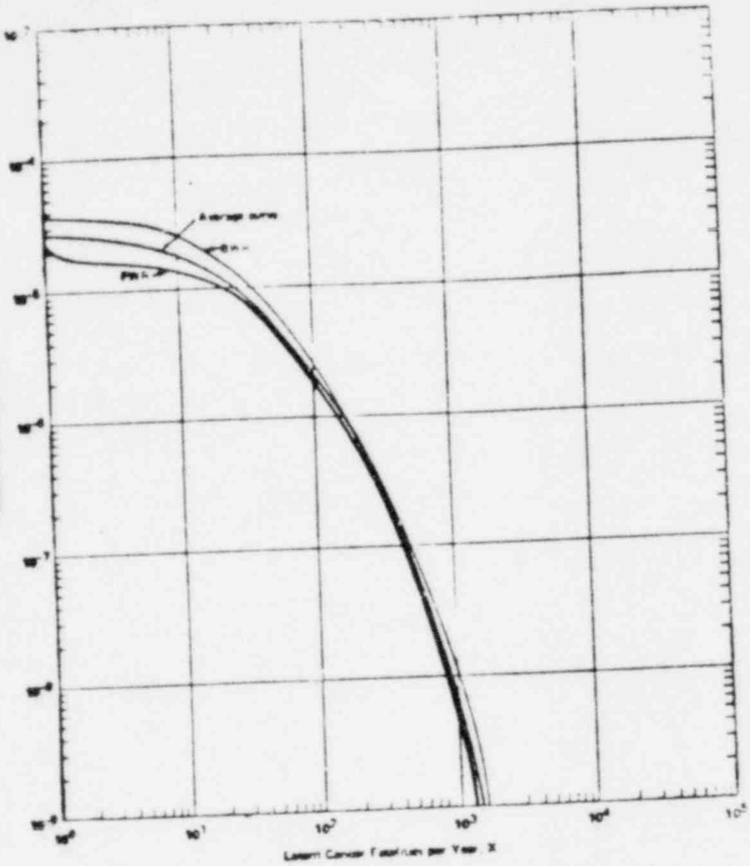

FIGURE 4.2 Probability Distribution for Latent Cancer Fatality Incidence per Reactor Year

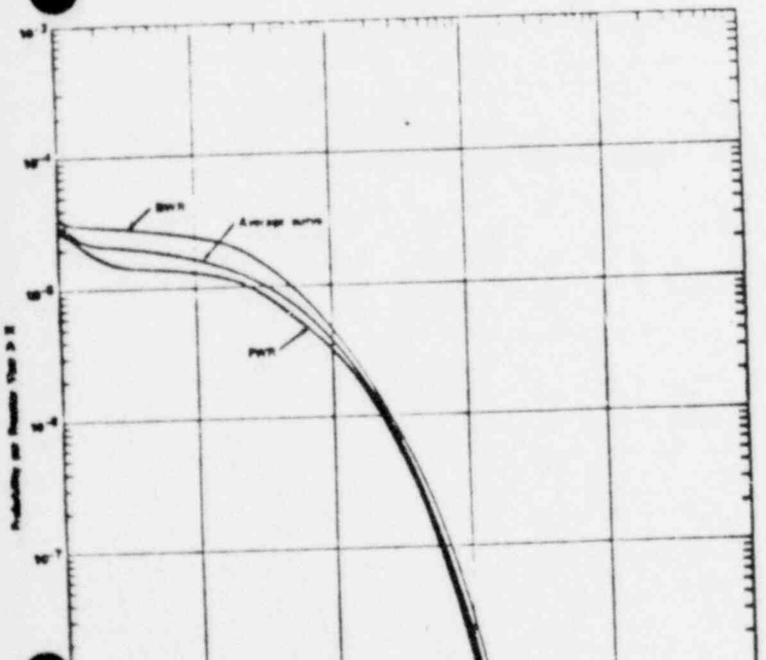

FIGURE 4.3 Probability Distribution for Incidence of Genetic Effects per Reactor Year

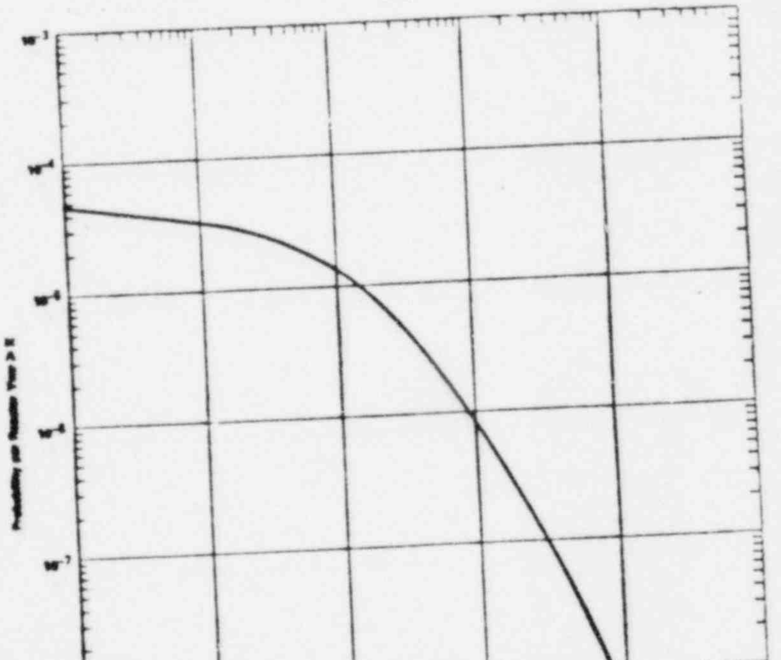

FIGURE 4.4 Probability Distribution for Property Damage per Reactor Year

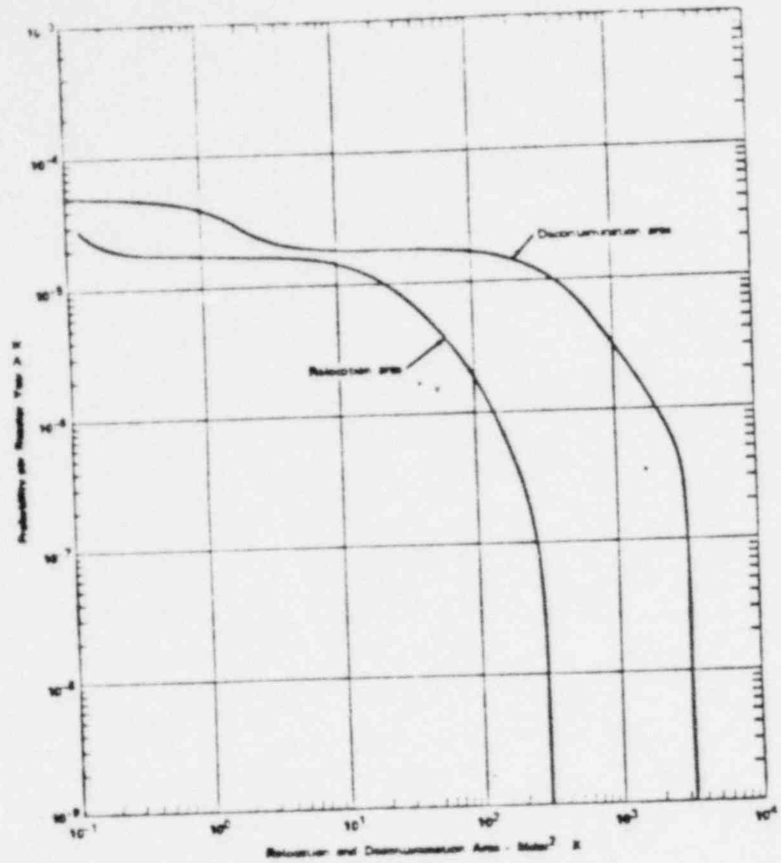

FIGURE 4.5 Probability Distribution for Relocation and Decontamination Area per Reactor Year

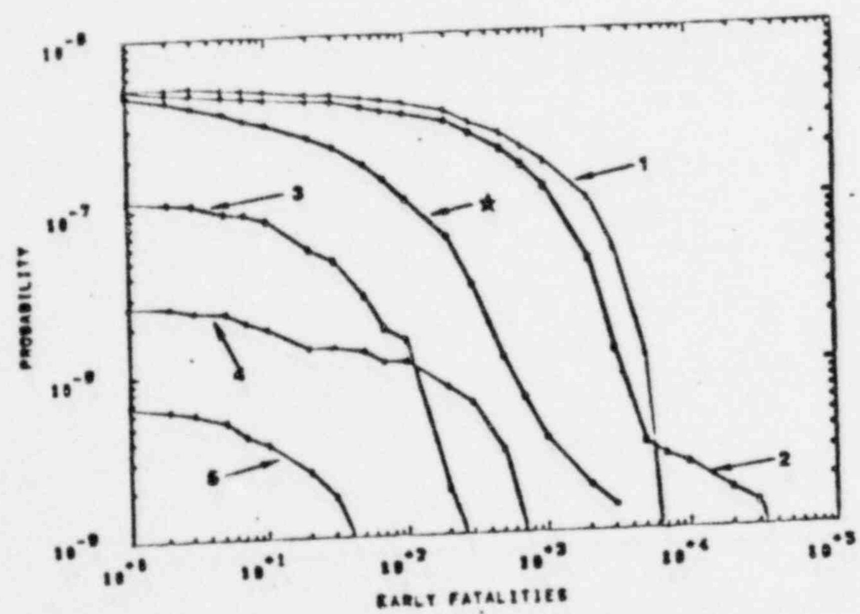

FIGURE 4.6: Log-Log plot of probability (per reactor-year versus early fatalities showing the dispersion of site specific CCDF's about the Reactor Safety Study CCDF.

1 = Indian Point (2985 MWt)
2 = Zion (3150 MWt)
3 = Palo Verde (3713 MWt)
4 = Millstone BWR (1956 MWt)
5 = San Onofre (1290 MWt)
★ = Reactor Safety Study

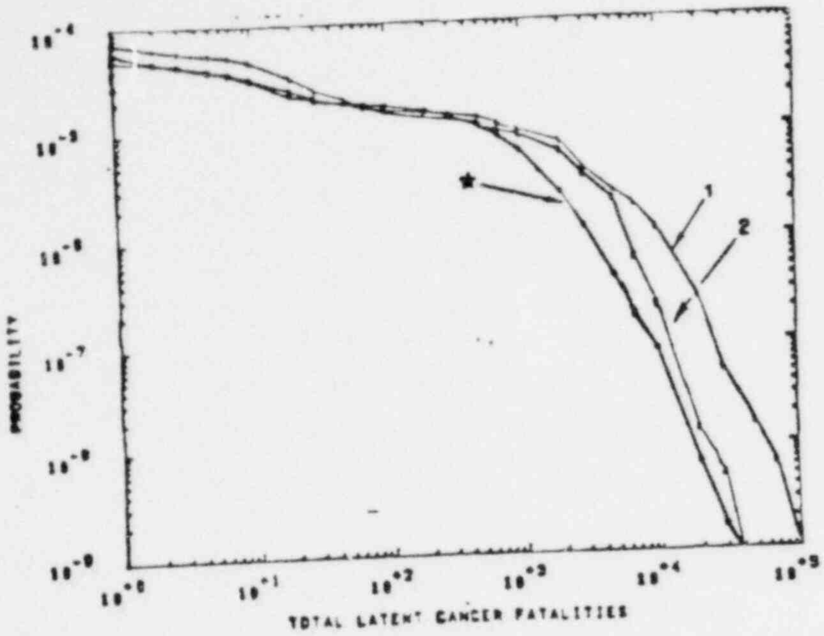

* REACTOR SAFETY STUDY
1. SURRY AT THE INDIAN POINT SITE
2. SURRY AT THE ZION SITE

FIGURE 4.7

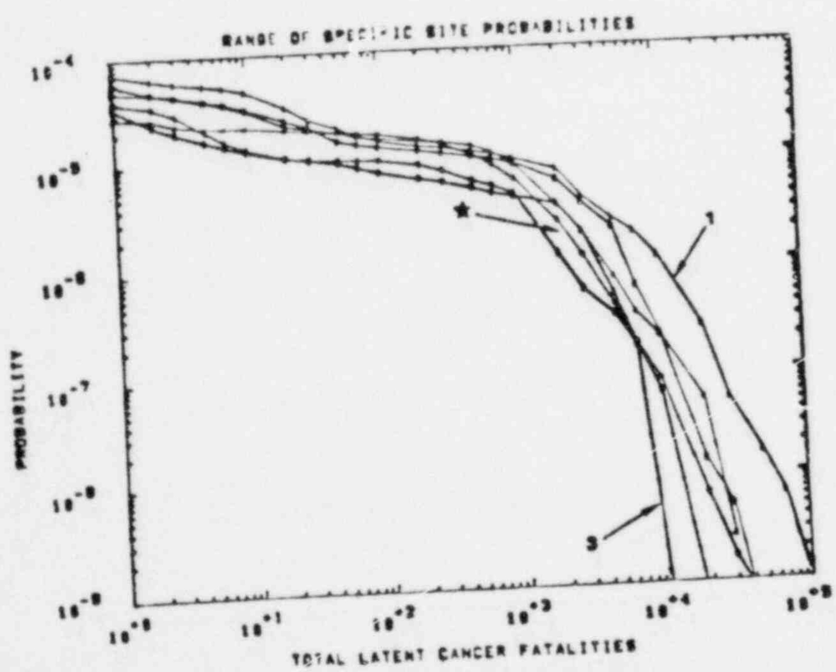

FIGURE 4.8: Log-Log plot of probability (per reactor-year) versus total latent cancer fatalities showing the dispersion of site specific CCDF's about the Reactor Safety Study CCDF.
1 = Indian Point (2895 MWt) 2 = Zion (3150 MWt)
3 = Palo Verde (3713 MWt) 4 = Millstone BWR (1956 MWt)
5 = San Onofre (1290 MWt) * = Reactor Safety Study

TABLE 4.1

EXPECTED VALUE OF THE RISK FOR THE

RESULTS PRESENTED IN WASH-1400

Consequence	Risk
Early Fatalities (average curve)	4.4×10^{-5} deaths/year
Latent Cancer Fatalities (total over 30 years) (average curve)	2.7×10^{-2} deaths/year
Genetic Effects (total over 30 years)	3.6×10^{-3} genetic effects/year
Property Damage	$\$5.5 \times 10^{3}$ \$/year
Relocation Area	7.5×10^{-4} mi^2/year
Decontamination Area	9×10^{-3} mi^2/year

TABLE 4.2

EXPECTED VALUE OF THE RISK FOR THE SURRY REACTOR PLACED ON SELECTED SITES AND SCALED TO THE POWER OF THE REACTOR AT THAT SITE

CURVE	RISK (DEATHS/YEAR) EARLY	RISK (DEATHS/YEAR) LATENT (TOTAL)
RSS (WASH-1400)*	4.4×10^{-5}	2.7×10^{-2}
RSS (FIG. 4.5)	4.0×10^{-5}	2.1×10^{-2}
San Onofre	4.9×10^{-8}	NOT AVAIL.
Indian Point	4.7×10^{-4}	4.96×10^{-2}
Zion	2.5×10^{-4}	3.5×10^{-2}
Millstone	3.0×10^{-6}	NOT AVAIL.
Palo Verde	4.3×10^{-6}	NOT AVAIL.

---

* This difference is attributed to updating the consequence modelling in the computer codes used to generate the curves.

TABLE 4.3

EXPECTED SOCIAL COST (RISK) OF EARLY DEATHS/YEAR AS A FUNCTION OF RISK AVERSION FOR 5 HYPOTHETICAL CASES COMPARED TO THE RSS (FROM FIG. 4.6)

CURVE	EXPECTED SOCIAL COST - RISK (Ec) (early deaths/yr)		
	$\alpha = 1$	$\alpha = 1\frac{1}{2}$	$\alpha = 2$
* RSS (FIG. 4.6)	4.0×10^{-5}	7.38×10^{-4}	2.1×10^{-2}
5 San Onofre	4.9×10^{-8}	2.7×10^{-7}	1.25×10^{-6}
1 Indian Point	4.7×10^{-4}	2.3×10^{-2}	1.2
2 Zion	2.5×10^{-4}	1.4×10^{-2}	1.3
4 Millstone	3.0×10^{-6}	4.7×10^{-5}	8.0×10^{-4}
3 Palo Verde	4.3×10^{-6}	4.1×10^{-5}	4.5×10^{-4}

TABLE 4.4

EXPECTED SOCIAL COST (RISK) OF TOTAL LATENT DEATHS/YEAR AS A FUNCTION OF RISK AVERSION FOR TWO HYPOTHETICAL CASES COMPARED TO THE RSS (FROM FIG. 4.7)

CURVE	EXPECTED SOCIAL COST - RISK (Ec) (latent deaths/yr)		
	$\alpha = 1$	$\alpha = 1\frac{1}{2}$	$\alpha = 2$
* RSS (FIG. 4.6)	2.06×10^{-2}	0.89	53.0
1 Surry at I.P. (1)	4.96×10^{-2}	4.01	410
2 Surry at Zion (2)	3.5×10^{-2}	2.12	149

TABLE 4.5

COST PER DEATH AVERTED AS A FUNCTION OF RISK AVERSION FOR A $1 MILLION AND A $10 MILLION IMPROVEMENT OVER 30 YEARS

COST OF IMPROVEMENT	CURVE (1) TO RSS*		
CASE	$\alpha = 1$	$\alpha = 1\frac{1}{2}$	$\alpha = 2$
I. $1 million	1.15×10^6	$10,000	$93.00
II. $10 million	$\$1.15 \times 10^7$	$106,000	$933.00

COST OF IMPROVEMENT	CURVE (2) TO RSS*		
	$\alpha = 1$	$\alpha = 1\frac{1}{2}$	$\alpha = 2$
III. $1 million	$\$2.31 \times 10^6$	$27,100	$347.00
IV. $10 million	$\$2.31 \times 10^7$	$271,000	$3,472

* FIGURE 4.7

TABLE 4.6

MAXIMUM AMOUNT OF EXPENDITURE FOR A $\$2 \times 10^6$ COST PER DEATH AVERTED OVER 30 YEARS AS A FUNCTION OF RISK AVERSION

CASE*	MAXIMUM EXPENDITURE ($)		
	$\alpha = 1$	$\alpha = 1\frac{1}{2}$	$\alpha = 2$
CURVE (1) TO RSS	1.7×10^6	1.9×10^8	2.14×10^{10}
CURVE (2) TO RSS	0.9×10^6	0.70×10^8	0.56×10^{10}

* Figure 4.7

TABLE 4.7

COST PER DEATH AVERTED AS A FUNCTION OF UNCERTAINTY FOR A $1 MILLION AND A $10 MILLION IMPROVEMENT OVER 30 YEARS

CASE	COST OF IMPROVEMENT	CURVE (1) RSS* UNCERTAINTY IN RISK		
		0	+10	+100
I.	$1 Million	$\$1.5 \times 10^6$	1.15×10^5	1.15×10^4
II	$10 Million	$\$1.5 \times 10^7$	1.15×10^6	1.15×10^5

	COST OF IMPROVEMENT	CURVE (2) RSS* UNCERTAINTY IN RISK		
		0	+10	+100
III.	$1 Million	$\$2.31 \times 10^6$	2.31×10^5	2.32×10^4
IV.	$10 Million	$\$2.31 \times 10^7$	2.31×10^6	2.31×10^5

* Figure 4.7

TABLE 4.8

MAXIMUM AMOUNT OF EXPENDITURE FOR A $\$2 \times 10^6$ COST PER DEATH AVERTED OVER 30 YEARS AS A FUNCTION OF UNCERTAINTY

CASE*	MAXIMUM EXPENDITURE (Uncertainty in Risk)		
	0	+10	+100
CURVE (1)→RSS	1.7×10^6	1.7×10^7	1.7×10^8
CURVE (2)→RSS	0.9×10^6	0.9×10^7	0.9×10^8

* Figure 4.7

ASSUMPTIONS FOR COAL RISK STUDY

1. HEALTH EFFCTS BASED UPON GROUND LEVEL POPULATION EXPOSURES TO SULFATE CONCENTRATION

2. EFFECTS OF NITRATES, HEAVY METALS and HEAVY METAL SULFITES IGNORED

3. LINEAR DAMAGE FUNCTION

4. LONG RANGE TRANSPORT OF SULFATE IGNORED

5. 1000 MWe plant @ 38% THERMAL EFFIC.
   75% Plant Factor – 3% Sulfur Coal.

6. 305 Meter Stack (8.2 m dia)
   16 m/s exit velocity

7. 4.15 kg/sec total sulfur effluent rate.

8. Meteorological Data from Pitts. Int'l Airport.

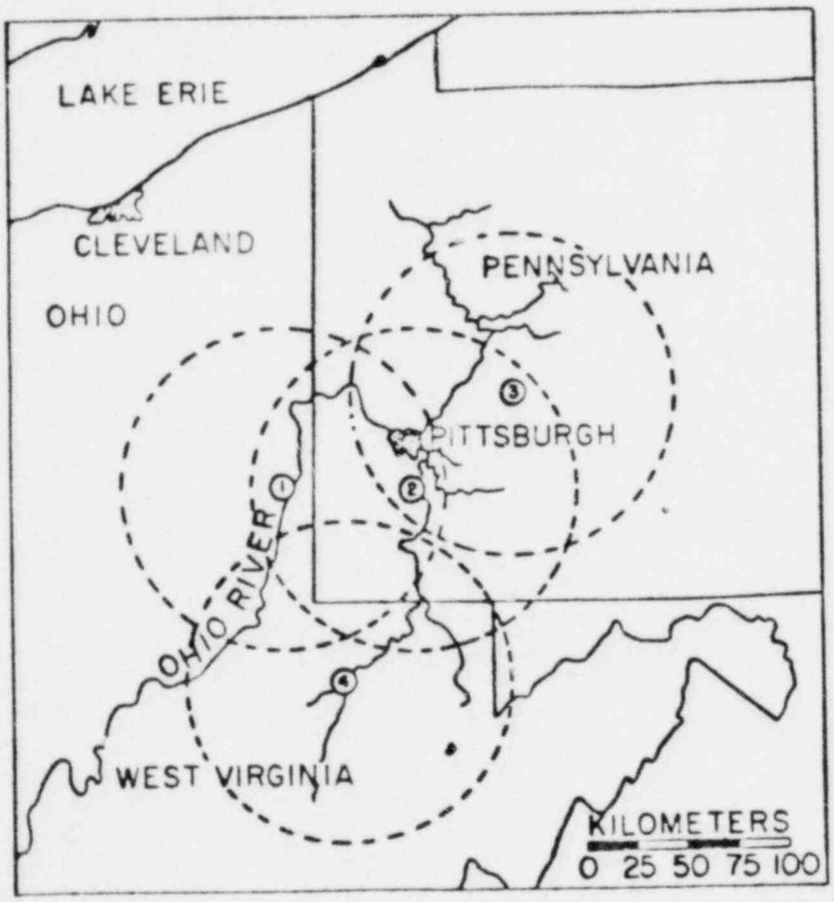

Figure 4. Locations used as the sites of the four identical hypothetical 1000-MW(e) power plants in our quantitative example.

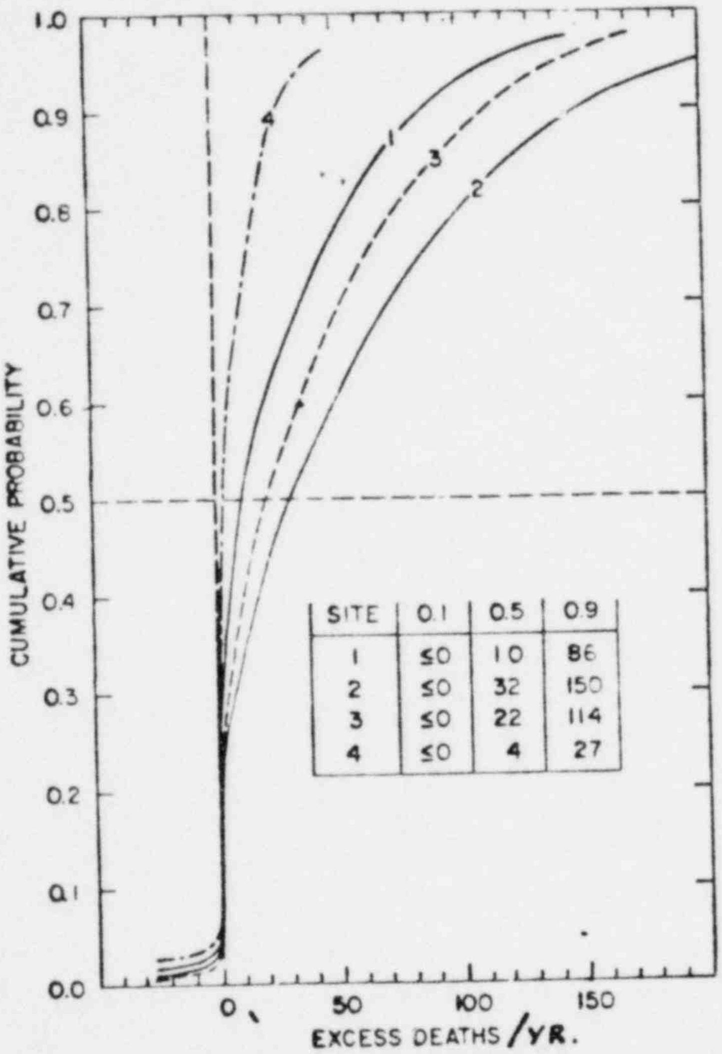

SITE	0.1	0.5	0.9
1	≤0	10	86
2	≤0	32	150
3	≤0	22	114
4	≤0	4	27

Cumulative distributions for the probability density functions shown in Figure 8. Plants are uncontrolled hypothetical 1000-MW(e) coal-fired plants burning 3% sulfur coal, located as shown in Figure 4. A first-order estimate of the effect of control can be obtained by multiplying the horizontal axis by $(1-\eta)$, where η is the fractional sulfur removal efficiency. Limiting assumptions and plant operating parameters are provided in the text.

FIGURE 5.1

TABLE 5.1

SOCIETAL AND INDIVIDUAL RISK FOR
A 1000 MWe COAL PLANT ON FOUR
PENNSYLVANIA SITES

SITE*	1	2	3	4
Population in 80 km. radius.	2.8×10^6	3.3×10^6	2.9×10^6	6.2×10^5
SOCIETAL RISK (deaths/yr.)				
Expected Value	24	49	34	10
Upper 90% Confidence Bound	86	150	114	27
AVERAGE INDIVIDUAL RISK ($yr.^{-1}$)				
Expected Value	0.86×10^{-5}	1.5×10^{-5}	1.2×10^{-5}	1.6×10^{-5}
Upper 90% Confidence Bound	3.1×10^{-5}	4.5×10^{-5}	3.9×10^{-5}	4.3×10^{-5}

*FIGURE 5.1

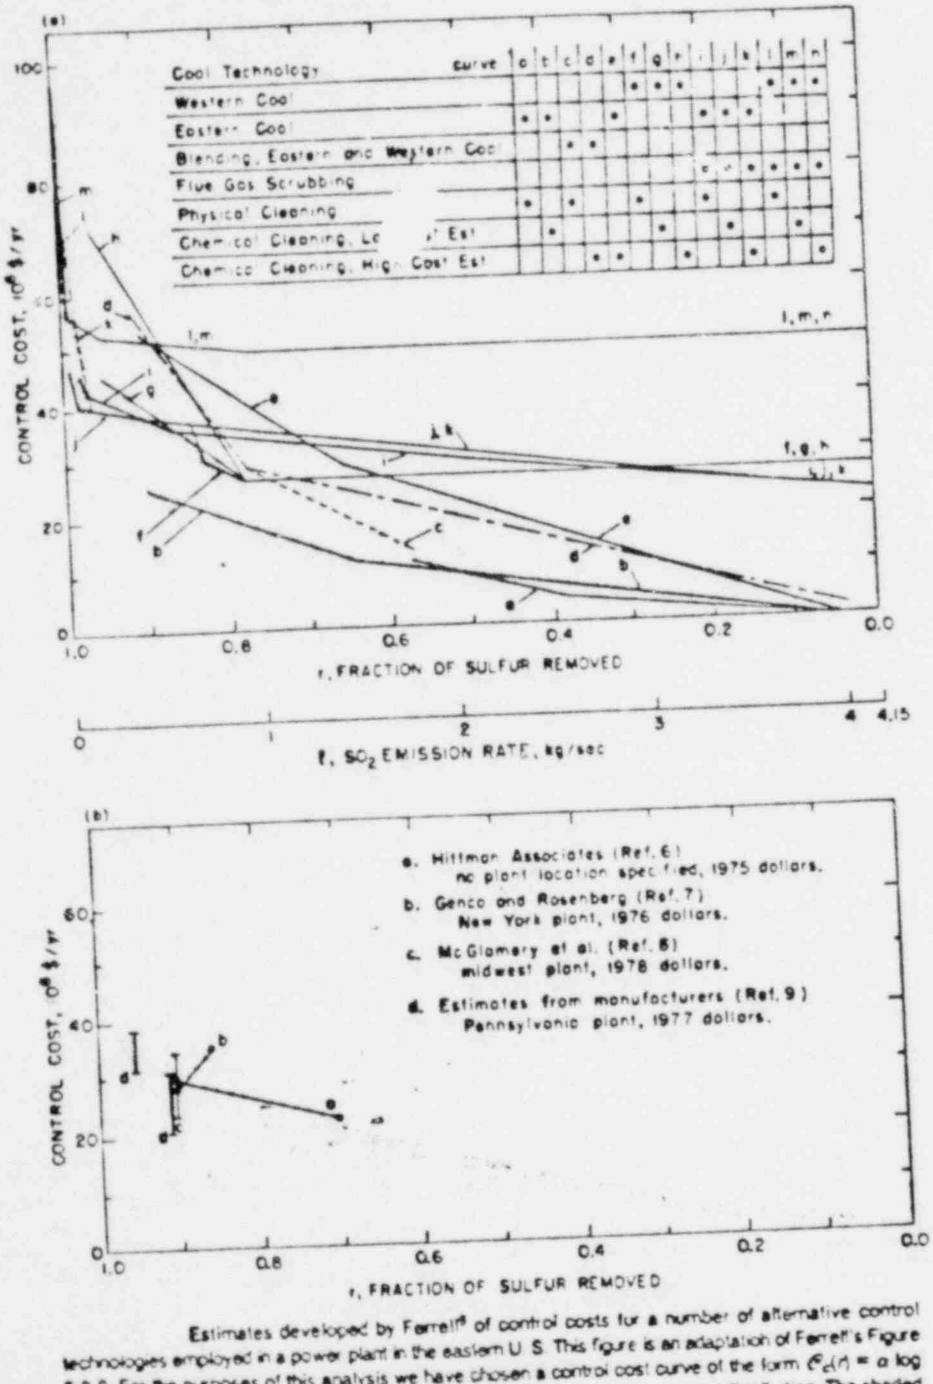

Estimates developed by Ferrell[8] of control costs for a number of alternative control technologies employed in a power plant in the eastern U.S. This figure is an adaptation of Ferrell's Figure 6.2-9. For the purposes of this analysis we have chosen a control cost curve of the form $\mathscr{C}_c(r) = a \log 1/(1-r)$. Our uncertainty in the parameter a has been represented by a normal distribution. The shaded region corresponds to $24 \leq a \leq 42$, which represents the 90% confidence interval for our subjective estimate of a. B. comparison of the 90% confidence interval of the set of control cost curves used in this analysis with a number of recent published and unpublished estimates of such costs.

FIGURE 5.2

TABLE 5.2

FRACTIONAL REDUCTION, COST AND COST PER DEATH AVERTED FOR VARIOUS SULFUR CONTROL STRATEGIES

STRATEGY*	FRACTION OF SULFUR REMOVED	COST ($/yr)	COST PER DEATH AVERTED ($/death averted)
		AT EXPECTED VALUE	
1 ⟶ 4	0.58	$15x10^6$	$1.07x10^6$
3 ⟶ 4	0.71	$18x10^6$	$0.75x10^6$
2 ⟶ 4	0.80	$22x10^6$	$0.56x10^6$
		AT 90% UPPER CONFIDENCE BOUND	
1 ⟶ 4	0.68	$18x10^6$	$0.30x10^6$
3 ⟶ 4	0.76	$20x10^6$	$0.22x10^6$
2 ⟶ 4	0.82	$23x10^6$	$0.19x10^6$

* Reducing the societal risk from the plant at sites 1, 2 and 3 to that of site 4.

TABLE 5.3

FRACTIONAL REDUCTION, COST AND COST
PER DEATH AVERTED FOR PLANT #4
FOR VARIOUS SULFUR CONTROL STRATEGIES

DEATHS/YR AVERTED	FRACTION OF SULFUR REMOVED	COST ($/yr.)	COST PER DEATH AVERTED ($/death averted)
	AT EXPECTED VALUE*		
4	.40	$8x10^6$	$2.0x10^6$
5	.50	$10x10^6$	$2.0x10^6$
7	.70	$18x10^6$	$2.5x10^6$
9	.90	$38x10^6$	$4.2x10^6$
	AT 90% UPPER CONFIDENCE BOUND*		
8.3	.31	$5x10^6$	$0.60x10^6$
13.5	.50	$10x10^6$	$0.75x10^6$
20	.75	$20x10^6$	$1.0x10^6$
25	.92	$40x10^6$	$1.6x10^6$
26	.96	$52x10^6$	$2.0x10^6$

* Expected value is 10 deaths/yr, at 90% upper confidence bound the risk is 27 deaths/yr.

TABLE 5.4

FRACTIONAL REDUCTION, COST AND COST PER DEATH AVERTED FOR VARIOUS SULFUR CONTROL STRATEGIES

STRATEGY*	FRACTION OF SULFUR REMOVED	COST ($/yr)	COST PER DEATH AVERTED ($/death averted)
AT EXPECTED VALUE			
1 → 4 i	0.80	22×10^6	1.2×10^6
3 → 4 i	0.85	26×10^6	0.9×10^6
2 → 4 i	0.89	37×10^6	0.84×10^6
AT 90% UPPER CONFIDENCE BOUND			
1 → 4 i	.99	80×10^6	$.94 \times 10^6$
3 → 4 i	.99	80×10^6	$.70 \times 10^6$
2 → 4 i	.99	80×10^6	$.54 \times 10^6$

* Reducing the societal risk from the plant represented by curves 1, 2 and 3 to that of an improved curve 4. At expected value, the improved curve 4 is 5 deaths/year, at the 90% upper confidence bound the improved curve 4 risk is 1 death/year.

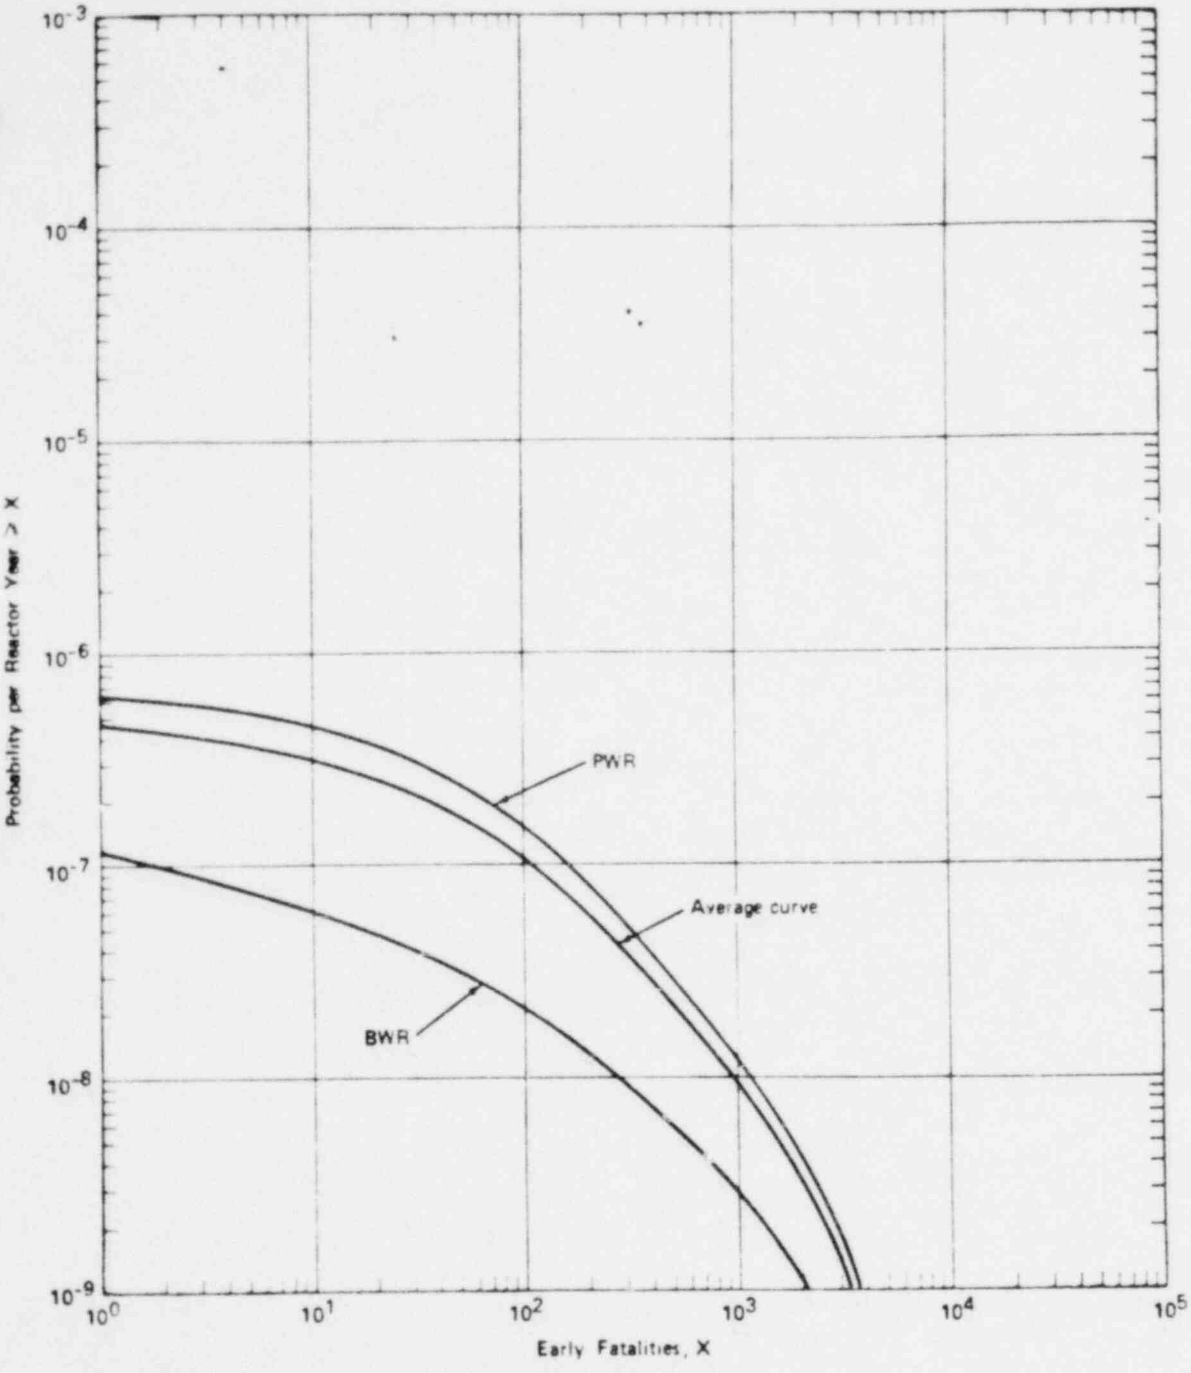

FIGURE 4.1 Probability Distribution for Early Fatalities per Reactor Year

Note: Approximate uncertainties are estimated to be represented by factors of 1/4 and 4 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.

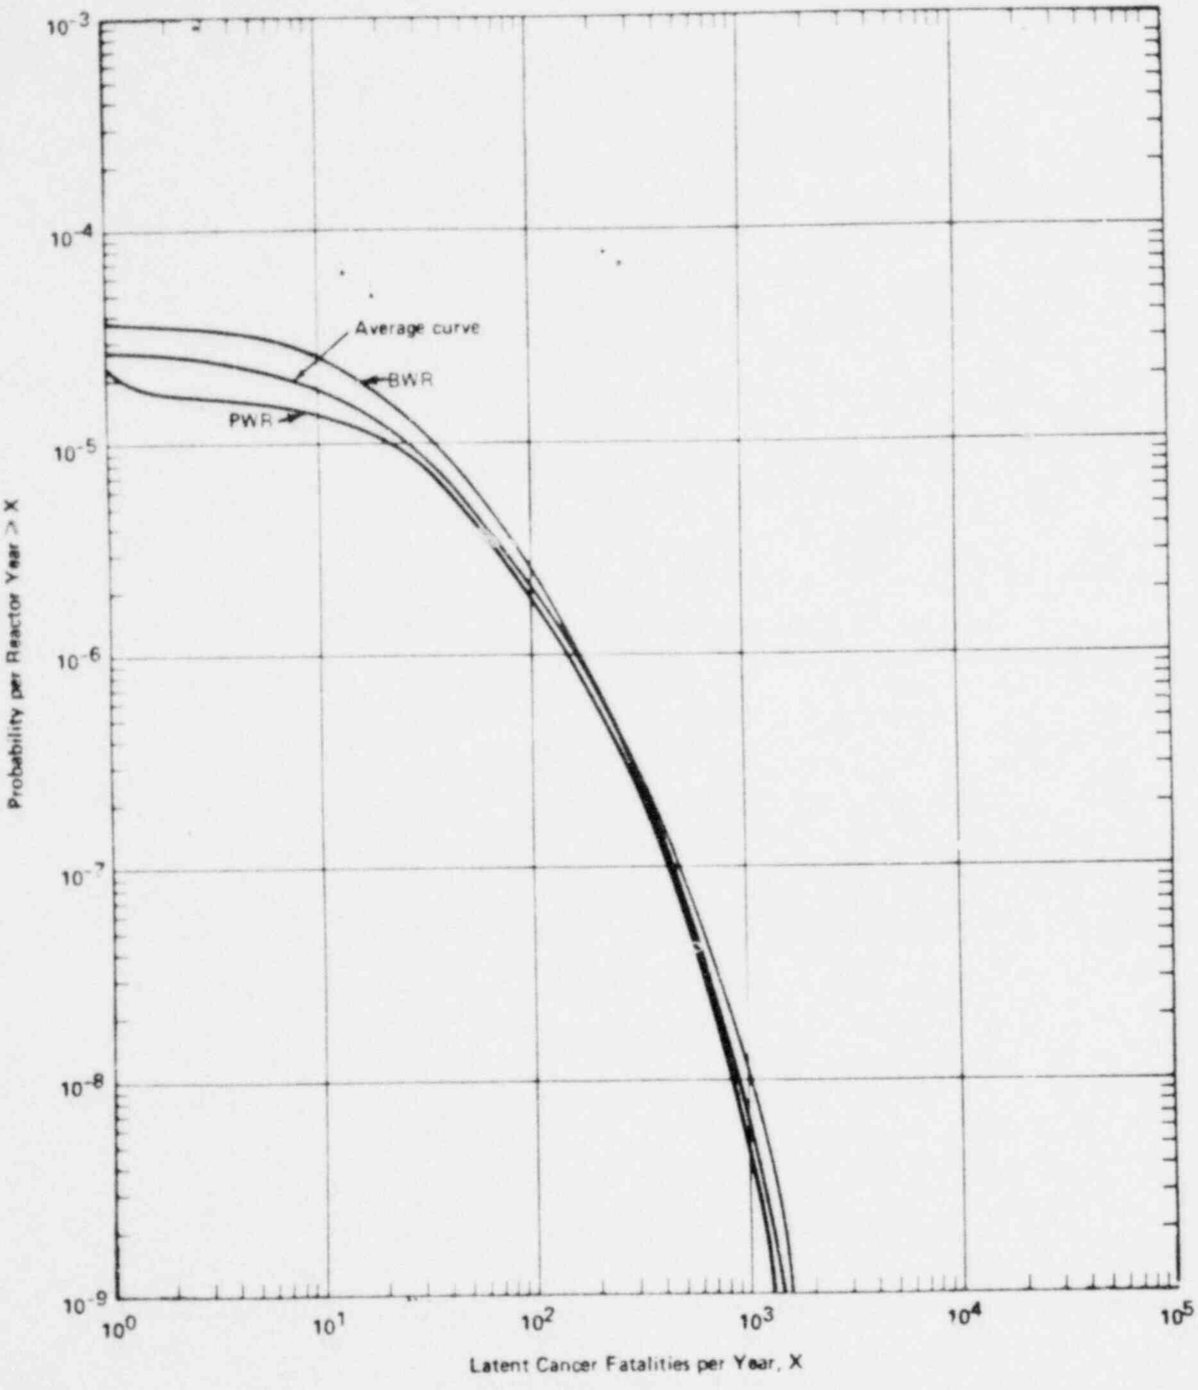

FIGURE 4.2 Probability Distribution for Latent Cancer Fatality Incidence per Reactor Year

Note: Approximate uncertainties are estimated to be represented by factors of 1/6 and 3 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.

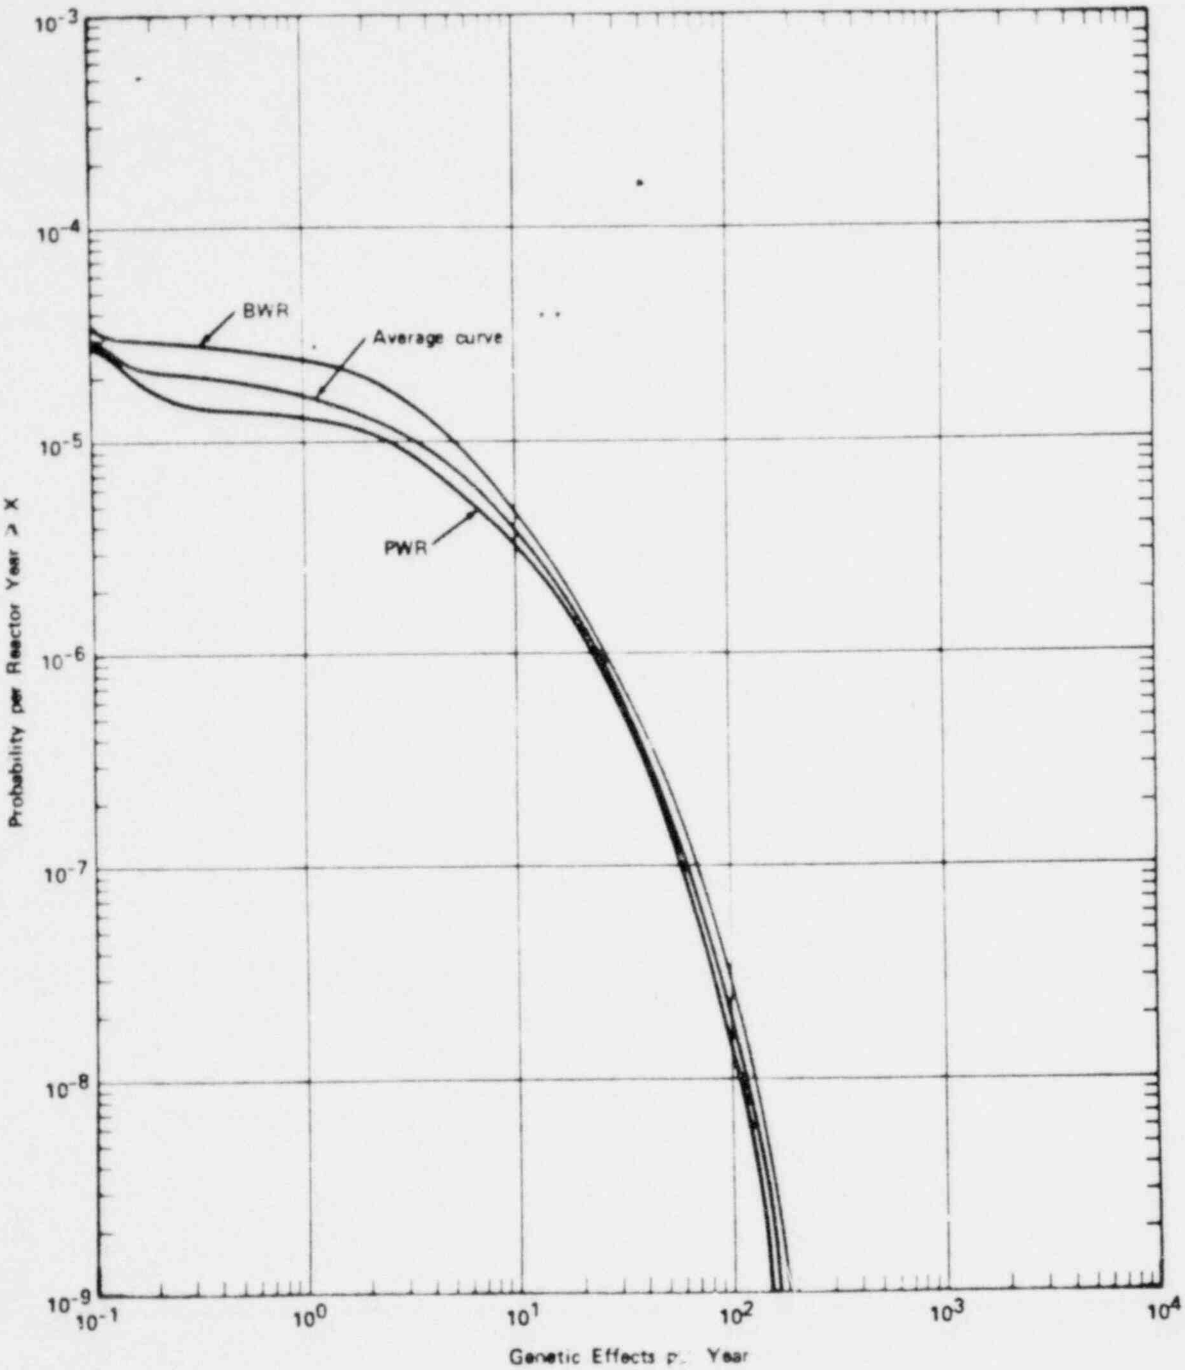

FIGURE 4.3 Probability Distribution for Incidence of Genetic Effects per Reactor Year

Note: Approximate uncertainties are estimated to be represented by factors of 1/3 and 6 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.

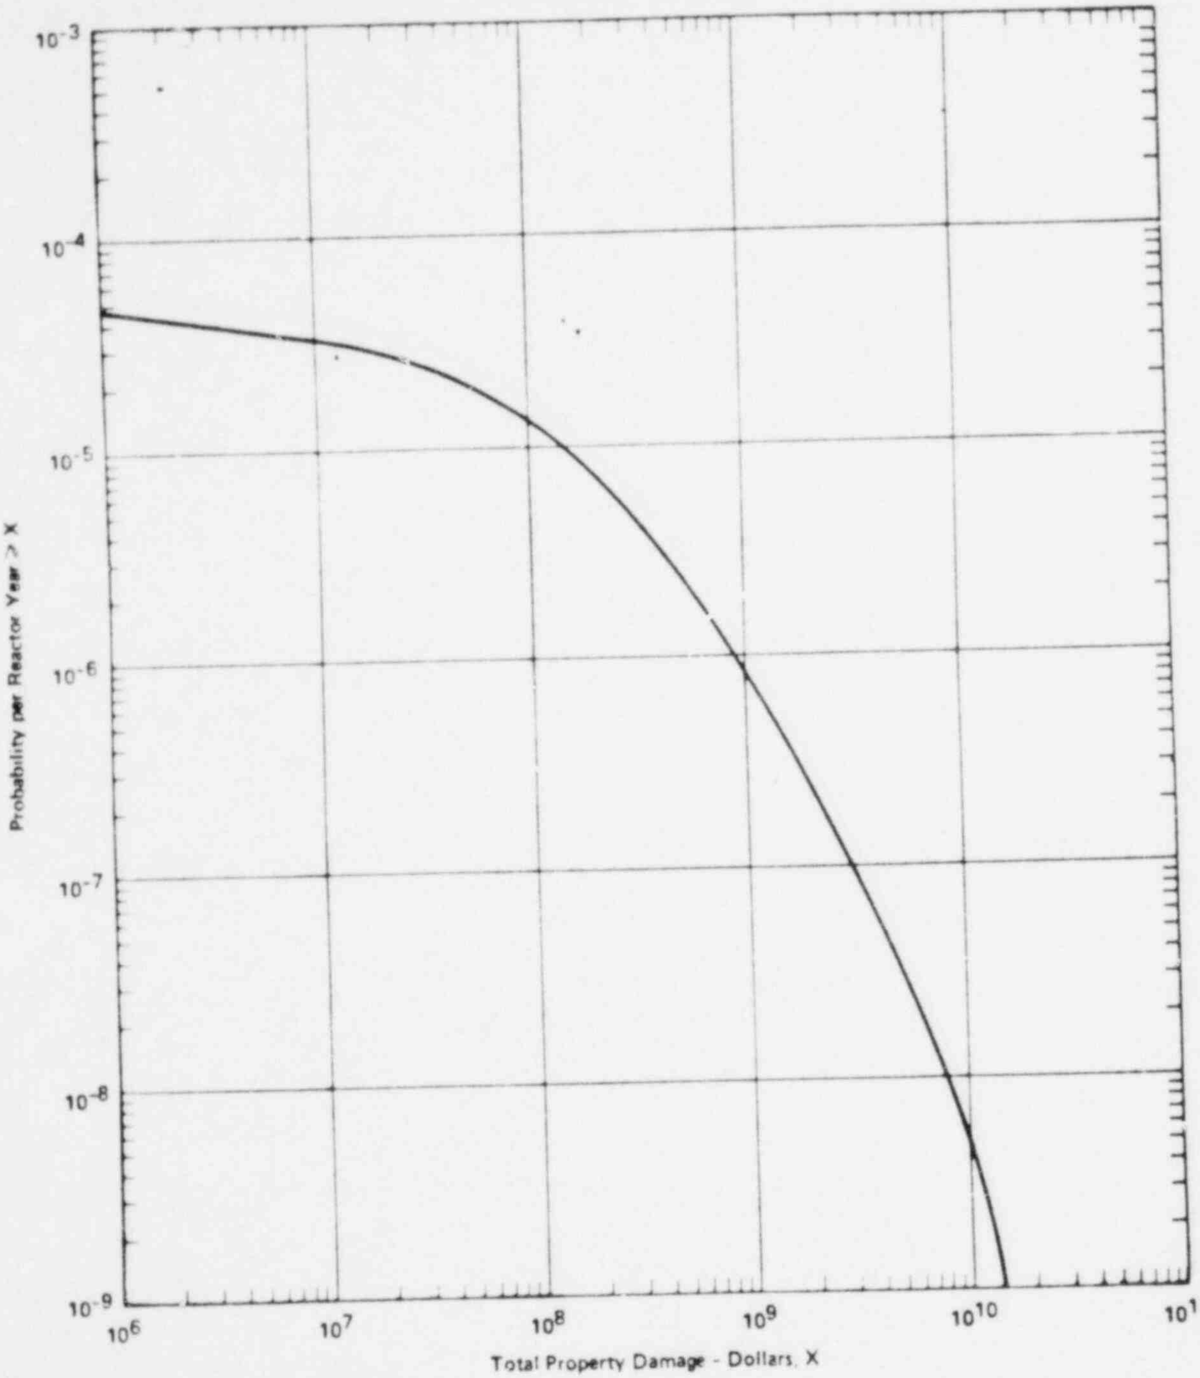

FIGURE 4.4 Probability Distribution for Property Damage per Reactor Year

Note: Approximate uncertainties are estimated to be represented by factors of 1/5 and 2 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.

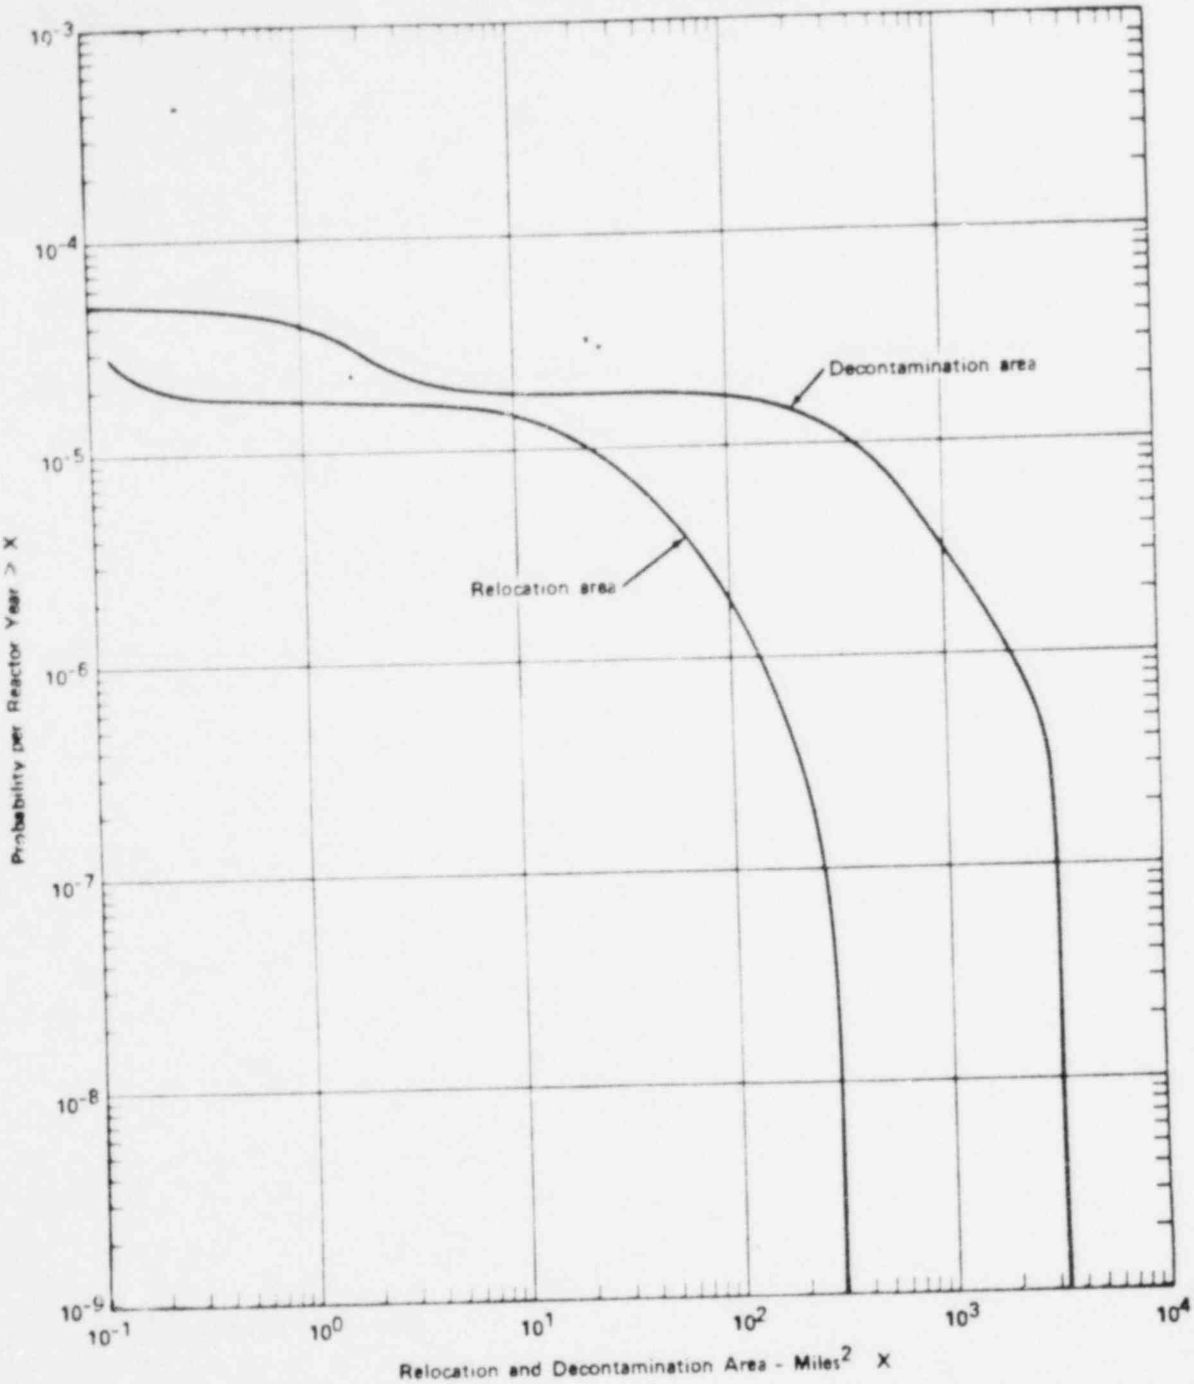

FIGURE 4.5 Probability Distribution for Relocation and Decontamination Area per Reactor Year

Note: Approximate uncertainties are estimated to be represented by factors of 1/5 and 2 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.

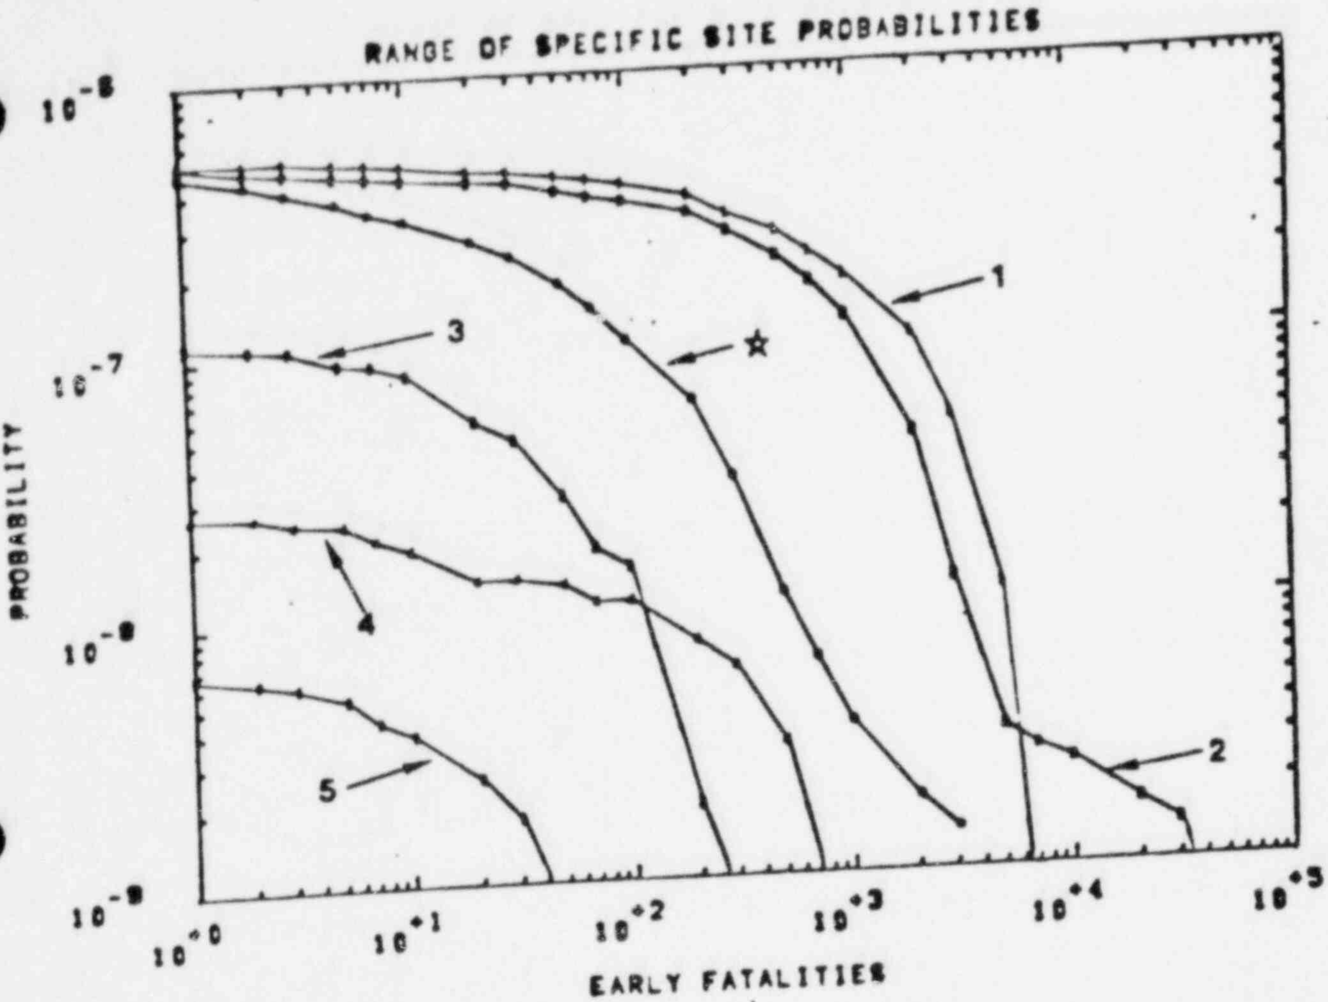

FIGURE 4.6: Log-Log plot of probability (per reactor-year versus early fatalities showing the dispersion of site specific CCDF's about the Reactor Safety Study CCDF.

1 = Indian Point (2985 MWt)
2 = Zion (3150 MWt)
3 = Palo Verde (3713 MWt)
4 = Millstone BWR (1956 MWt)
5 = San Onofre (1290 MWt)
☆ = Reactor Safety Study

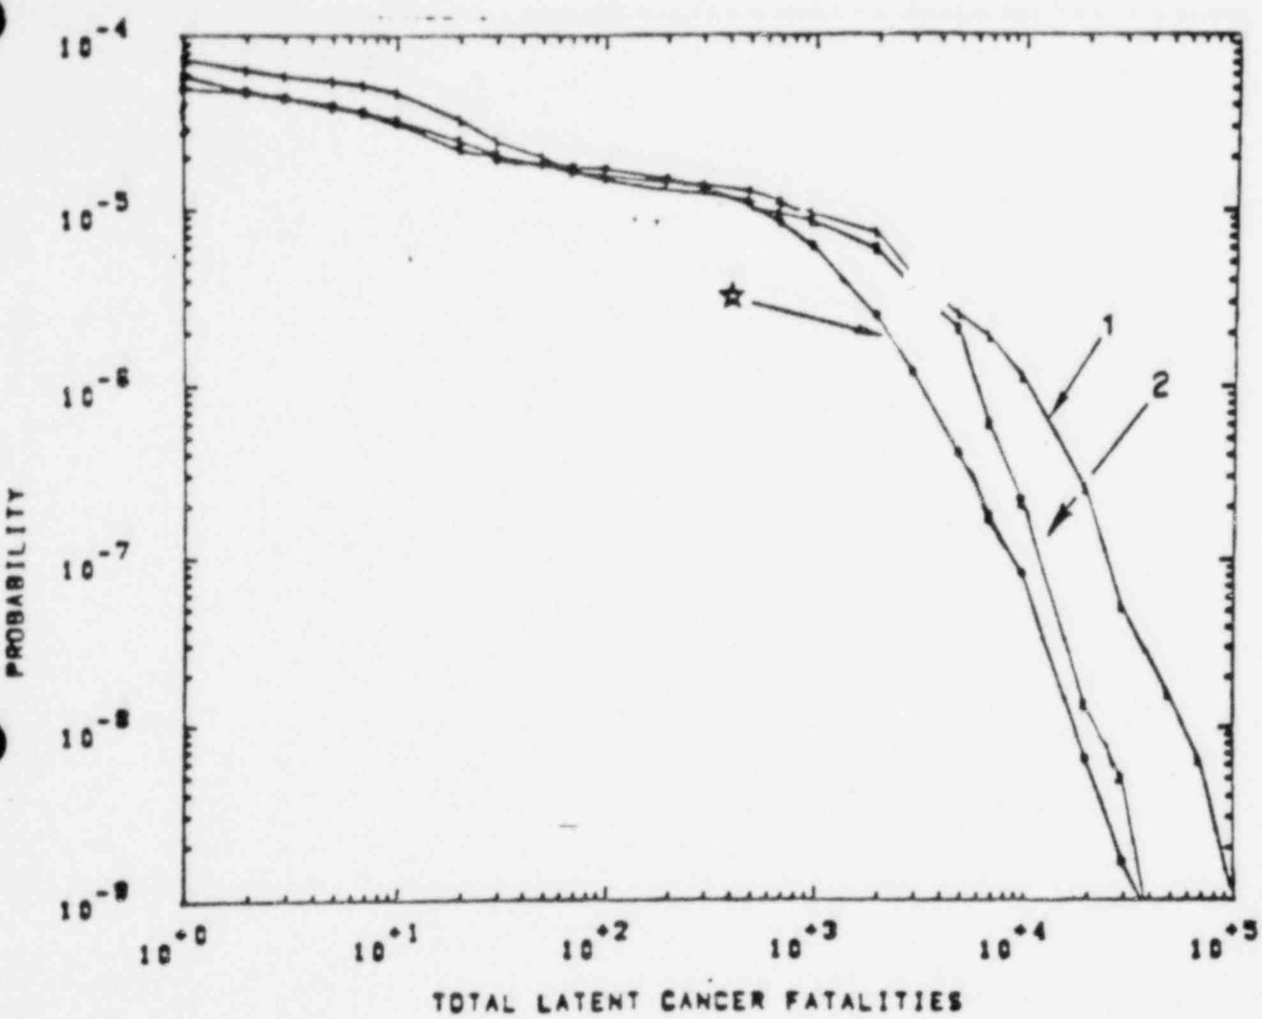

* REACTOR SAFETY STUDY

1. SURRY AT THE INDIAN POINT SITE
2. SURRY AT THE ZION SITE

FIGURE 4.7

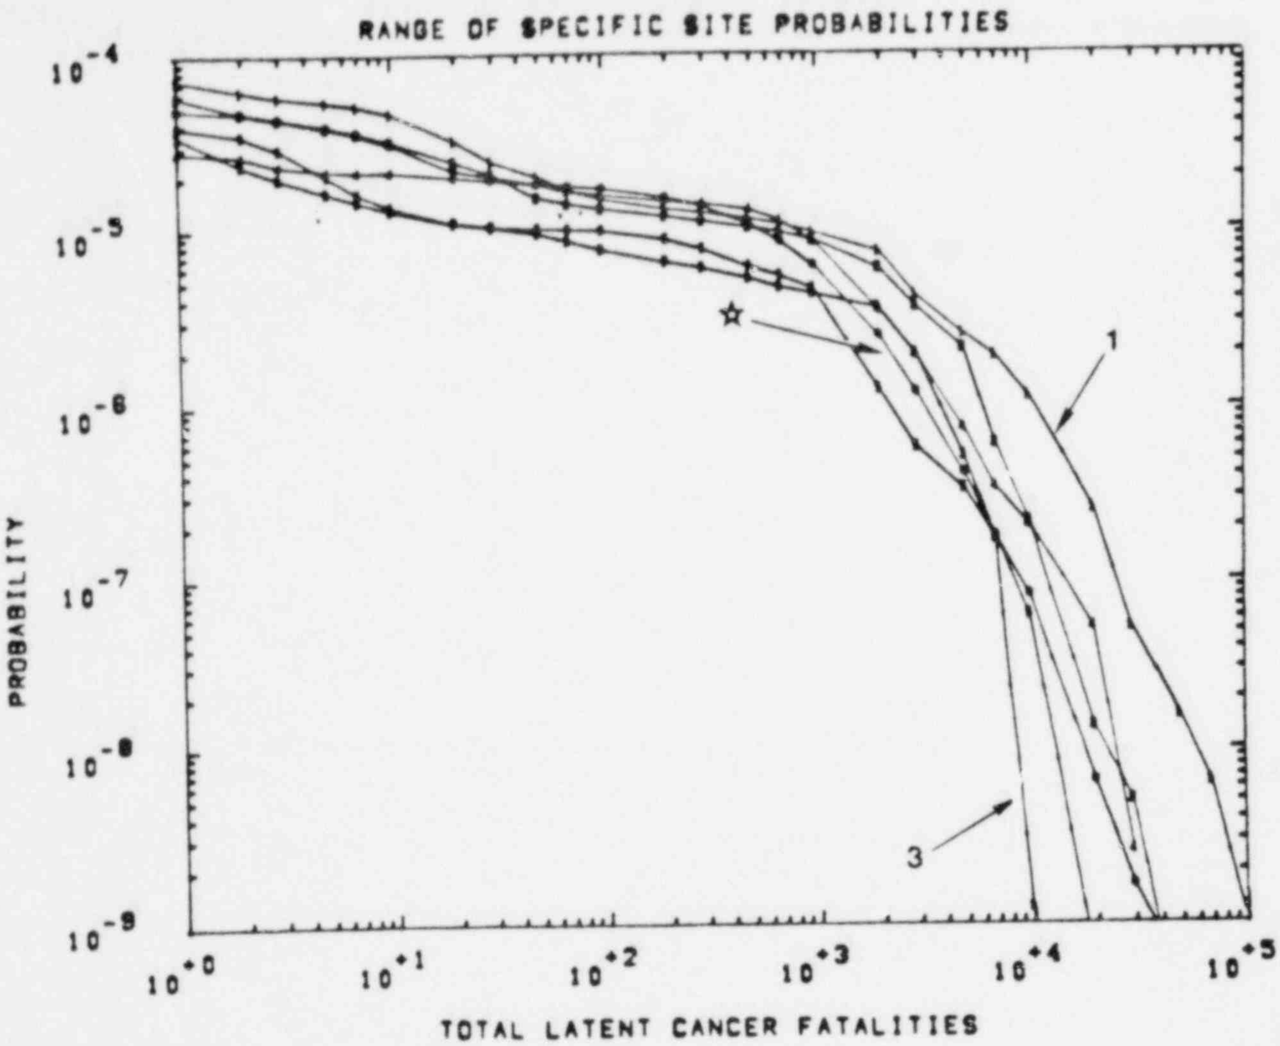

FIGURE 4.8: Log-Log plot of probability (per reactor-year) versus total latent cancer fatalities showing the dispersion of site specific CCDF's about the Reactor Safety Study CCDF.

1 = Indian Point (2895 MWt)
2 = Zion (3150 MWt)
3 = Palo Verde (3713 MWt)
4 = Millstone BWR (1956 MWt)
5 = San Onofre (1290 MWt)
* = Reactor Safety Study

~~POWERS~~ CORKS

CORE DEBRIS BEHAVIOR

PROGRAM OBJECTIVE: TO DEVELOP THE TECHNOLOGY ASSOCIATED WITH POST-ACCIDENT DEBRIS BEHAVIOR TO THE POINT THAT ADEQUATE ASSESSMENTS CAN BE MADE AS TO CONTAINMENT SUBSEQUENT TO A CORE DISRUPTION OR DEGRADATION

CORE DEBRIS BEHAVIOR

PROGRAM

COMPREHENSIVE

FORMATION ⇨ REMELT

EXPT. ⇨ ANALYSIS ⇨ MODELS

INTERNATIONAL

EURATOM FRANCE

JAPAN GERMANY

UK

CORE DEBRIS BEHAVIOR

REGIMES

FORMATION

SUB-DRYOUT

DRYOUT

POST-DRYOUT

STEEL MELT/MIGRATION

UO_2 MELT/MIGRATION,

VESSEL, CONTAINMENT ATTACK

CORE DEBRIS BEHAVIOR

Controlling Parameters

Initial Bed Characteristics

- Particle Size/Shape
- Particle Size Distribution
- Bed Geometry
- Stratification
- Presence of Steel
- Bed Depth
- Bed Packing

Cooling Available

- Through Flow
  - Natural
  - Forced
  - Gas Added
- U-Flow
  - Adiabatic Lower Boundary
  - Bottom Cooling

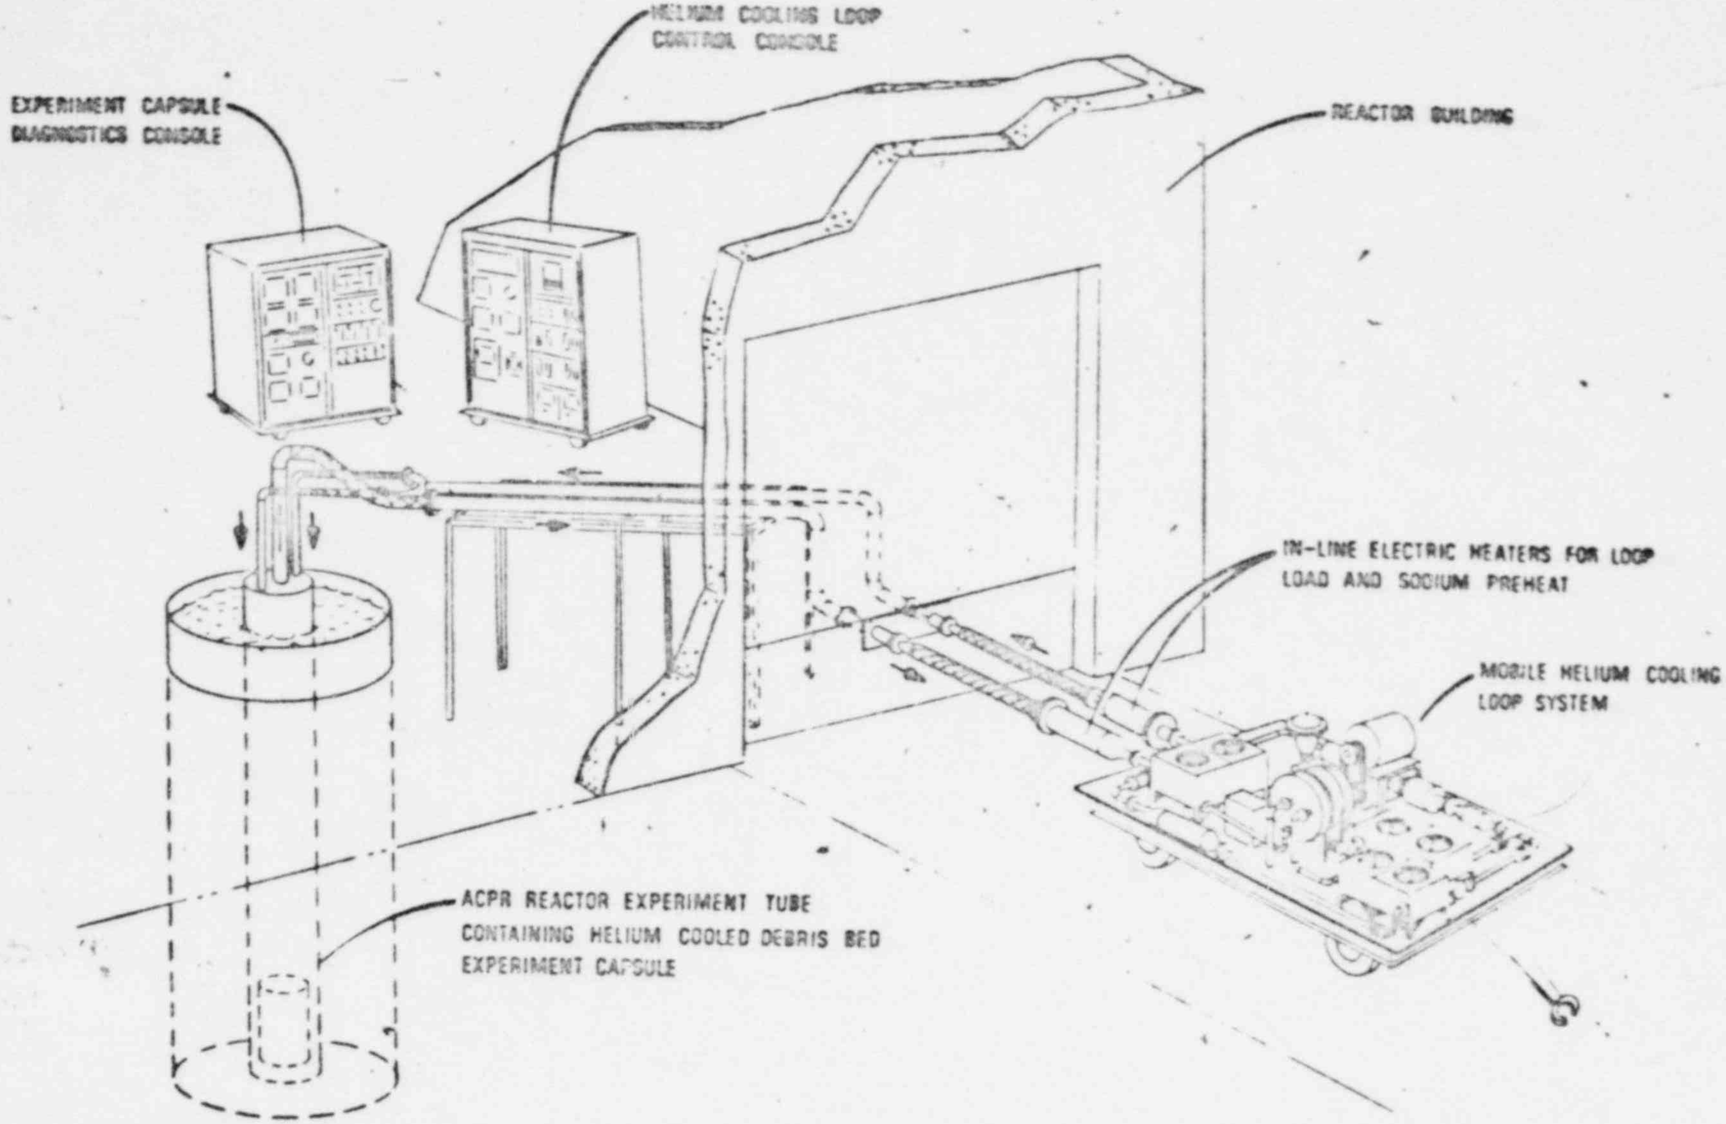

POST-ACCIDENT HEAT REMOVAL DEBRIS BED EXPERIMENT INSTRUMENTATION AT SANDIA LABORATORIES

REF. 14

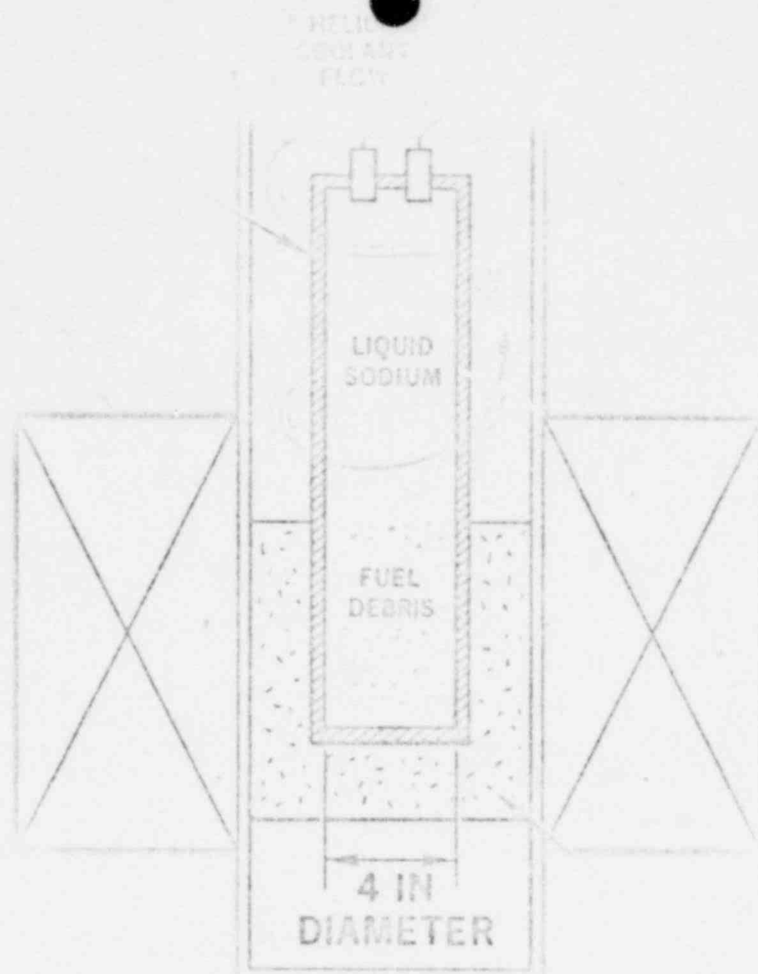

POST-ACCIDENT HEAT REMOVAL DEBRIS BED EXPERIMENT
AT SANDIA LABORATORIES

SKETCH SHOWS CUTAWAY SIDE VIEW DIAGRAM OF ANNULAR CORE PULSE REACTOR CORE WITH EXPERIMENTAL AREA.

TENTATIVE D-SERIES TEST PLAN

Test	Depth (cm)	Schedule Date	Full Size Distr.	Stratification	Added Steel	Bottom Cooling	Regime
D1	(~6)	Done					Sub-Dryout
D2	(~11)	Done					Dryout
D3	(~16)	Done					Dryout
D4	(8)	Done					
D6	(11)	80		X			Dryout
D5	(11)	80					Post-Dryout
D7	(6)	81					Dryout
D8	(8)	81	X				Dryout
D9	(16)	81					Dryout
D10	(8)	82	X		X		Dryout
D11	(16)	82	X			X	UO_2 Melt
D12	(8)	83	X	X	X		Dryout
D13	(11)	83	X			X	UO_2 Melt
D14	(11)	83	X		X	X	Steel/Melt
D15	(16)	84	X	X	X	X	Steel/Melt
D16	(16)	84	X			X	Extended UO_2 Melt

D4 EXPERIMENT

- FISSION-HEATED UO_2-SODIUM PARTICLE BED
- OBSERVATIONS

  DRYOUT FLUX WITH COLD OVERLYING SODIUM 4.5 TIMES SMALLER THAN DRYOUT WITH HOT OVERLYING SODIUM.

  A DISTURBANCE (SIMILAR TO THAT SEEN IN D2 BUT NOT IN D3) SEEN WITH HOT OVERLYING SODIUM ONLY

  STAGNANT BED THERMAL CONDUCTIVITY HIGHER AFTER DISTURBANCE

  PRE-DISTURBED BED DRYOUT DATA AGREES WITH SANDIA DRYOUT AND SERIES CONDUCTION MODELS

- INTERPRETATIONS

  COLD OVERLYING SODIUM SUPPRESSES CHANNEL FORMATION IN SHALLOW BEDS AND PREVENTS ENHANCED HEAT REMOVAL

  BED DISTURBANCES SEEM TO BE A SHALLOW-BED HOT-SODIUM EFFECT WHICH PRODUCE CHANNELS AND ENHANCED HEAT REMOVAL

  DISTURBANCE LEFT PERMANENT CHANNELS (EVEN AFTER COOLING)

CORE DEBRIS BEHAVIOR
DEBRIS BED EXPERIMENT-6

OBJECTIVE: TO DETERMINE THE EFFECT OF PARTICLE STRATIFICATION ON DEBRIS BED COOLABILITY

APPROACH: DUPLICATE DEPTH, DIAMETER, PACKING AND TEMPERATURE OF D2 BUT WITH STRATIFIED RATHER THAN UNIFORM BED

FEATURES: DEPTH -- 11 CM   DIAMETER -- 10 CM

ADDITIONAL SAFETY FEATURES (ISOLATION VALVING)

IMPROVED PRESSURE TRANSDUCER

IMPROVED ULTRASONIC THERMOMETRY

BETTER CHARACTERIZATION OF DOWNWARD AND RADIAL HEAT LOSSES

IMPROVED NEUTRON SHIELD DESIGN

CORE DEBRIS BEHAVIOR

DEBRIS BED EXPERIMENT-5

OBJECTIVES: · EXPLORATION OF THE POST-DRYOUT REGIME

DEVELOPMENT OF HIGH TEMPERATURE EXPERIMENT PACKAGE DESIGN

FEATURES: DEPTH ~11 CM   DIAMETER ~10 CM

SODIUM VOID FRACTION ≃ 43

MAXIMUM BED TEMPERATURE -- 1800° C

O_2 CONTENT < 50 PPM

APPLICATION OF DOUBLE CONTAINMENT

APPLICATION OF CRUCIBLE DESIGN

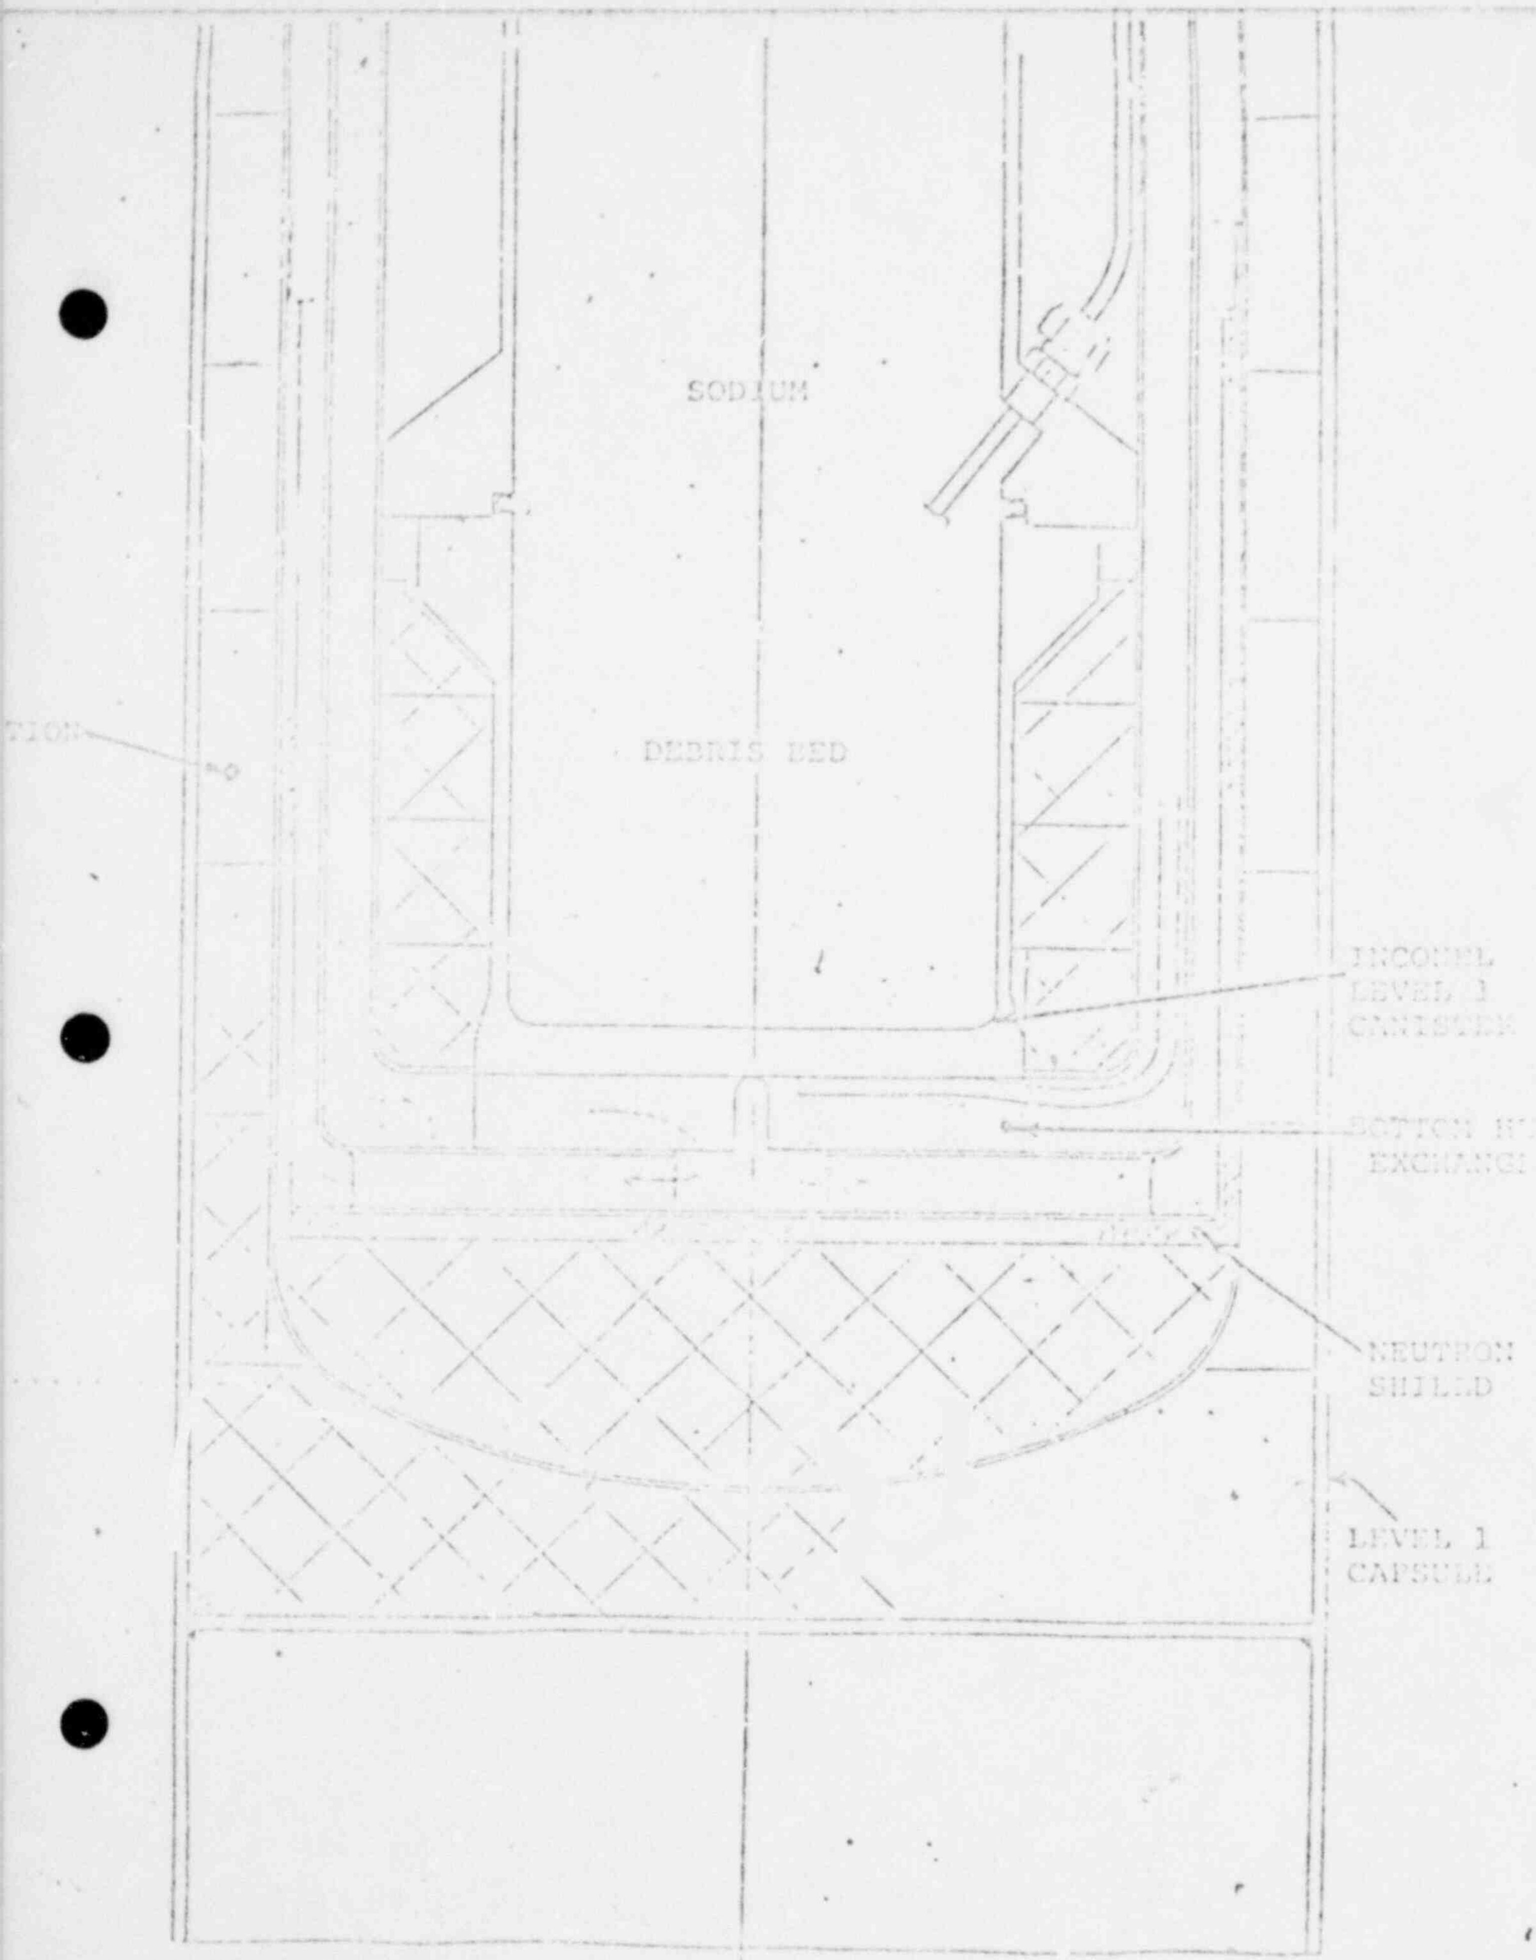

Non-Melt Bottom Cooled Capsule

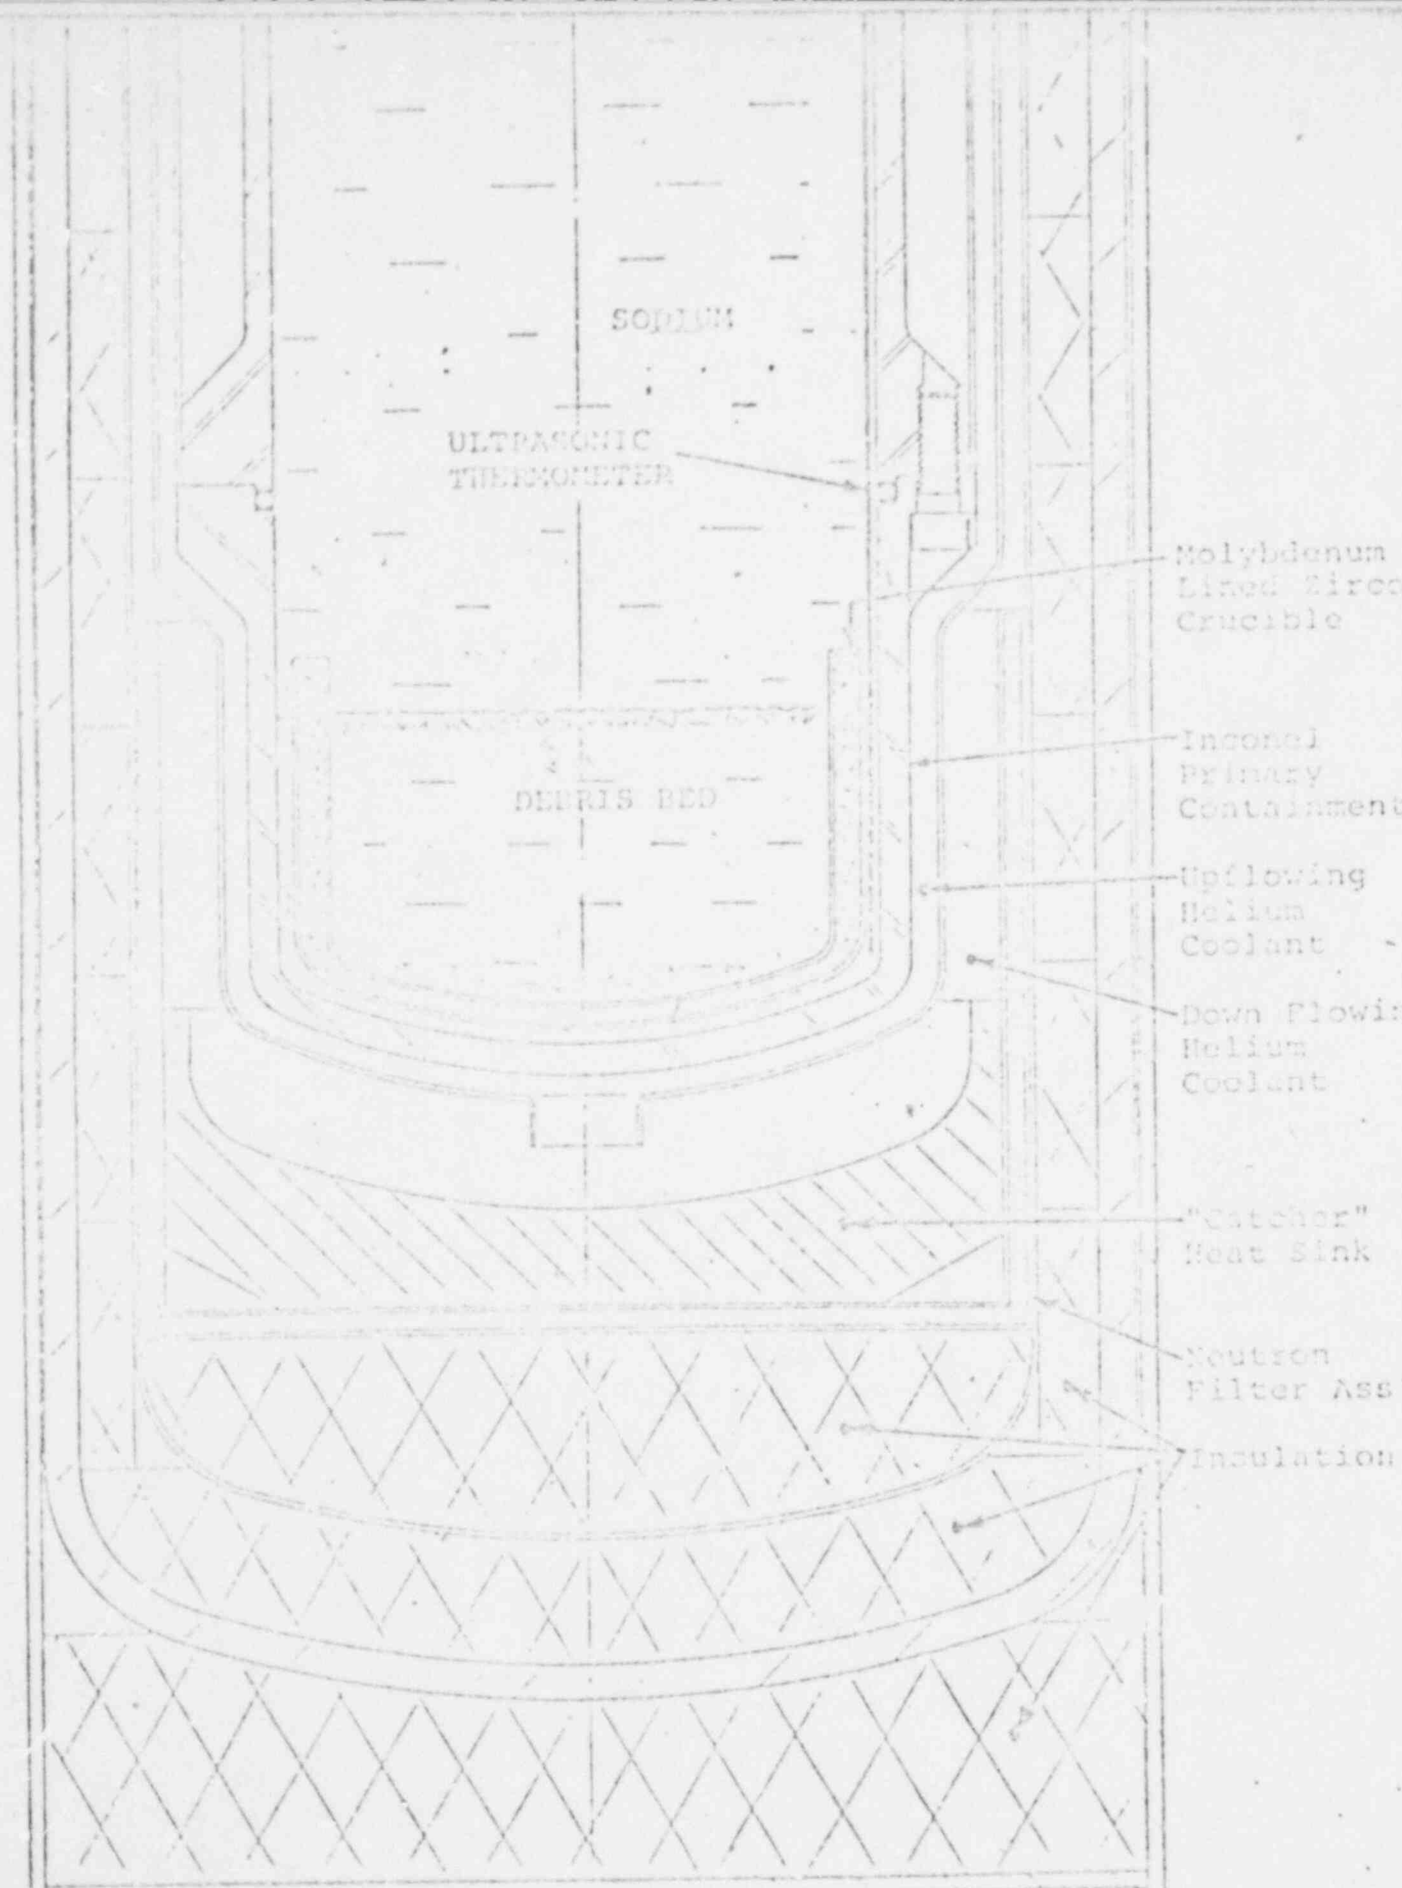

Preliminary Design of the Incipient

Melt Experiment Apparatus

PARTICLE BED MODELS DEVELOPED

- NON-CHANNELED BED LAMINAR/TURBULENT GRAVITATIONAL/-CAPILLARY DRYOUT MODEL. PREDICTS LOWER DRYOUT FLUX AND LESS VOID FRACTION DEPENDENCE FOR LARGE PARTICLES THAN EARLIER MODELS
- DOWNWARD BOILING MODEL FOR BOTTOM-COOLED BEDS. PREDICTS 3-4 TIMES INCREASE IN DRYOUT POWER FOR 10-CM HIGH BEDS WITH BOTTOM COOLING
- POST-DRYOUT MODEL FOR NON-CHANNELED BEDS
- DRYOUT MODEL BASED ON FLOODING. AGREES WITH TURBULENT LIMIT OF L/T G/C MODEL
- DRYOUT MODEL FOR VERY DEEP BEDS WITH BOTTOM ENTRY OF COOLANT
- CHANNEL PENETRATION MODEL TO EXPLAIN INCREASE IN DRYOUT FLUX FOR SHALLOW BEDS WITH HOT OVERLYING SODIUM

D4 DRYOUT BEHAVIOR

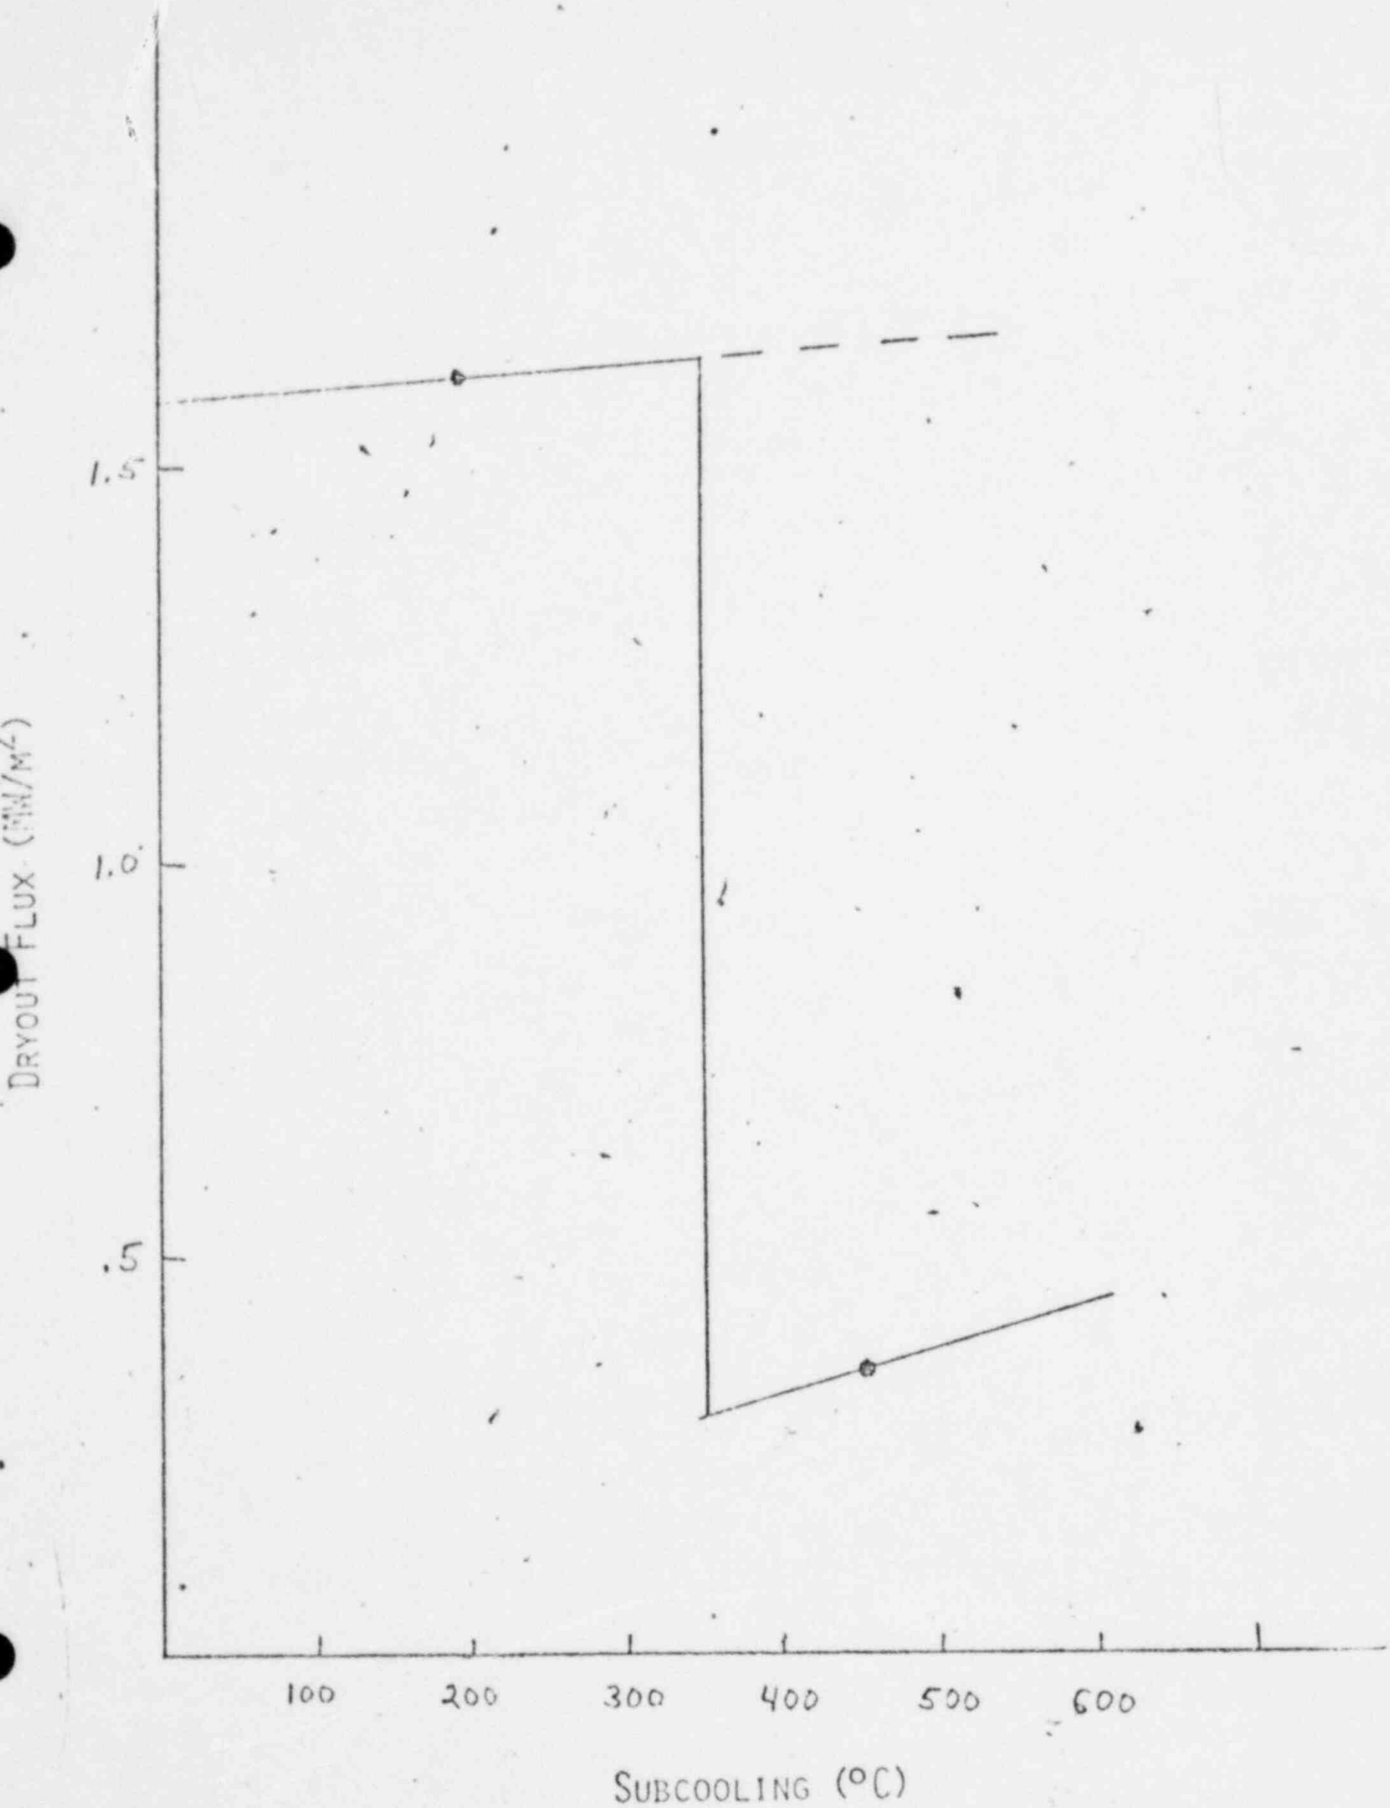

SODIUM

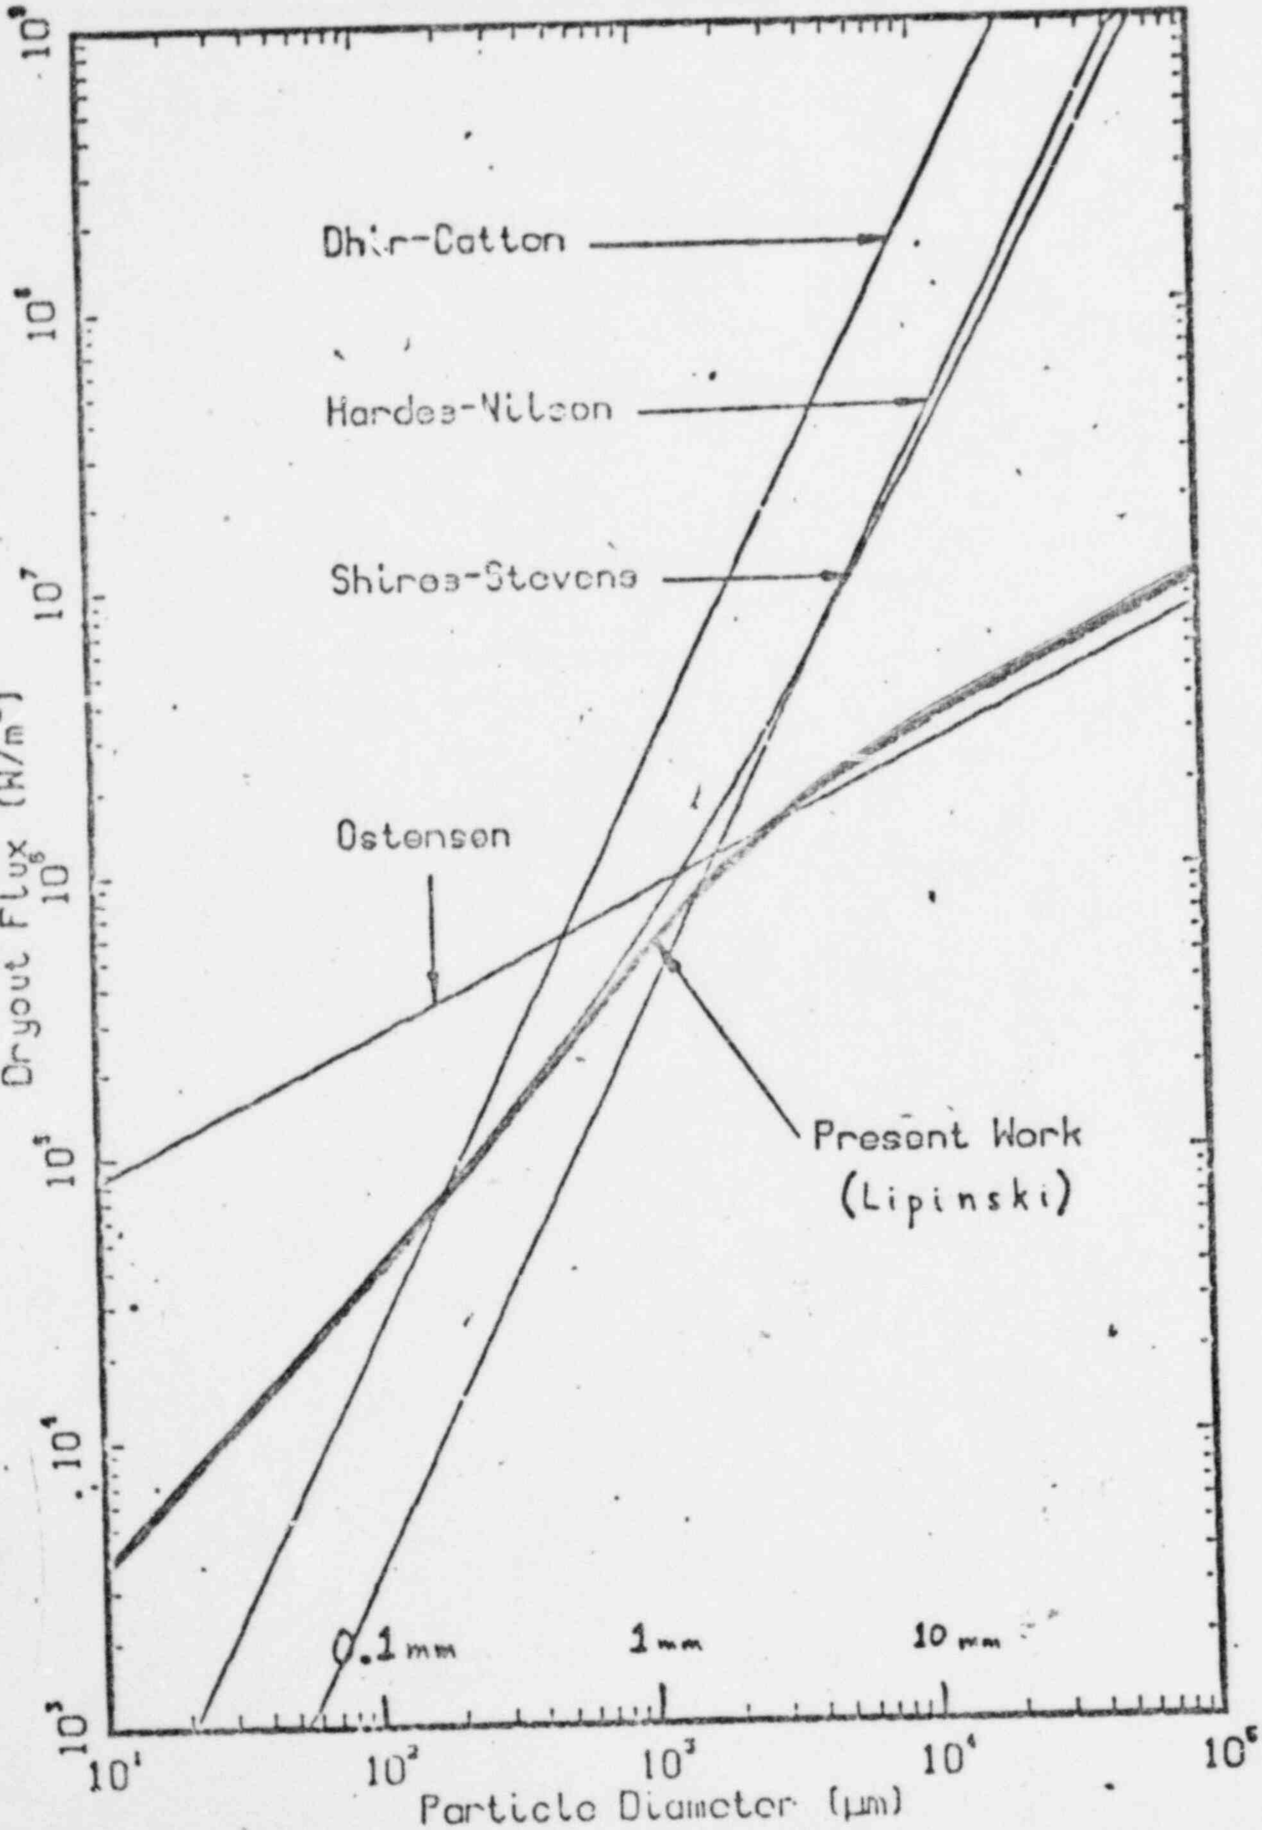

Dhir-Catton
Hardee-Nilson
Shires-Stevens
Ostensen
Present Work
(Lipinski)
0.1mm
1mm
10mm
Dryout Flux (W/m²)
Particle Diameter (μm)

CORE DEBRIS BEHAVIOR

MP -- OUT-OF-PILE EXPERIMENTATION

UO_2/MgO STUDIES

ULTRASONIC THERMOMETRY DEVELOPMENT

UO_2/STEEL MELT CONTAINMENT STUDIES

STEEL RELOCATION, VAPORIZATION, CONDENSATION AND WETTING STUDIES TO 2670°C

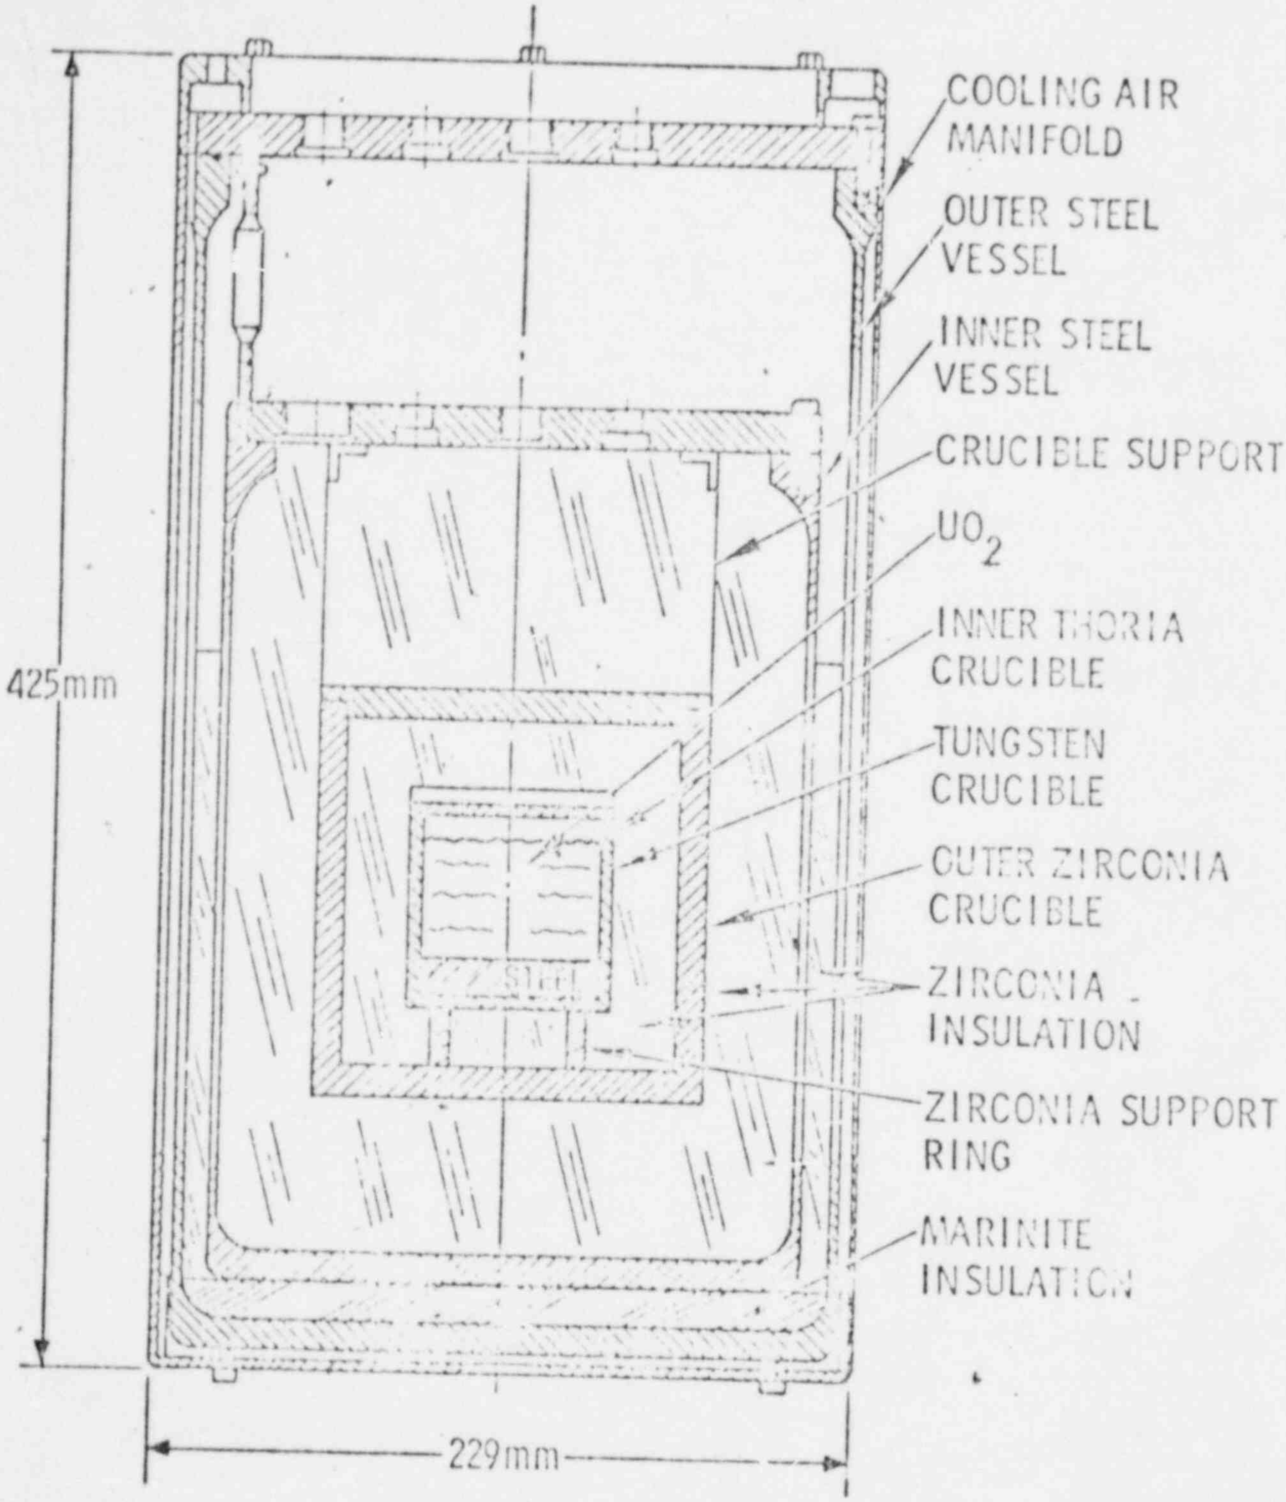
COOLING AIR MANIFOLD
OUTER STEEL VESSEL
INNER STEEL VESSEL
CRUCIBLE SUPPORT
UO_2
INNER THORIA CRUCIBLE
TUNGSTEN CRUCIBLE
OUTER ZIRCONIA CRUCIBLE
ZIRCONIA INSULATION
ZIRCONIA SUPPORT RING
MARINITE INSULATION
STEEL
425mm
229mm

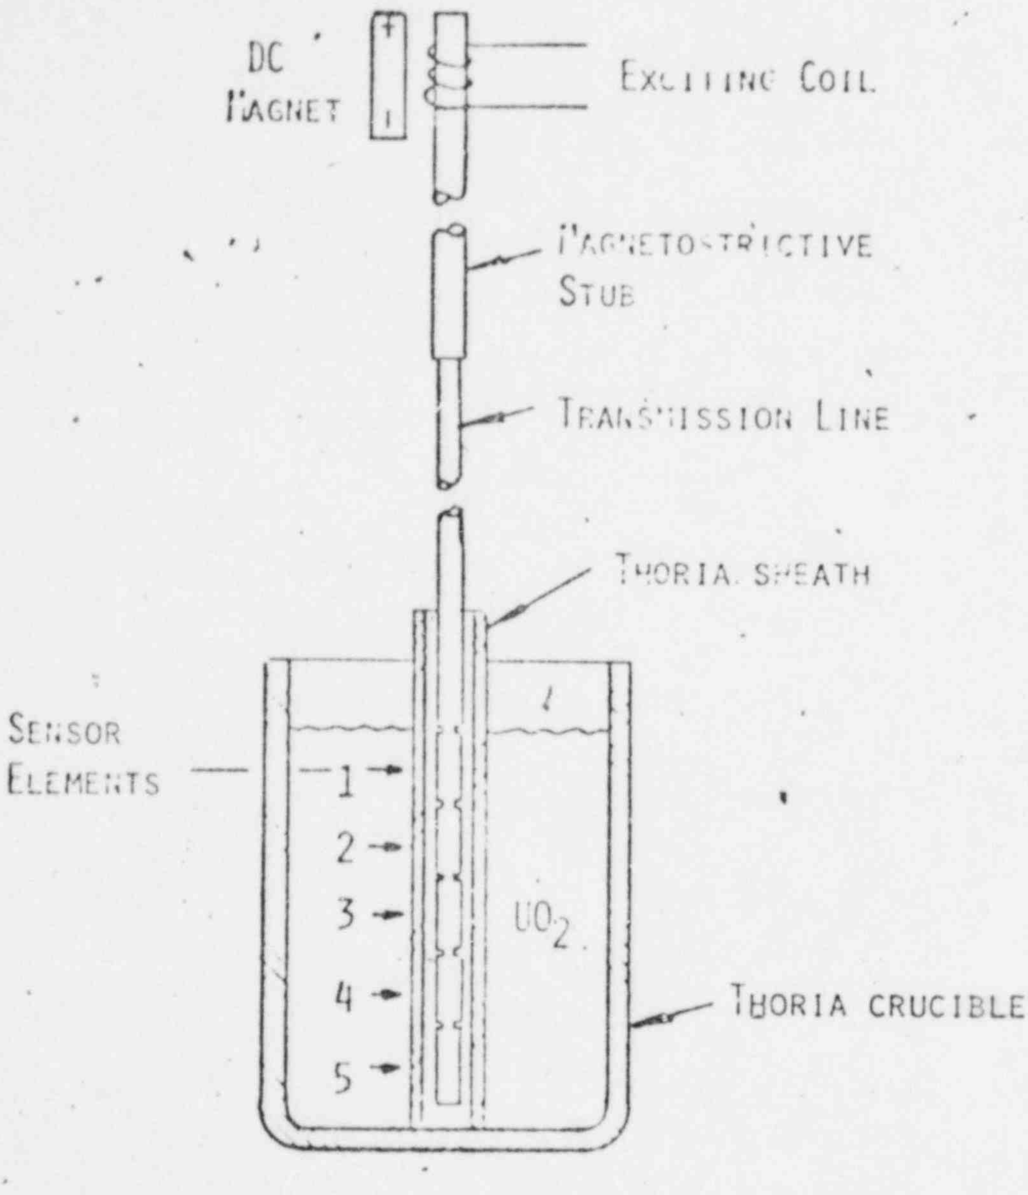

5-ELEMENT ULTRASONIC THERMOMETER
IN UO_2 (SCHEMATIC)

CORE DEBRIS BEHAVIOR

MOLTEN POOL EXPERIMENT-4

- OBJECTIVE: INVESTIVATE THE INTERACTION OF INTRINSICALLY HEATED UO_2 WITH MgO

FEATURES:

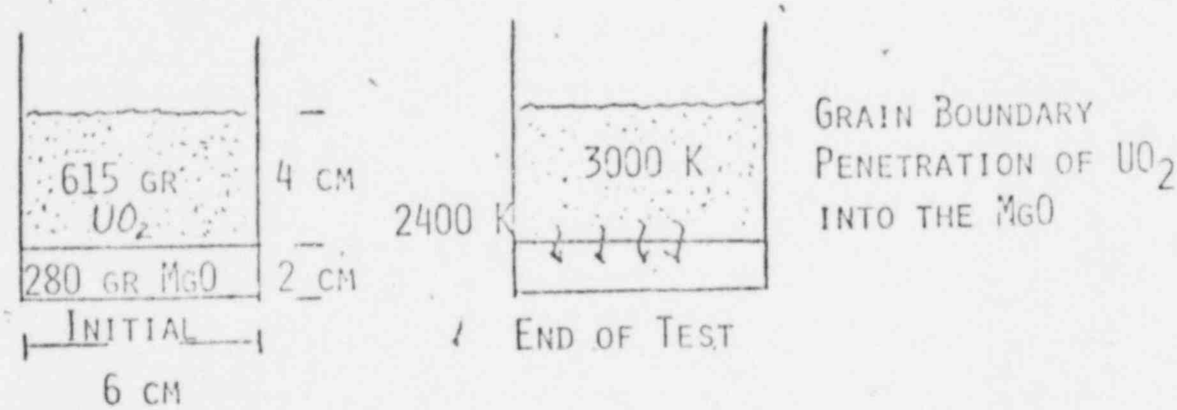

- RESULTS: USED ULTRASONIC THERMOMETRY TO MAP TEMPERATURES IN UO_2 AND MgO → DETERMINATION OF HEAT FLUX DISTRIBUTIONS

  POST-TEST EXAMINATION NOT YET STARTED--RESULTS WILL BE COMPARED WITH CORRELATIONS DEVELOPED IN OUT-OF-PILE FURANCE TESTS.

CORE DEBRIS BEHAVIOR

MOLTEN POOL EXPERIMENT-5

OBJECTIVE: INVESTIGATE STEEL MIGRATION
AND AGGLOMERATION ⟶
EFFECT ON EFFECTIVE BED
THERMAL CONDUCTIVITY
DIRECT SUPPORT OF D10
AND D12

FEATURES:

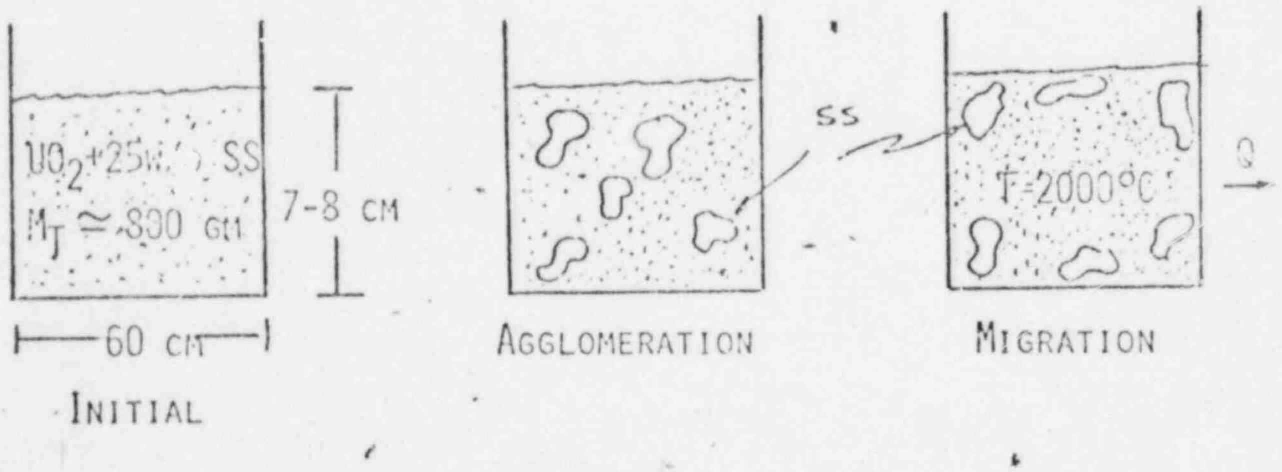

CORE DEBRIS BEHAVIOR

MOLTEN POOL EXPERIMENT-6

OBJECTIVE: INVESTIGATE THERMAL BEHAVIOR FOLLOWING DRYOUT AND THE TRANSITION FROM A PARTICULATE BED TO A MOLTEN POOL

COMPLEMENTS D5 AND PROVIDES DIRECT SUPPORT FOR D11, D13 AND D16

FEATURES:

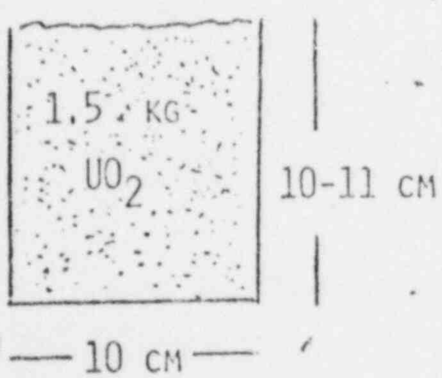

INITIAL

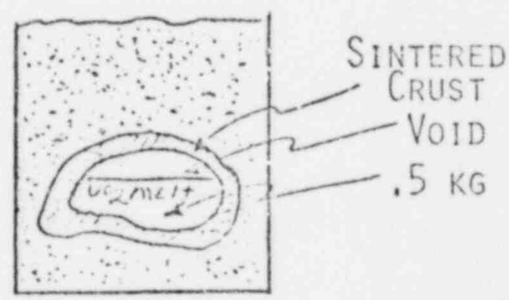

CORE DEBRIS BEHAVIOR
MOLTEN POOL EXPERIMENT-7

OBJECTIVE: INVESTIGATE STEEL AGGLIMERATION
AND MIGRATION MECHANISMS
AT HIGH TEMPERATURES
(MIGRATION VAPOR PRESSURE DRIVEN?)
DETERMINE HEAT FLUX DISTRIBUTION
SUPPORTS D10, D12, D14, D15

FEATURES:

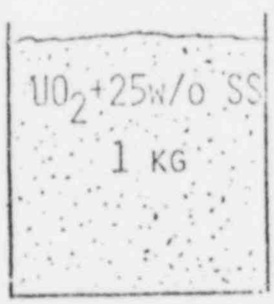

INITIAL

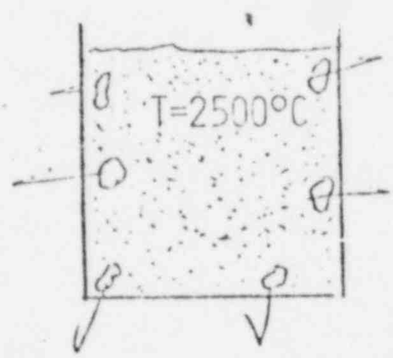

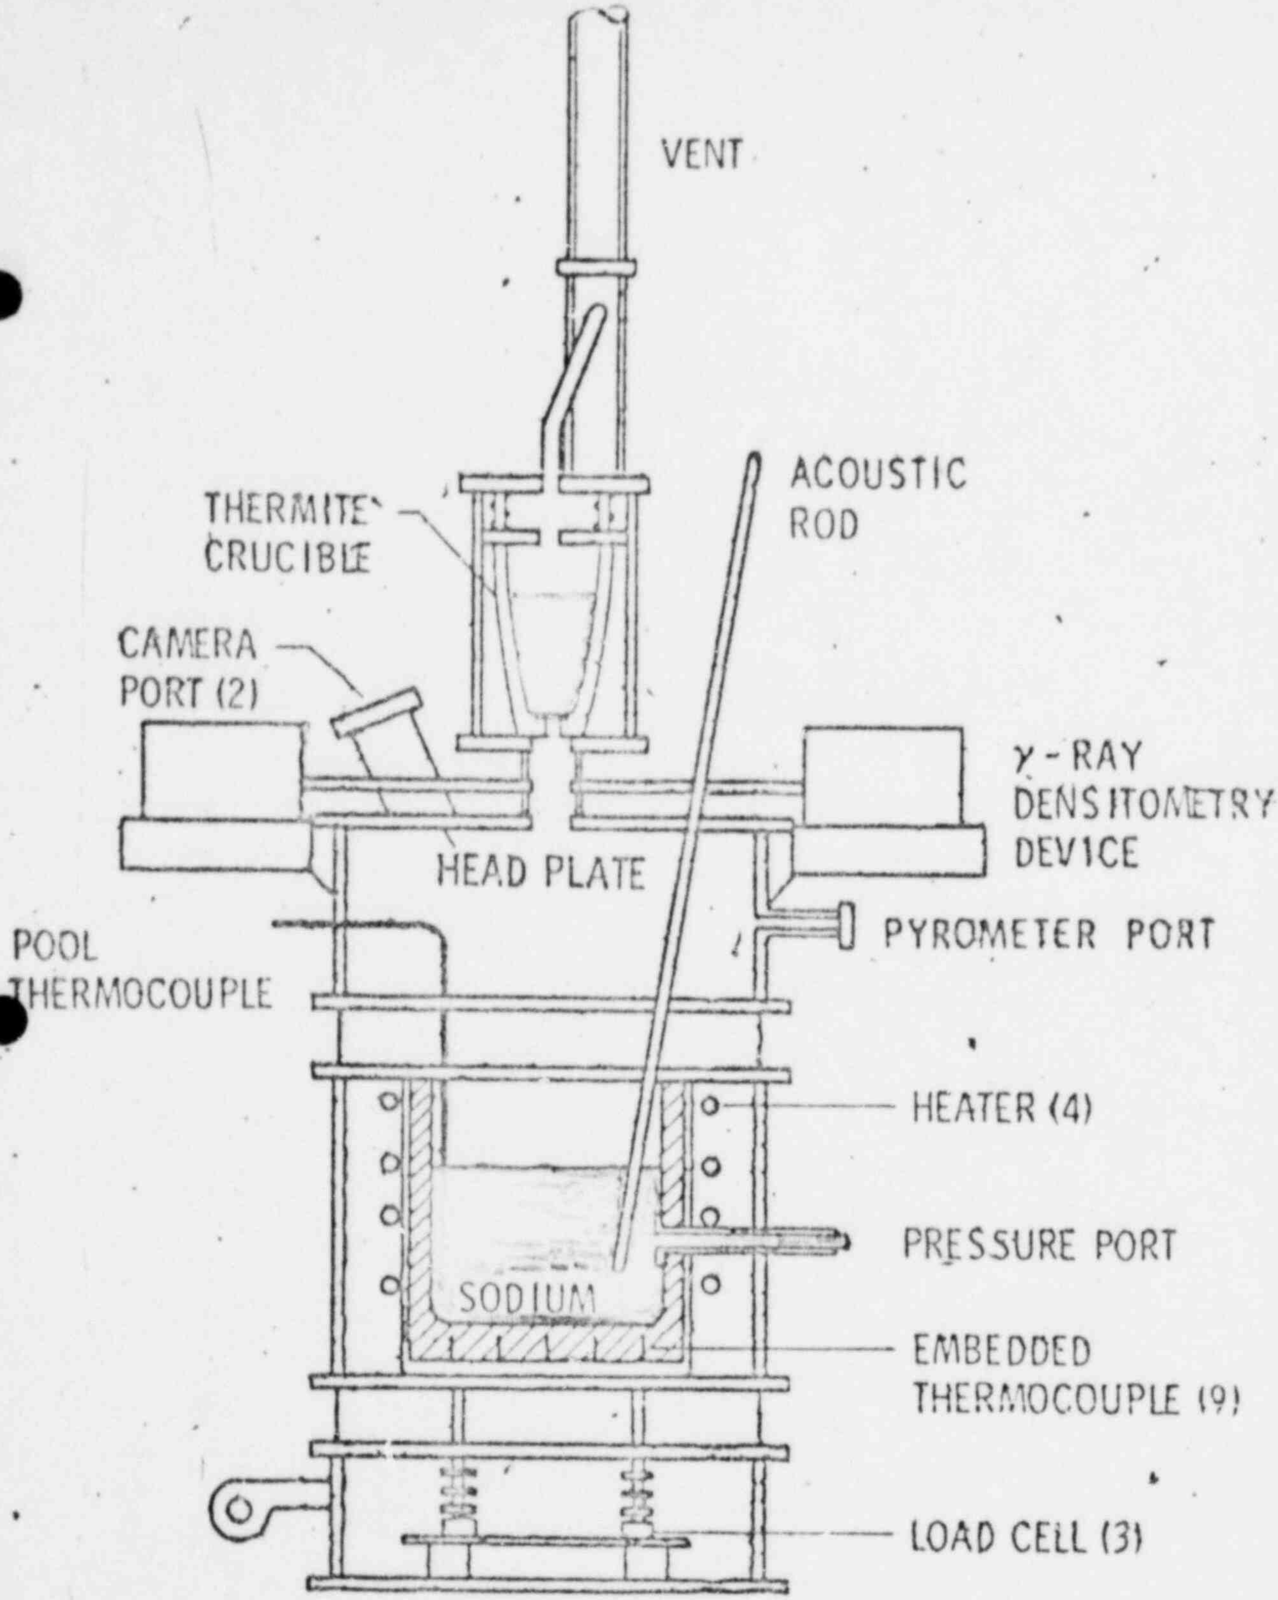
VENT
ACOUSTIC
ROD
THERMITE
CRUCIBLE
CAMERA
PORT (2)
γ - RAY
DENSITOMETRY
DEVICE
HEAD PLATE
PYROMETER PORT
POOL
THERMOCOUPLE
HEATER (4)
PRESSURE PORT
SODIUM
EMBEDDED
THERMOCOUPLE (9)
LOAD CELL (3)

DEBRIS BED STRUCTURE FROM
FRAG 3

FRAGMENTATION

STATUS

COMPLETED PLANNED TEST SERIES ON ~20 KG SCALE

- 5 EA; MELT INTO NA
- 2 EA; NA ONTO MELT
- 1 EA; NA ONTO MELT IN PRESENCE OF CONCRETE

RESULTS

EXCELLENT DATA ON

- PARTICLE SIZE
- PARTICLE SIZE DISTRIBUTION
- BED STRATIFICATION
- BED POROSITY
- BED FORMATION

OBSERVATIONS

A) 1ST ON STRATIFICATION

B) PARTICLE SIZE AND DISTRIBUTION NOT SIGNIFICANTLY DIFFERENT FROM SMALL SCALE RESULTS OF ANL AND OTHERS

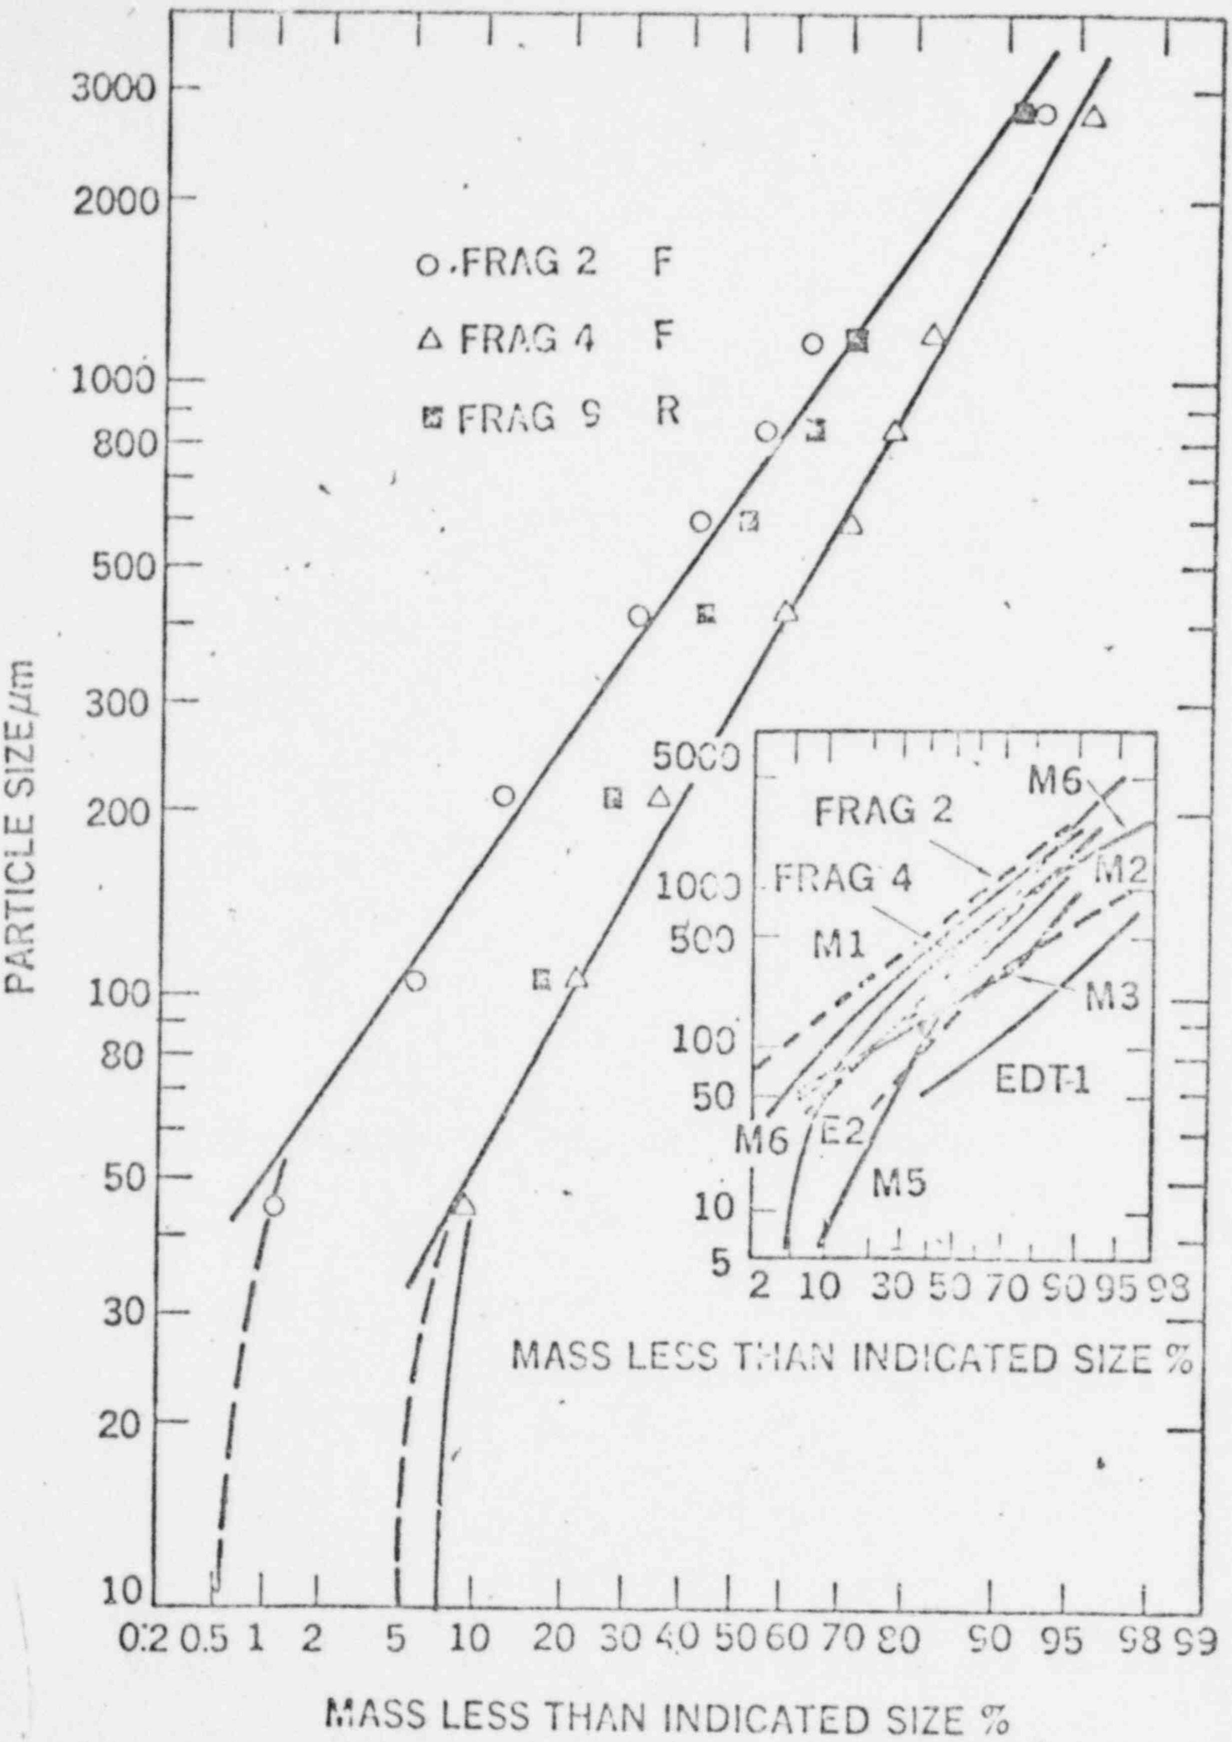

PARTICLE SIZE μm
MASS LESS THAN INDICATED SIZE %
FRAG 2 F
FRAG 4 F
FRAG 9 R
3000
2000
1000
800
500
300
200
100
80
50
30
20
10
0.2 0.5 1 2 5 10 20 30 40 50 60 70 80 90 95 98 99
5000
1000
500
100
50
10
5
2 10 30 50 70 90 95 98
FRAG 2
FRAG 4
M1
M6
M2
M3
EDT-1
E2
M5
MASS LESS THAN INDICATED SIZE %

	Forward: Melt Into Sodium					Reversed: Sodium Into Melt		
	Frag 1	Frag 2	Frag 4	Frag 5	Frag 6	Frag 7	Frag 11	Frag 12
Steam Size MM	13	51	51	51	51	102	102	102 Concrete Bottom Crucible
Melt	Fe Al$_2$O$_3$ (44%)	Fe Al$_2$O$_3$ (44%)	UO$_2$-ZrO$_2$ (70%) Stainless Steel	UO$_2$-ZrO$_2$ (70%) Stainless Steel	UO$_2$-ZrO$_2$ (70%) Stainless Steel	Fe Al$_2$O$_3$ (44%)	UO$_2$-ZrO$_2$ (70%) Stainless Steel	UO$_2$-ZrO$_2$ (70%) Stainless Steel
Melt Mass KG	12	13	20	20	20	14	21	21
Sodium KG	23	23	23	23	23	23	23	23
Sodium Temperature °C	520	480	420	250	690	500	510	520
Max Observed Pressure Bar	2	>20	--	2	7	--	--	16
Mechanical Damage	No	Yes	No	No	No	No	No (Movement of Apparatus)	No (Movement of Apparatus)

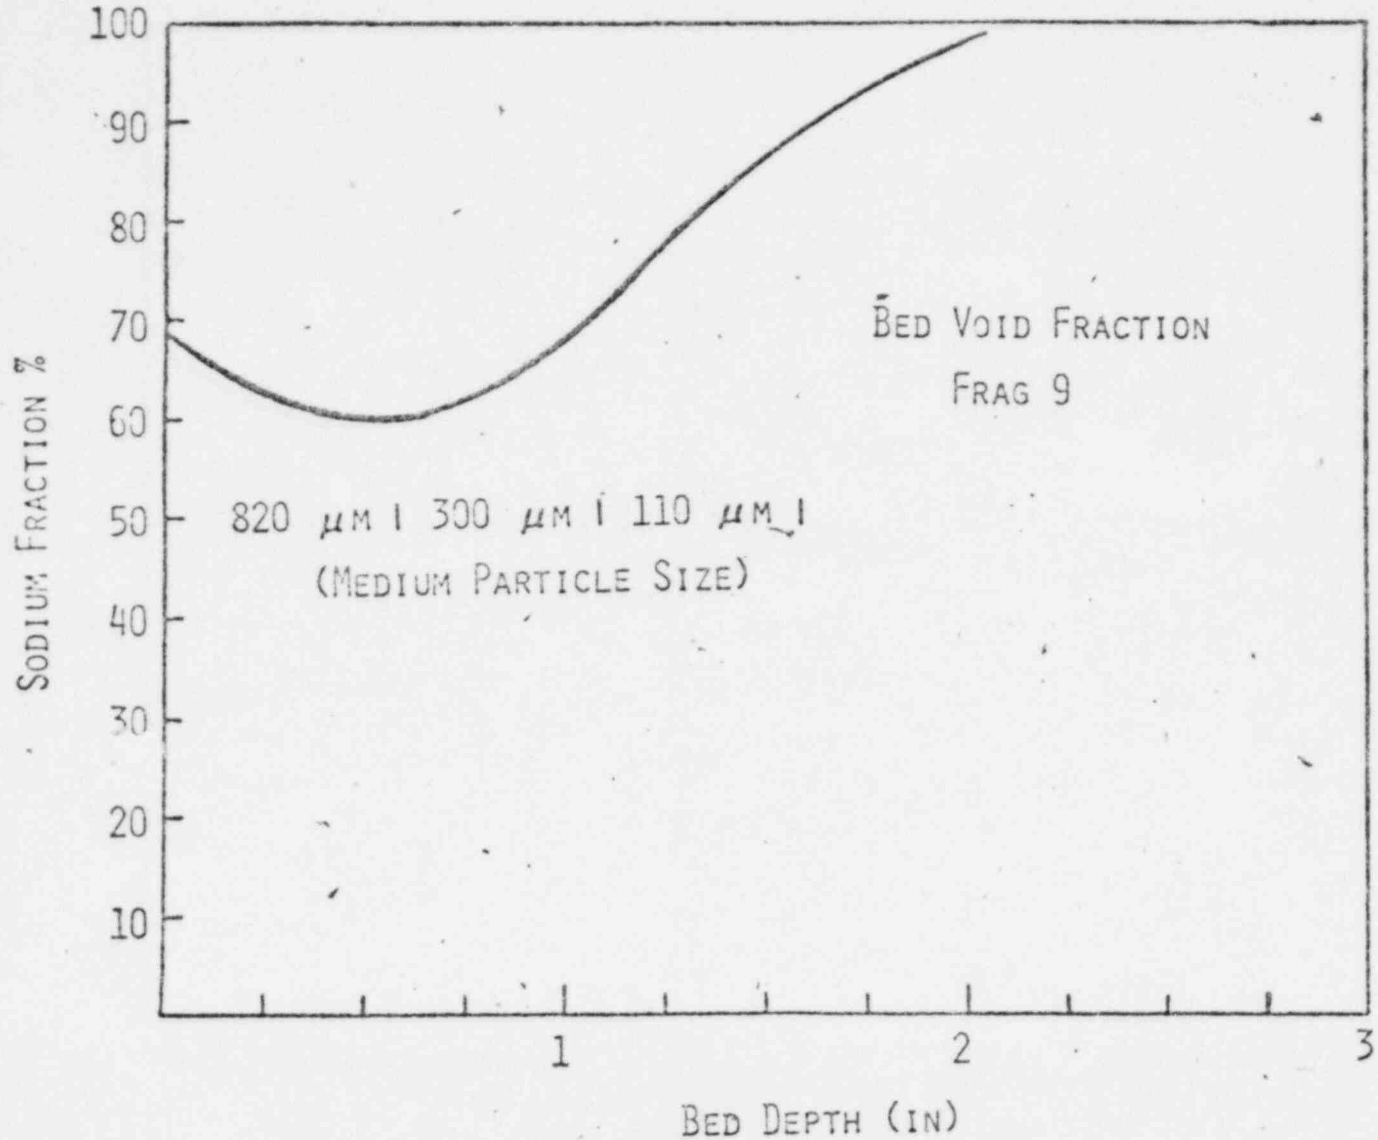

100
90
80
70
60
50
40
30
20
10
Sodium Fraction %
1
2
3
Bed Depth (in)
Bed Void Fraction
Frag 9
820 μm | 300 μm | 110 μm |
(Medium Particle Size)

CLAUSER

Accident Delineation Study

A comprehensive, systematic delineation of
LMFBR accident sequences
undertaken at

Sandia National Laboratories

Staff: D. C. Williams, J. A. Sholtis,* F. W. Sciacca,**
E. R. Copus, P. J. McDaniel

Presented to the ACRS by M. J. Clauser 6/30/80

*USAF **Energy, Inc., Albuquerque, New Mexico

LMFBR Accident Delineation Study

- Comprehensive delineation of entire LMFBR accident sequences:
  accident initiation, accident phenomenology, post-accident phenomenology

- Investigate applicability of event trees and fault trees to LMFBRs

- Initially based on CRBR, alternative designs now being examined

- Event trees and fault trees have been constructed, and branch-point likelihood has been estimated for the purpose of:

  - Developing the methodology
  - Delineating the plausible accident sequences
  - Determining the dominant sequences
  - Identifying the key phenomena and uncertainties in sequences

- Outcomes:

  - Basis for prioritizing research, design and development efforts
  - Basis for assessing relative safety of different components and designs
  - Help establish LMFBR licensing criteria such as design base accidents

LMFBR Accident Delineation Study - Status

Recent accomplishments

Phase 1 completed, report drafted

Qualitative delineation with event trees in three areas

- Accident initiation
- Accident phenomenology
- Post-accident phenomenology

Likelihood estimates made

Dominant pathways established

Fault trees established and partially quantified for CRBR engineered safety systems

Current Activities - Phase II - Quantification

Fault tree quantification

Studies with mechanistic systems codes

SAS
SIMMER
BRENDA
SSC
CONTAIN

Alternative containment design review

Accident Delineation Study -- Result

Principle result of Phase I:

Comprehensive, systematic delineation of entire LMFBR accident sequences

ACCIDENT DELINEATION STUDY ORGANIZATION

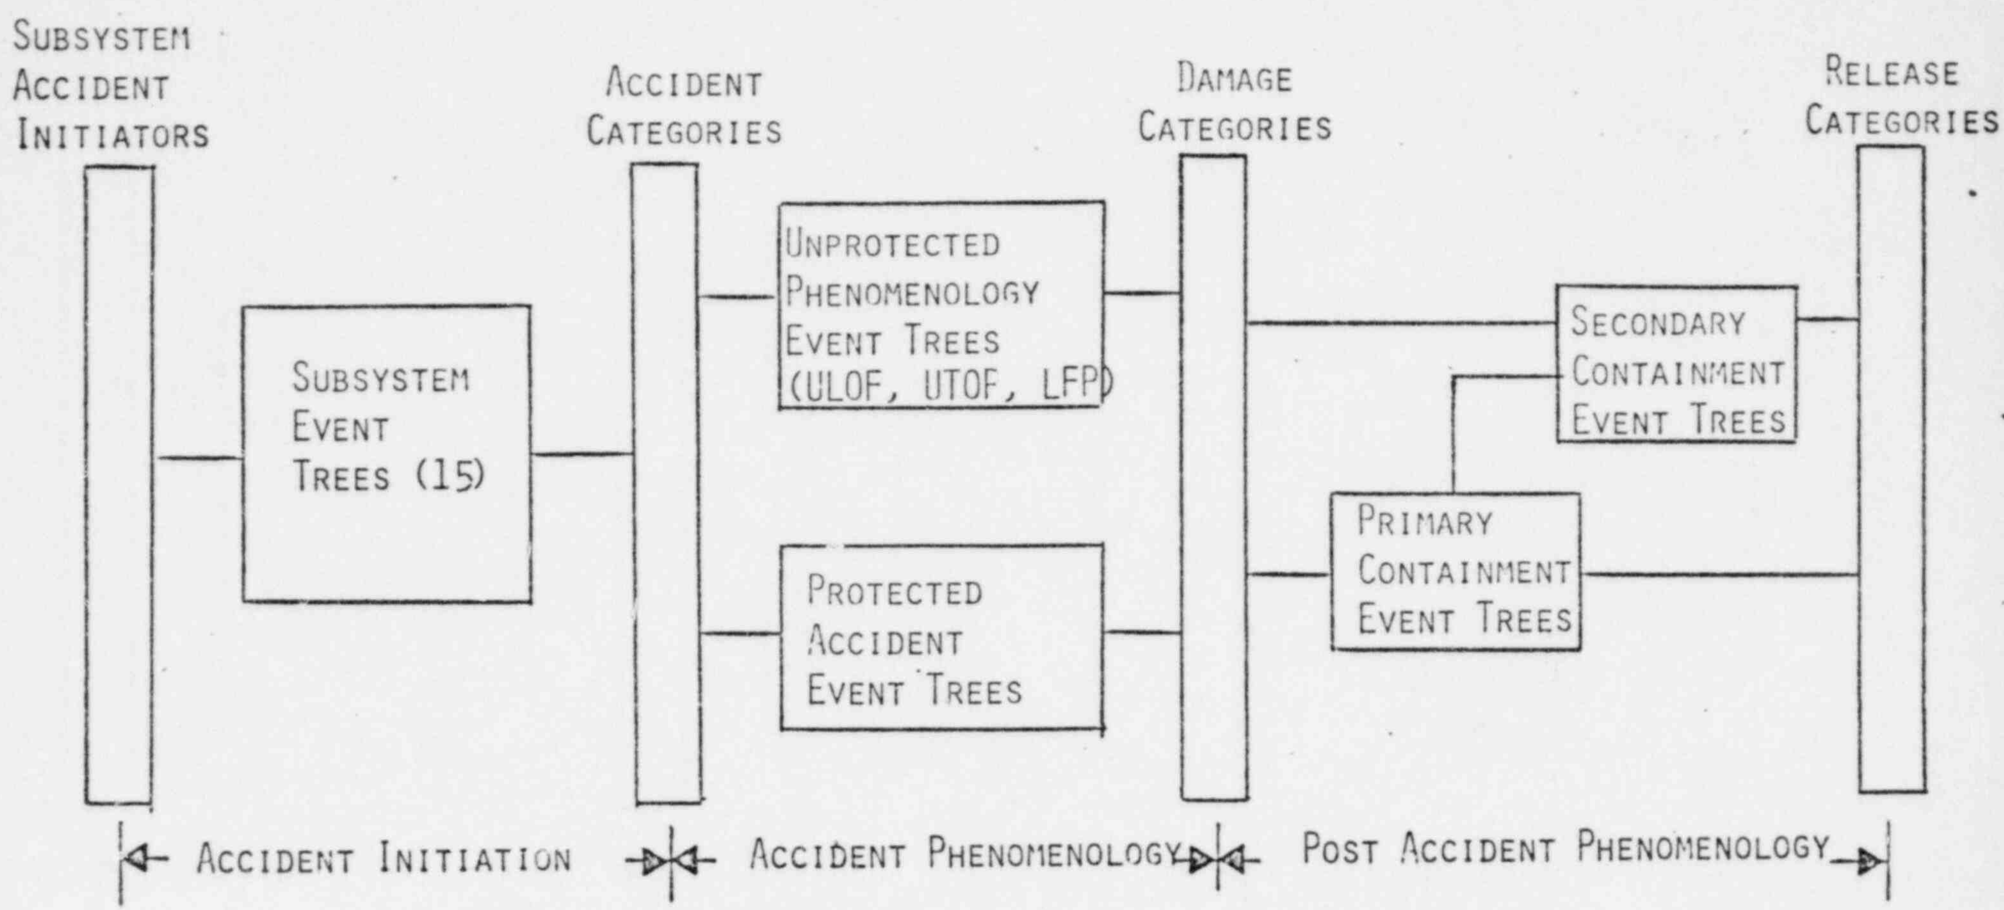

ENGINEERED SAFETY SYSTEMS' EVENT TREE

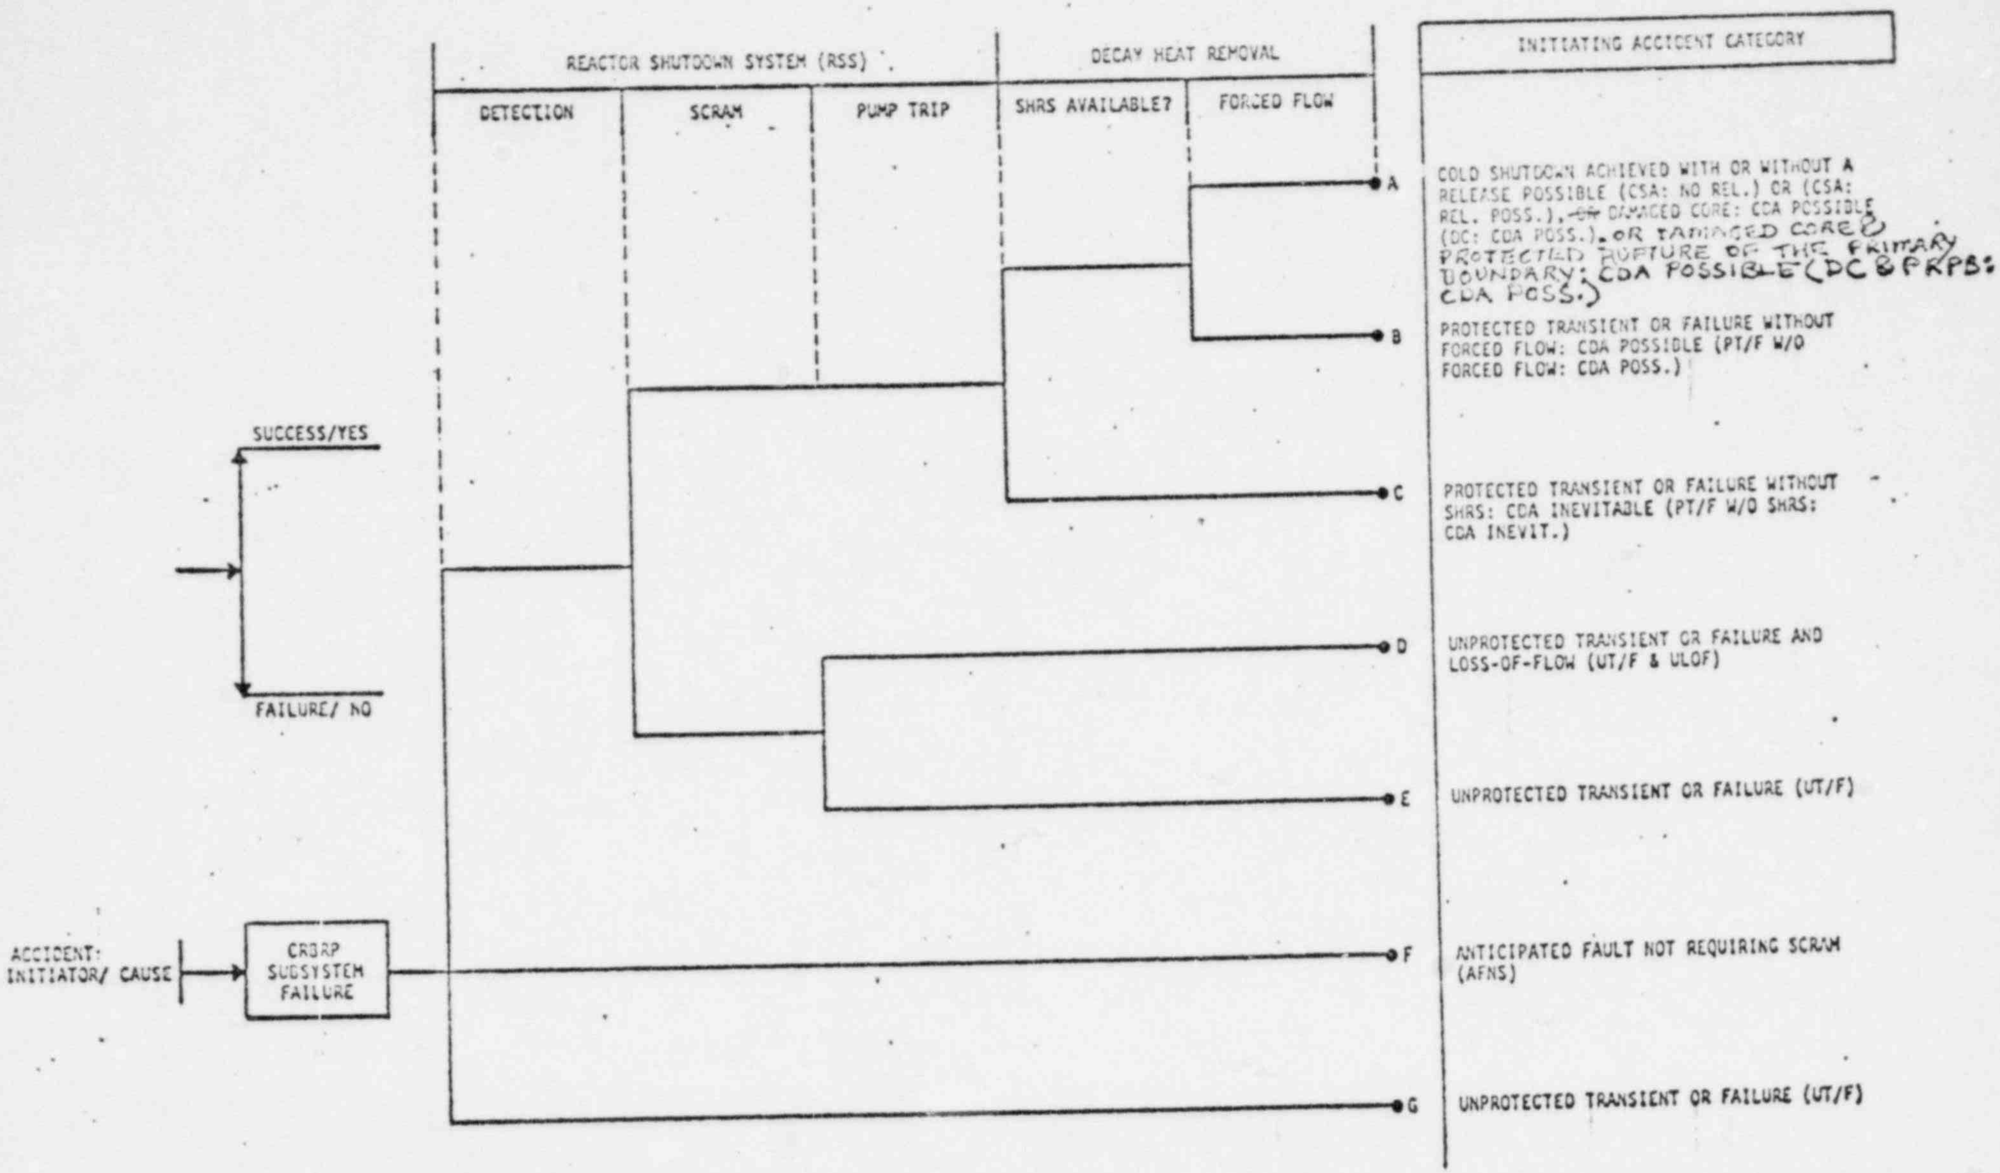

ULOF in Homogeneous Core

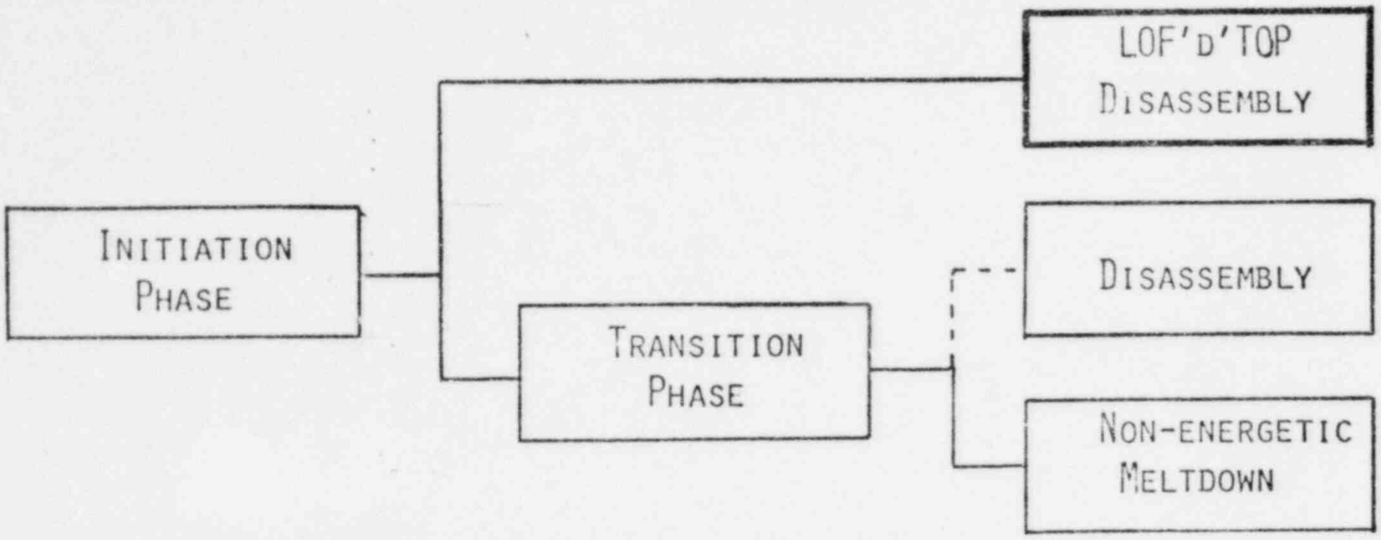

ULOF in Heterogeneous Core

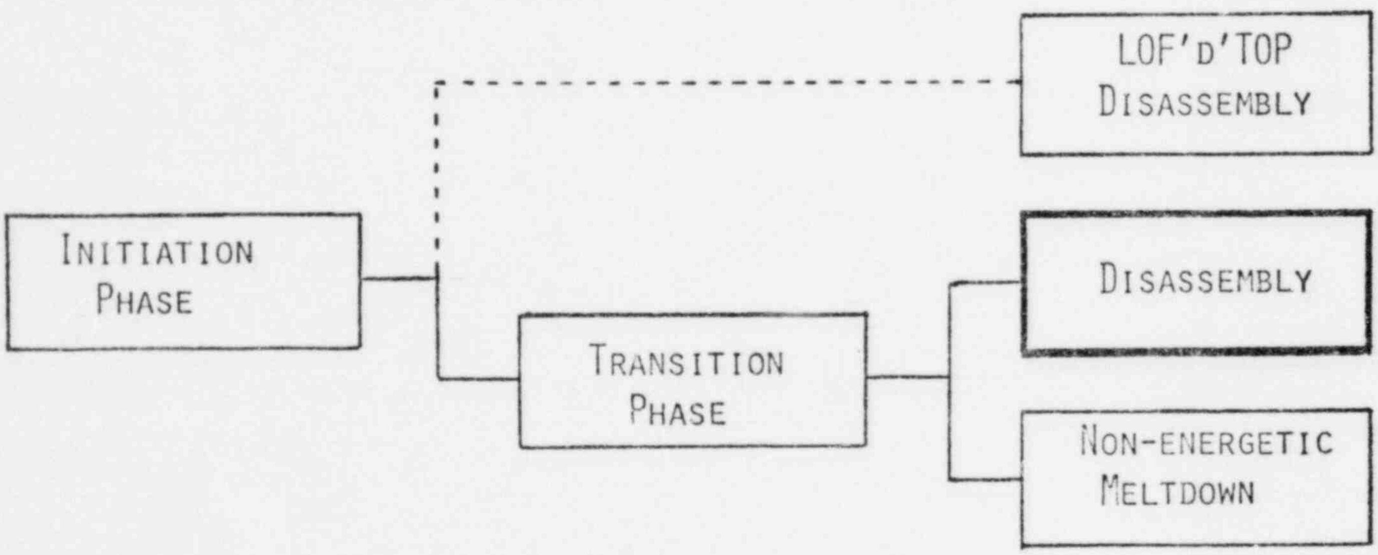

Protected CDA

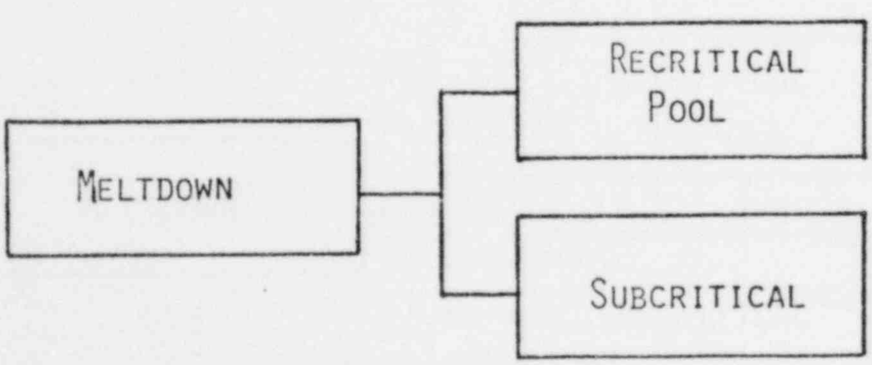

Accident Delineation Study - Phase I Conclusion

1 - Protected Accidents

-- Protected CDAs more frequent than unprotected CDAs (CRBRP Safety Study)

-- Protected accident consequences may be as severe as unprotected

-- Therefore protected accidents may constitute a greater risk to the public

-- Uncertainties in protected accident phenomenology are large

-- Recommendation: devote considerably more research and development effort to protected accidents

-- Principal questions:
   Core coolability
   Recriticallity

ACCIDENT DELINEATION STUDY - PHASE I CONCLUSION

2 - LOW RAMP-RATE UTOPS ($\lesssim$ 30¢/s)

-- LOW RAMP-RATE UTOPS MUCH MORE FREQUENT THAN HIGH RAMP-RATE UTOPS

-- LOW RAMP-RATE UTOPS MAY HAVE CONSEQUENCES COMPARABLE TO HIGH RAMP-RATE UTOPS

-- THEREFORE LOW RAMP-RATE UTOPS MAY CONSTITUTE A GREATER RISK TO THE PUBLIC

-- UNCERTAINTIES IN LOW RAMP-RATE UTOP PHENOMENOLOGY ARE LARGE

-- HOWEVER UTOPS AS A WHOLE CONSTITUTE RELATIVELY SMALL RISK

-- RECOMMENDATION: DEVOTE SOME MORE RESEARCH EFFORT TO LOW RAMP-RATE UTOP RESEARCH

Accident Delineation Study - Phase I Conclusion

3 - Local Fault Propagation (LFP) Accidents

-- LFP initiators (e.g., single-pin failures) occur quite frequently

-- Propagation appears unlikely, but uncertain

-- Risk from LFP accidents cannot yet be disregarded

Accident Delineation Study - Phase I Conclusion

4 - Containment

-- Containment reduces probability of atmospheric release by roughly one or two orders of magnitude

-- Half of all LMFBR CDAs may result in basemat failure

-- Consequences of basemat failure are far less severe than RCB failure

-- Containment substantially mitigates the effects of a CDA

CLAUSER

CONTAIN

An integrated reactor–containment systems code under development at

Sandia National Laboratories

Staff: Jay P. Odom,* Michael E. Senglaub, David K. Rudeen*

Presented to the ACRS by M. J. Clauser 6/30/80

*Science Applications, Inc., Albuquerque, New Mexico

CONTAIN

-- AN INTEGRATED, REACTOR-CONTAINMENT SYSTEMS CODE

- TO ANALYZE POST-ACCIDENT SEQUENCES FOLLOWING RELEASE OF MATERIAL FROM THE PRIMARY VESSEL THROUGH RELEASE FROM SECONDARY CONTAINMENT

-- FOR ALL TYPES OF ACCIDENTS

-- FOR ALL TYPES OF REACTORS

-- FOR ALL TYPES OF CONTAINMENT

-- MODELS FOR MASS AND ENERGY GENERATION AND TRANSPORT PROCESSES

-- DETERMINES PRESSURE, TEMPERATURE, LOCATION AND STATE OF FISSION PRODUCTS

CONTAIN FEATURES

- STATE-OF-THE-ART PHYSICS MODELS
  - -- GENERAL CAVITY DEBRIS-POOL MODEL (SINTER)
  - -- MULTI-COMPONENT, SECTIONAL AEROSOL MODEL (MAEROS)
  - -- GENERAL, DETAILED FISSION PRODUCT DECAY AND TRANSPORT
- MODULAR STRUCTURE
  - -- MODELS READILY UPDATED
  - -- PHYSICS READILY ALTERED
- FIRST VERSION IS OPERATIONAL

CONTAINMENT ANALYSIS

Cell Interaction	COMPARE, CACECO
Cell Atmosphere	
Aerosol Behavior	HAARM-3, AEROSIM, AEROSOL
Sodium Pool Burning	SOFIRE-2, NABRUND, CACECO, FRASC
Sodium Spray Burning	SOMIX-2, SPRAY-3
Materials Interactions	
Concrete-Sodium	CACECO
Concrete-Melt	INTER, WECHSL, CORCON
Concrete Thermal Response	USINT
Cavity Debris Pool	
Na Pool	CACECO
Debris Bed	Core Debris Behavior
Molten Pool Behavior	Program
Fission Product	
Inventory	ORIGEN, RIBD
Distribution	
Release	COMRADEX-II

CONTAIN Status (6/30/80)

Cavity Debris Pool Physics Models (SINTER)

Drift flux model	OPERATIONAL
Phase transitions	OPERATIONAL
Aerosol sources	OPERATIONAL
Fission product transport	UNDER DEVELOPMENT
Fission product decay	OPERATIONAL
Sodium-pool fires (SOFIRE II)	OPERATIONAL
Sodium-concrete interactions	UNDER DEVELOPMENT
Debris-coolant heat transfer	OPERATIONAL
Debris-concrete interactions	PLANNED
Concrete model (USINT)	UNDER DEVELOPMENT

CONTAIN Status (6/30/80)

Atmosphere Physics Models

Flow of mass and energy	OPERATIONAL
Two-phase thermodynamics (EOS)	OPERATIONAL
Aerosol transport (MAEROS)	OPERATIONAL
Fission product transport	UNDER DEVELOPMENT
Fission product decay	OPERATIONAL
Sodium spray fire (SPRAY)	OPERATIONAL
Surface heat transfer	OPERATIONAL
Surface condensation	OPERATIONAL
Engineered systems	UNDER DEVELOPMENT

CONTAIN Status (6/30/80)

General

Interim time step control	OPERATIONAL
Modular code structure	OPERATIONAL
Dynamic storage allocation	OPERATIONAL
INPUT	NEARLY OPERATIONAL
OUTPUT	OPERATIONAL
Plots	OPERATIONAL
Restart	OPERATIONAL

CONTAIN -- Z/IP Study

TMLB' steam spike

- -- Steam and Aerosol generation rate specified from March
- -- Multi-species aerosol model
- -- surface condensation and heat transfer
- -- Multi-cell, two-phase flow
- -- Fission product transport and decay
- -- Test of CONTAIN atmosphere calculation
- -- Investigated effect of gap conductivity

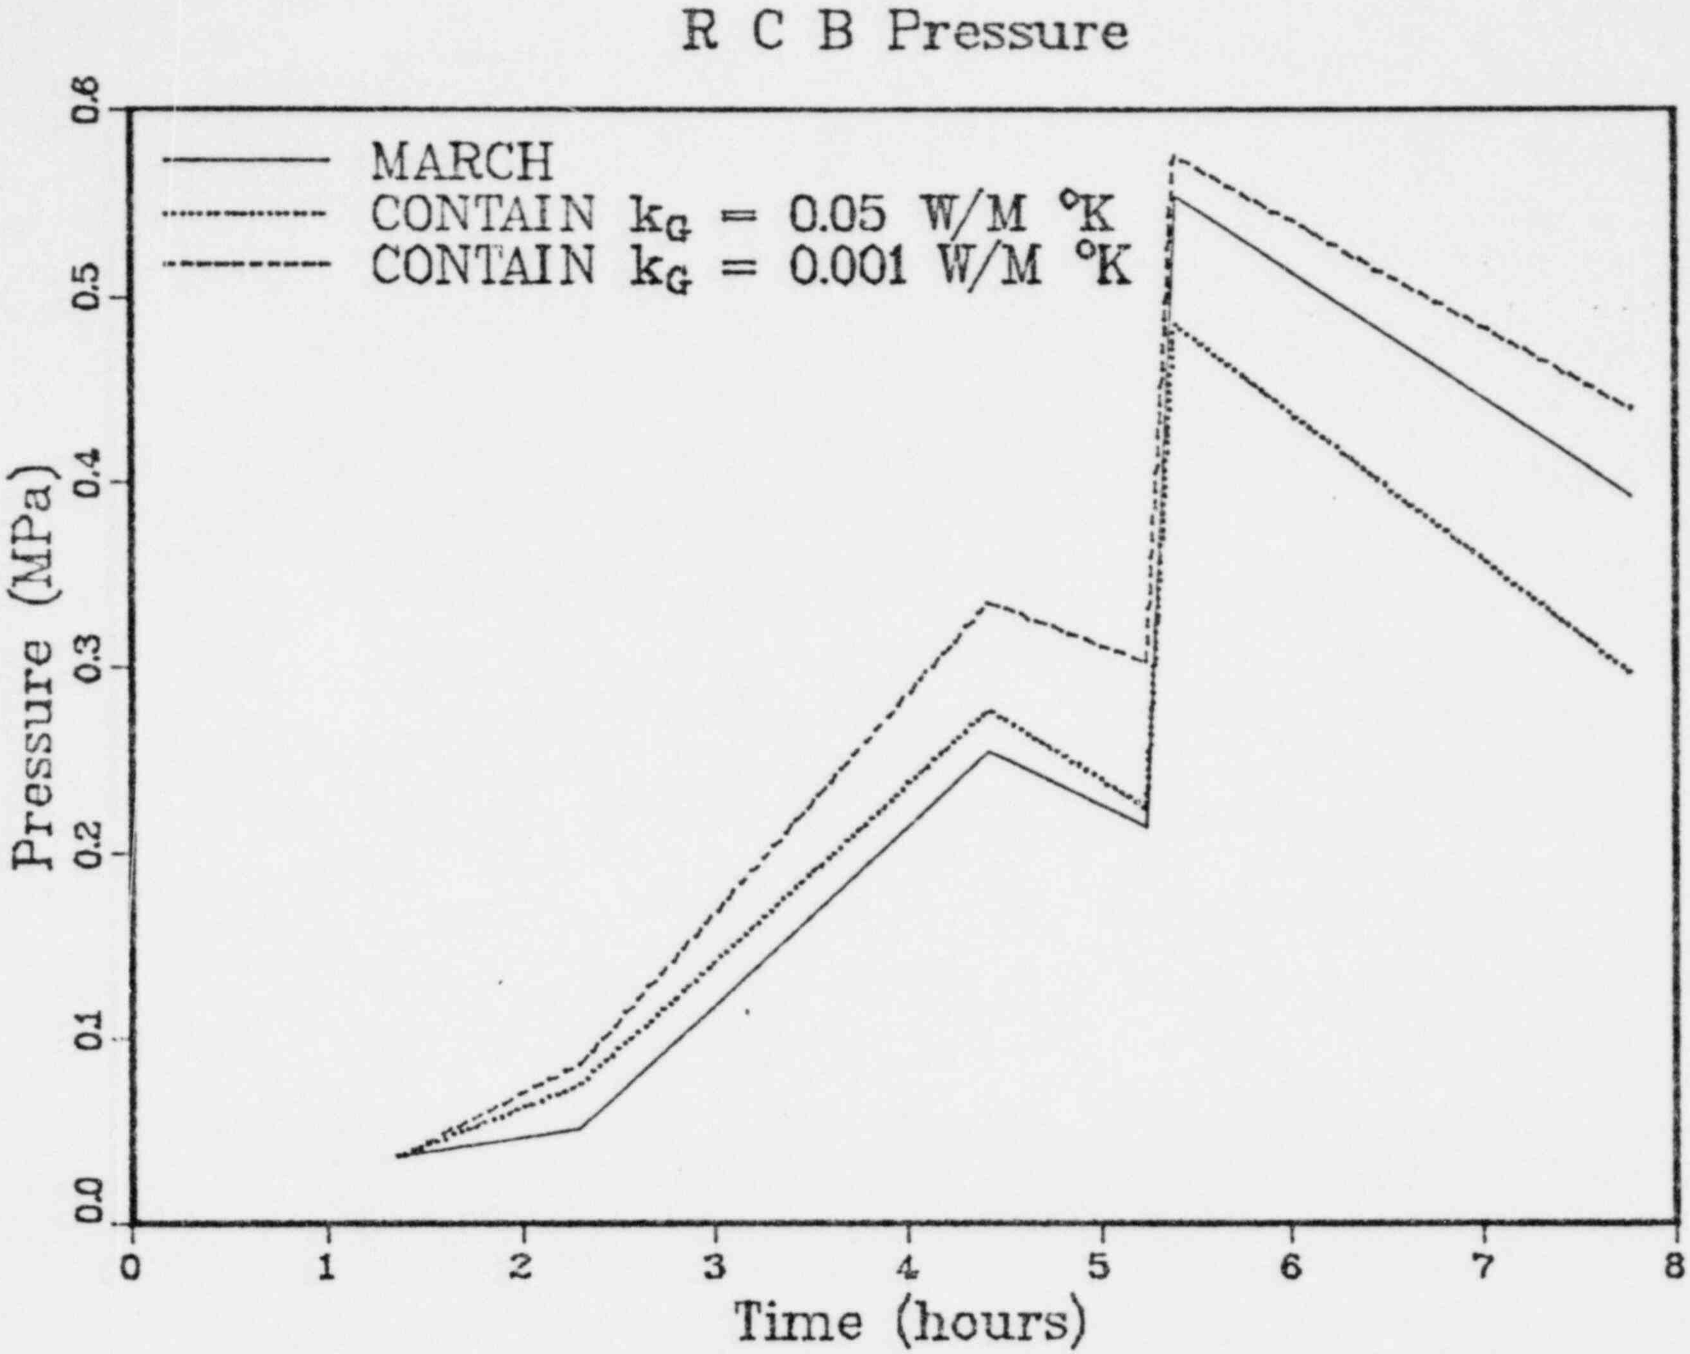

R C B Pressure
MARCH
CONTAIN kG = 0.05 W/M °K
CONTAIN kG = 0.001 W/M °K
Pressure (MPa)
0.0
0.1
0.2
0.3
0.4
0.5
0.6
Time (hours)
0
1
2
3
4
5
6
7
8

BOHL

LASL ANALYSIS OF STEAM EXPLOSIONS FOR THE Z/IP STUDY

Approach

- Use Sandia Experiments to Calibrate a 2-D Multiphase, Multicomponent Steam Explosion Expansion Model

- For Reactor Calculations, Use Heat Transfer Assumptions Consistent with Experiment Calibration

- Analyze Various Steam Explosion Expansions in Reactor Geometry Following from Assumed Premixed Interacting Configurations

LASL
LOS ALAMOS SCIENTIFIC LABORATORY

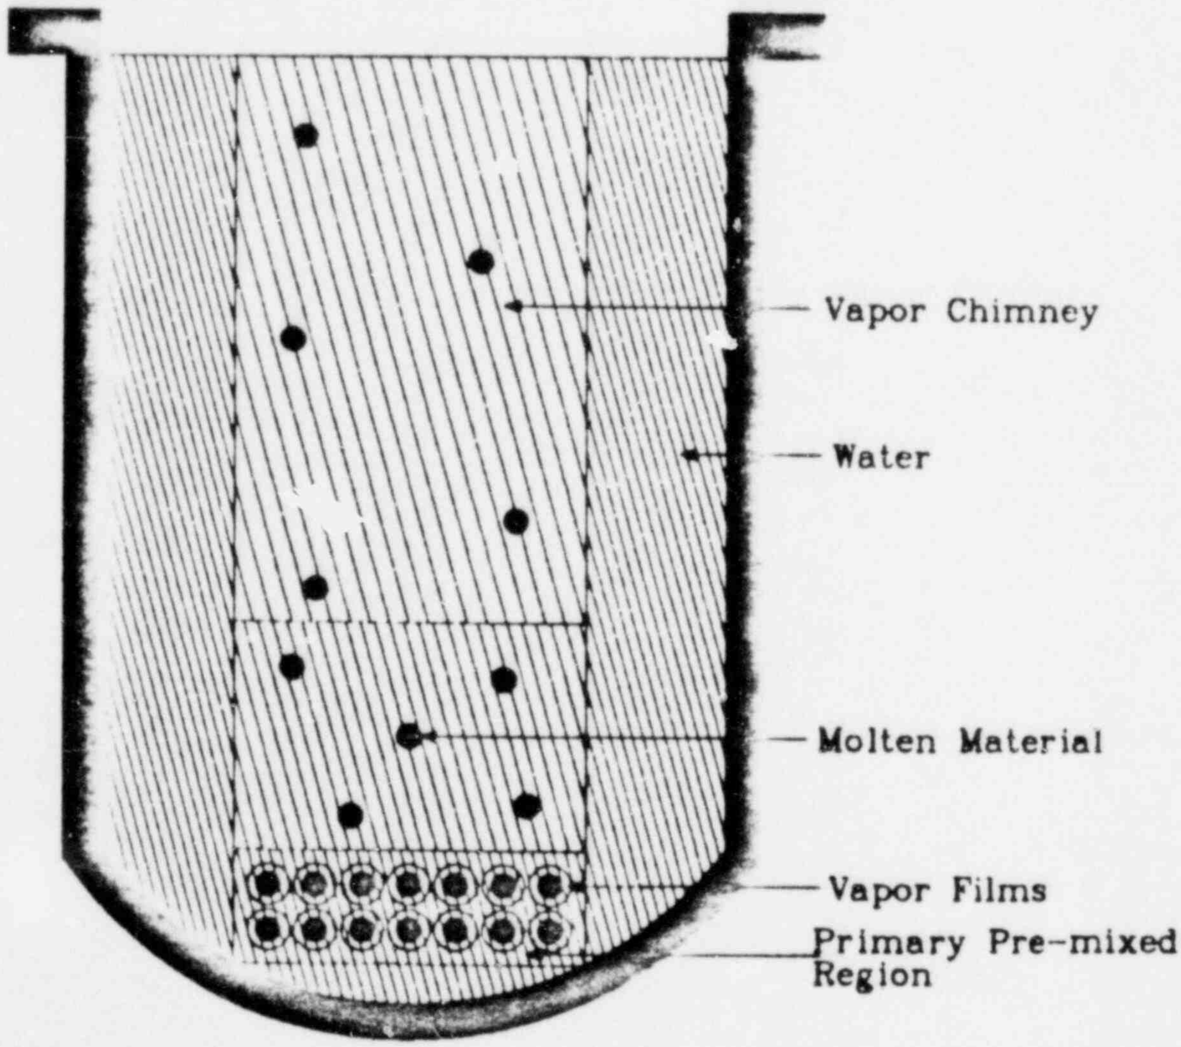

Vapor Chimney
Water
Molten Material
Vapor Films
Primary Pre-mixed Region

LASL ANALYSIS OF STEAM EXPLOSIONS FOR THE Z/IP STUDY

Findings–Experiment Calibration

- Key Assumptions –
  - Premixed Region – 0.1–m Radius
  - Vapor Chimney – 0.1–m Radius Above Premixed Region
  - 300 Micron Diameter Fuel Particle Size after Triggering
  - 50:25:25 (Fuel:Water:Vapor) Mixture Volume Fraction Ratios
- These Give Reasonable Agreement with Sandia Test 43 Pressure History (Rapid Rise to Near Critical and Rapid Decay) and Efficiency (0.5% Calculated, 0.43% Measured)
- Calibration Is Reasonable but Not Necessarily Unique

LOS ALAMOS SCIENTIFIC LABORATORY

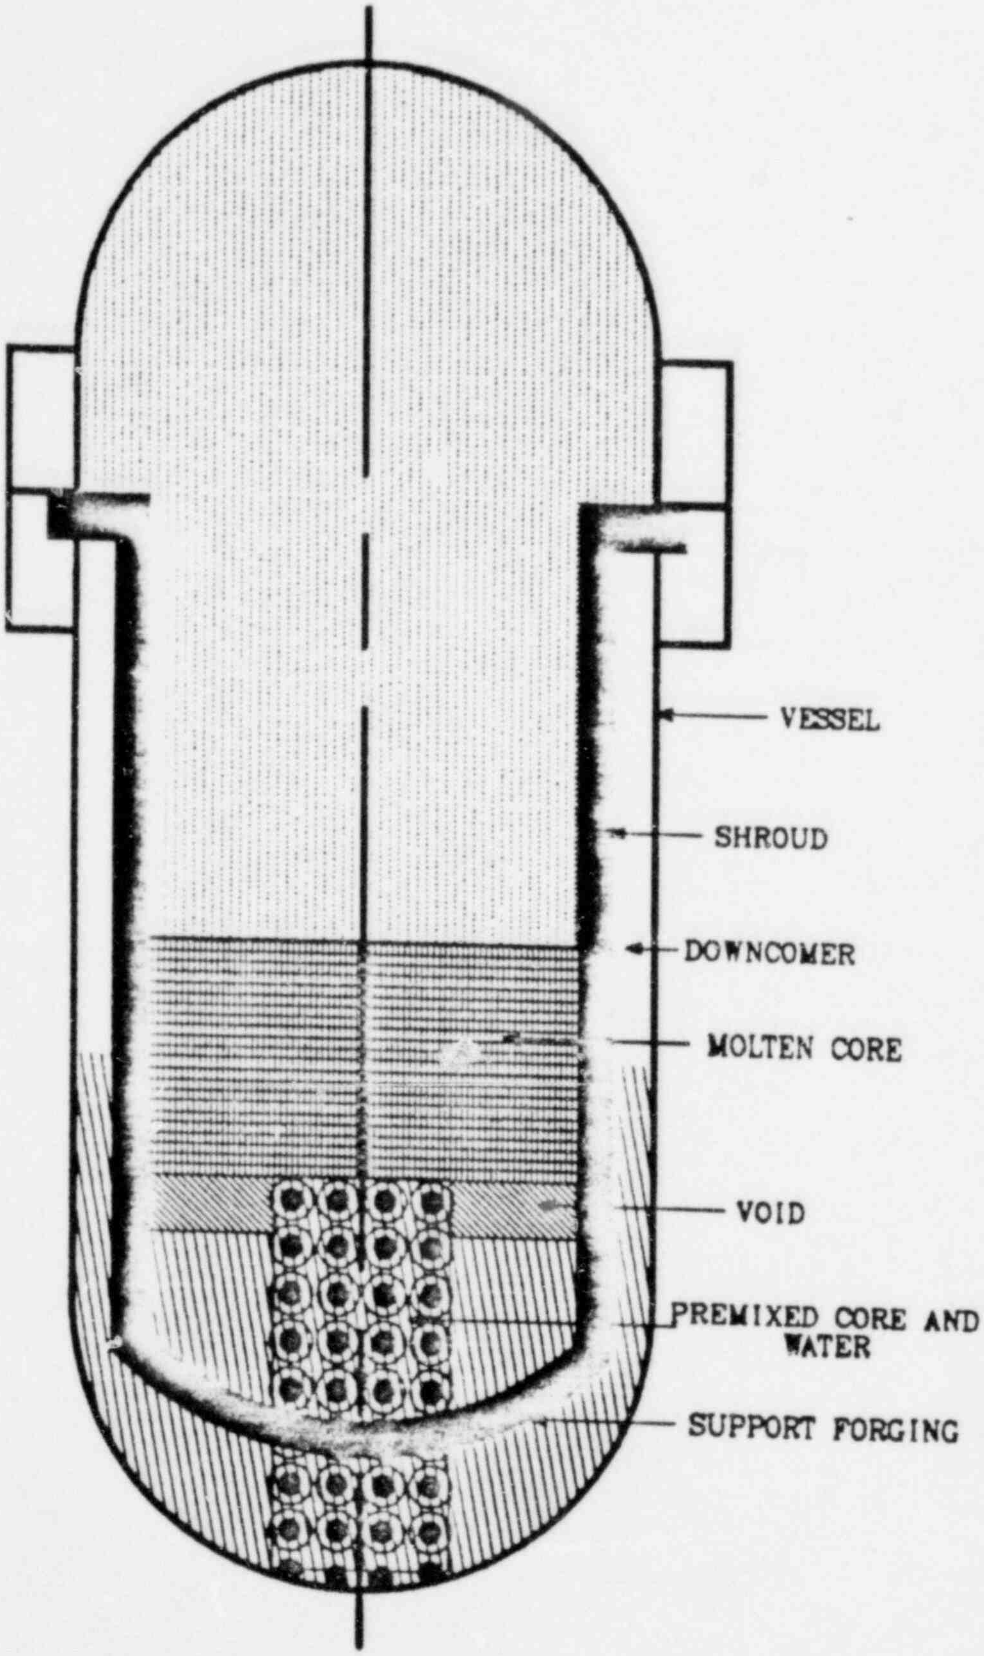
VESSEL
SHROUD
DOWNCOMER
MOLTEN CORE
VOID
PREMIXED CORE AND WATER
SUPPORT FORGING

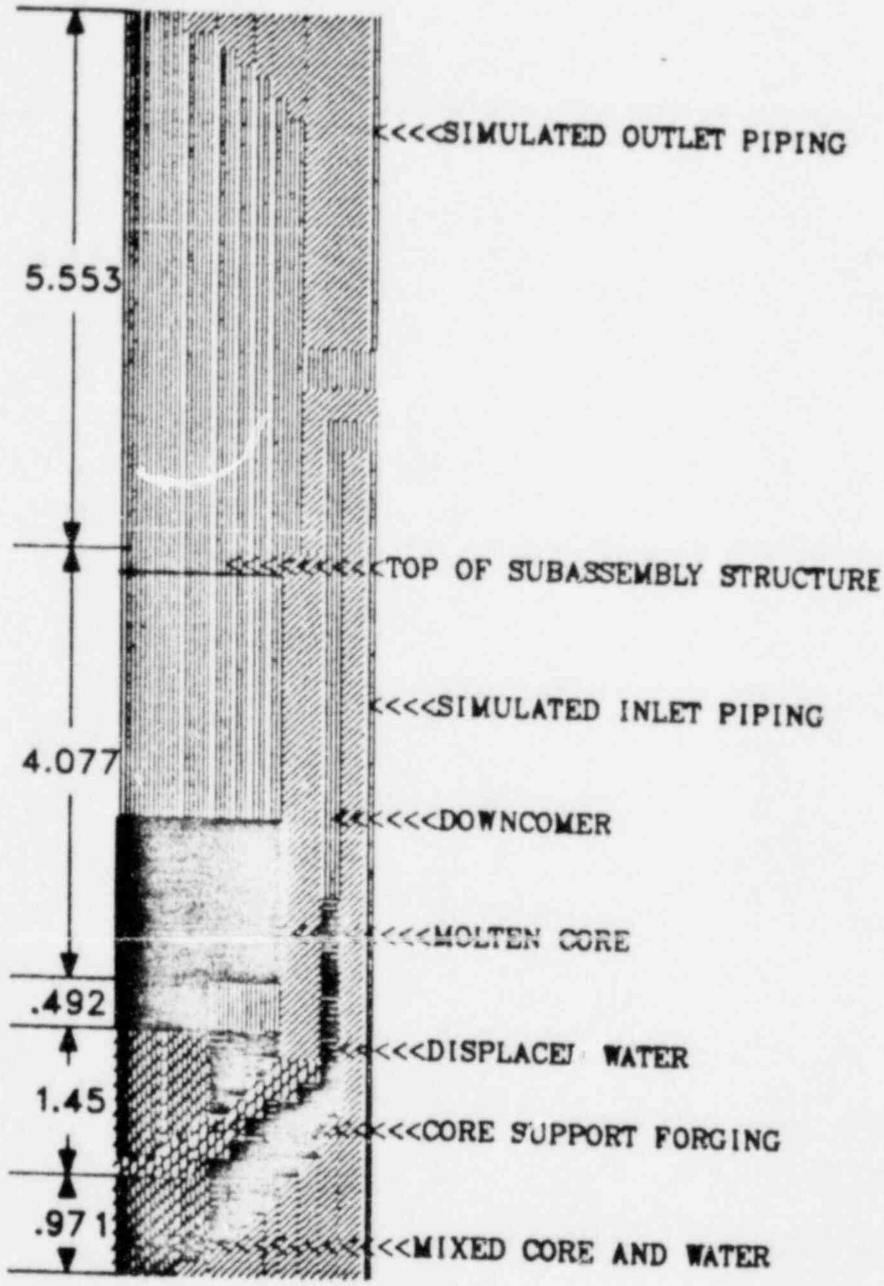

5.553
4.077
.492
1.45
.971
<<<<SIMULATED OUTLET PIPING
<<<<<<<<<<<TOP OF SUBASSEMBLY STRUCTURE
<<<<SIMULATED INLET PIPING
<<<<<<DOWNCOMER
<<<<<<MOLTEN CORE
<<<<<DISPLACED WATER
<<<<<<CORE SUPPORT FORGING
<<<<<<<<<<<MIXED CORE AND WATER

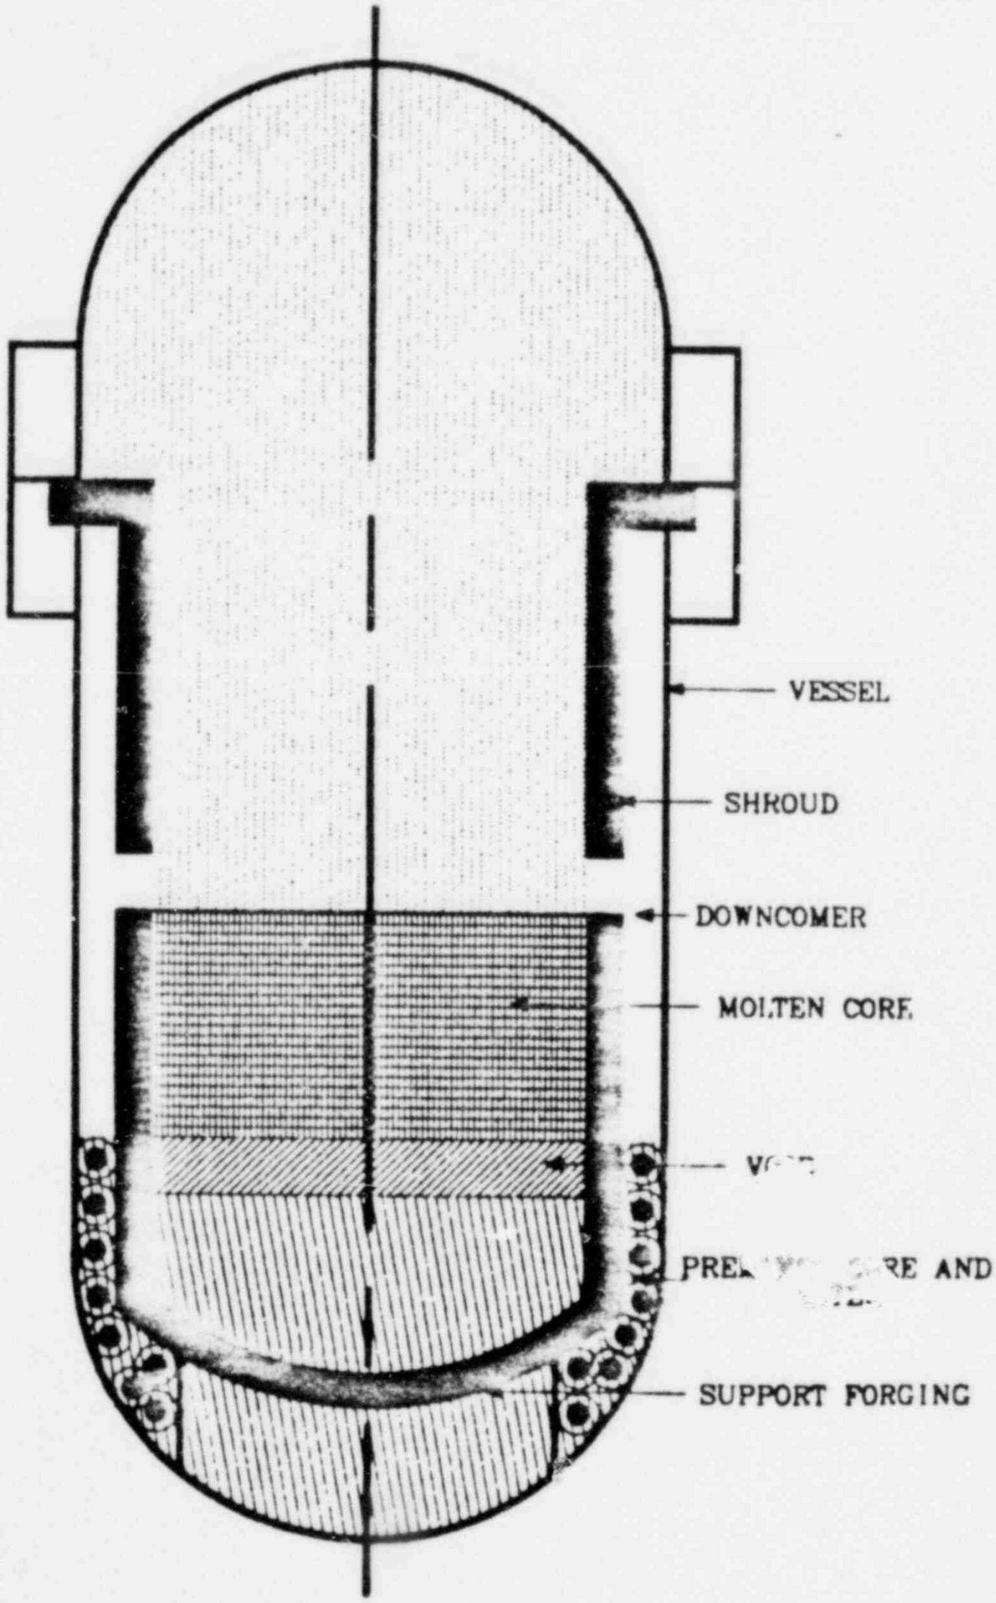
VESSEL
SHROUD
DOWNCOMER
MOLTEN CORE
SUPPORT FORGING

LASL A' \LYSIS OF STEAM EXPLOSIONS FOR THE Z/IP STUDY

Assumptions - Zion Analyses

- Same Heat Transfer Assumptions as for Experiment Calibration
- Pouring Mode of Mixing with 10-20% of Molten Core Materials Premixed with Water/Steam
- Overlying Molten Core Precludes Vapor Chimney and Provides Inertial Constraint
- No Vessel Internal Structures

LOS ALAMOS SCIENTIFIC LABORATORY

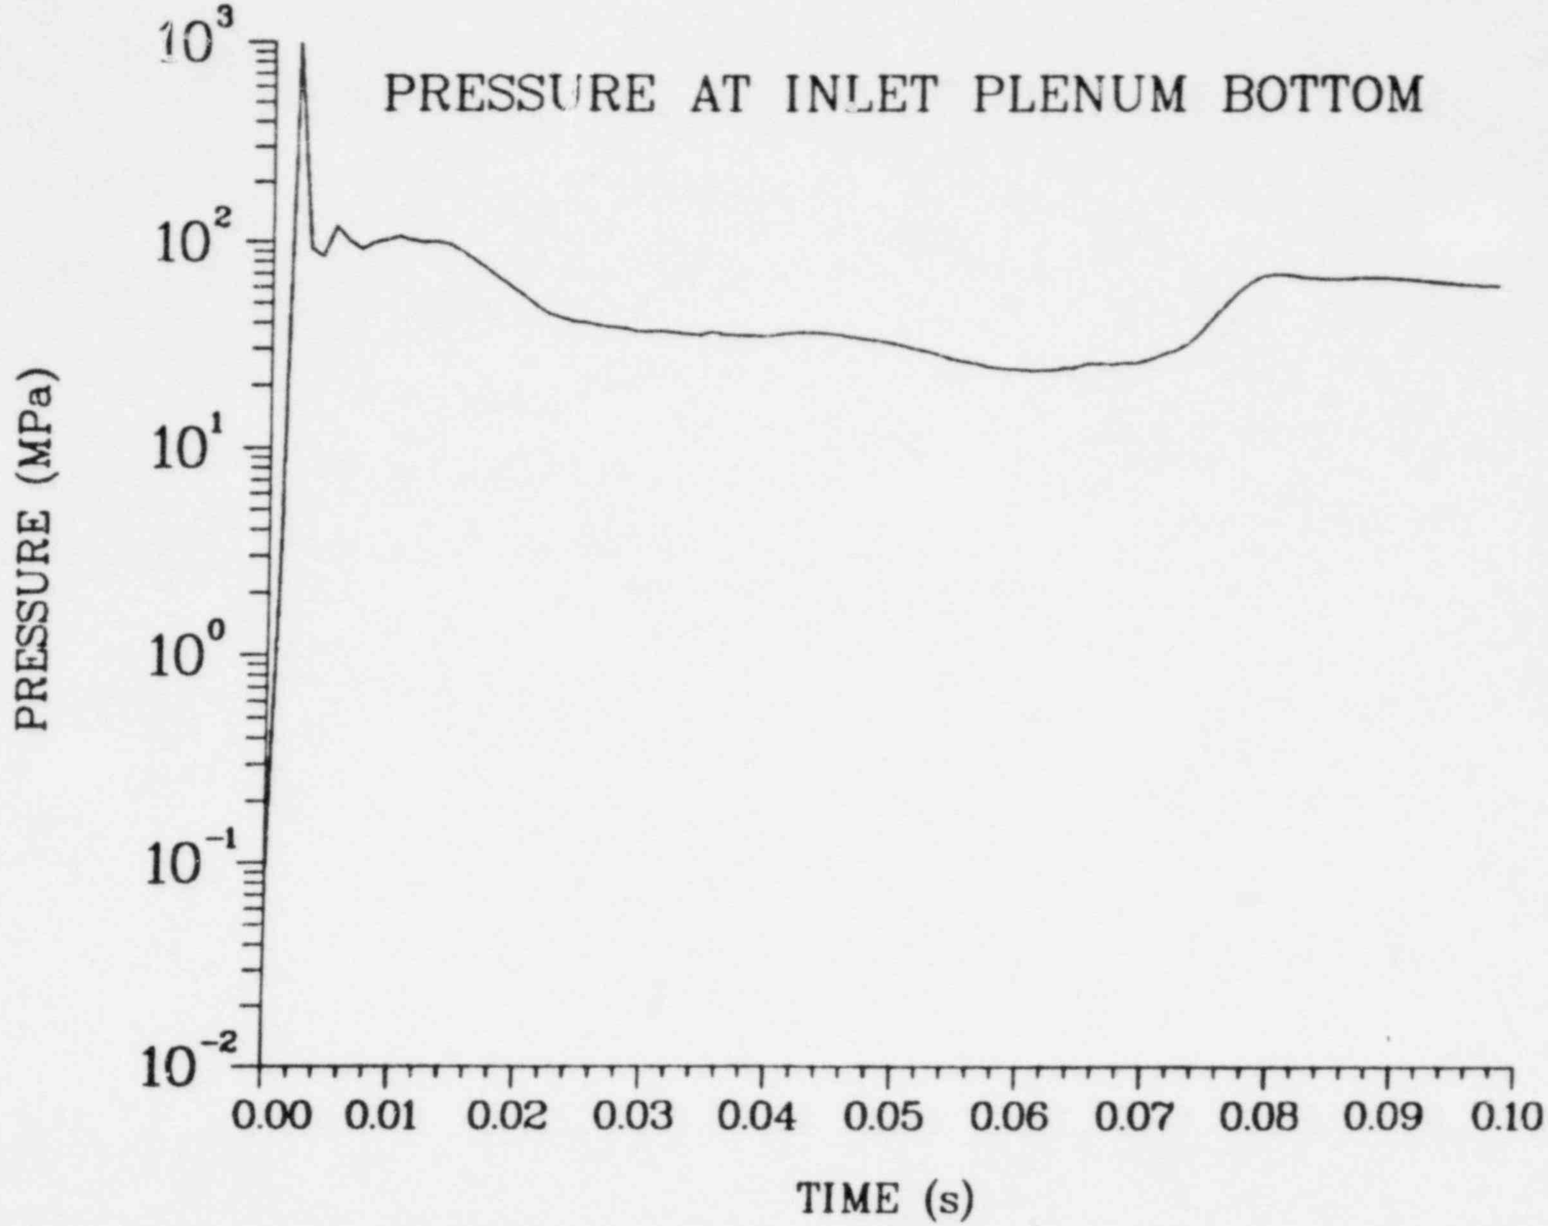
PRESSURE AT INLET PLENUM BOTTOM
PRESSURE (MPa)
10^3
10^2
10^1
10^0
10^-1
10^-2
0.00 0.01 0.02 0.03 0.04 0.05 0.06 0.07 0.08 0.09 0.10
TIME (s)

LASL ANALYSIS OF STEAM EXPLOSIONS FOR THE Z/IP STUDY

Findings - Zion Analyses - 1

- Inertial Constraint Lengthens Expansion Time and Increases Efficiency Relative to Experiment Simulation

- Given the Assumed Initial Configuration, Upward Directed Fuel Kinetic Energies of 1000 - 2000 MJ Seem Likely

- Better Quantification of Containment Failure Likelihood Must Consider the Core Melt Sequence, Incoherence in Fuel Dynamic Loading, and Structural Accommodation

LOS ALAMOS SCIENTIFIC LABORATORY

LASL ANALYSIS OF STEAM EXPLOSIONS FOR THE Z/IP STUDY

Findings - Zion Analyses - 2

- Two-Dimensional Behavior Strongly Influences Loading Dynamics
- Loadings Biased Towards Apex Decrease Likelihood of Large Missiles
- Lower Head Failures Likely Prior to Any Upper Head Failure -- Decreasing the Likelihood of Energetic Missile Generation in the Upward Direction
- Eventual Verification of Lower Probability for Containment Threats from Steam Explosions Is Likely

LASL
LOS ALAMOS SCIENTIFIC LABORATORY

WALKER

ELEVATED TEMPERATURE DESIGN ASSESSMENT

1.0 Introduction

The goal of this program at Sandia has been to improve our understanding of structural deformation and failure of engineering alloys under elevated temperature service conditions. In the course of this, we can provide input and guidance to the on-going development of elevated temperature design codes.

This program has involved tasks in three areas: structural analysis, materials testing and analysis, and non-destructive examination. Work in these three areas has been interconnected. For instance, mechanical test data are generated as inputs for structural analysis, and microstructural examinations give a basis for explaining and predicting alloy stress-strain and fracture behavior. Since the size of this program (FY 80 funding of $417K) does not permit a large effort in any one of these subjects, emphasis has been placed on coordinating the work of the task areas where practical. In addition, research areas have been selected which tend to complement, not duplicate, existing DOE-funded work for elevated temperature design.

2.0 Structural Analysis

Structural analysis activities are divided into three main areas: 1) numerical analysis of non-Sandia large scale component tests; 2) calculations in support of biaxial creep-fatigue experiments; 3) development of improved materials models for predicting plasticity and creep deformations.

2.1 Final reporting is being completed on time-dependent buckling calculations which compare analytical and experimental results from a Japanese PNC pipe elbow experiment. Analytical results were generated using two alternative existing finite elements, an inexpensive simplified element and a sophisticated double curved shell element, both from the MARC code. Calculations using the curved shell element demonstrated the need for using actual, as opposed to handbook, materials properties for these calculations.

2.2 Structural analyses have been reported for two geometries of the 316 stainless steel hollow tubular specimen used for biaxial creep-fatigue testing, Figure 1. These calculations address questions of possible specimen buckling under compressive strains and examines strains in the shoulder fillet region during monotonic and cyclic loadings. Results consist of specimen strains and displacements for monotonic compressive loading at 20°C and cyclic loading at 593°C; runs were made with and without internal pressure. Calculations for uniaxial compressive loading of both 0.635 mm and 1.27 mm wall thickness specimens predict no significant buckling of either specimen up to an axial strain of 5%. Experiments conducted on the specimens showed buckling limit strains of 2% and 5% for the thin-walled and thick-walled specimens, respectively. Cyclic deformation calculations predicted a small strain range increase in the near-shoulder region of the specimen, which would increase probability of failure in that location. Failure locations for fatigue experiments run to date, however, show no tendency for failure near the shoulder fillet, and it is possible that cyclic hardening, which was not accounted for in the calculations, acts to smooth out the strain distribution.

2.3 A task recently included in the structural analysis area of this program is to improve materials creep-plasticity deformation models used in finite element calculations. Analyses currently do not realistically handle cyclic hardening in fatigue or combined creep and plasticity. Cyclic hardening behavior will be added

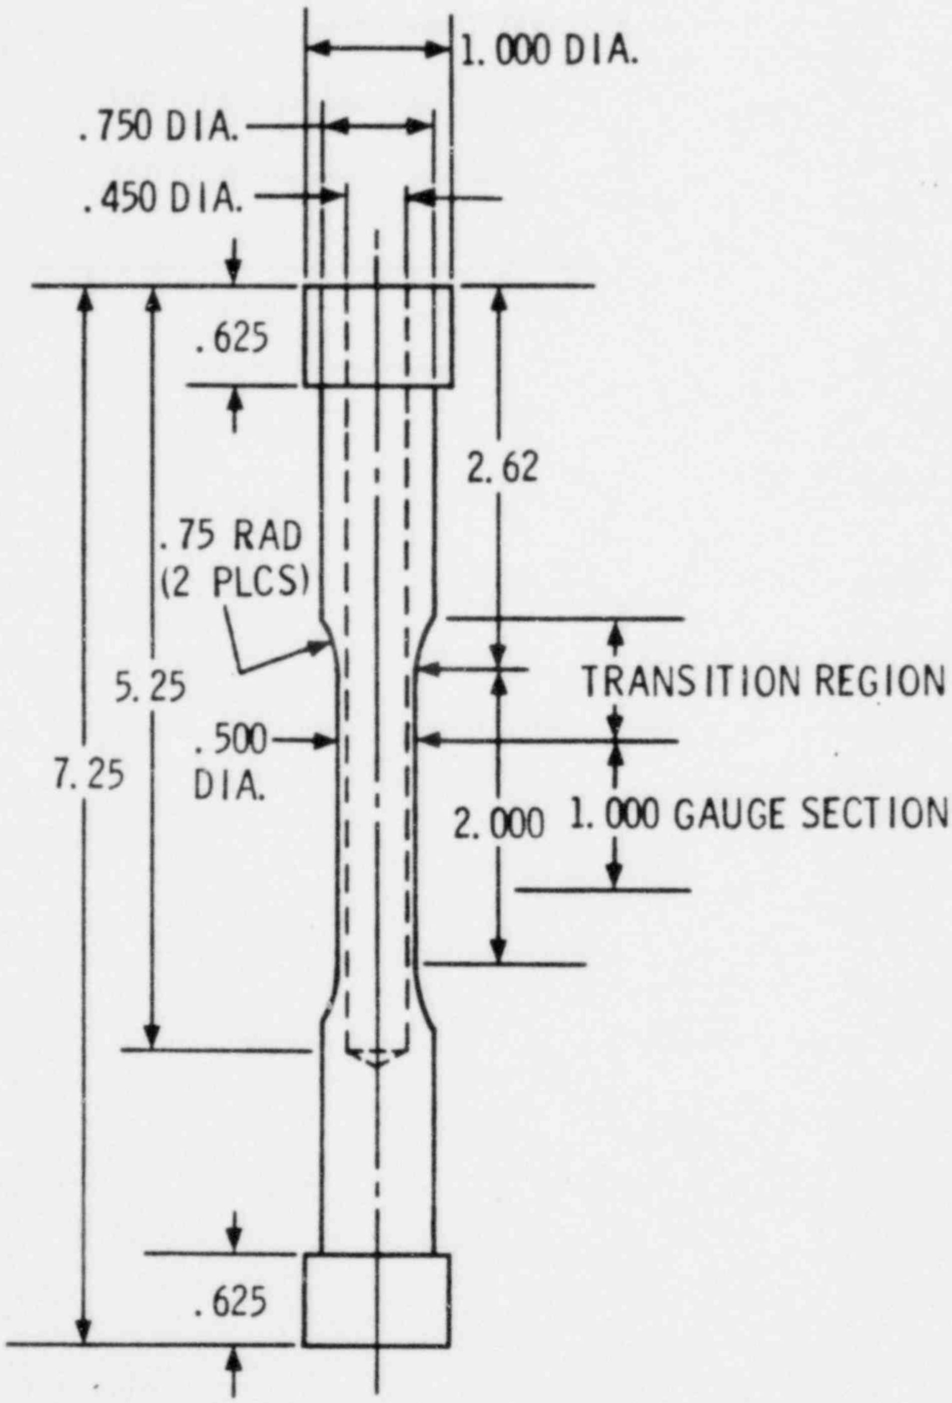

FIGURE 1. MUTIAXIAL SPECIMEN GEOMETRY.
(ALL DIMENSIONS IN INCHES)

to the MARC code, and the ability of several unified creep-plasticity models to predict experimentally measured behavior will be assessed. Both unified and biaxial analyses and experiments will be performed.

3.0 Non-Destructive Examination (NDE)

NDE studies were focused on attempts to measure microstructural changes which occur concurrently with creep-fatigue deformation. Such changes, although not necessarily associated with the specific damage process, could potentially be correlated with damage or fraction of life. Positron annihilation and ultrasonic techniques were investigated and compared with microstructural examinations by transmission electron microscopy. Neither of these techniques was found to provide a practical monitor of creep-fatigue damage in 316 stainless steel at elevated temperature.

A major conclusion was that microstructural changes in general do not correlate well with creep-fatigue life and that emphasis should be placed on reliable detection of small cracks. The use of crack detection methods for in-service NDE requires a methodology for dealing with them when they are detected: it becomes necessary to be able to predict critical crack lengths for overload failure and crack growth rates in creep-fatigue loading. This fracture mechanics approach has been successfully applied in many structures, including light water reactors, where stress conditions are linear-elastic. However, fracture mechanics is not well established either theoretically or experimentally for the fully plastic, time-dependent deformation loadings which may exist in structural components of advanced reactors. Some work in this area is being performed by Westinghouse Advanced Reactors Division to develop a leak-before-break design criteria to prevent sudden overload failures of flawed components.

As described in the following section, Sandia has initiated studies of fatigue crack growth under fully plastic conditions with and without a creep damage component.

4.0 Experimental Programs on Creep-Fatigue Behavior

The majority of research in this program is being conducted on 316 stainless steel. There are four general areas of interest: fatigue under biaxial loading, deformation modeling, crack growth studies, and metallurgical analysis of deformed samples.

4.1 A facility has been completed which can conduct elevated temperature fatigue testing of hollow tubular biaxial specimens with internal pressure and axial push-pull loading. The test specimen geometry is shown in Figure 1.

An initial series of thirteen low cycle fatigue tests at 593°C has been completed. The failure results are plotted in Figure as axial plastic strain range versus cycles to failure. Results from two sets of round bar uniaxial fatigue tests are also included. The tubular specimen test data are well-behaved and exhibit the typical Coffin-Manson low cycle fatigue life behavior. Comparison with the round bar results show that both the thin and thick-walled tubular specimens tend to fail much earlier at a given strain range, and the thin-walled samples fail before the thick walled. It is believed that these differences are caused in part by the increased number of cycles required for crack propagation through the thicker walled tube and the solid bar specimens. Data discussed below indicate that crack propagation in low cycle fatigue of 316 stainless steel comprises a relatively large fraction of total fatigue life.

The internally pressurized, non-hold period tests in Figure 2 show no change in life from the zero hoop stress tests. However, two samples with one minute tensile hold period are decreased in fatigue cyclic life by a factor of two. Additional data reduction on these tests are being carried out to investigate cyclic hardening and diametral ratchetting behavior.

4.2 Testing for deformation modeling includes both uniaxial and biaxial experiments to generate stress-strain-time flow properties. In addition, several transient stress relaxation studies are to be run to generate materials parameters to fit unified creep-plasticity deformation models. This area of study interacts closely with structural analysis code development activities.

4.3 Experiments to measure crack growth during fully plastic straining typical of low cycle fatigue are being performed in two ways. First, local crack growth rates defined by fatigue striation spacings have been measured on fracture surfaces of failed specimens. These striation spacings are measured in the scanning electron microscope as a function of crack length, and these data are integrated to give plots of growth rate versus cycle number and of total cycles to propagate (Np). Knowing Np and the number of cycles to failure, one can back calculate the number of cycles to initiate the fatigue crack. An example of this is shown in Figure 3 which indicates fraction of life spent in crack initiation versus plastic strain range for two reactor structural alloys--Incoloy 800 and 316 stainless steel. Data for 316 stainless steel are still incomplete; but the trend indicates considerably different behavior for the two alloys: at a given strain range, Incoloy 800 spends a much larger fraction of its cyclic life initiating fatigue cracks; whereas 316 stainless steel, particularly at large strain ranges typical of most materials

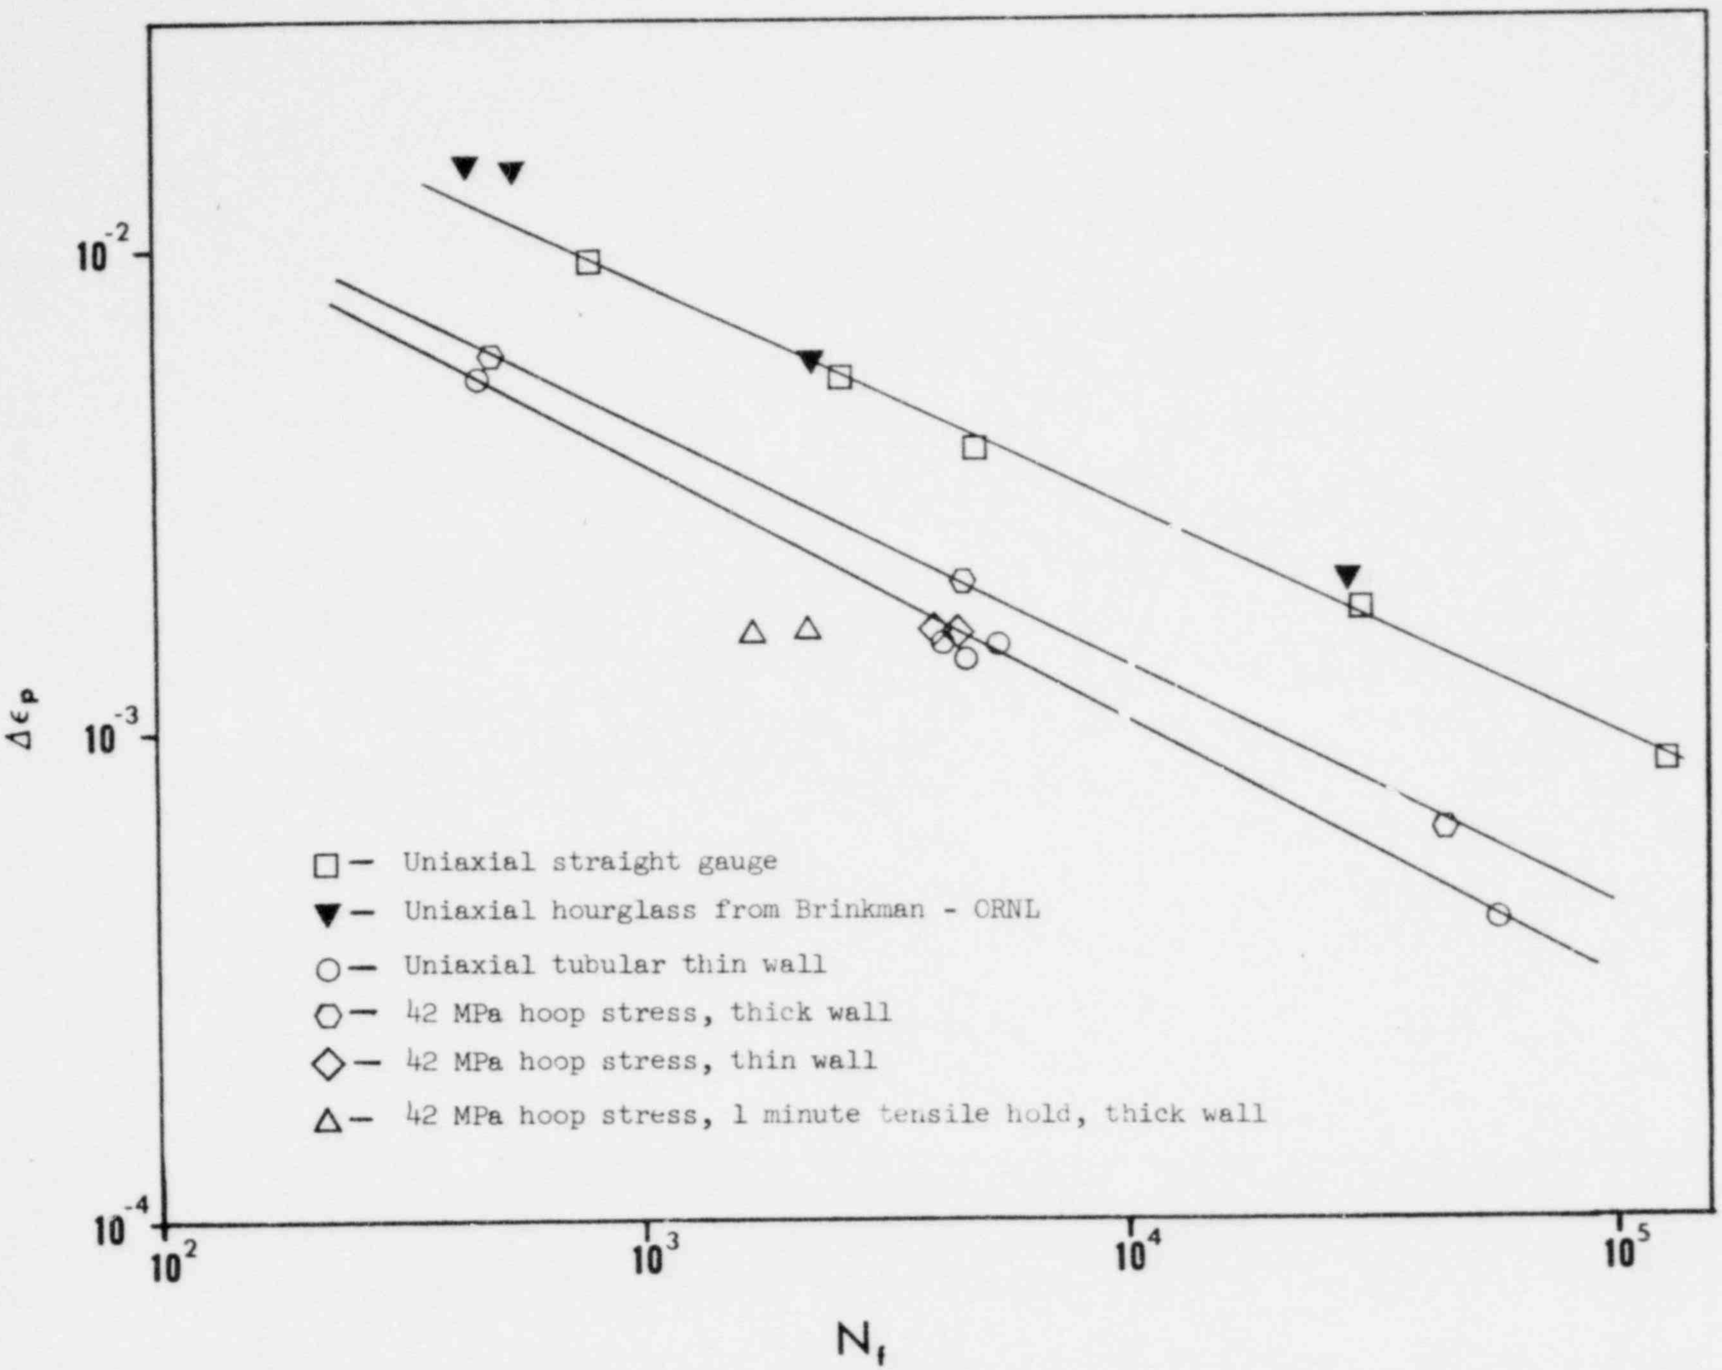

Figure 2. Low Cycle Fatigue Results for Biaxial Loadings

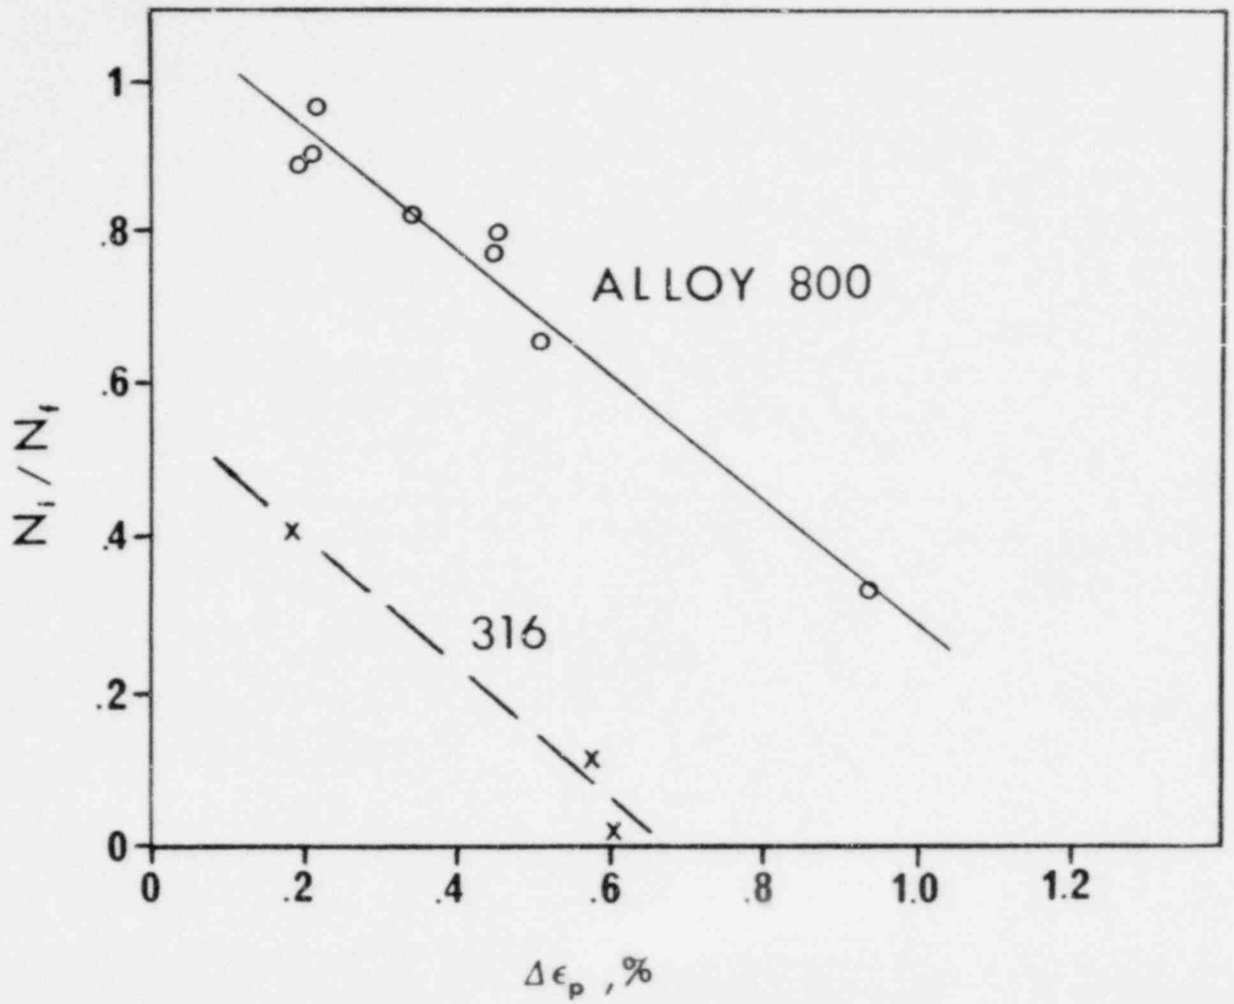

Figure 3. Fraction of Low Cycle Fatigue Life Spent in Crack Initiation.

testing, initiates quickly and spends virtually all of its life in crack propagation. The implications of this on design rule development for the two alloys is being assessed.

Additional crack growth rate experiments are being conducted using single edge notch specimens of rectangular cross-section and monitoring crack growth both optically and by potential drop techniques. This method allows growth rate measurements even when striations are not formed--as in tensile hold period tests where crack propagation is intergranular. These studies will indicate growth rates in 316 stainless steel under fully plastic straining with and without superimposed creep damage. Two goals of these measurements are 1) to investigate mechanical parameters which control crack growth rates for fully plastic loading (analogous to ΔK for linear-elastic loading) and 2) to characterize growth rates in combined creep-fatigue loading. In addition, a study will be conducted on pre-conditioned specimens with metallurgical structures more typical of mid-life service.

4.4 Metallurgical analysis of deformed specimens by transmission electron microscopy (TEM) has been an on-going task in this program. Its goals are to relate observed microstructural changes to measured flow properties and fracture processes.

An important finding of these investigations has been that microstructural changes, particularly carbide precipitation in 316 stainless steel, occur during cycling at elevated temperature which never occur by simple aging processes or steady-state creep deformation. Hold periods combined with cycling loading particularly increase precipitation of fine carbides at 593°C, and for long test times these apparently increase cyclic hardening behavior.

TEM examinations were completed on a 2.25 Cr-1 Mo steel specimen obtained from C. R. Brinkman at Oak Ridge National Laboratory. The specimen was cycled at 538°C at the low strain range of 0.1% with a 0.05 h compressive hold period. It ran under these conditions for 8124 hours at which time the strain range was increased to 0.40% and cycling was continued to failure. Of particular interest to Sandia were the dislocation and precipitate substructures. Earlier work at Sandia has shown that cycling at 593°C results in significant changes in carbide precipitation that do not occur during thermal aging or creep deformation. Examination of the ORNL sample showed that only thermal aging effects had occurred and that the strain range-temperature combination used in the test did not result in added carbide precipitation. In addition, the dislocation density observed in the foils was very low considering the 3733 cycles at $\Delta\epsilon$ = 0.4% given the specimen prior to failure. These observations indicate that accelerated testing of 2.25 Cr-1 Mo steel using higher-than-service temperatures needs to be re-evaluated.

WALKER

IMPROVED TEST TECHNOLOGY - 1980

FACILITIES

- LARGE MELT FACILITY
- ACRR ADVANCED MODES
- HOT CELL FACILITIES
- DATA ACQUISITION AND DISPLAY SYSTEM
- SODIUM SUPPORT FACILITY

DIAGNOSTICS

- CODED APERTURE IMAGING SYSTEM
- IN CORE FUEL MOTION SYSTEM
- ULTRASONIC THERMOMETRY
- AEROSOL SAMPLING
- OPTICAL FUEL MOTION

TEST HARDWARE

- ACRR NA LOOP
- BOTTOM COOLED PAHR CAPSULE DSN
- IN PILE TRANSITION PHASE

RECENT TESTS CONDUCTED - SNL

Accident Energetics

ACRR

- Prompt burst energetics - 13S & 14S (2)
- Equation of state (1)
- Fuel disruption - 2.1, 2.2, 2.3, 2.4 (4)
- High ramp rate disassembly - HRR 1, 2, 3, 4 (4)
- Coded aperture imaging (2)
- In core fuel motion detection - 7 pin, 17 pin (2)

Core Debris Behavior

ACRR

- Debris bed PAHR - D-4 (1)
- Molten pool PAHR - MP-4 (1)

Molten Core Technology

- Large melt facility - LMF 1 (1)
- MgO-UO_2 compatibility (15)
- High Al cement - HAC-2 (1)
- UO_2-concrete code comparison tests - CC1, CC2 (2)

Sodium Concrete

- Na-magnetite concrete interaction - T-14, 15 (2)
- Faulted liner tests - FLT 1, 2, 3, 4 (4)

FOREIGN COLLABORATIONS

ACRR-CABRI EXCHANGE AGREEMENT

CABRI:

- SNL & LASL STAFF AT CADARACHE
- PRE & POST TEST CALCULATIONS
- DIAGNOSTICS - SNL STAFF

ACRR:

- SG CARBIDE PBE TESTS - NRC/KFK
- HRR DISASSEMBLY TESTS - NRC/KFK
- ACRR EOS TESTS - NRC/KFK
- FD-4 FUEL DISRUPTION - NRC/KFK

CORE DEBRIS COOLABILITY - 1980-1983

- NRC/EURATOM/JAPAN SPONSORSHIP
- 10 ACRR DEBRIS BED TESTS
- 12 ACRR MOLTEN POOL TESTS

TRANSITION PHASE

- TECHNICAL DEFINITION STAGE
- LARGE MELT FACILITY
- NRC/KFK

FUEL MOTION

- In core diagnostics - BR-2 MOL 7C
- IEFM - UKAEA, KfK

STAFF

- Dr. Wolfgang Breitung - KfK
- Dr. Frank Briscoe - UKAEA
- Mr. Michel Schwarz - CEA
- (Dr. Guenter Fieg - KfK)
- (Mr. Katsuro Takahashi - PNC)

ARSR PROGRAM MILESTONES

ACCIDENT DELINEATION

ACCIDENT ENERGETICS

- Prompt Burst Work Potential
- Equation of State
- Fuel Disruption

POST ACCIDENT CONTAINMENT

- Debris Coolability
- Materials Interaction
- CONTAIN Code

POWERS

Ex-vessel Core Debris and Coolant Interactions with Reactor Cavity Materials

D. A. Powers

Presentation to the ACRS on June 30, 1980
at Los Angeles, CA

MODES OF
EX-VESSEL INTERACTIONS

- *MOLTEN SODIUM ONLY*
- *MOLTEN SODIUM and FRAGMENTED CORE DEBRIS*
- *MOLTEN SODIUM and MOLTEN CORE DEBRIS*
- *MOLTEN CORE DEBRIS ONLY*

WHY EX-VESSEL INTERACTIONS ARE OF CONCERN

* *INTERACTIONS THREATEN CONTAINMENT*

* *INTERACTIONS PROVIDE A RADIOACTIVE SOURCE TERM*

THREATS TO CONTAINMENT

- GAS GENERATION
- FLAMMABLE PRODUCTION
- ENERGY TRANSPORT
- BASEMAT PENETRATION

SOURCE TERM

* RADIOACTIVE AEROSOLS :

  - SPARGING
  - VAPORIZATION
  - CHEMICAL TRANSPORT

* AEROSOLS ALSO THREATEN MITIGATION SYSTEMS :

  - COAT AIR COOLERS
  - CLOG FILTERS
  - PLUG VENTS

TECHNICAL ISSUES RAISED BY EX-VESSEL INTERACTIONS

CORE CATCHER SYSTEM		*UNPROTECTED PLANT*		*FILTERED VENT SYSTEM*
		⇓		⇓
		GAS GENERATION	⇐	GAS GENERATION
⇓		H_2, CO, CH_4 PRODUCTION	⇐	H_2, CO, CH_4 PRODUCTION
UPWARD HEAT TRANSFER	⇒	UPWARD HEAT TRANSFER		
AEROSOL GENERATION	⇒	AEROSOL GENERATION	⇐	AEROSOL GENERATION
BASEMAT EROSION	⇒	BASEMAT EROSION		

SODIUM/CONCRETE INTERACTIONS

SAFETY QUESTIONS

- *MAGNITUDE OF THE PHENOMENA*
  - *GAS GENERATION*
  - *AEROSOL GENERATION*
  - *EROSION RATE*
- *IS THE INTERACTION SELF-LIMITING?*
- *DO LINERS HAVE TO BE ENGINEERED SAFETY FEATURES?*

Sodium - Concrete Interactions
Program Accomplishments in FY80

1) Interaction of sodium with magnetite concrete characterized in two large-scale tests

2) Phase stability studies of the Na-O-H system

3) Hydrology of sodium flow through faulted liners studied

4) USINT model of concrete behavior declared the state-of-the-art model in a code comparison effort at General Electric

5) Chemical hypothesis of sodium/concrete interaction formulated into a mathematical model

POSSIBLE MECHANISMS FOR LIMITING INTERACTIONS

- DEPLETION OF THE SODIUM
- DEPLETION OF CONCRETE & WATER
- REACTION PRODUCT BARRIER

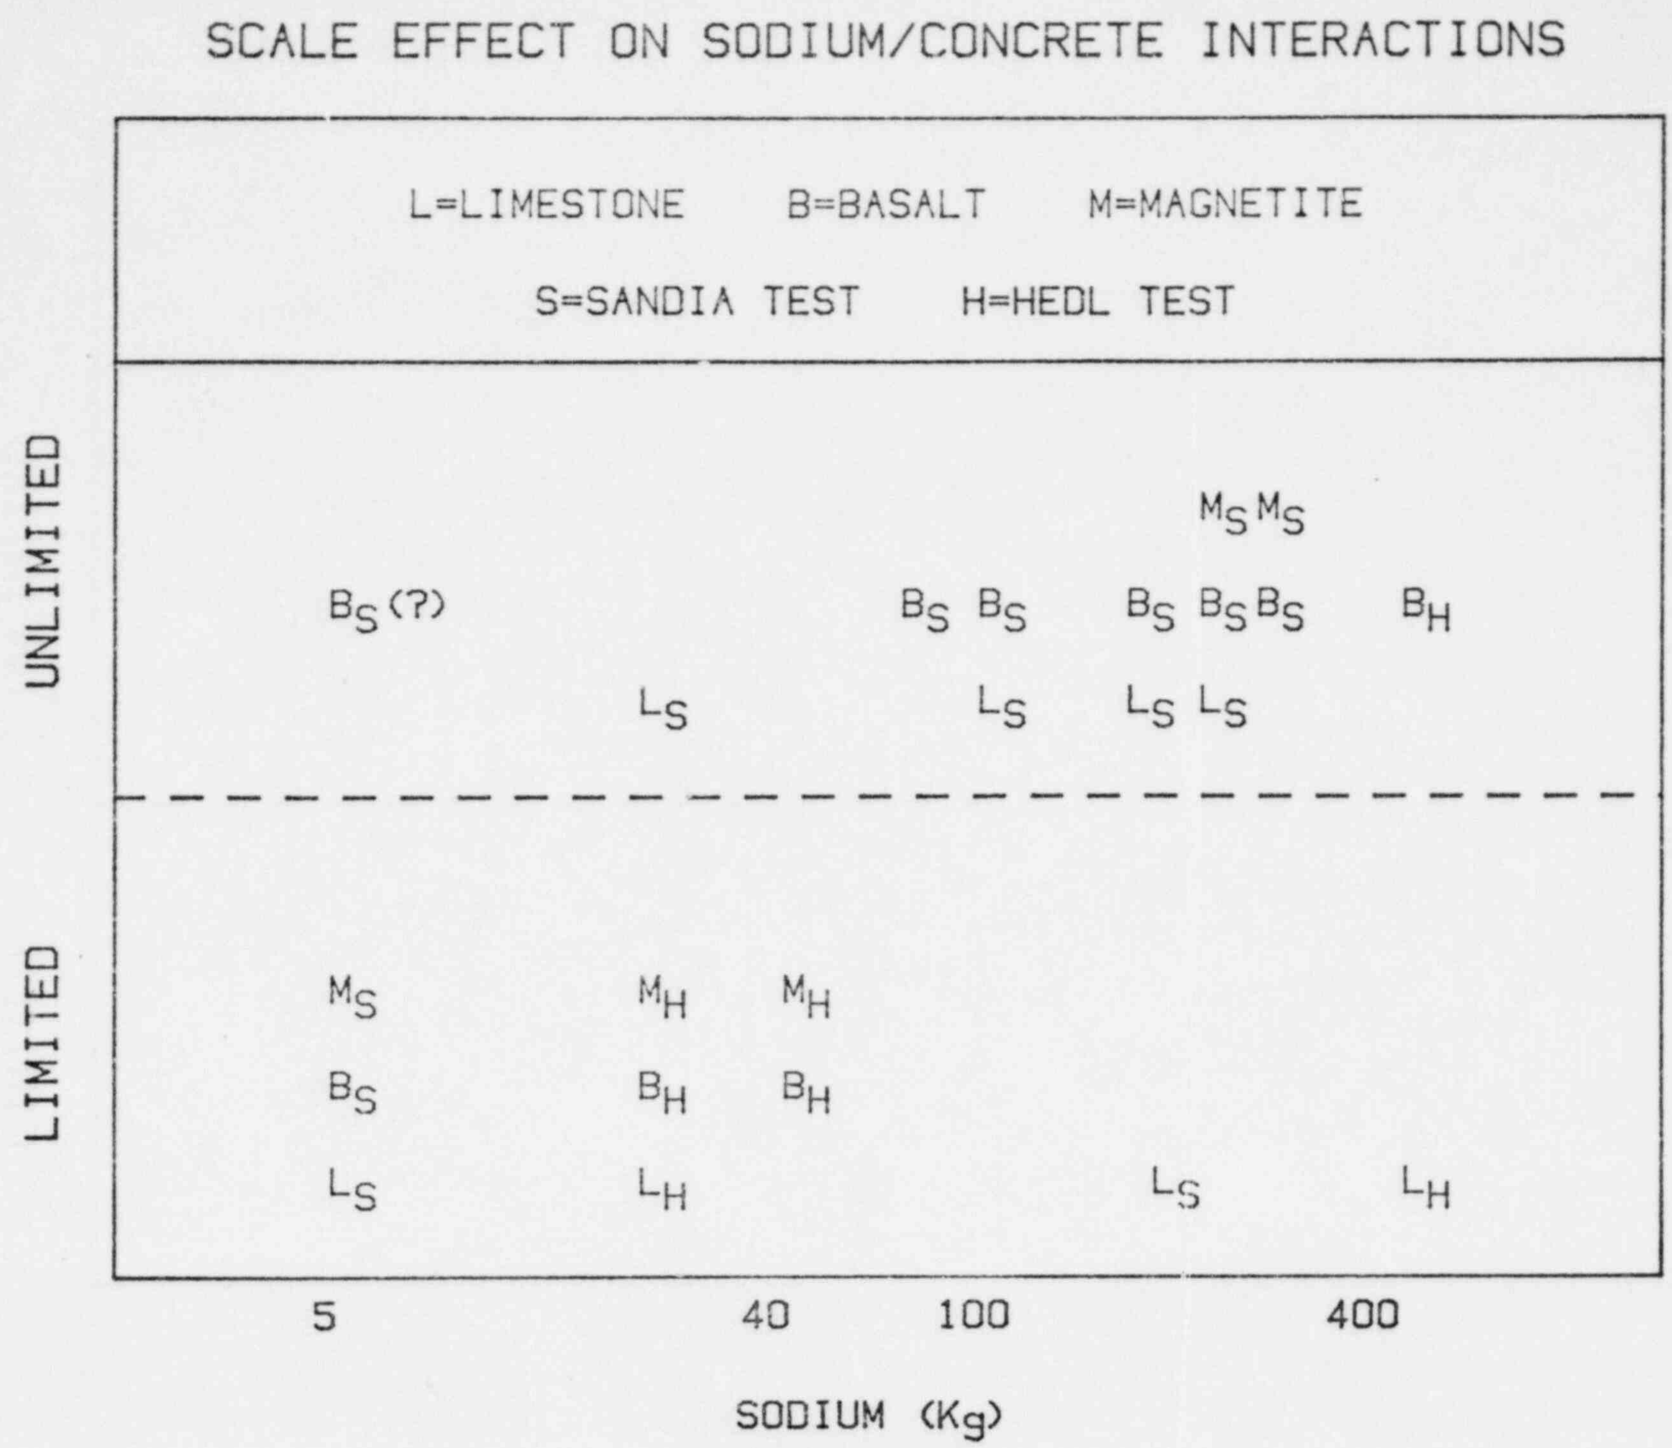

SCALE EFFECT ON SODIUM/CONCRETE INTERACTIONS
L=LIMESTONE B=BASALT M=MAGNETITE
S=SANDIA TEST H=HEDL TEST
UNLIMITED
LIMITED
5
40
100
400
SODIUM (Kg)

MELT TEMP. DURING INTERACTION WITH CONCRETE

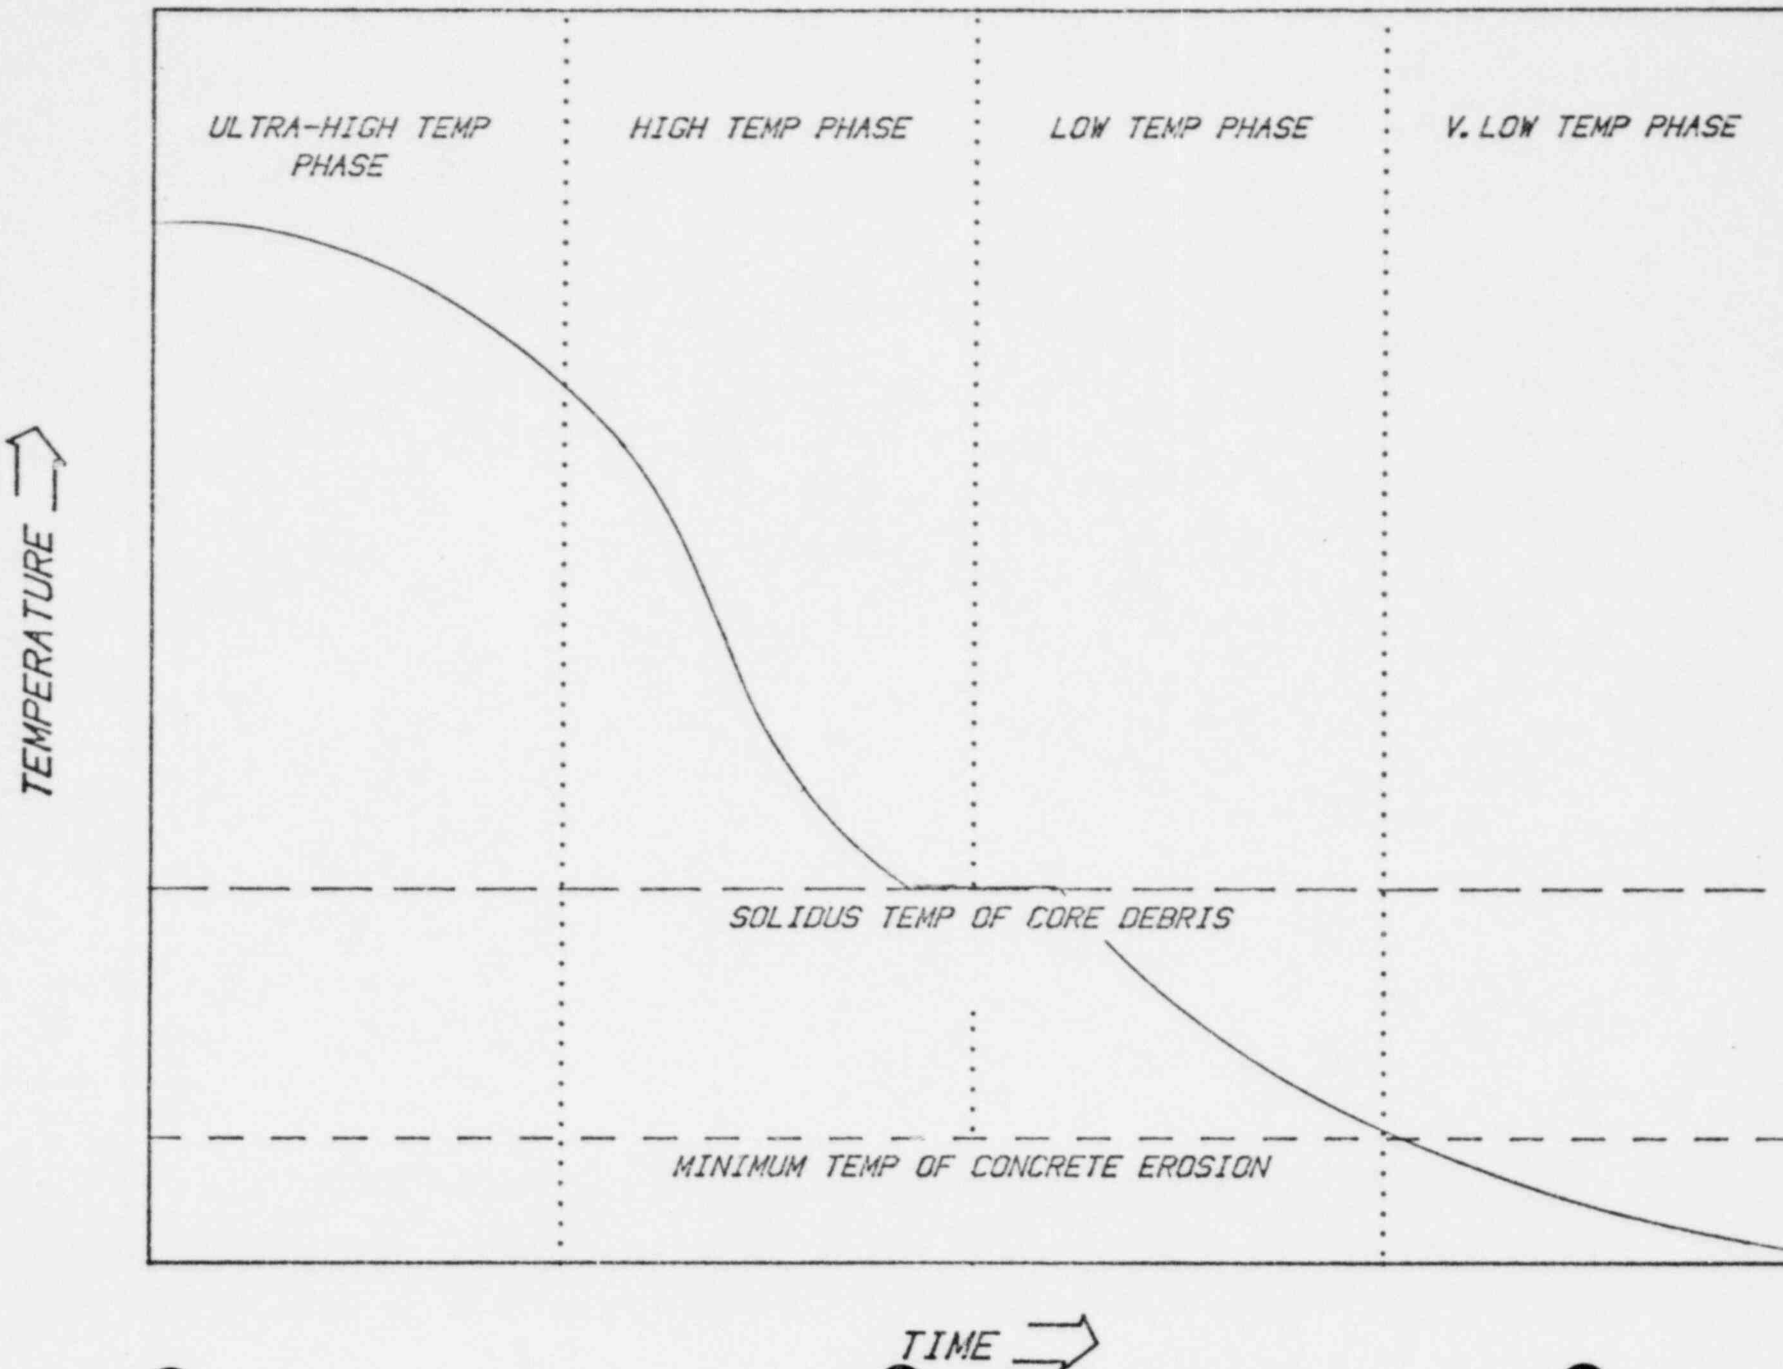

MELT TEMP. DURING INTERACTION WITH CONCRETE

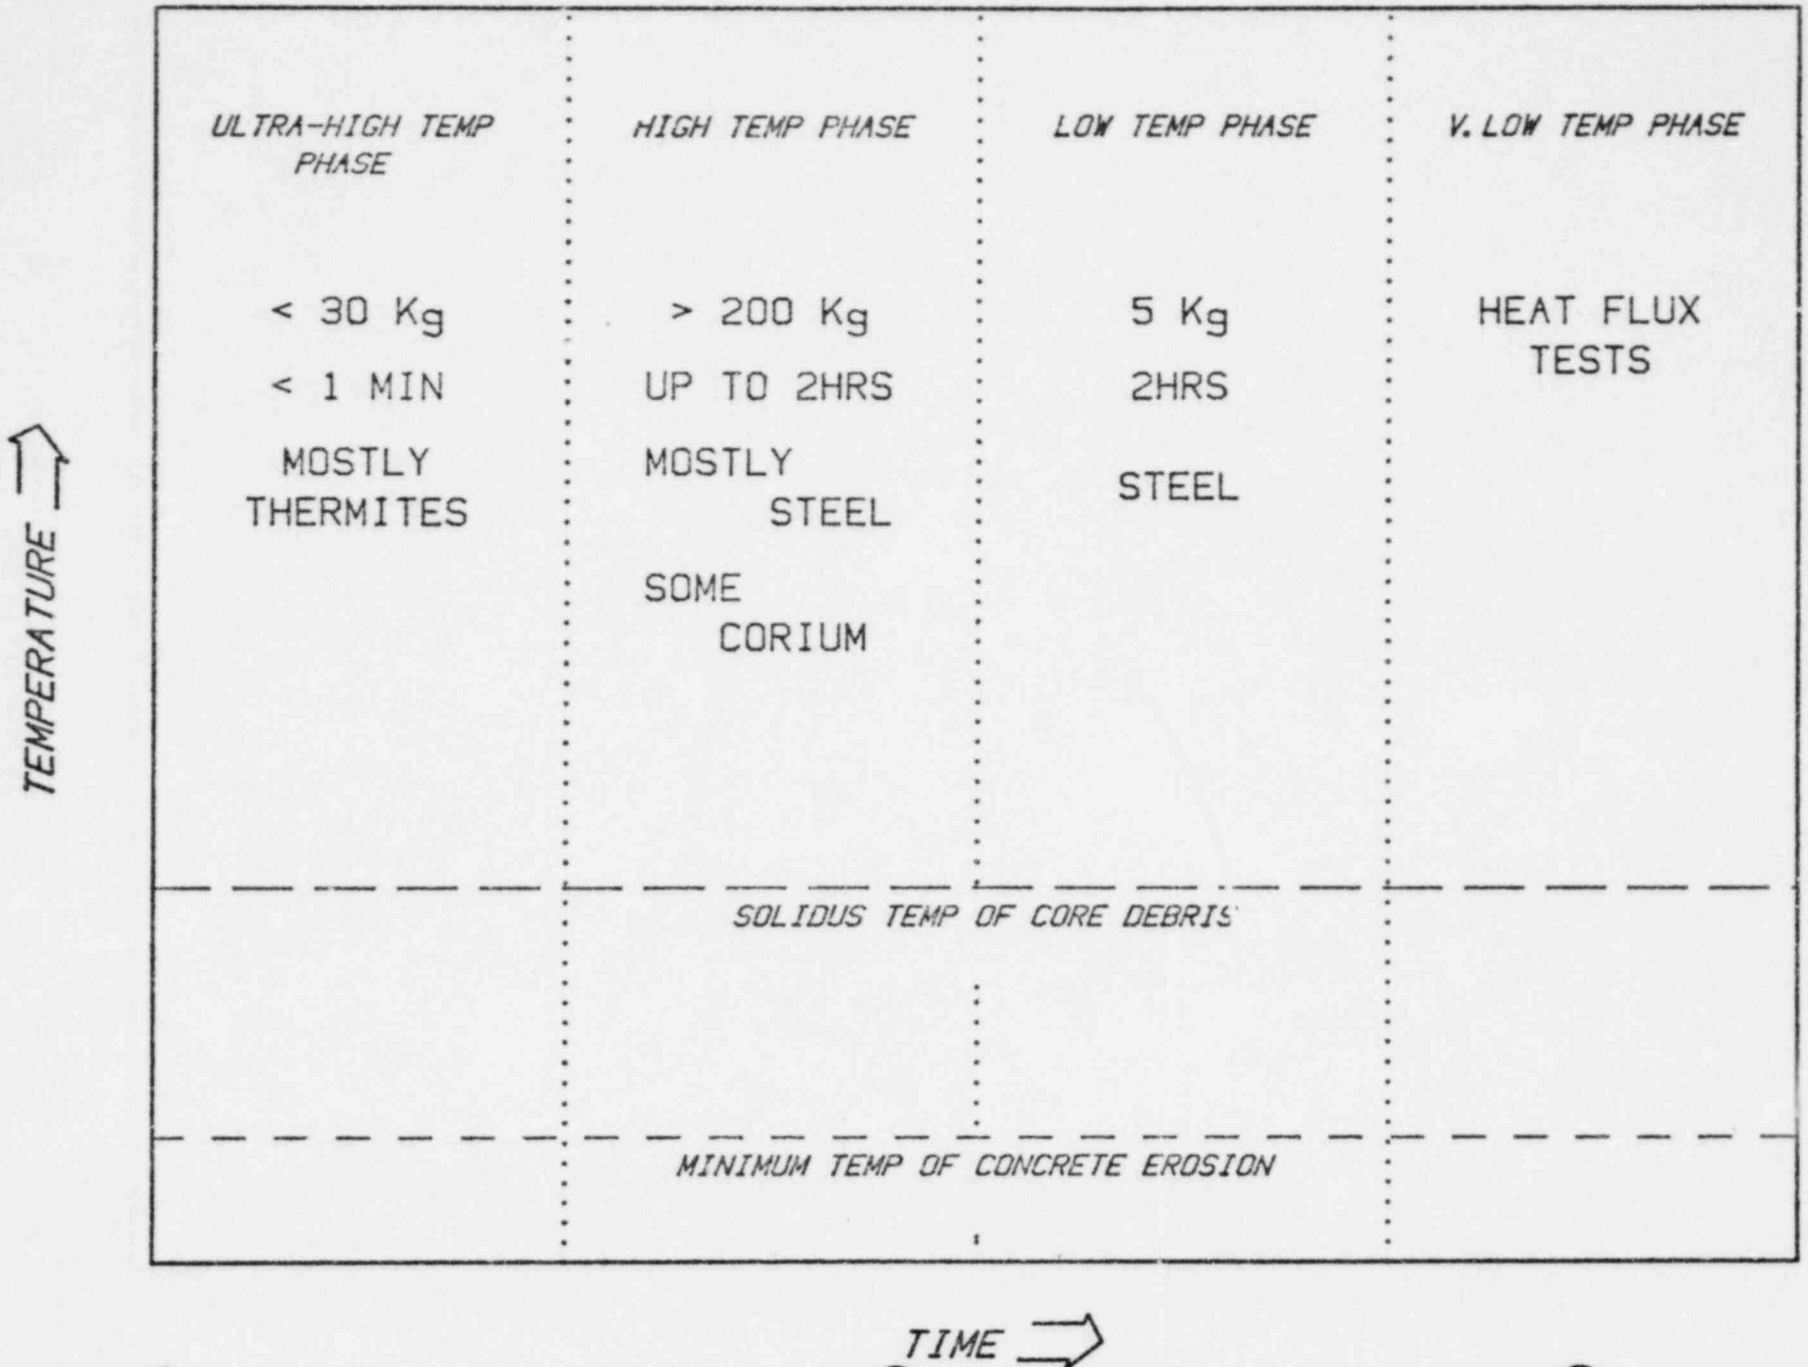

SUMMARY OF EXPERIMENTS

ULTRA-HIGH TEMPS:

- TRANSIENT (<1 min)
- SMALL SCALE (<30 kg)
- THERMITE SIMULANTS

HIGH TEMPS:

- TRANSIENT & SUSTAINED
- LARGE AND SMALL SCALE
- SOME CORIUM TESTS

LOW TEMPS:

- SMALL SCALE
- SUSTAINED STEEL
- FEW TESTS

VERY LOW TEMPS:

- PREDICTIVE, UNVERIFIED
- CODE AVAILABLE

Debris - Concrete Interactions
Program Accomplishments FY 80

1) CODE COMPARISON TESTS RUN

2) SUSTAINED HOT SOLID DEBRIS TESTS

3) CORCON DEVELOPMENT AND USED IN THE CODE COMPARISON EFFORT

4) AEROSOL SOURCE TERM MODEL DEVELOPED

CODE COMPARISON EFFORT

* TWO SUSTAINED TEST RUN.

* 'BLIND' PREDICTIONS BY THE CODE DEVELOPERS ARE UNDER WAY
  - CORCON (SANDIA)
  - WECHSL (KfK)
  - KAVERN (KWU)

* GAS AND AEROSOL GENERATION ARE NOT INCLUDED IN THE TRIALS

MODEL FAILINGS

- GAS GENERATION IS UNDERESTIMATED SIGNIFICANTLY
- AEROSOL GENERATION IS NOT PREDICTED
- FREEZING MAY BE PREDICTED TOO SOON
- SENSITIVE TO CONCRETE PROPERTIES.

LONG-TERM SUSTAINED HOT SOLID INTERACTIONS

- *DETERMINE THE ULTIMATE EXTENT OF GAS GENERATION AND BASEMAT EROSION.*

- *SCALING IS CRITICAL*

  - *VOLUME HEATING*
  - *SURFACE HEAT LOSSES*

- *DURATION OF TESTS IMPORTANT*

Debris - Sacrificial Material
Program Accomplishments FY 80

1) SUSTAINED INTERACTION TESTS WITH CANDIDATE MATERIALS CONDUCTED

2) HIGH ALUMINA CEMENT IS AN ATTRACTIVE MAT'L

3) UO_2-MgO CHEMICAL INTERACTION STUDIED

4) SOME EARLY UO_2-MgO INTERACTION MODEL DEVELOPMENT

5) ATTEMPTED LARGE-SCALE (210 KG) FUEL-MgO INTERACTION TEST

MODES OF COOLANT INTERACTIONS WITH CORE DEBRIS

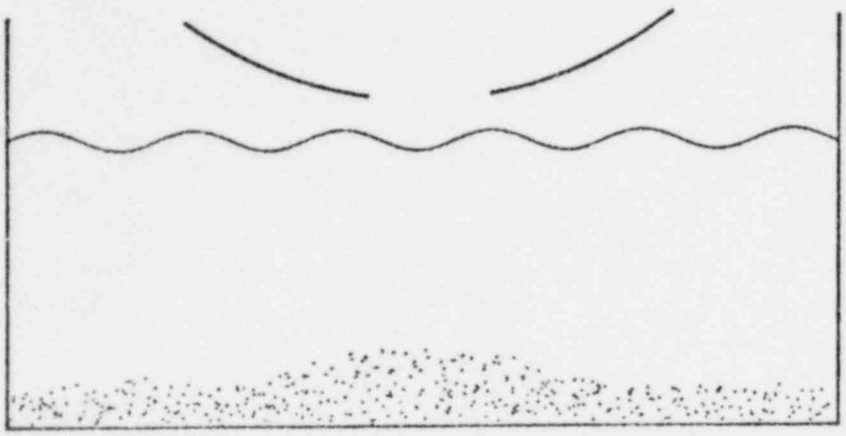

DEBRIS BED UNDER A MOLTEN Na POUL

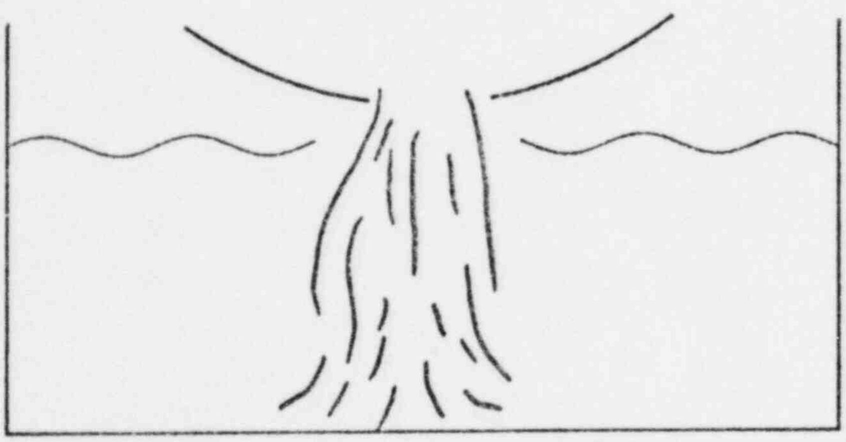

DEBRIS STREAMING INTO A Na POOL

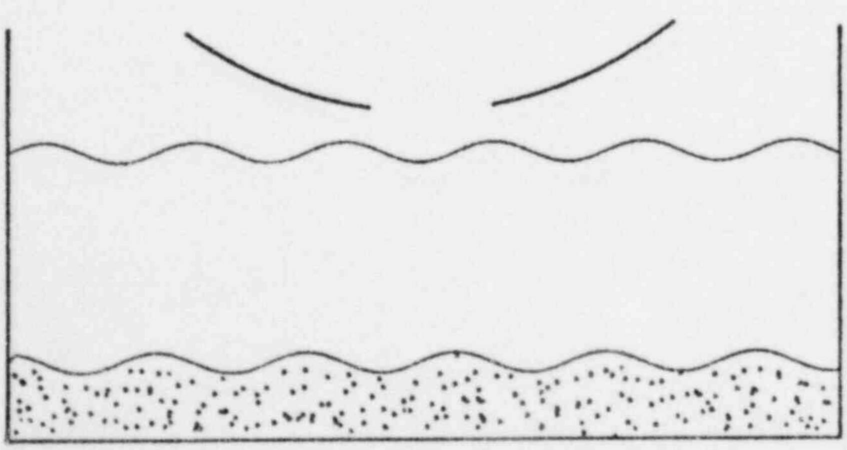

MOLTEN Na OVER MOLTEN CORE DEBRIS

EXPERIMENTAL PLAN

WATER INJECTED ONTO A
SUSTAINED MELT ON MgO

WATER INJECTED ONTO A
SUSTAINED MELT ON CONCRETE

SODIUM INJECTED ONTO A
SUSTAINED MELT

CRITICAL PARAMETERS :

- MELT DEPTH
- MELT TEMPERATURE
- COOLANT FLOW RATE
- COOLANT VELOCITY